Multiwinner Voting with Priority Candidates

MSc Thesis *(Afstudeerscriptie)*

written by

Philemon L. Huising
(born December 21st, 1993 in Wageningen, The Netherlands)

under the supervision of **Prof. dr. Ulle Endriss**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
July 11, 2023

Members of the Thesis Committee:
Prof. dr. Ulle Endriss (supervisor)
Dr. Malvin Gattinger (chair)
Dr. Davide Grossi
Simon Rey, MSc

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis is dedicated to the axiomatic study of multiwinner voting where certain kinds of candidates are to receive preferential treatment. We introduce a priority model for multiwinner voting, which extends the standard model by distinguishing between priority and non-priority candidates and incorporates a quota specifying the minimum number of priority candidates to be elected. Suitable rules are obtained by adapting rules from the standard setting by reserving seats for priority candidates, where reserved seats may be filled first, last, or at the latest possible moment this must be done in order to guarantee that the quota is met. These rules are studied axiomatically. We consider the standard axioms of anonymity and neutrality, as well as axioms designed to guarantee that priority candidates are not adversely affected. Furthermore, we focus on two kinds of elections, excellence-based elections, in which the highest-quality candidates are to be elected, and those elections in which proportional representation is important. In the excellence-based setting we show that there is a trade-off between the level of preferential treatment given and the quality of the elected committee. In the context of proportional representation we identify one particularly suitable rule, the adaptation of Phragmén's Sequential Rule that restricts attention to priority candidates at the latest possible moment.

Contents

Chapter 1

Introduction

1.1 Introduction and motivation

Multiwinner voting concerns the problem of how a group can decide, given a set of alternatives, which ones to select based on the preferences of its members. The real-life scenarios that fit this problem are many and varied. Representative governmental bodies such as parliaments are elected by citizens, shareholders vote for boards of publicly traded companies, and residential communities choose neighbourhood councils. Where the foregoing examples concern the election of representative bodies, multiwinner voting is also found in cases where the most qualified candidates are to be selected based on the votes of experts; for example, where an admissions committee selects students, judges determine the shortlist for a prize, or a limited number of job applicants are selected for an interview. Besides these more conventional contexts, applications are found in a wide range of technical domains, for example in blockchain technology (Cevallos & Stewart, 2021) and genetic programming (Faliszewski et al., 2017a).

The above examples are instances of *multiwinner elections*. A multiwinner election consists of a set of candidates, a set of voters, each of which has preferences regarding the candidates, and a desired committee size k. The objective is to choose a committee of size k based on the voters' preferences. In order to do so, a *multiwinner voting rule* is used. Choosing an appropriate rule for a multiwinner election is no simple task. Accordingly, the central line of research in this field has been the analysis of rules with respect to the desirable properties they exhibit, so-called *axiomatic analysis*.

In this thesis we are concerned with multiwinner elections in which certain kinds of candidates are to be afforded preferential treatment. Numerous examples highlight the prevalence of such scenarios. Various countries implement mandatory gender quotas in public boards and representative bodies, diversity requirements are applied to shortlists for prizes, and admissions processes for selective educational programs may prioritise applicants who are systemically disadvantaged. Many of these instances will be complex, concerning, e.g., several different types of candidates, each with different requirements regarding preferential treatment.

Little research has been done on preferential treatment in multiwinner voting, so we present a simple model as a first step. We consider an *approval-based* model, in which voters specify a subset of the candidates of which they approve.[1] The model, which we call *multiwinner voting with priority candidates*, augments the standard framework in two ways. Firstly, the set of candidates is divided into *priority* and *non-priority candidates*. Secondly, an election specifies a *quota* detailing the minimum number of priority candidates to be elected, which we denote q. The main aim of this thesis is twofold: we are looking to identify axioms that encode desirable properties in the priority model, and to find rules that satisfy these axioms.

There are roughly two sources of axioms. On the one hand we have axioms that are unique to the priority setting, for example, those that guarantee that a candidate is never worse off as a priority candidate than as a non-priority candidate. On the other hand we have axioms that are derived from those studied in the standard setting. Regarding the latter, forcing the election of q priority candidates will often mean that standard axioms are not satisfiable. Consequently, we define weakenings that allow us to reason about important properties in the context of the priority model.

We consider two different kinds of elections for which different axioms and rules are appropriate. The first kind are *excellence-based* elections, where the goal is to elect the k candidates of the highest quality. The second kind are elections concerning *proportional representation*, where a committee is to be elected that represents the distribution of preferences of the voters.

The rules that we consider for the priority model are adaptations of those for the standard model. The central idea behind these adaptations is that the q seats that are reserved for priority candidates can either be filled first or last. We will see that this distinction has a significant impact on the properties that rules possess.

1.2 Outline

In Chapter 2 we formally introduce the priority model and the rules that will be the focus of the rest of this thesis. To accomplish this, we first present the standard model and the corresponding rules that are thereafter adapted to the priority setting. The discussion of related work relies on an understanding of the priority model and is thus presented at the end of the second chapter. Chapter 3 first considers the fundamental

[1] The alternative is to use ordinal ballots, which require voters to rank candidates. For the comparative merits and demerits of approval voting see Laslier and Sanver (2010). Other, secondary variations concern restrictions on the size of the ballots. We do not consider those here.

axioms of anonymity and neutrality. Next we present a number of axioms capturing the basic requirement that candidates ought never to be worse off as priority candidates than as non-priority candidates under a priority rule. In Chapter 4 we first consider excellence-based elections in the priority context, introducing axioms that capture the fundamental nature of excellence-based elections as well as different degrees of preferential treatment for priority candidates. The second part of this chapter concerns weakenings of efficiency that are suited to the priority model. Chapter 5 investigates proportionality. We first introduce well-known axioms for the standard model that concern the representation required for cohesive groups of voters. We then consider how the notion of cohesiveness can be interpreted in the priority model, thus yielding several axioms in light of which the priority rules are analysed. For Chapters 3, 4 and 5, respectively, Tables 3.1, 4.1 and 5.1—each presented at the beginning of the corresponding chapter—provide an overview of the most important results with references to the corresponding proofs. Finally, in Chapter 6, the conclusion, we summarise and evaluate the work of the foregoing chapters and point towards future research.

1.3 Notation and prerequisites

For any set S, we will use $S[j]$ to denote all the subsets of S of size j, i.e., $S[j] = \{S' \subseteq S : |S'| = j\}$. We will sometimes use $[k]$ to denote the set $[k] = \{1, \ldots, k\}$ for $k \in \mathbb{Z}_+$. For any set S, we denote the powerset of S by $\mathcal{P}(S)$.

In several of the chapters we present tables summarising the results, i.e., which rules (do not) satisfy which axioms. In such tables, a number in square brackets refers to the corresponding proposition or corollary. If the corresponding result is a counterexample, this is indicated by prefixing the number with 'ex.'. If no reference is provided, then the result is proven in the text, rather than in a separate proof.

Chapter 2

Multiwinner voting with priority candidates

In this chapter, we introduce the fundamental components of approval-based multiwinner voting with priority candidates. The first section sets out the *standard model* and presents a number of corresponding rules. In the second section, to which the initiated reader may want to skip, we first introduce the *priority model*, which enriches the standard model with priority candidates and a quota detailing the number of priority candidates to be elected. Thereafter, we introduce the *priority multiwinner voting rules* with which we will be concerned in the following chapters. The assumptions made regarding the number of priority candidates are discussed in the third section. Finally, the fourth section presents a brief discussion of related work.

2.1 Approval-based multiwinner voting without priority candidates

In this section, we briefly cover the standard approval-based multiwinner voting model and a number of corresponding rules that will later be adapted to the priority setting. For a more extensive overview of approval-based multiwinner voting, including computational perspectives that are not considered here, see the recently published book by Lackner and Skowron (2023).

2.1.1 The standard model

Let C be a finite set of *candidates*, sometimes referred to as *alternatives*. For concrete examples it aids readability to have candidates named with letters of the alphabet, e.g., $C = \{a, b, c, d\}$; however, we will often want candidates to be indexed, in which case we will use $C = \{c_1, \ldots, c_m\}$. The number of candidates is always denoted $|C| = m$. Let $N = \{1, 2, \ldots, n\}$ be a finite set of $|N| = n$ *voters*. Each voter $i \in N$ submits an *approval ballot* (or *ballot*) $A_i \subseteq C$, also called i's *approval set*, specifying the candidates of which they approve.[1] An *approval profile* (or *profile*) is then a list $\boldsymbol{A} = (A_1, \dots, A_n)$ that specifies a ballot for each voter. If a voter $i \in N$ approves of a candidate c, i.e., $c \in A_i$, we say that i *supports* or *is a supporter of* c. We denote the set of all supporters of candidate c by N_c, i.e., $N_c = \{i \in N : c \in A_i\}$.

Any set $S \subseteq C$ of candidates is referred to as a *committee*. In multiwinner voting we are generally concerned with selecting a committee of a particular size, which we denote by k.[2] It is assumed that k is a positive integer no greater than m.

An *election* is a tuple $E = (C, \boldsymbol{A}, k)$ specifying a set of candidates, a profile, and a committee size. The set of voters can be deduced from the ballot profile $\boldsymbol{A}$, hence we often do not specify it explicitly.

Example 2.1. The owners of eight apartments in a block need to elect a homeowners' association to represent their interests and manage communal areas. Six of the owners declare their candidacy, $C = \{a, b, c, d, e, f\}$, and all owners submit an approval ballot:

$$\begin{array}{llll} A_1 = \{a, b, c\} & A_2 = \{b, c\} & A_3 = \{b, c\} & A_4 = \{b, c\} \\ A_5 = \{d, e\} & A_6 = \{d, e\} & A_7 = \{f\} & A_8 = \{f\}. \end{array}$$

For clarity we will often choose to represent the same profile as:

$$1 \times \{a, b, c\} \qquad 3 \times \{b, c\} \qquad 2 \times \{d, e\} \qquad 2 \times \{f\}.$$

This is particularly apt if we do not care about voters' identities, i.e., if we only care how often each approval set occurs.[3] Sometimes we may, for simplicity, present a profile in such a way, while still wanting to refer to individual voters. In such cases, we use the following convention to assign approval sets to the voters moving from left to right. In this case, the first voter, 1, has approval set $\{a, b, c\}$, the next three voters, 2, 3, and 4, agree on approval set $\{b, c\}$, etc. Finally, if we do not explicitly specify a set of candidates when presenting an election, one may assume that the candidates are exactly those that occur on some ballot, i.e., $C = \bigcup_{i \in N} A_i$.

An *approval-based multiwinner voting rule*—or simply *rule*, since we consider only approval ballots—is a function $\mathcal{R}$ that assigns to every election $E = (C, \boldsymbol{A}, k)$ a set $\mathcal{R}(E)$ of one or more *winning committees*,

[1] We are implicitly assuming that voters' preferences are *dichotomous*: for any candidate, a voter either approves or disapproves of that candidate.

[2] Though most of the work on multiwinner voting concerns *fixed-size committees*, in some situations it makes more sense to allow winning committees of *at most* size k. For an overview of such rules and their applications see, e.g., the works by Kilgour (2016) or Faliszewski et al. (2020).

[3] This is essentially the requirement encoded in the axiom of *anonymity*, which is covered in the next chapter (Chapter 3).

which are all size-k subsets of C. That is, $\emptyset \subset \mathcal{R}(E) \subseteq C[k]$. Sometimes it will be convenient to write $\mathcal{R}(C, \boldsymbol{A}, k)$ instead of $\mathcal{R}(E)$, omitting a pair of parentheses for clarity, and when C and k are clear from the context, we may simply write $\mathcal{R}(\boldsymbol{A})$. If a rule returns multiple winning committees, we say that these committees are *tied*. Given a winning committee W, we call the winning candidates of which voter i approves, i.e., the candidates in $A_i \cap W$, the *representatives* of i.

If a rule always returns a single winning committee, then it is said to be *resolute*, and otherwise it is *irresolute*. In most applications only a single winning committee is desired. In such cases, if an irresolute rule is used, ties are usually broken using a predetermined strict linear order on the size-k committees, also called a *tiebreaking order*. However, some of the resolute rules we consider, e.g., Phragmén's Sequential Rule and the Method of Equal Shares, make use of a predetermined strict linear order over candidates rather than committees. These rules add candidates to the winning committee in rounds and must have a way of choosing between tied candidates in any round. It should be noted that such rules yield irresolute versions if we simply consider all the possible ways of breaking ties between candidates.

Having defined the standard model, we now move on to rules for this setting.

2.1.2 Non-priority rules

The rules that we consider for the priority setting are all adaptations of existing rules. Accordingly, we present the relevant rules for the standard setting before moving on to the priority setting. Given the large and varied range of multiwinner voting scenarios, it is no wonder that a great many rules have been suggested.[4] Here we simply define the relevant rules, for the most part without discussing their properties. We distinguish between *scoring rules*, considered first, which return committees maximising a score, and *sequential rules*, which add candidates in rounds. For the following definitions, we assume that an election $E = (C, \boldsymbol{A}, k)$ and a corresponding set of voters N are given.

Scoring rules

Many of the rules considered work by assigning real-valued scores to committees and returning the size-k committees with maximal scores. Accordingly, in keeping with Kilgour and Marshall (2012), we formally define a *score* or *scoring function* as a function sc that takes a profile-committee pair $(\boldsymbol{A}, S)$ and returns a real number $\mathrm{sc}(\boldsymbol{A}, S)$.[5] When it is

[4] For an explicit categorisation of different kinds of multiwinner settings and a brief discussion of corresponding rules, see the works by Faliszewski et al. (2017c) and Elkind et al. (2017).

[5] This definition of scoring functions is not to be confused with the *committee scoring functions* presented by Elkind et al. (2017), which concern ordinal ballots and

clear from the context which profile is being considered, we will leave it out of the notation, writing $\text{sc}(S)$ instead of $\text{sc}(\boldsymbol{A}, S)$. Furthermore, when we are concerned with the score of a singleton committee $\{c\}$, we will write $\text{sc}(c)$ instead of $\text{sc}(\{c\})$. A *scoring rule* is a rule that returns all and only the committees with maximal score for some scoring function.

A candidate's *approval score* is simply the number of voters that approve of them. When only one winner is to be elected, it would seem that the only viable voting rule is that which elects the candidate with maximal approval score. A natural extension of this idea to the multiwinner setting yields *Multiwinner Approval Voting*.

Rule 1: Multiwinner Approval Voting, AV

The *approval score*, or AV-score, of a committee $S \subseteq C$ given the profile $\boldsymbol{A}$ is

$$\text{sc}_{\text{AV}}(S) = \sum_{i \in N} |A_i \cap S|.$$

AV then returns the committees that maximise the AV-score. That is,

$$\text{AV}(E) = \operatorname*{arg\,max}_{S \in C[k]} \text{sc}_{\text{AV}}(S).$$

Intuitively, we can view AV as maximising the total voter satisfaction, if a voter's satisfaction depends only on and is directly proportional to the number of representatives they have. If we instead take a voter's satisfaction to depend proportionally on the percentage of their ballot that is elected, we end up with a different rule (introduced by Kilgour, 2010, and discussed in detail by Brams and Kilgour, 2015):

Rule 2: Satisfaction Approval Voting, SAV

The *satisfaction approval score*, or SAV-score, of a committee $S \subseteq C$ given the profile $\boldsymbol{A}$ is defined as

$$\text{sc}_{\text{SAV}}(S) = \sum_{i \in N} \frac{|A_i \cap S|}{|A_i|},$$

where we take ${}^{|A_i \cap S|}/_{|A_i|} = 0$ whenever $A_i = \emptyset$. SAV returns the committees that maximise the SAV-score.

Both AV and SAV are instances of *candidate-wise scoring rules* (Kilgour & Marshall, 2012). A score, sc, is called *candidate-wise* if for any committee $S \subseteq C$, we have $\text{sc}(S) = \sum_{c \in S} \text{sc}(c)$. That is, the score of a committee can be computed by simply adding the scores of its constituent

generalise the notion of a positional scoring rule from single-winner elections to the multiwinner scenario.

candidates.[6] A *candidate-wise scoring rule* is then simply a scoring rule based on a candidate-wise score. Thus, candidate-wise scoring rules simply return size-k committees that consist of k of the highest-scoring candidates. To see that AV and SAV are indeed candidate-wise, note the following.

$$\begin{aligned}
\mathrm{sc}_{\mathrm{AV}}(S) &= \sum_{i\in N} |A_i \cap S| & \mathrm{sc}_{\mathrm{SAV}}(S) &= \sum_{i\in N} \frac{|A_i \cap S|}{|A_i|} \\
&= \sum_{i\in N}\sum_{c\in S} |A_i \cap \{c\}| & &= \sum_{i\in N}\sum_{c\in S} \frac{|A_i \cap \{c\}|}{|A_i|} \\
&= \sum_{c\in S}\sum_{i\in N} |A_i \cap \{c\}| & &= \sum_{c\in S}\sum_{i\in N} \frac{|A_i \cap \{c\}|}{|A_i|} \\
&= \sum_{c\in S} \mathrm{sc}_{\mathrm{AV}}(c) & &= \sum_{c\in S} \mathrm{sc}_{\mathrm{SAV}}(c)
\end{aligned}$$

We now consider two scoring rules that are not candidate-wise. The first, *Proportional Approval Voting*, intuitively captures the idea that there is a diminishing return on the increase in satisfaction that a voter experiences for each added representative.

Rule 3: Proportional Approval Voting, PAV

Let $h(x)$ be the xth harmonic number, i.e., $h(x) = \sum_{i=1}^{x} 1/i$ for $x \in \mathbb{Z}_+$, and let $h(0) = 0$. The PAV-score of a committee S given profile $\boldsymbol{A}$ is

$$\mathrm{sc}_{\mathrm{PAV}}(S) = \sum_{i\in N} h(|A_i \cap S|),$$

and PAV elects those size-k committees that maximise this score.

The following rule, *Approval Chamberlin-Courant*, can be seen as encoding an extreme interpretation of diminishing returns, where a voter is taken as fully satisfied if they have at least one representative and not satisfied if they do not. That is, the satisfaction of a voter with at least one representative does not increase if they receive more representatives. This is a somewhat unnatural interpretation of the rule. More intuitively, Approval Chamberlin-Courant returns committees that give as many voters as possible a representative.

Rule 4: Approval Chamberlin-Courant, CC

The CC-score of a committee S is defined as

$$\mathrm{sc}_{\mathrm{CC}}(S) = |\{i \in N : A_i \cap S \neq \emptyset\}|.$$

[6] Such scoring functions are also called *additive*, since for any two disjoint committees S_1 and S_2 it holds that $\mathrm{sc}(S_1 \cup S_2) = \mathrm{sc}(S_1) + \mathrm{sc}(S_2)$.

CC then returns all committees that maximise this score.

PAV and CC are both explicitly discussed by Thorvald N. Thiele (1895) in the context of a class of rules now known as *Thiele methods*.[7]

Rule 5: w-Thiele Method, w-Thiele

We call a non-decreasing function $w : \mathbb{N} \to \mathbb{R}$ with $w(0) = 0$ a *Thiele function*. Given a Thiele function w, the corresponding w-score of a committee S given a profile $\boldsymbol{A}$ is

$$\mathrm{sc}_w(S) = \sum_{i \in N} w(|A_i \cap S|).$$

w-Thiele then returns all committees that maximise the w-score.

Of the rules thus far considered, only SAV is not a Thiele method. This is because a voter's contribution to the SAV-score of a committee is a function not just of the number of representatives that voter gets, but also of the size of their ballot. To see that AV, CC and PAV are Thiele methods, note that they correspond to the following Thiele functions:

$$w_{\mathrm{AV}}(x) = x \qquad w_{\mathrm{CC}}(x) = \min(x, 1) \qquad w_{\mathrm{PAV}}(x) = h(x).$$

We have now considered all the scoring rules and move on to the sequential rules.

Sequential rules

Informally, we call a rule *sequential* if it works in rounds, adding a candidate in each round until a committee of the desired size has been reached.[8] Sequential rules are generally computationally tractable. As such, a number of sequential approximations of computationally hard rules have been considered, e.g., for CC and PAV, which are both NP-hard (Procaccia et al., 2008; Skowron et al., 2016). However, recent years have seen a number of sequential rules (re-)emerge that are interesting in their own right, satisfying strong axiomatic properties while retaining computational tractability.

We first consider the sequential analogues of scoring rules, which we call *sequential scoring rules*. Given some scoring function, we can define a corresponding sequential rule that in each of k rounds greedily selects the candidate that increases the score of the (provisional) committee by the most. Thus, for any scoring rule $\mathcal{R}$, we have the following sequential analogue.

[7] In the context of ordinal ballots, Chamberlin and Courant (1983) introduce a rule now known as the Chamberlin-Courant rule. Approval Chamberlin-Courant can be seen as an adaptation of this rule to the approval setting, hence its name.

[8] There exist other rules commonly called sequential that start with the set of all candidates and remove candidates in rounds to yield a size-k committee. We do not consider such rules here.

Rule 6: Sequential $\mathcal{R}$, seq-$\mathcal{R}$

Given the scoring function sc on which $\mathcal{R}$ is based, seq-$\mathcal{R}$ works as follows. Start with an empty committee $W_0 = \emptyset$. For each round $r \in [k]$ we set $W_r = W_{r-1} \cup \{c\}$ for the thus far unelected candidate $c \notin W_{r-1}$ that maximises $\mathrm{sc}(W_{r-1} \cup \{c\})$. If there are multiple such candidates, a predetermined tiebreaking order over candidates is used. The rule returns $W = W_k$.

The sequential analogues of the Thiele methods—e.g., seq-AV, seq-CC and seq-PAV—are commonly referred to as *sequential Thiele methods*. Given the above definition, seq-CC elects, in any round, the candidate that increases the CC-score of the provisional committee by the most; i.e., it elects the candidate that represents the greatest number of voters that do not yet have a representative. This process is suboptimal with regards to maximising the CC-score, hence seq-CC may return a committee that is not a winning committee according to CC (see Example 2.6).

In contrast, seq-AV can be viewed as fundamentally the same as AV. In each round, seq-AV adds a remaining candidate with the highest approval score to the committee. Consequently, the resulting committee maximises the approval score and is a winning committee according to AV. In order to make AV resolute, ties must be broken between *committees*, while seq-AV makes use of a tiebreaking order between *individual candidates*. There is no unambiguously optimal method for extending (ordinal) preferences over candidates to preferences over sets of candidates or vice versa (Barberà et al., 2004), hence we cannot strictly say that AV and seq-AV are equivalent. Nevertheless, all the ways of breaking ties between candidates will yield, using seq-AV, all the committees returned by AV. In keeping with the literature, we have treated AV as a scoring rule, but it will sometimes be helpful to consider it as a sequential rule, i.e., as seq-AV.

The above reasoning concerning the relationship between AV and seq-AV also applies to SAV and seq-SAV (which is not a sequential Thiele method). In fact, the sequential version of any candidate-wise scoring rule will always select k of the highest-scoring candidates, and thus yield a committee that maximises the score. And all the winning committees according to the scoring rule can be obtained by executing the sequential rule with all possible ways of breaking ties between candidates.[9]

We now consider two sequential rules that have received a lot of attention in recent years, neither of which is naturally interpreted as a sequential scoring rule: *Phragmén's Sequential Rule* (seq-Phragmén) and the *Method of Equal Shares* (MES). We start with the former.

[9] Given some committee W that maximises the (candidate-wise) score, any tiebreaking order in which all the members of W are ranked higher than all other candidates will yield W in combination with the sequential version of the rule.

Phragmén gave, towards the end of the 19th century, a number of different formulations of his sequential rule—Janson (2018) provides a detailed overview. Here, we take the formulation used by Brill et al. (2023) as fundamental. In this formulation, elected candidates are thought of as incurring a *load* among their supporters and the goal is to elect a committee for which the corresponding distribution of the loads over the voters is as balanced as possible. We thus first define what a *load distribution* is.

Definition 2.1 (Load distribution)**.** Given an election $E = (C, \boldsymbol{A}, k)$, a *load distribution* is a two-dimensional array $x = (x_{i,c})_{i \in N, c \in C}$, where $x_{i,c}$ is the load voter i receives from candidate c. The total load of voter i is then $\bar{x}_i = \sum_{c \in C} x_{i,c}$. A load distribution must meet the following requirements.

(i) $x_{i,c} \geq 0$ for all $i \in N$ and $c \in C$

(ii) $x_{i,c} = 0$ for any $i \in N$ and $c \notin A_i$

(iii) $\sum_{i \in N} \bar{x}_i = k$

(iv) $\sum_{i \in N} x_{i,c} \in \{0, 1\}$ for every $c \in C$

Condition *(ii)* guarantees that the load due to candidate c is distributed only among supporters of c. Conditions *(iii)* and *(iv)* guarantee that there is a size-k committee corresponding to the load distribution x, namely, $\{c \in C : \sum_{i \in N} x_{i,c} = 1\}$.

Phragmén's Sequential Rule starts with an empty committee and greedily adds candidates that minimise the resulting maximal voter load.

Rule 7: Phragmén's Sequential Rule, seq-Phragmén

The rule starts with an empty committee $W_0 = \emptyset$. In each of k rounds a new candidate c is added and one unit of load is distributed among the voters in N_c, to yield a size-k winning committee. Let $\bar{x}_i^r$ denote the total load incurred by voter i at the end of round r. Every voter $i \in N$ starts with a load of 0, i.e., $\bar{x}_i^0 = 0$. In each round, a voter maintains their previously accrued load, which may be added to only if the elected candidate is approved by the voter. We thus have $\bar{x}_i^r \geq \bar{x}_i^{r-1}$ for each $i \in N$ and $\bar{x}_i^r = \bar{x}_i^{r-1}$ when the candidate c elected in round r is not approved of by i, i.e., $c \notin A_i$. In any round $r \in [k]$, the candidate c and loads $\bar{x}_i^r$ are chosen that respect the above conditions and that minimise the resulting maximum voter load $\max_{i \in N} \bar{x}_i^r$. The committee after round r is then $W_r = W_{r-1} \cup \{c\}$. If there are multiple such candidates, ties are broken according to a predetermined order. After the kth round $W = W_k$ is returned.

It is easy to see that the resulting load distribution meets the requirements set out in Definition 2.1. Note that the above description of the rule assumes that there are at least k candidates with positive approval. In fact, often the assumption is made that all candidates have positive approval. When working with seq-Phragmén we also make this assumption, even though it will often be implausible in the priority setting. We come back to this later on. It should also be noted that sometimes, e.g., in the work of Peters and Skowron (2020), the load incurred due to the election of a candidate is set at n/k, rather than 1. This makes no difference to the working of the rule, but may, as we shall see, make it easier to work with in certain contexts.

We introduce two further (equivalent) formulations of seq-Phragmén: a *continuous formulation* which is intuitively the clearest, and a *discrete formulation*, which will be helpful in our proofs. To make clear that we have one fundamental formulation—that given in the statement of the rule (Rule 7)—we present these as lemmas.

Lemma 2.1 (The continuous formulation of seq-Phragmén)**.** *The rule corresponding to the following description is equivalent to seq-Phragmén.*

Voters start with a budget of 0, which increases continuously with time, such that at time t a voter's budget is t (if they have not spent any of it). As soon as a group of voters that unanimously approve of some candidate c (i.e., N_c) have a combined budget of 1, this candidate is elected and the budgets of the voters in N_c are set to 0. The voters who do not approve of the elected candidate keep their budget. This process is repeated, breaking ties between candidates according to a fixed tiebreaking order, until k candidates have been elected.

For an example of work that utilises the continuous formulation, see the article by Peters and Skowron (2020). Where the continuous formulation is intuitively simplest, the following discrete formulation, already presented by Phragmén (1899), clearly demonstrates how the rule is computed.

Lemma 2.2 (The discrete formulation of seq-Phragmén)**.** *The rule corresponding to the following description is equivalent to seq-Phragmén.*

In round $r \in [k]$, for any thus far unelected candidate $c \notin W_{r-1}$, let

$$\ell_r(c) = \frac{1 + \sum_{i \in N_c} \bar{x}_i^{r-1}}{|N_c|}.$$

Then the maximum load $\ell_r = \max_{i \in N} \bar{x}_i^r$ is

$$\ell_r = \min_{c \in C \setminus W_{r-1}} \ell_r(c),$$

and a candidate c for which $\ell_r(c)$ is minimal, using tiebreaking if necessary, is elected. If c is elected in round r, then the new loads are given by

$$\bar{x}_i^r = \begin{cases} \ell_r(c) & \textit{if } i \in N_c \\ \bar{x}_i^{r-1} & \textit{otherwise.} \end{cases}$$

Brill et al. (2023) also present the discrete formulation as a lemma (Lemma 4.5) and provide a corresponding proof. There is an optimisation-based rule corresponding to seq-Phragmén, called *Phragmén's Leximax Rule.*[10] In contrast to seq-Phragmén, Phragmén's Leximax Rule fails to satisfy a desirable property called *committee monotonicity* and it is NP-hard, hence we do not consider it here.[11]

The *Method of Equal Shares*, first introduced by Peters and Skowron (2020), endows every voter with the same starting budget, which can be used to elect candidates. Accordingly, it intuitively captures the idea that every voter ought to be able to determine an equal share of the winning committee. Roughly, the idea is that in any round, the candidate is elected that requires the smallest payment per supporting voter.

Rule 8: Method of Equal Shares, MES

MES has two phases. During the first phase, we start with an empty committee $W_0 = \emptyset$ and add a candidate in each of *at most* k rounds. If less than k candidates are elected during the first phase, the second phase is used to complete the committee.

Phase 1: Initially, each voter $i \in N$ receives an initial budget of $b_i(1) = 1$, where $b_i(r)$ denotes i's budget immediately before the start of round r. The *price*, to be paid by voters, of electing a candidate is $p = {}^{n}/_{k}$. In round r, a candidate $c \notin W_{r-1}$ is *x-affordable* if

$$\sum_{i \in N_c} \min(x, b_i(r)) \geq p,$$

where $x \in \mathbb{R}_{\geq 0}$. That is, the supporters of c can raise the funds required to elect this candidate among themselves without any of them paying more than x. If no candidate is x-affordable for any x—i.e., if no candidate is *affordable*—then phase 1 terminates. Otherwise, a candidate $c \notin W_{r-1}$ is elected that is x-affordable for minimal x. We then update the voters' budgets as follows:

$$b_i(r+1) = \begin{cases} b_i(r) - x & \text{if } i \in N_c \text{ and } b_i(r) \geq x \\ 0 & \text{if } i \in N_c \text{ and } b_i(r) < x \\ b_i(r) & \text{if } i \notin N_c \end{cases}$$

[10] This rule does not only return committees that minimise the maximum voter load. It also employs a specific kind of tiebreaking which is essential for the satisfaction of certain properties such as the proportionality requirement *proportional justified representation* (Brill et al., 2023).

[11] See the work by Brill et al. (2023) for more on Phragmén's Leximax Rule. The original work is by Phragmén (1896) and a modern overview is given by Janson (2018). We do not consider other variants, such as one that minimises the variance of the loads, since these do not fare well with regards to the proportionality axioms with which we will be concerned. For more on these rules also see the mentioned works.

Note that since $b_i(1) = 1$ for each $i \in N$ and $p = n/k$, there is exactly enough budget in total to elect k candidates. If there are no more affordable candidates and less than k candidates have been elected, we move to phase 2.

Phase 2: MES satisfies most interesting properties, for example the proportionality axiom *extended justified representation* (see Chapter 5), regardless of how the remaining candidates are elected (Peters & Skowron, 2020). Nevertheless, to guarantee that certain axioms (though not any that we consider) such as *priceability* are satisfied, it is recommended to complete the committee using seq-Phragmén, where voters have starting budgets equal to the budget they have left at the end of Phase 1. It is essential that in this case the load incurred by electing a candidate is the same as the price to be paid for a candidate in Phase 1.

We note that, as presented by Lackner and Skowron (2023), MES may also be formulated such that the price of a candidate is set at $p = 1$ and the initial budget of each voter $i \in N$ is $b_i(1) = k/n$. In that case, the load incurred for each candidate selected in Phase 2 using seq-Phragmén will be 1, as specified in our definition (Rule 7).

Having covered the standard model and the relevant rules for this setting, we are ready to move on to the priority setting.

2.2 Approval-based multiwinner voting with priority candidates

Analogously to the previous section, we first introduce the *priority model*, a novel enrichment of the standard model, and thereafter we define a number of new *priority rules* with which we will be concerned in the following chapters.

2.2.1 The priority model

The priority model enriches the standard model in two ways. Firstly, we specify a subset of the set of candidates $C^+ \subseteq C$ that consists of the *priority candidates*. This also yields a set $C^- = C \setminus C^+$ of *non-priority candidates*. When introducing a set of candidates or profile, we will often distinguish the priority candidates by representing them with an overline. For example, in the set of candidates $C = \{\overline{a}, \overline{b}, c, d\}$, $\overline{a}$ and $\overline{b}$ are priority candidates, while c and d are non-priority candidates. When we refer to a candidate's *status*, we mean their classification as a priority or non-priority candidate. Thus, e.g., two candidates have the same status if they are either both priority candidates or both non-priority candidates.[12]

[12] This vocabulary is particularly useful in situations where we consider the consequences of changing a candidate's status. In such cases we will leave out the overlines

Extending the new notation to arbitrary committees, we denote the set of priority candidates that belong to a committee $S \subseteq C$ by $S^+ = S \cap C^+$. Similarly, the set of non-priority candidates that belong to S is denoted $S^- = S \cap C^- = S \setminus S^+$.

The second addition to the standard model is the *priority quota* (or *quota*) q. The quota represents the number of 'reserved seats' on a committee; i.e., q specifies the number of seats that are to be filled with priority candidates. Consequently, we require that q be a non-negative integer of at most k.[13] We will be concerned with committees that always elect at least q priority candidates—such committees are said to *respect the quota*—hence, we also require that q be no greater than $|C^+|$, the number of priority candidates. This restriction, which might equally be formulated as an assumption concerning the number of priority candidates (i.e., that there are at least q of them) is necessary if the priority rules are to be well-defined. However, as we shall see, some of the considered rules require the stronger assumption that there are enough priority candidates to fill the committee, i.e., $|C^+| \geq k$. We discuss these assumptions in detail in Section 2.3, after the relevant rules have been introduced.

The two additions to the standard model are clearly reflected in the definition of a *priority election*, or simply *election*, which is a tuple $E = (C, C^+, \boldsymbol{A}, k, q)$. A *priority approval-based multiwinner voting rule*, or simply *rule*, is then a function $\mathcal{R}$ that takes as input a priority election E and returns one or more size-k winning committees, i.e., $\emptyset \subset \mathcal{R}(E) \subseteq C[k]$. A rule is said to *respect the quota*, if it returns only committees that respect the quota, i.e., that have at least q priority candidates. Resoluteness for priority rules is defined exactly as in the standard setting.

We use 'election' and 'rule' for both the standard model and the priority model. Usually, the context will make clear which setting is being considered. When this is not the case we will explicitly use 'non-priority' and 'priority' to distinguish between the two settings.

Armed with this vocabulary, we now note that each priority election $E = (C, C^+, \boldsymbol{A}, k, q)$ has an *underlying* non-priority election $E' = (C, \boldsymbol{A}, k)$. Consequently, each non-priority rule $\mathcal{R}$ can be seen as a priority rule that takes a priority election $E = (C, C^+, \boldsymbol{A}, k, q)$ and returns $\mathcal{R}(C, \boldsymbol{A}, k)$, i.e., all the winning committees for the underlying non-priority election.

Example 2.2. A company must elect a standing committee of size $k = 3$ that organises social events for employees. In order to guarantee that junior employees are also represented, it is required that at least

when presenting a profile to avoid confusion.

[13] It is interesting to consider whether priority rules reduce to the non-priority rules on which they are based in the extreme cases where $q = 0$ or $q = k$, i.e., where none or all of the seats are required to be filled by priority candidates.

one member of the committee is junior, i.e., $q = 1$. The set of candidates is $C = \{\overline{a}, \overline{b}, c, d\}$ and the set of priority candidates—in this case the set of junior candidates—is $C^+ = \{\overline{a}, \overline{b}\}$. The following ballots are cast.

$$5 \times \{\overline{a}\} \qquad 5 \times \{\overline{a}, \overline{b}, d\} \qquad 10 \times \{\overline{a}, c\} \qquad 5 \times \{c, d\}$$

If we consider the underlying election, it would seem that the most reasonable outcome is $\{\overline{a}, c, d\}$.[14] Though this committee meets the requirement of having $q = 1$ priority candidates, we will see that some priority rules will give different results (Example 2.3).

Having set out the priority model, we can now introduce the priority rules with which we will be concerned in the following chapters.

2.2.2 Priority rules

We first introduce *dual election rules* and after that consider adaptations of the sequential rules to the priority setting.

Dual election rules

Given a non-priority rule $\mathcal{R}$ and priority election $E = (C, C^+, \boldsymbol{A}, k, q)$, we might think to run two different elections using $\mathcal{R}$, one for priority candidates and one for non-priority candidates. This would result in a committee of exactly q priority candidates (and $k - q$ non-priority candidates). However, since we are here concerned with giving some kind of preferential treatment to priority candidates, it makes more sense to allow all committees that have *at least* q priority candidates.[15] In this case we might have a *quota election*, in which q *reserved* seats are to be filled with priority candidates, and an *open election*, in which $k - q$ candidates are to be elected and *all* (remaining) candidates, including any (remaining) priority candidates, are eligible. This gives rise to the concept of *dual election rules*. We have two kinds of dual election rules, corresponding to whether we fill the reserved seats first or last.

Rule 9: Reserved-first dual election rule based on $\mathcal{R}$

Given a priority election $E = (C, C^+, \boldsymbol{A}, k, q)$, the *reserved-first dual election rule based on non-priority rule $\mathcal{R}$*, denoted $D^f_{\mathcal{R}}$, works as

[14] With the exception of CC and seq-CC, all non-priority rules considered return $\{\overline{a}, c, d\}$ as the sole winning committee. CC also returns $\{\overline{a}, c, d\}$, but as one of three winning committees. Whether or not seq-CC returns $\{\overline{a}, c, d\}$ depends only on tiebreaking (alphabetic tiebreaking means $\{\overline{a}, \overline{b}, c\}$ is elected).

[15] The setting in which there is a strict desired distribution of different types of candidates, considered by, e.g., Bredereck et al. (2018) and Celis et al. (2018), is discussed in the section on related work at the end of this chapter.

follows. $W \in D^f_{\mathcal{R}}(E)$ if and only if there exist $Q \in \mathcal{R}(C^+, \boldsymbol{A}, q)$ and $O \in \mathcal{R}(C \setminus Q, \boldsymbol{A}, k-q)$ with $W = Q \cup O$. We refer to the elections $(C^+, \boldsymbol{A}, q)$ and $(C \setminus Q, \boldsymbol{A}, k-q)$ as the *quota election* and *open election*, respectively.

Note that we used the original profile $\boldsymbol{A}$ to specify both the quota election and the open election. Strictly speaking, we ought to consider a restriction of this profile, since it contains candidates that are not members of the considered candidate set. Formally, this is achieved by taking the intersection of each ballot with the relevant candidate set; however, we will often leave this implicit to avoid clutter.

Note that the above definition relies on the requirement that $|C^+| \geq q$, since otherwise there is no guarantee that there will be enough eligible candidates for the quota election.[16]

Reversing the order of the quota election and the open election yields another rule.

Rule 10: Reserved-last dual election rule based on $\mathcal{R}$

Given a priority election $E = (C, C^+, \boldsymbol{A}, k, q)$, the *reserved-last dual election rule based on non-priority rule* $\mathcal{R}$, denoted $D^\ell_{\mathcal{R}}$, works as follows. $W \in D^\ell_{\mathcal{R}}(E)$ if and only if there exist $O \in \mathcal{R}(C, \boldsymbol{A}, k-q)$ and $Q \in \mathcal{R}(C^+ \setminus O, \boldsymbol{A}, q)$ with $W = O \cup Q$. We refer to the elections $(C, \boldsymbol{A}, k-q)$ and $(C^+ \setminus O, \boldsymbol{A}, q)$ as the *open election* and *quota election*, respectively.

Where the definition of $D^f_{\mathcal{R}}$ relied on the assumption that $|C^+| \geq q$, the definition of $D^\ell_{\mathcal{R}}$ requires the stronger assumption that there are at least k priority candidates, i.e., $|C^+| \geq k$. This is because we must guarantee that there will be q unelected priority candidates *after the open election* in which $k-q$ priority candidates could be elected. As we shall see, there are other rules that require this additional assumption, which is discussed in Section 2.3.

Dual election rules are particularly interesting if the used non-priority rule is a candidate-wise scoring rule. In this light, we consider the dual election rules based on AV and SAV, i.e., D^f_{AV}, D^ℓ_{AV}, D^f_{SAV}, and D^ℓ_{SAV}. To see that the reserved-first and reserved-last rules are distinct, consider the following example.

Example 2.3. Consider again the election E from Example 2.2 with the set of candidates $C = \{\overline{a}, \overline{b}, c, d\}$, committee size $k = 3$ and quota $q = 1$. For both AV and SAV the reserved-first dual election rule yields $\{\overline{a}\}$ and $\{c, d\}$ as the unique winning committees for the quota and open elections, respectively. That is, $D^f_{\mathrm{AV}}(E) = D^f_{\mathrm{SAV}}(E) = \{\overline{a}, c, d\}$.

[16] Remember that non-priority rules are defined only for cases in which the committee size is no greater than the number of available candidates.

In contrast, for both AV and SAV, the reserved-last dual election rule yields $\{\overline{a}, c\}$ and $\{\overline{b}\}$ as the unique winning committees for the quota and open elections, respectively. That is, $D^{\ell}_{\mathrm{AV}}(E) = D^{\ell}_{\mathrm{SAV}}(E) = \{\overline{a}, \overline{b}, c\}$.

Intuitively, if the quota election is first, then $\overline{a}$ takes the reserved seat and $\overline{b}$ is outcompeted in the open election. However, if the open election is first, then $\overline{a}$ takes an unreserved seat, which allows $\overline{b}$ to take the reserved seat in the following quota election.

Note that, contrary to the claim that we must generally assume $|C^+| \geq k$, in the above example the candidate set contains fewer than k priority candidates, i.e., $|C^+| < k$. This is not a problem, since the reason for the assumption is that it might be the case that there are insufficiently many priority candidates for the quota election and this is not the case here.

The following lemma captures the fact that D^f_{sc}, the reserved-first dual election rule based on the non-priority rule maximising the candidate-wise score sc, returns exactly those committees with maximal scores among those that respect the quota. Intuitively, the reason is the following. If q priority candidates must be elected and we want to achieve the highest-possible score, we must always elect the highest-scoring q priority candidates. This is what happens in the quota election. After that, the $k - q$ highest-scoring remaining candidates are elected.

Lemma 2.3. *Let* sc *be a candidate-wise score.* D^f_{sc} *returns all and only the committees that maximise the score among those that respect the quota. That is,*

$$D^f_{\mathrm{sc}}(E) = \operatorname*{arg\,max}_{W \in C[k],\, |W^+| \geq q} \mathrm{sc}(W).$$

Proof. Denote the non-priority scoring rule corresponding to candidate-wise score sc by $\mathcal{R}_{\mathrm{sc}}$. We first show that any winning committee according to $\mathcal{R}_{\mathrm{sc}}$ has a maximal score among the committees with q or more priority candidates. Assuming $|C^+| \geq q$, the quota election clearly guarantees that q priority candidates are elected. Now let $W_1 \in D^f_{\mathrm{sc}}(E)$ be a winning committee and let $W_2 \in C[k]$ with $|W_2^+| \geq q$ be arbitrary. We show that $\mathrm{sc}(W_1) \geq \mathrm{sc}(W_2)$. There must be sets $Q_1 \in \mathcal{R}_{\mathrm{sc}}(C^+, \boldsymbol{A}, q)$ and $O_1 \in \mathcal{R}_{\mathrm{sc}}(C \setminus Q_1, \boldsymbol{A}, k - q)$ with $W_1 = Q_1 \cup O_1$. Now let $Q_2 \subseteq W_2^+$ consist of q of the highest-scoring priority candidates from W_2 such that $Q_1 \cap W_2 \subseteq Q_2$. This is guaranteed to be possible since Q_1 consists of q of the highest-scoring priority candidates simpliciter (i.e., not only from W_1, but from C). We then have $\mathrm{sc}(Q_1) \geq \mathrm{sc}(Q_2)$. Furthermore, since $Q_1 \cap W_2 \subseteq Q_2$, we have $O_2 = W_2 \setminus Q_2 \subseteq C \setminus Q_1$. But then, since $O_1 \in \mathcal{R}_{\mathrm{sc}}(C \setminus Q_1, \boldsymbol{A}, k - q)$, we must have $\mathrm{sc}(O_1) \geq \mathrm{sc}(O_2)$. Since $\mathrm{sc}(Q_1) \geq \mathrm{sc}(Q_2)$, it follows that $\mathrm{sc}(W_1) \geq \mathrm{sc}(W_2)$.

We now show that any committee $W \in C[k]$ that maximises the score among the committees with q or more priority candidates is a D^f_{sc} winning committee. Let $Q \subseteq W^+$ consist of q of the highest-scoring priority candidates from W. It follows that these must be q of the highest-scoring priority candidates from C simpliciter. That is, $Q \in \mathcal{R}_{\mathrm{sc}}(C^+, \boldsymbol{A}, k)$. Now suppose $W \setminus Q \notin \mathcal{R}_{\mathrm{sc}}(C \setminus Q, \boldsymbol{A}, k-q)$. It follows that there exist some $c \in W \setminus Q$ and $c' \in C \setminus Q$ such that $\mathrm{sc}(c') > \mathrm{sc}(c)$. But then $(W \setminus \{c\}) \cup \{c'\}$ respects the quota and has a higher score than W, which contradicts our assumption to the contrary. $\square$

Example 2.3 shows that for AV and SAV the reserved-first and reserved-last rules are distinct, hence Lemma 2.3 does not hold for the reserved-first rules. It should not come as a surprise that the other scoring rules considered fail to yield such a lemma for either their reserved-first or reserved-last adaptation. For candidate-wise scoring rules, the contribution of an elected candidate to the score of the final committee is simply their contribution to the score of the subcommittee in which they were elected. Similarly, the contribution of a voter to the score is simply the sum of their contributions to the scores of the subcommittees. This is not the case for CC or PAV.

In fact, Example 2.3 suffices to show that a lemma that is analogous to Lemma 2.3 does not hold for D^f_{CC} or D^ℓ_{CC}, since $\{\overline{a}, b, d\}$ receives the maximal CC-score of 25 but is not a winning committee according to either rule. Though the example considered does not show this, the same is true for D^f_{PAV} or D^ℓ_{PAV}.[17]

We now move on to consider priority adaptations of the sequential rules.

Sequential priority rules

Assuming we want to elect only committees that respect the quota, the natural way to adapt a sequential rule would be to restrict the candidate pool to priority candidates for q of the k rounds. Analogous to the dual election rules, we first consider two ways in which q seats might be reserved for priority candidates: the first q seats may be reserved, or the last q seats may be reserved.

Given a sequential rule $\mathcal{R}$, we refer to the priority rule that results from restricting the pool of eligible candidates to priority candidates in the first q rounds as *reserved-first* $\mathcal{R}$.

Rule 11: Reserved-first $\mathcal{R}$, rf-$\mathcal{R}$

Given a sequential non-priority rule $\mathcal{R}$, rf-$\mathcal{R}$ works as follows. Run

[17] In Example 2.3, both dual rules based on PAV return the committee that maximises the PAV-score, $\{\overline{a}, c, d\}$. However, it is not difficult to come up with an example where these rules return a committee that does not maximise the PAV-score. For example, if we have a profile where three voters approve $\{\overline{a}, \overline{b}\}$ and two voters approve $\{c\}$, with $k = 2$ and $q = 1$. Then $\{\overline{a}, \overline{b}\}$ is the winning committee according to both dual rules, while both $\{\overline{a}, c\}$ and $\{\overline{b}, c\}$ have a higher PAV-score.

the algorithm for $\mathcal{R}$ as usual, but in the first q rounds restrict the candidate pool to the remaining priority candidates. That is, for any round $1 \leq r \leq q$ the candidate pool is $C_r = C^+ \setminus W_{r-1}$, while for any round $q+1 \leq r \leq k$ we have $C_r = C \setminus W_{r-1}$ as before.

Note that, generally, the above definition, analogously to the reserved-first dual election rules, requires that $|C^+| \geq q$ in order to be well-defined. Taking the last q seats as reserved yields *reserved-last* $\mathcal{R}$.

Rule 12: Reserved-last $\mathcal{R}$, rl-$\mathcal{R}$

Given a sequential non-priority rule $\mathcal{R}$, rl-$\mathcal{R}$ works as follows. Run the algorithm for $\mathcal{R}$ as usual, but in the last q rounds, restrict the candidate pool to the remaining priority candidates. That is, for any round $1 \leq r \leq k-q$ we have the candidate pool $C_r = C \setminus W_{r-1}$ as before, while for any round $k-q+1 \leq r \leq k$ the candidate pool is $C^+ \setminus W_{r-1}$.

Similarly to the reserved-last dual election rules, the reserved-last sequential rules require the assumption that $|C^+| \geq k$, since, in order to elect k candidates, there must be at least q remaining priority candidates that can be elected in the last q rounds.

The sequential versions of the scoring rules considered—seq-AV, seq-SAV, seq-CC and seq-PAV—are well-defined and intuitively clear. However, for (all formulations of) seq-Phragmén and MES, the reserved-first adaptation is either ill-defined or yields a clearly undesirable rule. The following example illustrates this. In each case the problems stem fundamentally from the fact that 'worse' candidates may be elected before 'better' ones.

Example 2.4. Let $C = \{\overline{a}, b, c\}$, $k = 2$, and $q = 1$, and consider the following profile:

$$1 \times \{\overline{a}, b\} \qquad 9 \times \{b\} \qquad 1 \times \{c\}.$$

We first consider the reserved-first rule based on the discrete formulation of seq-Phragmén. In the first round a priority candidate must be elected, hence $\overline{a}$ is elected, resulting in a load of 1 for voter 1, i.e., $\overline{x}_1^1 = 1$. Since

$$\frac{1 + \sum_{i \in N_b} \overline{x}_i^1}{|N_b|} = \frac{1+1}{10} = 0.2$$

is minimal, we elect b in round 2 and set $\overline{x}_i^2 = 0.2$ for $i \in [10]$. But then we have a situation in which an initially distributed load (for $\overline{a}$) is redistributed, strictly decreasing the maximum load. This is made worse by the fact that the load distributed due to $\overline{a}$ is shouldered by voters that do not approve of this candidate.

Now consider the continuous formulation of seq-Phragmén. The first candidate elected must be a priority candidate, i.e., $\overline{a}$. This happens at time t, when voter 1 has a budget of 1. At this point all other voters also have a budget of 1. But that means that b and c are tied in that they are both affordable, while b, who became affordable at an earlier moment, is clearly a better choice (also in the spirit of seq-Phragmén).[18]

We now consider seq-Phragmén in terms of its fundamental definition (Rule 7). Electing $\overline{a}$ in the first round results in a load of 1 for voter 1. In the second round, the maximum voter load will remain 1 regardless of whether we elect b or c, while, as stated above, b is clearly the right choice in the spirit of seq-Phragmén.

The problem with MES arises from the fact that the first phase might end in a restricted round because there are not enough priority candidates that enjoy sufficient support, while there are non-priority candidates that are affordable. In this case the price will be $p = {}^{n}\!/_{k} = {}^{11}\!/_{2} = 5.5$. Since the only eligible candidate in the first round is $\overline{a}$ and they only have one supporter (with an initial budget of 1), there is no affordable eligible candidate, hence phase 1 ends while b is affordable.

Inspired by Phragmén's Leximax Rule, we might refine the reserved-first adaptation of seq-Phragmén (as defined) as follows. Rather than looking only at the resulting maximal load when deciding which candidate to elect, we can break ties between candidates by looking at the next highest loads that result from their election. In the above example, b would be preferred over c in the second round. In both cases the maximal voter load after round 2 will be 1; however, electing b would result in a next-highest voter load of ${}^{1}\!/_{9}$ (for voters 2 through 10) compared to 1 if c were elected (for voter 11).[19] Though this gives a well-defined rule in the spirit of seq-Phragmén, we defer its investigation to future research.

It is important to realise that any assumptions concerning a non-priority rule will carry over to its priority adaptations. For example, if, as we have done, it is assumed in the context of seq-Phragmén that there are no unapproved candidates, then rl-seq-Phragmén also requires this assumption. We briefly discuss the assumption that all candidates have positive approval in the context of the priority model in Section 2.3.

We saw in Lemma 2.3 that in the case of dual election rules with candidate-wise scoring functions, the reserved-first approach comes clos-

[18] It may be tempting to think that we could break ties by considering which candidate became affordable first. This will not do because, generally, the time a candidate became affordable cannot be taken as a measure of how suitable they are in later rounds. For example, the group of voters responsible for a candidate's becoming affordable early on may be given relatively many representatives in the following rounds, hence a candidate that became affordable later on, supported by a different group of voters, might be preferable.

[19] For a formal treatment of how this tiebreaking would work, see the definitions of the *leximax ordering* and Phragmén's Leximax Rule by Brill et al. (2023).

est to preserving the original rule (in this case maximising the score) while meeting the requirement that q priority candidates be elected. Though the reserved-first approach for seq-Phragmén and MES is complicated, there is another, intuitively promising option if we are concerned with filling the quota while, as best as possible, maintaining the desirable properties of the original rules. When executing a sequential rule $\mathcal{R}$, we can keep track of how many priority candidates have been elected and restrict the candidate pool to the remaining priority candidates at the *latest possible moment* that this must be done in order to guarantee that q priority candidates will be elected. We refer to this as *reserved-latest* $\mathcal{R}$.

Rule 13: Reserved-latest $\mathcal{R}$, late-$\mathcal{R}$

Let $\mathcal{R}$ be any of the described non-priority sequential rules, and let $p(r)$ denote the number of priority candidates that have been elected in the rounds preceding r, i.e., $p(r) = |W^+_{r-1}|$. Execute $\mathcal{R}$ as normal, checking at the start of each round r whether $q - p(r) > k - r$, i.e., whether the number of rounds remaining is smaller than the number of priority candidates that must still be elected in order to make sure the quota is respected. From the first round for which this is the case onward, the candidate-pool is restricted to the remaining priority candidates. I.e., if ℓ is the first round r for which $q - p(r) > k - r$, then for any round $\ell \leq r' \leq k$ we have the candidate pool $C_{r'} = C^+ \setminus W_{r'-1}$.

Note that neither seq-Phragmén nor MES give rise to any problems here. Intuitively, this is because a 'better' candidate is never elected after a 'worse' candidate. Also note that the above definition does not require the strong assumption that $|C^+| \geq k$.

We now briefly discuss a number of general observations concerning reserved-latest rules that are independent of the (non-priority) rules on which they are based.

First off, at most the last q rounds will be restricted. In any of the first $k-q$ rounds, there will be at least q rounds remaining. But then for each of these rounds there are at least as many rounds left (q or more) as priority candidates that need to be elected to respect the quota (at most q), hence these rounds cannot be restricted. When the last q rounds are restricted, the algorithm runs exactly as the corresponding reserved-last rule.

Similarly, it is clear that if the restriction to priority candidates does not come into play at any point, then the algorithm works exactly as the original non-priority rule. In any case, for both the reserved-last and reserved-latest rules, the first $k - q$ rounds are identical to those of the original rule (since these are guaranteed to be unrestricted).

If the restriction to priority candidates becomes active at any point, then *exactly* q priority candidates are elected. The restriction first comes into force in the first round r for which $q - p(r) > k - r$, from which it follows that $p(r) = q - k + r - 1$. Since a priority candidate is elected in each of the last $k - r + 1$ rounds, the total number of priority candidates elected is $p(r)+(k-r+1) = q-k+r-1+k-r+1 = q$. Intuitively, it is clear that this must be the case, since we are restricting the candidate pool to priority candidates at the *latest possible moment* that this can be done in order to guarantee the quota is respected. Taking the contrapositive, we also know that if more than q priority candidates have been elected, it cannot be the case that the restriction came into force.

Finally, we can say something about *when* the restriction comes into force. First off, it cannot be the case that the restriction comes into force and that q or more priority candidates have been elected before that.[20] Secondly, if exactly (at least) $x < q$ priority candidates are elected before the first restricted round, then exactly (at least) the first $k-q+x$ rounds are unrestricted.

Interestingly, the reserved-first and reserved-latest adaptations of sequential versions of candidate-wise rules, e.g., seq-AV and seq-SAV, are equivalent. The underlying intuition is that for candidate-wise scoring rules, the reserved-first rule also aims to deviate from the original rule only as much as necessary to respect the quota. First, the highest-scoring q priority candidates are elected, guaranteeing that the quota is met with the strongest possible candidates. After that, the highest-scoring remaining candidates are elected.[21]

Proposition 2.4. *If $\mathcal{R}$ is a candidate-wise scoring rule, then rf-seq-$\mathcal{R}$ and late-seq-$\mathcal{R}$ are equivalent.*

Proof. Let $\mathcal{R}$ be a candidate-wise scoring rule, let E be an arbitrary priority election, and let $W_{\text{rf}} = \text{rf-seq-}\mathcal{R}(E)$ and $W_{\text{late}} = \text{late-seq-}\mathcal{R}(E)$. We make a case distinction in whether the restriction to priority candidates came into force in the execution of late-seq-$\mathcal{R}$ or not. When we say that the candidates from some set are the highest-scoring, we mean this to take tiebreaking into consideration.

Suppose the restriction did not come into effect. Then W_{late} simply consists of the k highest-scoring candidates. Now let W^q_{late} be the q highest-scoring priority candidates in W_{late}. These are the highest-scoring priority candidates simpliciter and will thus be elected in the first q rounds of rf-seq-$\mathcal{R}$. But then $W_{\text{late}} \setminus W^q_{\text{late}}$ consists of the highest-scoring candidates that are not elected in the first q rounds of rf-seq-$\mathcal{R}$. Thus, $W_{\text{rf}} = W_{\text{late}}$.

[20] The previous paragraph already tells us that the restriction did not come into force if more than q priority candidates were elected. *A fortiori*, this holds if more than q priority candidates are elected in unrestricted rounds. Here, however, we additionally claim that it cannot be the case that the restriction came into force after (exactly) q priority candidates were elected in unrestricted rounds.

[21] This is in contrast with the reserved-last rule, in which the reserved-seats are not necessarily filled with the strongest priority candidates.

Suppose the restriction did come into effect. Then we must have $|W^+_{\text{late}}| = q$. The priority candidates in W_{late} are the q highest-scoring and are thus elected in the first q rounds of rf-seq-$\mathcal{R}$. Analogously to the previous case, we can then conclude that $W_{\text{rf}} = W_{\text{late}}$. □

We have already seen that D^f_{sc}, where sc is a candidate-wise scoring function, returns committees that respect the quota with maximal score (Lemma 2.3). Thus, D^f_{sc} is as close as we can get to the original rule if we want to make sure q priority candidates are elected. Recall that the same intuition motivates reserved-latest sequential rules, hence it should not come as a surprise that these two approaches coincide for candidate-wise rules. This is not the case for seq-CC, as the following example shows.

Example 2.5. Consider the election with the candidate set $C = \{\overline{a}, \overline{b}, \overline{c}, d, e\}$, committee size $k = 3$, quota $q = 1$, and the profile

$$4 \times \{\overline{a}, d\} \qquad 3 \times \{\overline{b}\} \qquad 1 \times \{d\} \qquad 2 \times \{e\}.$$

It is easy to check that rf-seq-CC, rl-seq-CC and late-seq-CC all give different outcomes, namely, $\{\overline{a}, \overline{b}, e\}$, $\{\overline{b}, \overline{c}, d\}$ and $\{\overline{b}, d, e\}$, respectively.

We have argued that, intuitively, the reserved-latest adaptation attempts to come as close as possible to the original rule, and in the above example, the committee returned by late-seq-CC has a maximal CC-score. In this light, it might be tempting to think that late-seq-CC will always yield a better CC-score than rf-seq-CC or rl-seq-CC. The following example shows that this is not the case.

Example 2.6. Consider the election with $k = 3$, $q = 2$ and the profile

$$7 \times \{\overline{a}\} \quad 2 \times \{\overline{b}, \overline{c}\} \quad 3 \times \{\overline{b}, e\} \quad 1 \times \{\overline{c}\} \quad 2 \times \{\overline{d}\} \quad 3 \times \{\overline{d}, e\}.$$

Both rf-seq-CC and rl-seq-CC return the committee $\{\overline{a}, \overline{b}, \overline{d}\}$, which has the maximal CC-score of 17, while late-seq-CC returns the seq-CC winning committee $\{\overline{a}, \overline{c}, e\}$, which has a CC-score of 15.

Note that the above example also shows that seq-CC differs from CC (as was claimed before), and that priority adaptations of seq-CC may actually perform better than seq-CC.[22] Example 2.6 also shows that rl-seq-Phragmén and late-seq-Phragmén differ. The former gives the outcome $\{\overline{a}, \overline{b}, \overline{d}\}$, while the latter returns the seq-Phragmén winning committee, which is $\{\overline{a}, \overline{b}, e\}$ or $\{\overline{a}, \overline{d}, e\}$ depending on tiebreaking.

[22] These two facts are fundamentally related. Greedily selecting candidates may ultimately harm the CC-score. However, in the priority setting it is possible that a (priority) candidate that does not maximally increase the score must be elected, making it possible to elect an ultimately more representative committee.

The following lemma gives some insight into how the priority adaptations of seq-Phragmén relate to the original rule, and allows us to apply the discrete and continuous formulation of seq-Phragmén to these adaptations.

Lemma 2.5. *Let $E = (C, C^+, \boldsymbol{A}, k, q)$ be a priority election and let $W =$ late-seq-Phragmén(E) be the late-seq-Phragmén winning committee. Now consider the non-priority election $E' = (C', \boldsymbol{A}, k)$, where $C' = C^+ \cup W$ consists of the winning candidates from W and any remaining priority candidates. That is, C' is obtained by removing from C the non-priority candidates that are not selected by late-seq-Phragmén. Then it is the case that $W =$ seq-Phragmén(E'). The same holds for rl-seq-Phragmén.*

Proof. Suppose that the first ℓ rounds of late-seq-Phragmén on E are unrestricted. Then these are identical to the first ℓ rounds of seq-Phragmén on $(C, \boldsymbol{A}, k)$, which in turn are identical to the first rounds of seq-Phragmén on $(C', \boldsymbol{A}, k)$. This relies on the fact that we can remove any unelected candidates from the candidate set (and formally also from the ballots) without making a difference to the execution of seq-Phragmén. To see that this is the case, simply note that we choose the candidates who at each step minimise the resulting maximum load. Removing candidates that do not do this does not change that, since the loads resulting from a candidate depend solely on the voters that approve of them and their current loads. From round $\ell + 1$ on, in late-seq-Phragmén, we simply run seq-Phragmén considering only the priority candidates. But then, by the same reasoning as above, these rounds will be identical to rounds $\ell + 1$ to k of seq-Phragmén on $(C', \boldsymbol{A}, k)$. The same reasoning applies to rl-seq-Phragmén. $\square$

It follows directly that we can use the continuous and discrete formulations (Lemma 2.1 and Lemma 2.2) to describe both rl-seq-Phragmén and late-seq-Phragmén; the only difference lying in the restriction to priority candidates from the specified round onward. Thus, from here on, we regard these lemmas as applying to rl-seq-Phragmén and late-seq-Phragmén as well.

We have now introduced all of the rules with which we will be concerned in the following chapters. The next section discusses in more detail the assumptions made.

2.3 Assumptions

We have seen that priority rules require assumptions regarding the number of priority candidates. In particular, all of the priority rules require the assumption that $|C^+| \geq q$ in order to be well-defined. Assuming there is cause to give preferential treatment to priority candidates in the first place, this assumption will often not be satisfied in practice. This, however, need not be a problem. When less than the desired number of priority candidates are available, we can simply set $q = |C^+|$. For all the rules considered—i.e., rules that respect the quota—this will result in all priority candidates being elected, which seems desirable when less than the desired number of priority candidates are available.

More problematic is the assumption, necessary for the well-definedness of the reserved-last rules, that there are sufficient priority candidates to fill the entire committee, i.e., $|C^+| \geq k$. Again, the underlying reasons for giving preferential treatment to certain candidates in the first place will often make it make it very unlikely that a large number of such candidates is available. In certain cases there are easy adaptations that deal with this problem. E.g., for candidate-wise scoring rules the committee can simply be completed with the highest-scoring remaining candidates. For other rules, however, this is less trivial. In any case, it should be clear that, though it is likely that the assumption is not met, this will often not be a problem.[23] The ill-definedness of the reserved-last rules in the absence of the assumption that $|C^+| \geq k$ comes, as we have seen, from the possibility that there may not be enough priority candidates left for the reserved seats, which are filled last. Since candidates are elected, roughly, from 'best' to 'worst', it is in practice unlikely, given that priority candidates require some sort of preferential treatment, that many will be elected in unrestricted rounds. Nevertheless, in the future, these rules ought to be refined so that they do not require this assumption. In favour of simplicity, we do not do so in this thesis.

Finally, we mention again briefly the assumption concerning seq-Phragmén that all candidates are approved of by some voter. Again, it is likely that in practice this assumption will be violated, given the fact that priority treatment is necessary in the first place. There are ways to refine seq-Phragmén such that these assumptions need not be made, but this makes the presentation and handling of the rule in proofs more difficult. For this reason we stick with the assumption here.

Having introduced the priority model and the rules with which we will be concerned in this thesis, we can now briefly consider related work.

2.4 Related work

Besides excellence-based elections and elections concerning proportional representation, both of which were discussed in the introduction, elections have also been considered where the emphasis is on representing smaller groups of voters.[24] In the extreme case this requires what is called *diversity*. The goal here is to elect a committee that represents as many voters as possible. Examples include selecting locations for a

[23] This is also illustrated by the examples that do not satisfy this assumption, e.g., Examples 2.3, 2.5, and 2.6.

[24] In excellence-based elections candidates with the most support are elected, thus favouring bigger groups of voters. Proportional representation prescribes a linear relation between the size of a group of voters and the representation they require.

given number of facilities such as fire stations, choosing the products to be advertised by an online store, and deciding which movies will be offered for in-flight entertainment (Elkind et al., 2017; Faliszewski et al., 2017c). This is very different from the setting we consider, in which an exogenously given quota is to be respected whatever the voters' preferences may be. Such a quota will often be motivated by considerations of diversity with regards to the types of candidates elected; however, this is entirely separate to the notion of diversity with regards to the voters that are represented. In fact, though we do not do so here, it would be possible to consider the above described notion of diversity with regards to voter representation in the context of the priority model (much like we have done for excellence-based elections and proportionality). For more work on the notion of diversity described above, see the accompanying footnote.[25]

Lang and Skowron (2018) present a framework in which candidates have external attributes, such as gender, political affiliation, or academic faculty. For each value that an external attribute can take, e.g., the value *woman* for the attribute *gender*, the ideal proportion of the committee that instantiates this value is specified. The goal is then to select a fixed-size committee that comes as close as possible to the desired distribution of attributes. This generalises settings such as *bi-apportionment*, where each candidate has a political affiliation and a region they represent (see Pukelsheim, 2017). The desired distribution may be based on votes or given exogenously. In any case, this model takes into account only external constraints on the attributes of candidates, wherever these may come from. Work has been done on combining this model with multi-winner voting rules (Bredereck et al., 2018; Celis et al., 2018), though attention has been restricted to rules that provide a ranking over committees. Essentially, we can interpret our priority model as one in which there is a single attribute: priority status. However, in the priority model we are not concerned with approximating as best as possible some *fixed* ideal distribution; i.e., we are not considering committees with *exactly* q candidates. Rather, we have a requirement concerning the minimal representation of priority candidates, which we assume can always be satisfied. This also allows us to consider adaptations of rules, such as seq-Phragmén, which are more problematic in the above setting as they do not naturally yield a ranking of committees.

The problem of *school choice* is a *matching* problem in which students are to be assigned to schools based on their preferences and those of the schools. *Controlled school choice*, in which schools reserve seats for certain kinds of students, for example for students of certain ethnic backgrounds or from within a walk zone, has been widely studied since the first mechanism-design approach to school choice, credited to Abdulkadiroğlu and Sönmez (2003). This has been done both from the per-

[25] See, e.g., Faliszewski et al., 2017b; Faliszewski and Talmon, 2018; Jaworski and Skowron, 2022; Lackner and Skowron, 2018.

spective of matching mechanisms as a whole (Abdulkadiroğlu & Sönmez, 2003; Doğan, 2016; Hafalir et al., 2013; Kojima, 2012), as well as from the perspective of individual schools (Doğan et al., 2021; Dur et al., 2018; Dur et al., 2020; Echenique & Yenmez, 2015). In the mechanisms studied, based on *deferred acceptance* (Gale & Shapley, 1962), the main distinction is between whether schools fill reserved seats first or last, a distinction that is central to the definition of our priority rules as well.[26] It is shown that this choice has significant consequences. Given that schools' preferences are modelled as rankings over students, we can draw a parallel with candidate-wise multiwinner scoring rules, which yield a ranking over candidates. In this light, the first half of Chapter 4 is analogous to this line of research in school choice. A significant difference is that in school choice, schools' rankings over students are given exogenously, while in the priority model rankings over the candidates must be deduced from voters' preferences. Consequently, whereas in school choice rankings over students are strict, the analogous rankings over candidates in the priority model may be weak. Furthermore, we consider properties such as efficiency, which concern voters' satisfaction and have no direct counterpart in matching.

In this chapter we introduced the priority model and the corresponding rules with which we will be concerned. We briefly discussed the significant assumptions that we make and gave an overview of related work. The following chapters explore desirable properties in the priority setting and use these to analyse the rules here defined.

[26] Other *reservation sequences* are considered in the matching literature, but these have yet to be clearly characterised.

Chapter 3

Anonymity, neutrality and priority treatment for priority candidates

In the foregoing chapter we introduced the priority model along with a number of priority rules. In this chapter we consider some basic properties of such rules. The first section considers (adaptations of) the standard axioms *anonymity* and *neutrality*. The following section introduces a number of axioms that capture the basic requirement that priority rules may not (inadvertently) disadvantage priority candidates. Table 3.1 gives an overview of which rules (fail to) satisfy the properties considered in this chapter. The abbreviations PW, GW, CPW, and CGW, refer to the axioms *possible win*, *guaranteed win*, *committee possible win*, and *committee guaranteed win*, respectively.

	axiom			
rule	PW	GW	CPW	CGW
D^f_{sc}	✓[3.1]	✓[3.2]	✗[ex.3.2]	✗[ex.3.2]
D^ℓ_{sc}	✓[3.3.1]	✓[3.4]	✓[3.3]	✓[3.5]
rf-seq-CC	-	✓	-	✗[ex.3.3]
rl-seq-CC	-	✓[3.6.1]	-	✓[3.6]
late-seq-CC	-	✓[3.7]	-	✗[ex.3.3]
rl-seq-Phragmén	-	✓[3.6.1]	-	✓[3.6]
late-seq-Phragmén	-	✓[3.7]	-	✗[ex.3.3]
rl-MES	-	✓[3.6.1]	-	✓[3.6]
late-MES	-	✓[3.7]	-	✗[ex.3.3]

Table 3.1: We have used sc as a placeholder for any candidate-wise scoring function. Since all considered rules satisfy *respect of quota*, we do not include it. Furthermore, we do not include *anonymity* or *priority neutrality* since the results mirror the non-priority setting. Lastly, as for resolute rules PW and GW (respectively, CPW and CGW) are equivalent, we only show the results for GW (respectively, CGW).

3.1 Anonymity and neutrality

Anonymity is generally considered one of the basic requirements of a mechanism in the social choice literature (Arrow et al., 2002).[1] Informally, anonymity requires that voters be treated equally; i.e., that it only matters which ballots are submitted and that it does not matter which voters submit them. We will need the following notation for the formal definition. Given a profile $\boldsymbol{A}$ and a permutation $\pi : N \to N$ on the voters, let $\boldsymbol{A} \circ \pi = (A_{\pi(1)}, \ldots, A_{\pi(n)})$. That is, $\boldsymbol{A} \circ \pi$ is the profile in which each voter $i \in N$ submits the ballot that voter $\pi(i)$ submits in $\boldsymbol{A}$.[2]

Axiom 1: Anonymity

A rule $\mathcal{R}$ satisfies *anonymity* if for any election $(C, \boldsymbol{A}, k)$ and permutation $\pi : N \to N$ on the voters, it is the case that $\mathcal{R}(C, \boldsymbol{A} \circ \pi, k) = \mathcal{R}(C, \boldsymbol{A}, k)$.

The above is the standard definition of anonymity for (multiwinner) voting and applies directly to the priority setting as well. We note that all of the (non-priority and priority) rules considered here satisfy anonymity.

Neutrality, another basic requirement in the social choice literature, requires that candidates be treated equally (Arrow et al., 2002). Normally, this is captured by the demand that changing candidates' names results in a corresponding change in the winning committees. In order to capture this formally we need the following notation. Given a profile $\boldsymbol{A}$ and a permutation $\pi : C \to C$ on the candidates, we define $\pi^* \circ \boldsymbol{A} = (\pi^*(A_1), \ldots, \pi^*(A_n))$, where $\pi^*(X) = \{\pi(c) : c \in X\}$.[3] A rule $\mathcal{R}$ is then said to be *neutral* if for any election $(C, \boldsymbol{A}, k)$ and permutation $\pi : C \to C$ we have $\mathcal{R}(C, \pi \circ \boldsymbol{A}, k) = \pi^*(\mathcal{R}(C, \boldsymbol{A}, k))$, where $\pi^*(\mathcal{F}) = \{\pi^*(S) : S \in \mathcal{F}\}$ for any family of committees $\mathcal{F} \subseteq \mathcal{P}(C)$.

When giving some kind of preferential treatment to priority candidates, we may expect to violate neutrality. After all, priority candidates are purposefully treated differently to non-priority candidates. The following example illustrates this.

Example 3.1. Consider the election E with set of candidates $C = \{\overline{a}, b, c\}$, committee size $k = 2$, quota $q = 1$ and the profile $\boldsymbol{A}$:

$$A_1 = \{\overline{a}, b\} \qquad\qquad A_2 = \{\overline{a}\}.$$

[1] For anonymity in the context of more recent developments, such as judgement aggregation, sequential voting and voting rule verification, see, for example, Brandt et al. (2016) or Endriss (2017).

[2] I use the notation of Lackner and Skowron (2023) who treat $\boldsymbol{A}$ as a function from N to $\mathcal{P}(C)$. The composition $\boldsymbol{A} \circ \pi$ is then also a function from N to $\mathcal{P}(C)$.

[3] See footnote 2. If π^* is a bijection on $\mathcal{P}(C)$ and $\boldsymbol{A}$ a function from N to $\mathcal{P}(C)$, then $\pi^* \circ \boldsymbol{A}$ is also a function from N to $\mathcal{P}(C)$.

We then have $D^f_{\text{AV}}(E) = \{\overline{a}, b\}$. Now consider the permutation $\pi : C \to C$ with $\pi(\overline{a}) = b$, $\pi(b) = c$ and $\pi(c) = \overline{a}$. We then have

$$\begin{aligned} D^f_{\text{AV}}(\pi^* \circ \boldsymbol{A}) &= D^f_{\text{AV}}(\{b, c\}, \{b\}) \\ &= \{\overline{a}, b\} \\ &\neq \{b, c\} \\ &= \pi^*(D^f_{\text{AV}}(E)). \end{aligned}$$

Nevertheless, it seems that D^f_{AV} is as neutral as can be given its priority treatment. A natural first attempt at a notion of neutrality for the priority setting might require that any two candidates of the same status are treated equally. That is, we might keep the definition of neutrality as is, except that we restrict attention to permutations $\pi : C \to C$ for which $\pi(c) = c'$ implies that c and c' have the same status. We refer to this as *restricted neutrality*. Though intuitive, restricted neutrality is too weak. Priority neutrality ought to entail that when we switch the names of candidates c and c' *and switch their status as well*, c' is elected if and only if c was originally elected. We thus arrive at the following notion of neutrality.

Axiom 2: Priority neutrality

A rule $\mathcal{R}$ satisfies *priority neutrality* if for any election $E = (C, C^+, \boldsymbol{A}, k, q)$ and permutation $\pi : C \to C$ we have

$$\mathcal{R}(C, \pi^*(C^+), \pi^* \circ \boldsymbol{A}, k, q) = \pi^*(\mathcal{R}(E)).$$

To illustrate that our definition of neutrality prohibits non-neutral treatment of candidates that is not motivated by priority treatment we consider an example. Let $\mathcal{R}$ be the rule that is as AV except that when candidate a is a priority candidate, only non-priority candidates are elected. Thus, when a is a non-priority candidate, k of the most-approved candidates are elected, and when a is a priority candidate, k of the most-approved non-priority candidates are elected. This is clearly a violation of priority neutrality; however, it is not a violation of the earlier form of neutrality discussed: since no priority candidates are elected whenever a is a priority candidate, all priority candidates are treated equally in that case.

To make this a little clearer, we can formalise it. We know that when a is not a priority candidate, $\mathcal{R}$ is equivalent to AV. AV is neutral and, *a fortiori*, also satisfies restricted neutrality, hence when a is a priority candidate, restricted neutrality is respected. Now let E be any election for which $a \in C^+$. Then any $W \in \mathcal{R}(E)$ consists of k of the most-approved non-priority candidates. Now let $\pi : C \to C$ be a permutation that never assigns to a candidate another candidate with a differing status and let $E_\pi = (C, C^+, \pi^* \circ \boldsymbol{A}, k, q)$. We know that $\mathcal{R}(E_\pi)$ consists of committees of the most-approved non-priority candidates of C. But these are exactly $\{\pi^*(W) : W \in \mathcal{R}(E)\}$. Thus $\mathcal{R}$ satisfies restricted neutrality.

To see that $\mathcal{R}$ does not satisfy priority neutrality, consider the election E with $C = \{a, b\}$, $C^+ = \{a\}$, $k = 1$, $q = 0$, and two voters: $A_1 = \{a\}$ and $A_2 = \emptyset$. Since $a \in C^+$, we have $\mathcal{R}(E) = \{b\}$. Now let π be the non-identity permutation on C and consider the election $E_\pi = (C, \pi^*(C^+), \pi^* \circ \boldsymbol{A}, k, q)$. We have $\mathcal{R}(E_\pi) = \{b\}$, since now b is the most-approved candidate and $a \notin \pi^*(C^+)$. Since $\pi^*(\mathcal{R}(E)) = \{a\}$, this violates priority neutrality.

It is easy to see that the dual election rules, which are not resolute, satisfy priority neutrality. In contrast, the sequential rules, which are resolute, violate priority neutrality since they employ a tiebreaking order over candidates. Note that when $q = 0$, the priority rules that we consider reduce to the underlying non-priority rules. In such cases, priority neutrality also reduces to neutrality. It follows, since the non-priority sequential rules considered are not neutral, that their priority adaptations fail priority neutrality.[4] Thus, analogously to the non-priority setting, there is a trade-off between neutrality and resoluteness.

Having looked at the fundamental properties of anonymity and neutrality, we now move on to consider properties that explicitly concern the preferential treatment given to priority candidates.

3.2 Priority treatment for priority candidates

We are concerned with rules that will guarantee that at least q priority candidates are elected. We already defined *respect of quota* in the previous chapter, but state it formally here.

Axiom 3: Respect of quota, RoQ

A committee $W \in C[k]$, given an election $E = (C, C^+, \boldsymbol{A}, k, q)$, is said to *respect the quota* if the existence of an unelected priority candidate $c \in C^+ \setminus W$ implies $|W^+| \geq q$. A rule $\mathcal{R}$ satisfies *respect of quota* if it returns only committees that respect the quota.

The above axiom requires that whenever there are at least q priority candidates, at least q such candidates are elected, and otherwise, if there are fewer than q priority candidates, all priority candidates will be elected. Given our assumption that $|C^+| \geq q$, this is equivalent to the requirement that q priority candidates be elected. Nevertheless, we have chosen to formulate the axiom so that it does not depend on this assumption.

[4] The same reasoning applies when $q = k$, for then a sequential priority rule can be seen as simply executing the underlying non-priority rule on the profile obtained by removing all the non-priority candidates.

The following axioms that we consider are concerned with the minimal requirement that candidates are not worse off as priority candidates than as non-priority candidates. Such a requirement can also be seen as a kind of incentive-compatibility. For example, in some situations priority candidates may have to register in order to be officially recognised as priority candidates. If in such cases it cannot be guaranteed that being considered a priority candidate will not harm one's chances of being selected, priority candidates may be hesitant to communicate their actual priority status.

This idea gives rise to a number of different axioms, which vary according to how we interpret "harm one's chances". Perhaps the two most obvious requirements are that if a non-priority candidate wins in some (all) outcomes, then they should also win in some (all) outcomes when their status is changed to priority. We formalise these requirements as follows.

Axiom 4: Possible win, PW

A rule $\mathcal{R}$ satisfies *possible win* if for any election $E = (C, C^+, \boldsymbol{A}, k, q)$ and non-priority candidate $c \in C^-$, it holds that if there is some winning committee $W \in \mathcal{R}(E)$ such that $c \in W$, then for the election $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$, there must be a winning committee $W' \in \mathcal{R}(E')$ such that $c \in W'$.

Axiom 5: Guaranteed win, GW

A rule $\mathcal{R}$ satisfies *guaranteed win* if for any election $E = (C, C^+, \boldsymbol{A}, k, q)$ and non-priority candidate $c \in C^-$, it holds that if $c \in W$ for every $W \in \mathcal{R}(E)$, then $c \in W'$ for every $W' \in \mathcal{R}(E')$, where $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$.

Note that PW and GW are equivalent for resolute rules, for then there is only a single winning committee. We first consider the dual election rules based on candidate-wise scoring rules in relation to these axioms. In what follows sc is any candidate-wise scoring function and $\mathcal{R}_{\text{sc}}$ is the (non-priority) rule that returns committees maximising this score.

Proposition 3.1. *D^f_{sc}, i.e., the reserved-first dual election rule based on the candidate-wise scoring function* sc*, satisfies PW.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$ and $W \in D^f_{\text{sc}}(E)$ be an arbitrary election and winning committee, respectively. It follows that there exist sets $W_1 \in \mathcal{R}_{\text{sc}}(C^+, \boldsymbol{A}, q)$ and $W_2 \in \mathcal{R}_{\text{sc}}(C \setminus W_1, \boldsymbol{A}, k - q)$ with $W_1 \cup W_2 = W$. If there are no elected non-priority candidates, the claim holds vacuously, so let $c \in W^-$ be an arbitrary winning non-priority candidate (from which it follows that $c \in W_2$). Now consider the election $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$ in which c is a priority candidate. If

$W_1 \in \mathcal{R}_{\text{sc}}(C^+ \cup \{c\}, \boldsymbol{A}, q)$, then, since $W_2 \in \mathcal{R}_{\text{sc}}(C \setminus W_1, \boldsymbol{A}, k-q)$, we have $W \in D^f_{\text{sc}}(E')$. Otherwise, if $W_1 \notin \mathcal{R}_{\text{sc}}(C^+ \cup \{c\}, \boldsymbol{A}, q)$, then it must be the case that c has a strictly higher score than some member of W_1 and, consequently, c is a member of each winning committee in $\mathcal{R}_{\text{sc}}(C^+ \cup \{c\}, \boldsymbol{A}, q)$. Thus, in either case c is a member of some winning committee. □

Proposition 3.2. *D^f_{sc} satisfies GW.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$ and $c \in C^-$ be an arbitrary election and non-priority candidate, respectively, and define $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$. Now suppose c is a member of every winning committee $W \in D^f_{\text{sc}}(E)$ and, towards a contradiction, suppose that there is a committee $W' \in D^f_{\text{sc}}(E')$ with $c \notin W'$. It follows that there exist $W'_1 \in \mathcal{R}_{\text{sc}}(C^+ \cup \{c\}, \boldsymbol{A}, q)$ and $W'_2 \in \mathcal{R}_{\text{sc}}(C \setminus W'_1, \boldsymbol{A}, k-q)$ such that $W'_1 \cup W'_2 = W'$. Thus, W'_1 consists of q of the most-approved priority candidates in E', which, since $c \notin W'_1$, are also q of the most-approved priority candidates in E, i.e., $W'_1 \in \mathcal{R}_{\text{sc}}(C^+, \boldsymbol{A}, q)$. But then, since $W'_2 \in \mathcal{R}_{\text{sc}}(C \setminus W'_1, \boldsymbol{A}, k-q)$, we must have $W' = W'_1 \cup W'_2 \in D^f_{\text{sc}}(E)$. This contradicts the assumption that c is a member of every winning committee for E. □

As might be expected, the reserved-last dual election rule based on candidate-wise scoring function sc satisfies both PW and GW as well. In fact, D^ℓ_{sc} satisfies a property stronger than PW, which we call *committee possible win* (CPW). CPW requires that a winning committee remain a winning committee when the status of one of its non-priority candidates is changed to priority. *A fortiori*, a non-priority candidate that is a member of some winning committee will be a member of some, namely, the same, winning committee when their status is changed to priority.

Axiom 6: Committee possible win, CPW

A rule $\mathcal{R}$ satisfies *committee possible win* if for any election $E = (C, C^+, \boldsymbol{A}, k, q)$, winning committee $W \in \mathcal{R}(E)$ and non-priority candidate $c \in W^-$ it holds that $W \in \mathcal{R}(E')$, where $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k)$.

Proposition 3.3. *D^ℓ_{sc} satisfies CPW.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$ and $W \in D^\ell_{\text{sc}}(E)$ be an election and winning committee, respectively, and let $c \in W^-$ be any elected non-priority candidate (for the claim holds vacuously when $W^- = \emptyset$). It follows that there exist $W_1 \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k-q)$ and $W_2 \in \mathcal{R}_{\text{sc}}(C^+ \setminus W_1, \boldsymbol{A}, q)$ with $W_1 \cup W_2 = W$. Furthermore, it must be the case that $c \in W_1$. Now consider the election $E = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$. Since $c \in W_1$ we have $C^+ \setminus W_1 = (C^+ \cup \{c\}) \setminus W_1$. But then we have $W_1 \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k-q)$ and $W_2 \in \mathcal{R}_{\text{sc}}((C^+ \cup \{c\}) \setminus W_1, \boldsymbol{A}, q)$. That is, $W = W_1 \cup W_2 \in D^\ell_{\text{sc}}(E')$. □

Corollary 3.3.1. *D^{ℓ}_{sc} satisfies PW.*

Proposition 3.4. *D^{ℓ}_{sc} satisfies GW.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$ be an election and suppose $c \in W^-$ for every $W \in D^{\ell}_{\text{sc}}(E)$. That is, c is a non-priority candidate that is guaranteed to be selected. It follows that for any $W_1 \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k - q)$ and $W_2 \in \mathcal{R}_{\text{sc}}(C^+ \setminus W_1, \boldsymbol{A}, q)$, we have $c \in W_1$. It follows directly that $c \in W'$ for every winning committee $W' \in D^{\ell}_{\text{sc}}(E')$ where $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$. Informally, the open election for E and E' are identical, so if c is always elected in the former it must always be elected in the latter. $\square$

Where PW requires that a possibly winning candidate remain possibly winning when their status is changed to priority, CPW requires that a winning committee remain a winning committee when the status of one of its candidates is changed to priority. Analogously, we can define an axiom *committee guaranteed win*, which requires that if there is a unique winning committee, this committee remain the sole winning committee when the status of one of its candidates is changed to priority.

Axiom 7: Committee guaranteed win, CGW

A rule $\mathcal{R}$ satisfies *committee guaranteed win* if for any election $E = (C, C^+, \boldsymbol{A}, k, q)$ with a unique winning committee W that contains a non-priority candidate $c \in W^-$, it holds that W is the unique winning committee for $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k)$, i.e., $\mathcal{R}(E') = \{W\}$.

Note that, while CPW implies PW, CGW does not imply GW. This is because a candidate is guaranteed to be selected if they occur in all, possibly multiple, winning committees, whereas a particular committee is guaranteed to be the winning committee only if there are no other winning committees.

Proposition 3.5. *D^{ℓ}_{sc} satisfies CGW.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$ and $W \in D^{\ell}_{\text{sc}}(E)$ be an election and *unique* winning committee, respectively, and let $c \in W^-$ be an elected non-priority candidate (for the claim holds vacuously when $W^- = \emptyset$). Suppose, towards a contradiction, that there exists some $W' \neq W$ such that $W' \in D^{\ell}_{\text{sc}}(E')$, where $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$. It follows that there exist $W'_1 \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k - q)$ and $W'_2 \in \mathcal{R}_{\text{sc}}((C^+ \cup \{c\}) \setminus W'_1, \boldsymbol{A}, q)$ with $W'_1 \cup W'_2 = W'$. It must be the case that $c \in W'_1$, otherwise there would be some winning committee for E that does not contain c, which contradicts the assumption that W is the unique winning committee for E. But then we have $C^+ \setminus W'_1 = (C^+ \cup \{c\}) \setminus W'_1$, from which it follows that $W'_2 \in \mathcal{R}_{\text{sc}}(C^+ \setminus W'_1, \boldsymbol{A}, k - q)$. That is, $W'_1 \cup W'_2 = W' \in D^{\ell}_{\text{sc}}(E)$. Since $W' \neq W$, this contradicts our assumption that W is the unique winning committee for E, and we can conclude that there is no such W'. $\square$

PW and GW capture the fact that an individual candidate ought not to be worse off as a priority candidate than as a non-priority candidate. It is clear why these properties are desirable whenever priority treatment is in order. This is not as clear for CPW and CGW, as the following example, which shows that D^f_{sc} may violate both axioms, illustrates.

Example 3.2. This example works both if we take the approval score $\mathrm{sc}_{\mathrm{AV}}$ or the satisfaction approval score $\mathrm{sc}_{\mathrm{SAV}}$. Consider the election with the set of candidates $C = \{a, b, c\}$ of which a is the only priority candidate ($C^+ = \{a\}$), committee size $k = 2$, and quota $q = 1$. Furthermore, we have a two-voter profile:

$$A_1 = \{b, c\} \qquad\qquad A_2 = \{b\}.$$

Since a is the only priority candidate, they are elected in the quota election. Next, as b is the highest-scoring remaining candidate, b is elected in the open election. Thus, $\{a, b\}$ is the unique winning committee. Now consider what happens when b becomes a priority candidate: b wins the quota election and c is preferred to a in the following open election. Thus $\{b, c\}$ is the unique winning committee, which violates both CPW and CGW.

The above example illustrates why CPW and CGW ought not to be considered minimal requirements that any priority rule must meet. In order to respect the quota, q priority candidates must be elected. When a stronger priority candidate becomes available, it will often make sense that some weaker, previously elected priority candidate should lose their place if there is a stronger, previously unelected non-priority candidate available.

We now consider the sequential rules. Since PW and GW are equivalent for these (resolute) rules, and the same holds for CPW and CGW, we consider only GW and CGW. It should also be noted that in the context of resolute rules, CGW does imply GW.

Proposition 3.6. *rl-seq-CC, rl-seq-Phragmén and rl-MES satisfy CGW.*

Proof. Let $\mathcal{R}$ be any of the three rules. Now let $E = (C, C^+, \boldsymbol{A}, k, q)$ and $W \in \mathcal{R}(E)$ be an election and the corresponding winning committee, respectively. The desired result follows directly if $W^- = \emptyset$, so let $c \in W^-$ be an elected non-priority candidate. It follows that c must have been elected in one of the first $k - q$ unrestricted rounds. Now consider the election $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$. As the restricted rounds are fixed—i.e., exactly the last q rounds are restricted—the same candidates are eligible and selected at each round of $\mathcal{R}$ on E'. Note that this relies on the fact that a candidate's status is only relevant to whether they are eligible for the restricted rounds. $\square$

Corollary 3.6.1. *rl-seq-CC, rl-seq-Phragmén and rl-MES satisfy GW.*

The corollary follows immediately, since CGW implies GW when the considered rule is resolute. We now consider the reserved-latest rules.

Proposition 3.7. *late-seq-CC, late-seq-Phragmén and late-MES satisfy GW.*

Proof. This proof is almost entirely analogous to that of Proposition 3.6. The only difference is that in this case changing c's status to priority means that the restriction to priority candidates comes into force one round later in E' than in E, which means we cannot guarantee that the same committee is elected. Nevertheless, the unrestricted rounds of E, say there are ℓ, will be identical to the first ℓ rounds of E', hence c is elected. □

The fact that rf-seq-CC satisfies GW follows from similar reasoning.

Proposition 3.8. *rf-seq-CC satisfies GW.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$, $W \in \mathcal{R}(E)$ and $c \in W^-$ be an election, winning committee, and winning non-priority candidate, respectively. It follows that c must have been elected in one of the last $k-q$ (unrestricted) rounds. Now consider the election $E' = (C, C^+ \cup \{c\}, \boldsymbol{A}, k, q)$. Now, either c is elected in one of the restricted rounds, or, otherwise, they are elected in exactly the same round as before. □

As the above proofs suggest, late-seq-CC, late-seq-Phragmén, late-MES and rf-seq-CC do not satisfy CGW. This is shown by the following example.

> **Example 3.3.** Like in Example 3.2, let $C = \{a, b, c\}$ be the set of candidates, with a the only priority candidate, and let $k = 2$ and $q = 1$. This time we consider a different profile with three voters:
>
> $$1 \times \{a\} \qquad\qquad 3 \times \{b\} \qquad\qquad 2 \times \{c\}.$$
>
> We first consider late-seq-CC, late-seq-Phragmén and late-MES. These rules will elect b in the first round and a, being the only priority candidate, in the second (restricted) round.[5] That is, $\{a, b\}$ is the unique winning committee. If, however, b's status is changed to priority, then b is still elected in the first round and, since the quota has already been met, the restriction does not come into force in round 2 and c is elected instead of a. This violates CGW.
>
> For rf-seq-CC, a is originally elected in the first (restricted) round and b in the second round. However, when b is a priority candidate, they will be elected in the first round instead of a, in which case c will be elected in the second round, thus violating CGW.

[5] Note that for MES, since $p = n/k = 3$, b is the only affordable candidate. Consequently, a is chosen in Phase 2 using seq-Phragmén.

The above example supports the earlier claim that CPW and CGW are too strong to be considered minimal requirements for all priority rules. The same thing is going on here as in Example 3.2: the availability of a new, strong priority candidate means that we can replace a weaker priority candidate with a non-priority candidate without violating the quota.

Nevertheless, the above results show an interesting trend: the priority rules that are (intuitively) minimal adaptations designed to respect the quota, i.e., reserved-first and reserved-latest rules, do not satisfy CPW and CGW, while the reserved-last rules do. This suggests that CPW and CGW may be read as guaranteeing some kind of priority treatment. Both axioms guarantee that a priority candidate will not be disadvantaged even if some other winning non-priority candidate gets priority status.[6] This agrees with the general finding concerning dual rules discussed in the coming chapter, that the reserved-last rules implement a greater degree of priority treatment.

Having looked at some of the basic requirements that priority rules ought to meet, we now move on to consider axioms in the context of excellence-based elections.

[6] Actually, this also holds for non-priority candidates: a winning non-priority candidate's status being changed cannot mean a different non-priority candidate is not elected because of that. However, this also holds for some of the rules that violate CPW and CGW, e.g., the dual election rules based on candidate-wise scores. Thus, CPW and CGW can be thought of as effectively capturing a kind of priority treatment.

Chapter 4

Excellence-oriented axioms and efficiency

The first section of this chapter is concerned with excellence-based elections in the context of the priority model. In a non-priority context, excellence-based elections can be thought of a situations where voters are experts and the goal is to elect the best k candidates. In the priority setting, the challenge is how to go about doing this if we want to respect the quota and afford preferential treatment to priority candidates. In Subsection 4.1.1 we introduce the notion of excellence-based elections in a bit more detail, motivate the consideration of such elections in the priority model, and explain why we consider only the candidate-wise dual rules in this context. After that, in Subsection 4.1.2, we consider a number of relevant axioms and show how these may be used to characterise the candidate-wise dual election rules.

The second section of the chapter concerns *efficiency* in the context of the priority model. This is considered in this chapter because the rules (excluding PAV) that are concerned with proportional representation, considered in Chapter 5, are not efficient.

Table 4.1 specifies which of the rules (do not) satisfy the axioms considered in this chapter. The abbreviations AQ, LtQ, LtAQ, qPE and PE, stand for *active quota*, *limit to quota*, *limit to active quota*, *quota priority efficiency* and *priority efficiency*, respectively. When something holds for both AV and SAV, we will use (S)AV to avoid clutter. For example, we will use $D^f_{\text{(S)AV}}$ when what we are claiming holds for both D^f_{AV} and D^f_{SAV}.

4.1 Excellence-based elections with priority candidates

This section, which concerns excellence-based elections in the priority model, is split into two parts. We first consider excellence-based elections and corresponding rules in the non-priority context, arguing that the priority context is also relevant. In this subsection we also highlight

	axiom				
rule	AQ	LtQ	LtAQ	qPE	PE
D^f_{sc}	-	✓[4.3]	✓[4.3.1]	-	-
D^ℓ_{sc}	✓[4.4]	-	✓[4.4]	-	-
$D^f_{\text{(S)AV}}$	✗	✓[4.3]	✓[4.3.1]	✓[4.6]	✓
$D^\ell_{\text{(S)AV}}$	✓[4.4]	✗	✓[4.4]	✗[ex.4.2]	✓[4.7]

Table 4.1: The satisfaction table for the excellence-based axioms and efficiency axioms, where sc is any candidate-wise scoring function. We do not include *priority merit* and *non-priority merit*, since these are satisfied by all the rules considered. The dashes show that the satisfaction of the axiom depends on the particular scoring function used.

similarities between this setting and that of *controlled school choice* (Abdulkadiroğlu & Sönmez, 2003). In the second subsection, we introduce axioms for the priority setting inspired by those considered in controlled school choice and use these to characterise the candidate-wise dual election rules.

4.1.1 Excellence-based elections

In excellence-based elections, voters can be viewed as experts and the goal is to elect the k best candidates, where it is assumed that there are no interdependencies between candidates. Thus, two similar candidates ought, generally, to either both be elected or both be left unelected. This is in contrast with, for example, the goal of proportional representation (see Chapter 5), where we may want to avoid electing two similar candidates in favour of representing more voters.

An example of a kind of excellence-based election is shortlisting for a prize. For instance, in the case of the Oscars, experts, i.e., members of the Academy of Motion Picture Arts and Sciences, vote for candidates, i.e., movies, in order to create a shortlist of finalists. The eventual winner—the recipient of the Oscar—is then chosen from this shortlist.[1] Another example would be admissions to a programme with limited places on the basis of votes of some admissions committee.

To see that there might be reason to give preferential treatment to certain candidates in excellence-based elections, consider again the two examples given. When shortlisting books for an international prize, preferential treatment might be warranted to guarantee that the resulting shortlist does not neglect typically underrepresented languages or countries. Nevertheless, we are still concerned with shortlisting the best books, given the relevant restrictions. Similarly, a programme to which students are to be admitted might suffer from, for example, a gender im-

[1] This last step, in which a *single* winner is chosen, is not an instance of multiwinner voting. In fact, often (though not in the case of the Oscars), the winner will be chosen by a jury through deliberation, rather than voting.

balance. Consequently, applicants from underrepresented genders might be given some kind of preferential treatment. Here too, the goal is still to elect the best candidates given the restrictions.

We will refer to rules that are suitable to excellence-based elections as *excellence-oriented* rules. In the context of ordinal ballots, *committee monotonicity* has been suggested as a requirement for excellence-oriented rules (Barberà & Coelho, 2008; Elkind et al., 2017).[2] For resolute rules, committee monotonicity requires that any candidate that wins for committee size k, also wins for committee size $k+1$.[3]

Since there are no interdependencies between candidates, excellence-oriented rules ought, for committee size k, to elect the k best candidates. Increasing the committee size to $k+1$ should then simply result in the next best, i.e., the $(k+1)$st best, candidate being added to the committee. Committee monotonocity thus seems like a reasonable requirement. However, it also seems too weak to capture the idea that there can be no interdependencies between candidates. Consider, for example, seq-Phragmén, which satisfies committee monotonicity. Clearly, and by design, in any round, which candidate is chosen depends on the candidates chosen in the previous rounds. Since the other (non-priority) sequential rules considered violate committee monotonicity and seq-Phragmén clearly encodes dependencies between candidates, these rules are ill-suited to the excellence-based setting.

It goes beyond the scope of this thesis to discuss more suitable or additional axioms that capture excellence-oriented rules (in the non-priority setting).[4] However, we argue that (sequential) candidate-wise scoring rules are particularly well-suited to this setting. Not only are such rules committee monotone, but—restricting our attention to sequential rules for simplicity—in any round, whether or not a candidate is elected depends only on the scores of the other remaining candidates. That is, selecting the next candidate in no way depends on which candidates have already chosen.[5] Consequently, for the excellence-based setting, we consider the priority rules based on candidate-wise non-priority rules, i.e., the candidate-wise dual election rules.

[2] Barberà and Coelho (2008) refer to committee monotonicity as *enlargement consistency*.

[3] For irresolute rules the definition is slightly more involved. It requires that every winning committee for k has a superset that is winning for $k+1$, and every winning committee for $k+1$ has a subset that is winning for k. For a formal definition see the work of Elkind et al. (2017).

[4] One option that has been discussed with regards to choice problems more generally, though not explicitly in the context of excellence, is *gross substitutes* (Doğan et al., 2021; Echenique & Yenmez, 2015; Kojima, 2012). Interpreting this axiom in the context of multiwinner voting naturally yields the requirement that if a candidate is selected, they must still be selected if only a subset of the candidates is considered. Note that seq-Phragmén would not satisfy this property, while candidate-wise scoring rules would.

[5] This property has also been formalised in the context of choice rules by Doğan et al. (2021), who call it *irrelevance of accepted alternatives*.

Controlled school choice and excellence-based priority elections

The matching problem of *controlled school choice* is similar to that of excellence-based elections with priority candidates. Controlled school choice concerns the matching of students to schools, based on students' and schools' preferences, where schools (are required to) give preferential treatment to certain kinds of candidates, e.g., those from the catchment area or those of underrepresented ethnic backgrounds (Abdulkadiroğlu & Sönmez, 2003).[6] Schools' preferences are strict linear orders over the students and vice versa.

Most of the mechanisms studied in this setting are based on *deferred acceptance* (DA), first introduced by Gale and Shapley (1962), and work in rounds. In the first round, students apply to their top-ranked schools. Schools then use some choice rule to determine which students are tentatively accepted. The students that are rejected apply to their next highest-ranked school in the next round. Schools then use their choice rule to select which of the currently matched and newly applying students to tentatively accept (or reject). This continues until all students are matched to a school. We are here concerned with the choice rules that schools use to determine which students from the applicant pool to (tentatively) accept.

When there are no requirements regarding preferential treatment, it is assumed that schools simply accept the highest-ranked candidates. However, when a school needs to guarantee that a certain number of seats are reserved for priority candidates, there are several choices for doing so.[7] Two options, which are analogous to the priority rules considered in this thesis, have been extensively studied in the literature: filling reserved seats first, and filling them last.[8] The parallel with excellence-based elections comes from the assumption that schools will want to fill any seat with the highest-ranked, eligible candidate. That is, the goal is to accept the best students, while meeting the quota for priority students.

To draw out the parallel a little further, we can view candidate-wise scoring rules as generating a linear order over candidates (that allows for ties). The corresponding dual election rules can then be viewed as choice rules that select candidates according to this non-strict ranking, reserving either the first q seats or the last q seats for priority candidates.

The central axiom considered in matching is *stability*, which, roughly, requires that there do not exist a student and school where the student prefers the school to their current match and the school prefers

[6] This same framework is studied in different contexts as well, e.g., by Sayedahmed (2021), who considers the problem of matching refugees to host countries, where certain groups of refugees, such as those in war zones, are given priority.

[7] In the literature it is common to refer to 'minority' and 'majority' students. We stick with the more neutral terminology of 'priority' and 'non-priority' used in the rest of this thesis.

[8] Other options, such as alternating reserved and unreserved seats, are also studied, for example by Doğan et al. (2021); however, as of yet there are no results that give a thorough understanding of the effects of such *reservation sequences*.

the student to one of their current matches. Weakenings that take the preferential treatment of priority students into consideration and encode different degrees of preferential treatment have been studied by, e.g., Abdulkadiroğlu and Sönmez (2003), Hafalir et al. (2013), and Sayedahmed (2021). Though these properties concern the stability of matching mechanisms as a whole, rather than the choice rules used by schools, elements are easily translated into our current setting. In the following subsection we introduce a number of axioms inspired by those from the controlled school choice setting and characterise the dual election rules using them.

4.1.2 Excellence-based axioms

The election $E = (C, C^+, \boldsymbol{A}, k, q)$ and candidate-wise scoring function sc are taken as given for the definitions and proofs below. As before, the non-priority rule that simply maximises the score is denoted $\mathcal{R}_{\text{sc}}$.

In excellence-based elections we want to guarantee that the best candidates are chosen. In the priority setting, however, we will generally not be able to simply maximise the score of the winning committees due to restrictions regarding priority treatment—e.g., that q priority candidates must be elected. Priority treatment can be used to justify electing a priority candidate in favour of a higher-scoring non-priority candidate. However, it ought not to be the case that a priority candidate is left unelected in favour of a lower-scoring candidate (of any status). This is captured by the axiom *priority merit*.

Axiom 8: Priority merit, PM

A committee $W \in C[k]$ satisfies PM if $\text{sc}(c) \geq \text{sc}(c^+)$ for all $c \in W$ and $c^+ \in C^+ \setminus W$. A rule $\mathcal{R}$ satisfies PM if it only returns committees that satisfy PM.

As explained, an analogous axiom for non-priority candidates would not make sense in light of priority treatment.[9] However, it does make sense to require that a non-priority candidate cannot be elected in favour of a higher-scoring non-priority candidate. We call this requirement *non-priority merit*.

Axiom 9: Non-priority merit, NPM

A committee $W \in C[k]$ satisfies NPM if $\text{sc}(c_1^-) \geq \text{sc}(c_2^-)$ for all $c_1^- \in W^-$ and $c_2^- \in C^- \setminus W$. A rule $\mathcal{R}$ satisfies NPM if it only returns committees that satisfy NPM.

[9] Note that such an axiom combined with PM would simply enforce maximisation of the score, i.e., it would require $\mathcal{R}_{\text{sc}}$ to be used.

PM and NPM aim to enforce the selection of higher-scoring candidates when this does not contradict the desired priority treatment. In that light, they can also be viewed as fairness requirements in the context of excellence-based elections. They guarantee that, unless priority treatment is involved, candidates will be selected (or not) based only on their score.

We consider only rules that respect the quota, i.e., that satisfy RoQ. RoQ encodes a kind of priority treatment, since it guarantees that q priority candidates will be chosen, even if this means electing lower-scoring candidates. However, as Sayedahmed (2021) points out in her distinction between *representation* and *effective preferential treatment*, often preferential treatment for priority candidates aims to provide opportunities to priority candidates that they would not have otherwise. That is, we may want to use reserved seats to accommodate those priority candidates that may otherwise not be elected.

Consider, for example, the situation in which candidates from under-represented genders are to be given some kind of preferential treatment in a selection procedure for an educational programme. Suppose there is a priority candidate who would be elected even if there were no priority treatment, i.e., if the highest-scoring candidates simpliciter are chosen. This candidate will help to fill the quota. However, in essence, it might be argued, they did not receive any priority treatment, since they would have been elected anyway. Instead, a priority candidate that would otherwise not have been elected could make use of the reserved seat.

To capture this idea, we, following Sayedahmed (2021), first introduce the concept of *protected priority candidates*. Given a (winning) committee W, a priority candidate $c^+ \in W^+$ is said to be *protected* if their score is weakly less than that of any elected non-priority candidate $c^- \in W^-$. The set of protected priority candidates in W is thus

$$\{c^+ \in W^+ : \mathrm{sc}(c^-) \geq \mathrm{sc}(c^+) \text{ for each } c^- \in W^-\}.$$

We can then require that, whenever possible, q such priority candidates are to be elected:

Axiom 10: Active quota, AQ

A committee W satisfies AQ if the existence of an unelected priority candidate $c \in C^+ \setminus W$, implies that W counts at least q protected priority candidates, i.e.,

$$|\{c^+ \in W^+ : \mathrm{sc}(c^-) \geq \mathrm{sc}(c^+) \text{ for each } c^- \in W^-\}| \geq q.$$

A rule $\mathcal{R}$ satisfies AQ if its winning committees all satisfy AQ.

At a first glance this axiom seems too strong, for it requires that whenever fewer than q protected priority candidates have been elected, all priority candidates must be elected. To see that this is not as problematic as may seem, suppose, in violation of AQ, that there is an unelected priority candidate c^+ and that fewer than q protected priority candidates have been elected. Either c^+ is weakly lower-scoring than all elected non-priority candidates or not. In the former case, c^+ should have been assigned a reserved seat, since these are reserved for protected priority candidates. In the latter case, c^+ is better than some elected non-priority candidate, and should thus have been elected in their stead.[10]

It is essential that we do not require protected priority candidates to have a strictly lower score than elected non-priority candidates. To see this, consider a situation in which all candidates have the same approval score. Then AQ would require that no priority candidate is left unelected (even if there are more than k), which does not make sense. Use of the weak rather than strict inequality conceptually means, roughly, that a protected candidate is one that would *or might* (due to unfavourable tiebreaking) not have been elected in absence of priority treatment.

We now have two axioms, RoQ and AQ, that capture two different kinds of priority treatment. It is easy to see AQ implies RoQ, since the former requires that either all priority candidates are elected, or q protected priority candidates have been elected. In either case RoQ is satisfied. We can thus think of AQ as capturing a stronger form of priority treatment than RoQ.

Besides guaranteeing priority treatment for priority candidates, we may want to guarantee that a rule does not go any further in giving preferential treatment to priority candidates than is required by the agreed upon priority treatment. Again, this is required if we are concerned with excellence, but can also be seen as a fairness requirement. Non-priority candidates receive the guarantee that, other than having a low score, they may only be left unelected if this is necessary to meet the requirements of the specific kind of priority treatment considered. Accordingly, we have the following two axioms corresponding to RoQ and AQ, in that order.

Axiom 11: Limit to quota, LtQ

A committee W satisfies LtQ if $|W^+| > q$ implies $\mathrm{sc}(c) \geq \mathrm{sc}(c^-)$ for any $c \in W$ and $c^- \in C^- \setminus W$. A rule $\mathcal{R}$ satisfies LtQ if its winning committees all satisfy LtQ.

RoQ expresses the requirement that q priority candidates be elected. LtQ essentially says that if this minimal requirement can be met without affording any preferential treatment to priority candidates, then no such treatment should be given. To see that this is indeed the case, consider a violation of LtQ. In that case more than q priority candidates will be elected, yet at least one of these, say c^+, has a strictly lower score than

[10] Note that this latter case is also a violation of PM.

some unelected non-priority candidate c^-. But then c^-, being a higher-scoring candidate, might replace c^+, without this violating respect of quota.

Similarly, when we want to ensure that at least q protected priority candidates are elected, we may want to give every non-priority candidate the guarantee that at most q priority candidates with a worse score will be elected.

Axiom 12: Limit to active quota, LtAQ

A committee W satisfies LtAQ if for every non-priority candidate $c^- \in C^- \setminus W$ we have

$$|\{c^+ \in W^+ : \mathrm{sc}(c^-) > \mathrm{sc}(c^+)\}| \leq q.$$

A rule $\mathcal{R}$ satisfies LtAQ if its winning committees all satisfy LtAQ.

Note that the strict inequality is used in the above definition because more than q priority candidates with an equal score may be elected if the tiebreaking happens to be favourable to priority candidates. Since RoQ is a weaker requirement than AQ, limiting preferential treatment to RoQ is a stronger requirement than limiting it to AQ; that is, LtQ implies LtAQ. Formally, if a committee satisfies LtQ, then either q or fewer priority candidates are elected, or all elected priority candidates are weakly higher-scoring than all unelected non-priority candidates. In either case, for any unelected non-priority candidate, there are at most q elected priority candidates that are strictly lower-scoring.

Recall the reserved-first dual election rule based on candidate-wise scoring function sc, denoted D^f_{sc}. Each of its winning committees consists of a size-q $\mathcal{R}_{\mathrm{sc}}$ winning committee from among the priority candidates and a corresponding size-$(k-q)$ $\mathcal{R}_{\mathrm{sc}}$ winning committee from among the remaining candidates. Also recall the analogous reserved-last rule, denoted D^ℓ_{sc}, which starts with the open election. We are now in a position to characterise these rules. We first prove a number of useful lemmas. The first captures the fact that D^f_{sc} and D^ℓ_{sc} will always elect a candidate over a worse candidate of the same status.

Lemma 4.1. (*i*) *Given a winning committee $W \in D^f_{\mathrm{sc}}(E)$, the set W^+ consists of $|W^+|$ of the highest-scoring priority candidates and W^- consists of $|W^-|$ of the highest-scoring non-priority candidates.* (*ii*) *The same holds for D^ℓ_{sc}.*

Proof. For D^f_{sc}, this follows from Lemma 2.3. Suppose there is a winning committee $W \in D^f_{\mathrm{sc}}(E)$ and candidates $c \in W$ and $c' \in C \setminus W$ with $\mathrm{sc}(c') > \mathrm{sc}(c)$ that are of the same status. Then $(W \setminus \{c\}) \cup \{c'\}$ would still respect the quota and have a strictly greater score than W, which contradicts the fact that W, by Lemma 2.3, maximises the score amongst committees that respect the quota.

We now consider D^{ℓ}_{sc}. Let $W \in D^{\ell}_{\text{sc}}(E)$ be a winning committee. It follows that there are $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k-q)$ and $Q \in \mathcal{R}_{\text{sc}}(C^+ \setminus O, \boldsymbol{A}, q)$ such that $O \cup Q = W$. We first consider the priority candidates and then consider the non-priority candidates.

Suppose, towards a contradiction, that there are priority candidates $c_1^+ \in W^+$ and $c_2^+ \in C^+ \setminus W$ such that $\text{sc}(c_2^+) > \text{sc}(c_1^+)$. If $c_1^+ \in O$, then $(O \setminus \{c_1^+\}) \cup \{c_2^+\}$ has a strictly higher score than O, which contradicts $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k-q)$. Similarly, if $c \in Q$, replacing c with c' yields a committee with a higher score, which contradicts $Q \in \mathcal{R}_{\text{sc}}(C \setminus O, \boldsymbol{A}, q)$. Thus, each elected priority candidate scores at least as high as any unelected priority candidate.

Now suppose there are non-priority candidates $c_1^- \in W^-$ and $c_2^- \in C^- \setminus W$ with $\text{sc}(c_2^-) > \text{sc}(c_1^-)$. As c_1^- is not eligible for the quota election, we must have $c_1^- \in O$. It follows that $(O \setminus \{c_1^-\}) \cup \{c_2^-\}$ has a strictly higher score than O, which contradicts $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k-q)$. □

It follows directly from Lemma 4.1 that both D^{f}_{sc} and D^{ℓ}_{sc} satisfy NPM.

Lemma 4.2. *Both D^{f}_{sc} and D^{ℓ}_{sc} satisfy NPM.*

We can now characterise the candidate-wise reserved-first dual rules.

Proposition 4.3. *A rule $\mathcal{R}$ satisfies RoQ, LtQ, PM and NPM iff it is the reserved-first dual election rule based on the candidate-wise scoring function* sc, *i.e., D^{f}_{sc}.*

Proof. We first show that D^{f}_{sc} satisfies the axioms. We have already seen that D^{f}_{sc} satisfies RoQ and NPM (Lemma 4.2). To see that PM holds as well, let $W \in D^{f}_{\text{sc}}(E)$ be a winning committee. If there are candidates $c \in W$ and $c^+ \in C^+ \setminus W$ with $\text{sc}(c^+) > \text{sc}(c)$, then we could increase the score of W by replacing c with c^+, without decreasing the number of priority candidates. But, by Lemma 2.3, this is not possible.

Similarly, to see that LtQ is respected, suppose $|W^+| > q$ and suppose there are candidates $c \in W$ and $c^- \in C^- \setminus W$ with $\text{sc}(c^-) > \text{sc}(c)$. We could then increase the score of W by replacing c with c^-, while still having at least q priority candidates (exactly one less than before). By Lemma 2.3 this is not possible.

We now show the other direction. That is, we show that any committee satisfying the axioms is a D^{f}_{sc} winning committee. To that end, let W be such a committee. We show that it maximises the score among the committees that respect the quota, from which the desired result follows immediately through Lemma 2.3.

Let $c \in W$ and $c' \in C \setminus W$ be arbitrary such that $\text{sc}(c') > \text{sc}(c)$. In light of PM and NPM we can conclude that $c \in C^+$ and $c' \in C^-$. It follows, by LtQ, that $|W^+| = q$. But then we cannot improve on W's score without violating RoQ. □

Corollary 4.3.1. *D^{f}_{sc} satisfies LtAQ.*

Proof. This follows directly since LtQ implies LtAQ. □

The above characterisation is similar to the result presented by Hafalir et al. (2013) that the deferred acceptance algorithm for controlled school choice in which schools fill reserved seats first satisfies what they call *stability for minority reserves* or *affirmative action with minority reserves*. This axiom is essentially a weakening of stability that incorporates RoQ and LtQ.[11]

We now turn our attention to D^{ℓ}_{sc}.

Lemma 4.4. *D^{ℓ}_{sc} satisfies AQ, LtAQ, PM and NPM.*

Proof. Let $W \in D^{\ell}_{\text{sc}}(E)$ be a winning committee. It follows that there exist $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k - q)$ and $Q \in \mathcal{R}_{\text{sc}}(C^+ \setminus O, \boldsymbol{A}, q)$ with $O \cup Q = W$.

We know that W satisfies RoQ and NPM (Lemma 4.2). To see that it also satisfies PM, suppose, towards a contradiction, that there exist candidates $c \in W$ and $c^+ \in C^+ \setminus W$ with $\text{sc}(c^+) > \text{sc}(c)$. We then have that $c \in O$ contradicts $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k - q)$, while $c \in Q$ contradicts $Q \in \mathcal{R}_{\text{sc}}(C^+ \setminus O, \boldsymbol{A}, q)$.

To see that W satisfies AQ, it is sufficient to note that no candidate in Q has a strictly higher score than any candidate in O. Thus Q consists of q protected priority candidates.

Finally, for LtAQ, suppose there exists an unelected non-priority candidate $c^- \in C^- \setminus W$ such that there are more than q priority candidates with a strictly worse score. Given $|Q| = q$, there must then be a candidate $c^+ \in O^+$ such that $\text{sc}(c^-) > \text{sc}(c^+)$. But this contradicts $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k - q)$. □

Proposition 4.5. *A rule $\mathcal{R}$ satisfies* AQ, LtAQ, PM *and* NPM *iff it is the reserved-first dual election rule based on* sc, *i.e.,* D^{ℓ}_{sc}.

Proof. Given Lemma 4.4, we need only show that any committee that satisfies the axioms is a D^{ℓ}_{sc} winning committee. To that end, let $W \in C[k]$ be a committee satisfying AQ, LtAQ, PM and NPM. Let $Q \subseteq W^+$ consist of q of the lowest-scoring priority candidates from W and let $O = W \setminus Q$. It follows from AQ that for every priority candidate $c^+ \in Q$ we have $\text{sc}(c^-) \geq \text{sc}(c^+)$ for any $c^- \in W^-$.[12] It follows that any candidate from Q has a weakly worse score than any candidate from O. Now suppose, towards a contradiction, that there exist candidates $c \in O$ and $c' \in C \setminus O$ such that $\text{sc}(c') > \text{sc}(c)$. We know that $c' \notin Q$, hence $c' \in C \setminus W$. It follows, by PM and NPM that $c \in C^+$ and $c' \in C^-$. But then, since every candidate in Q has a score of at most $\text{sc}(c)$, there are at least $q+1$ priority candidates with a strictly worse score than c' in W. This contradicts LtAQ. We can thus conclude that $O \in \mathcal{R}_{\text{sc}}(C, \boldsymbol{A}, k - q)$.

[11] Sayedahmed (2021) calls this axiom *representation-stability*.

[12] This relies on the assumption that $|C^+| \geq k$. Also, note that this is vacuously true if all elected candidates are priority candidates.

By PM we know that there do not exist priority candidates $c_1^+ \in Q$ and $c_2^+ \in C^+ \setminus W$ with $\mathrm{sc}(c_2^+) > \mathrm{sc}(c^+1)$. But then $Q \in \mathcal{R}_{\mathrm{sc}}(C^+ \setminus O, \boldsymbol{A}, q)$, from which it follows that $W = O \cup Q \in D^{\ell}_{\mathrm{sc}}(E)$. □

The characterisation of D^{ℓ}_{sc} is analogous to the result of Sayedahmed (2021, Theorem 5) that the deferred acceptance algorithm for controlled school choice in which schools fill reserved seats last satisfies *protection-stability*, which is essentially a weakening of stability that incorporates AQ and LtAQ.

We know that $D^{f}_{\mathrm{(S)AV}}$ must violate AQ, since if it did not it would satisfy AQ, LtAQ, PM and NPM, and thus, by Proposition 4.5, be equivalent to $D^{\ell}_{\mathrm{(S)AV}}$, which we know is not the case (see Example 2.3). Similar reasoning in light of Proposition 4.3 shows that $D^{\ell}_{\mathrm{(S)AV}}$ violates LtQ.[13]

It might be tempting to conclude that any reserved-first dual election rule violates AQ . This is, however, not the case. The reasoning in the previous paragraph shows that if D^{f}_{sc} satisfies AQ, then it must be equivalent to D^{ℓ}_{sc}. The reserved-first and reserved-last dual election rules based on any trivial scoring function that assigns to each candidate the same score (regardless of the profile) will be equivalent and satisfy AQ, LtQ, LtAQ, PM and NPM. Analogous reasoning shows that we cannot conclude that D^{ℓ}_{sc} violates LtQ for every candidate-wise score sc.

A significant difference between the matching results and the results presented here is that in the former (exogenously given) schools' preferences regarding students are strict rankings and, accordingly, the rules considered for choosing which students from the pool to admit are resolute. In contrast, the candidate-wise scores considered here can be thought of as generating a weak ranking over candidates, and the corresponding dual election rules considered are irresolute. Though we are concerned with a setting in which there are voters, the above results can also be generalised to settings where a weak ranking over candidates is exogenously given (as opposed to being based on votes). The following section, however, concerns *efficiency*, a property that is inherently related to voters' satisfaction.

4.2 Efficiency

Pareto efficiency is a general and much-studied concept that concerns the comparison of outcomes given individuals' preferences. Informally, outcome α Pareto dominates outcome β if every individual weakly prefers α over β and for at least one individual this preference is strict. Pareto efficiency then requires that outcomes are not Pareto dominated. That is, Pareto efficiency requires that there is no way of making some individual better off without making a different individual worse off (Arrow et al., 2002).

[13] In fact, Example 2.3 shows directly that $D^{f}_{\mathrm{(S)AV}}$ violates AQ and $D^{\ell}_{\mathrm{(S)AV}}$ violates LtQ.

In order to apply this concept to multiwinner voting, we must define what it means for a voter to prefer one committee over another. That is, we need to extend voters' approval preferences on candidates to preferences on sets of candidates. The most common way to do this assumes that voters prefer committees in which they approve of more candidates.[14]

Axiom 13: Pareto efficiency

A committee W_1 *Pareto dominates* a committee W_2 if

(i) every voter approves of at least as many candidates in W_1 as in W_2; i.e., $|A_i \cap W_1| \geq |A_i \cap W_2|$ holds for every $i \in N$, and

(ii) at least one voter approves of strictly more candidates in W_1 than W_2; i.e., there is some $j \in N$ such that $|A_j \cap W_1| > |A_j \cap W_2|$.

We also say that W_1 is a *Pareto improvement* on W_2. A committee is *Pareto optimal* if it is not Pareto dominated by any committee of the same size, i.e., if it does not allow for a Pareto improvement.

A rule $\mathcal{R}$ satisfies Pareto efficiency if all its winning committees are Pareto optimal.

This definition of Pareto efficiency is sometimes referred to as *strong Pareto efficiency*. This is in contrast with *weak Pareto efficiency*, which requires that when some winning committee is Pareto dominated, then all committees that dominate it are also winning committees (Lackner & Skowron, 2023). None of the rules considered that violate (priority adaptations of) Pareto efficiency as we have defined it satisfy the (priority adaptations of) weak Pareto efficiency, hence we do not consider it.[15]

Of the rules that we have seen, AV, SAV and PAV satisfy Pareto efficiency (while CC satisfies weak Pareto efficiency). For these rules it is easy to see that a Pareto improvement results in an increase in the corresponding scores.[16] We use this fact repeatedly in what follows.

Pareto efficiency will often be too strong a requirement in the priority setting. In fact, as the following example shows, Pareto efficiency and

[14] See, for example, the work of Aziz and Monnot (2020) and Lackner and Skowron (2020, 2023). To the best of our knowledge, the only other set extension commonly used with dichotomous preferences takes a voter $i \in N$ to (weakly) prefer committee W_1 to committee W_2 if $A_i \cap W_2 \subseteq A_i \cap W_1$. This is used primarily in the context of *strategyproofness* (Peters, 2018). It should be noted that Pareto optimality as we define it implies this 'inclusion-type' Pareto optimality.

[15] An alternative definition of weak Pareto efficiency used by Lackner and Skowron (2020) requires that it is never the case that all winning committees are dominated by *the same* committee.

[16] Voters who approve of the same number of candidates contribute the same to the score, while voters who approve of strictly more candidates contribute strictly more to the score.

RoQ are not compatible.

Example 4.1. Consider the profile

$$2 \times \{\overline{a}, b, c\} \qquad\qquad 1 \times \{b, c\},$$

and suppose $k = 2$ and $q = 1$. The sole Pareto optimal committee, $\{b, c\}$, violates RoQ.

In order to account for priority treatment we consider a number of weakenings of Pareto efficiency. If we wish the quota to be respected, a natural way to weaken Pareto efficiency is to require that no Pareto improvements can be made without violating the quota.

Axiom 14: Quota priority efficiency, qPE

A committee W_1 *quota priority dominates*, or q-p-dominates, W_2 if

(i) W_1 Pareto dominates W_2, and

(ii) W_1 contains at least q priority candidates, i.e., $|W_1^+| \geq q$.

Quota priority optimality and *quota priority efficiency* are then defined analogously to Pareto optimality and efficiency (Axiom 13).

Note that qPE as stated here only makes sense if we assume that there are at least q priority candidates available. Otherwise, it is vacuously true for any committee $W \in C[k]$ that no Pareto improvement can be made that respects the quota (for there are no committees that respect the quota). Since we assume $|C^+| \geq q$, we do not adapt the definition of qPE to deal with cases in which $|C^+| < q$.

In a way, qPE captures the idea that priority treatment ought not to go further than guaranteeing respect of quota. That is, it requires of a rule that respects the quota that it come as close as possible to Pareto efficiency given this restriction. Accordingly, we find that qPE is satisfied by the reserved-first dual election rules based on AV and SAV.

Proposition 4.6. *$D^f_{\text{(S)AV}}$ satisfies qPE.*

Proof. Lemma 2.3 tells us that D^f_{sc} maximises the score of the winning committees among those that respect the quota. Since a Pareto improvement implies an increase in the (S)AV-score, we conclude that we cannot Pareto improve on any of the winning committees without violating RoQ. □

This result does not generalise to all candidate-wise scoring rules. Consider, for example, a kind of 'dictator AV': the dictator is one of the voters and the score of a committee is simply the size of its intersection with the dictator's ballot. This rule is candidate-wise, but scores need not increase with a Pareto improvement, since the dictator, though no worse off, need not be better off.

The reserved-last dual election rules based on AV and SAV do not satisfy qPE, as the following example shows.

Example 4.2. Consider the election with the set of candidates $C = \{\overline{a}, \overline{b}, c\}$, committee size $k = 2$, quota $q = 1$, and the following profile:

$$A_1 = \{\overline{a}, c\} \qquad\qquad A_2 = \{\overline{a}\}.$$

The unique winning committee according to $D^{\ell}_{\text{(S)AV}}$ is $\{\overline{a}, \overline{b}\}$, which is q-p-dominated by $\{\overline{a}, c\}$.

The intuition here will familiar by now: a strong priority candidate takes an unreserved seat, allowing a weaker priority candidate to make use of a reserved seat. But as a result, worse candidates have been elected than required in order to respect the quota.

When we set $q = 0$, qPE reduces to Pareto efficiency. Since the priority sequential rules reduce to the underlying non-priority sequential rules when $q = 0$, and none of these satisfy Pareto efficiency, none of the priority rules will satisfy qPE.

As we have seen, priority merit and non-priority merit combine to guarantee that the only way we can increase the score of a winning committee is by replacing priority candidates with non-priority candidates. This suggests a weaker form of priority efficiency that prohibits Pareto domination by committees with at least as many priority candidates.

Axiom 15: Priority efficiency, PE

A committee W_1 *priority dominates*, or p-dominates, W_2 if

(*i*) W_1 Pareto dominates W_2, and

(*ii*) W_1 contains at least as many priority candidates as W_2, i.e., $|W_1^+| \geq |W_2^+|$.

Priority optimality and *priority efficiency* are then defined as usual.

We first consider the relation between qPE and PE. Assuming that a committee W has at least q priority candidates, i.e., $|W^+| \geq q$, qPE implies PE, while the reverse does not hold. This is easy to see: if no Pareto improvement can be made with at least q priority candidates, then *a fortiori*, since $|W^+| \geq q$, no Pareto improvement can be made with more than $|W^+|$ priority candidates. Thus, since $D^{f}_{\text{(S)AV}}$ satisfies qPE and RoQ, it must also satisfy PE.

PE is more general than the above motivation suggests. The axiom captures a weakening of Pareto efficiency in light of strong priority treatment. That is, it can be understood as implicitly assuming that the number of elected priority candidates cannot be lower in light of priority treatment.

Proposition 4.7. $D^{\ell}_{\text{(S)AV}}$ *satisfies PE.*

Proof. Since $D^{\ell}_{\text{(S)AV}}$ satisfies PM and NPM, the only way we can increase the score of a winning committee W, is by replacing a priority candidate with a non-priority candidate. But then, as a Pareto improvement implies an increase in (S)AV-score, such an improvement requires a decrease in the number of priority candidates elected. $\square$

When we set $q = k$, the priority sequential rules effectively reduce to the non-priority rules where only the priority candidates are considered. Assuming that respect of quota is met, any Pareto improvement will constitute a priority improvement. Consequently, as none of the non-priority sequential rules are efficient, none of the sequential priority rules will satisfy PE.

We have seen that the dual election rules based on candidate-wise scores are suited to excellence-based elections, and that the reserved-last rules encode a greater degree of priority treatment at the cost of satisfying a stronger efficiency requirement. We also saw that the sequential rules are ill-suited to this setting and perform badly with regards to efficiency. The next chapter concerns proportionality, a goal to which the sequential rules are well-suited.

Chapter 5

Proportionality

Proportionality concerns the fair representation of groups of voters. The underlying idea is that groups of voters that are cohesive enough in their preferences deserve a number of representatives that is proportional to the size of the group. This chapter concerns proportionality in the context of priority treatment. The first section introduces the notion of proportionality in more detail along with a number of axioms for the non-priority setting. In the second section we briefly motivate the consideration of proportionality in the priority setting. Thereafter, we introduce adaptations of the non-priority axioms, showing which priority rules (do not) satisfy them. The results are summarised in Table 5.1 below. The abbreviations JR, PJR and EJR stand for *justified representation*, *proportional justified representation* and *extended justified reprsentation*, respectively. The prefix 'rp' abbreviates 'restricted priority' and the prefix 'p' abbreviates 'priority'. Thus, e.g., rp-EJR denotes *restricted priority extended justified representation*.

	axiom					
rule	rp-JR	rp-PJR	rp-EJR	p-JR	p-PJR	p-EJR
rf-seq-CC	✓[5.1.1]	✗	✗	✓[5.1]	✗	✗
rl-seq-CC	✓[5.1.1]	✗	✗	✓[5.1]	✗	✗
late-seq-CC	✓[5.1.1]	✗	✗	✓[5.1]	✗	✗
rl-seq-Phragmén	✓	✓	✗	✓[5.4]	✗[ex.5.1]	✗
late-seq-Phragmén	✓[5.3.1]	✓[5.3.1]	✗	✓[5.3.1]	✓[5.3]	✗

Table 5.1: An overview of the relevant rules and the proportionality axioms that they satisfy.

5.1 Proportionality in the non-priority model

Proportionality, as a general concept, requires that an elected committee reflects the distribution of voters' preferences. That is, proportional rules guarantee that (even small) groups of voters receive a level of representation that is proportional to their size. The topic of proportionality has been extensively studied in the context of *apportionment*, which concerns the distribution of seats, e.g., in a parliament, over political parties based on voters' preferences. Many countries have chosen to determine their parliaments in a proportional manner, decisions of such governing bodies being deemed legitimate only when the opinions of all citizens are taken into account, weighted according to their prevalence among the voting population.[1] Apportionment can be seen as a subdomain of multiwinner voting (Brill et al., 2018),[2] but, since proportional committees will correspond to fair committees in numerous other scenarios, recent years have seen the emergence of a great deal of research on proportionality in multiwinner voting simpliciter.

Peters and Skowron (2020) make a distinction between two fundamentally distinct forms of proportional representation, one concerning the *fair distribution of welfare*, i.e., the satisfaction of voters, and the other concerning the *fair distribution of voting power*. Essentially, the former is concerned with groups receiving a number of representatives proportional to their size, while the latter requires that groups can determine a fraction of the committee that is proportional to their size.[3] We are here concerned with proportionality with regards to the fair distribution of welfare, hence this is what we mean by 'proportionality' from here on.

The central concept in the study of proportionality in multiwinner voting is that of ℓ-cohesiveness.

Definition 5.1 (ℓ-cohesiveness)**.** For $\ell \geq 1$ and a profile $\boldsymbol{A}$, a group $V \subseteq N$ is ℓ-*cohesive* if

(*i*) $|V| \geq \ell \cdot {}^{n}/_{k}$, and

(*ii*) $\left|\bigcap_{i \in V} A_i\right| \geq \ell$

Intuitively, an ℓ-cohesive group, since it contains an ${}^{\ell}/_{k}$th fraction of the voters and agrees on ℓ candidates, ought to be given ℓ representatives in the winning committee. It seems natural to require that for any ℓ-cohesive group V, every voter $i \in V$ must have ℓ representatives in the winning committee. However, as shown by Aziz et al. (2018, Example 1), no rule meets this requirement.

[1] Many countries, e.g., the United States, do not have a system of proportional representation. It is important to note that while proportional representation is widely favoured as a form of democracy, its suitability remains a subject of debate. Critics contend, for example, that proportional representation stimulates the formation of splinter factions from larger parties, potentially compromising the stability of governments. It is also argued that proportional rules are difficult to understand. Any good introduction to electoral systems, such as Norris (1997), will cover these topics. For overviews including more formal material see the works by Balinksi and Young (2001) and Pukelsheim (2017).

[2] The underlying idea is that apportionment can be represented as multiwinner voting restricted to *party-list profiles*, which are profiles in which any two voters' ballots are either identical or disjoint.

[3] To see that these two concepts differ and can conflict, see the first example presented by Peters and Skowron (2020).

Different axioms result from different representation requirements for ℓ-cohesive groups, i.e., from different interpretations of what a 'representative' of a cohesive group is. The weakest axiom we consider, due to Aziz et al. (2017), called *justified representation*, requires that for every 1-cohesive group there is a voter from that group that has a representative.

Axiom 16: Justified representation, JR

A rule $\mathcal{R}$ satisfies *justified representation* if for any winning committee $W \in \mathcal{R}(E)$ and 1-cohesive group of voters V, there exists a voter $i \in V$ who is represented by at least one candidate in W, i.e., $A_i \cap W \neq \emptyset$.

This is a very weak requirement and, arguably, cannot be seen as a proportionality requirement at all, since even large cohesive groups (e.g., the set of all voters in a unanimous profile) only require a single representative in the elected committees (Lackner & Skowron, 2023). We nevertheless consider it because it is clearly a minimal requirement for a proportional rule and also naturally relates to the following stronger axioms.

Sánchez-Fernández et al. (2017) introduce a stronger axiom which requires that for any ℓ-cohesive group, there are at least ℓ elected candidates each of which represents at least one member of V.

Axiom 17: Proportional justified representation, PJR

A rule $\mathcal{R}$ satisfies *proportional justified representation* if for any winning committee $W \in \mathcal{R}(E)$ and ℓ-cohesive group of voters V, it is the case that $\left|\left(\bigcup_{i\in V} A_i\right) \cap W\right| \geq \ell$.

Finally, *extended justified representation*, introduced by Aziz et al. (2017), requires that for every ℓ-cohesive group, some member of the group has at least ℓ representatives.

Axiom 18: Extended justified representation, EJR

A rule $\mathcal{R}$ satisfies *extended justified representation* if for any winning committee $W \in \mathcal{R}(E)$ and ℓ-cohesive group of voters V, there is a voter $i \in N$ with ℓ or more representatives in W, i.e., $|A_i \cap W| \geq \ell$.

First, note that if we restrict PJR or EJR to 1-cohesive groups, we get JR. Furthermore, EJR implies PJR. We thus have a hierarchy of proportionality axioms, where JR is the weakest and EJR the strongest. Of the non-priority rules we have seen, only PAV and MES satisfy EJR

(Aziz et al., 2017; Peters & Skowron, 2020). Further, seq-Phragmén satisfies PJR (Brill et al., 2023) and both CC and seq-CC satisfy (only) JR (Aziz et al., 2017).

Having considered proportionality in the non-priority model, we are now ready to consider the priority model.

5.2 Proportionality in the priority model

Before presenting proportionality axioms for the priority setting and corresponding results, we briefly consider the relevance of proportionality in the priority setting.

We have already seen that proportionality is important in the context of apportionment, which can be seen as a restricted form of multiwinner voting. In practice, there are many apportionment settings in which priority treatment is to be afforded to certain candidates. For example, there are countries that reserve parliamentary seats for women.[4] Outside the apportionment setting, priority treatment is also relevant. Corporate boards, community organisations and academic institutions, to name a few, may desire preferential treatment for certain kinds of candidates, to promote, for example, the representation of junior employees, marginalised members of the community, and candidates from particular disciplines, respectively.

Given that proportionality is relevant in the context of the priority model, we now consider how to adapt the previously introduced axioms. We consider only priority rules that respect the quota. In essence, these rules guarantee that q priority candidates will be elected no matter the support that priority candidates enjoy among the voters. Though all sensible rules will take voters' preferences into account when filling the quota, in a worst-case scenario there are only $k-q$ seats that can be used to meet representation requirements for cohesive groups. Accordingly, we might adapt the notion of ℓ-cohesiveness so that it is relative not to the committee-size k, but to $k-q$.

Definition 5.2 (Restricted priority ℓ-cohesiveness, rp-ℓ-cohesiveness)**.** For $\ell \geq 1$ and a profile $\boldsymbol{A}$, a group $V \subseteq N$ is *rp-ℓ-cohesive* if

(i) $|V| \geq \ell \cdot {}^{n}/_{k-q}$, and

(ii) $\left|\bigcap_{i\in V} A_i\right| \geq \ell$

Replacing the notion of ℓ-cohesiveness with that of rp-ℓ-cohesiveness in the definitions of JR, PJR and EJR yields axioms for the priority

[4] A particularly striking example can be observed in the case of Rwanda, where 24 of the 80 seats in the lower house are reserved for women. At present, women hold 49 of the seats. For more information, see the gender quotas database of the International Institute for Democracy and Electoral Assistance, accessible at https://www.idea.int/data-tools/data/gender-quotas.

setting. We have *restricted priority justified representation* (rp-JR), *restricted priority proportional justified representation* (rp-PJR), and *restricted priority extended justified representation* (rp-EJR), respectively, which respect the hierarchy of the original axioms. When $q = 0$, an rp-ℓ-cohesive group is simply ℓ-cohesive. Coupled with the fact that the priority sequential rules reduce to their underlying non-priority rules when $q = 0$, we can conclude that if a non-priority rule does not satisfy JR, PJR, or EJR, none of its priority adaptations will satisfy the corresponding restricted priority adaptation.

The notion of rp-ℓ-cohesiveness is based on a worst-case scenario; that is, it assumes that only $k-q$ seats can be used to meet the representation requirements for cohesive groups. Consequently, any reserved-first dual election rule based on a rule satisfying JR, PJR or EJR, respectively, will satisfy the corresponding restricted priority adaptation. In fact, this will be the case regardless of how the quota is filled. The axioms based on rp-ℓ-cohesiveness thus seem too weak. After all, the quota may be filled in more or in less proportional ways. Nevertheless, the worst-case scenario is one which has to be taken into account.

Intuitively, a cohesive group can be afforded more representatives if the quota seats can be used to do so, i.e., if they approve of sufficiently many priority candidates. The more priority candidates a cohesive group approves of, the more seats can be used to offer them representatives. This intuition is captured in the following definition of z-ℓ-cohesiveness.

Definition 5.3 (z-ℓ-cohesiveness)**.** For $q \geq z \geq 0$ and $\ell \geq 1$, a group $V \subseteq N$ is *z-ℓ-cohesive* if

(i) $|V| \geq \ell \cdot n/_{k-q+z}$, or $\left|\bigcap_{i \in V} A_i^+\right| \geq \ell$ and $|V| \geq \ell \cdot n/k$, and

(ii) $\left|\bigcap_{i \in V} A_i\right| \geq \ell$, and

(iii) $\left|\bigcap_{i \in V} A_i^+\right| \geq z$.

When a group of voters does not unanimously approve of any priority candidates, i.e., $z = 0$, this reduces to rp-ℓ-cohesiveness, thus accounting for the worst-case scenario in which a cohesive group must have size $n/_{k-q}$ to claim a representative. By taking into consideration the number of priority candidates that a group unanimously approves of, we intuitively achieve two things. Firstly, cohesive groups that jointly approve of priority candidates can 'demand' more than $k-q$ seats, and secondly, they will require fewer voters per demanded seat. The fact that when $\left|\bigcap_{i \in V} A_i^+\right| \geq \ell$ we only require $|V| \geq \ell \cdot n/k$, reflects the intuition that when a cohesive group's proportionality requirements can be met with only priority candidates, all seats can be used to assign representatives and consequently there should be no further restrictions on the size of the group beyond those required in the non-priority setting. Note that similar reasoning applies when $z = q$. In that case we also only require $|V| \geq \ell \cdot n/k$ (even if $\ell > z$).

The new definition of cohesiveness induces three proportionality axioms that are analogous to JR, PJR and EJR, which we call *priority justified representation* (p-JR), *priority proportional justified representation* (p-PJR) and *priority extended justified representation* (p-EJR), respectively. Like for the restricted priority axioms, these are also analogously ordered in terms of strength: p-EJR implies p-PJR, which implies p-JR. Furthermore, since any rp-ℓ-cohesive group is a z-ℓ-cohesive group, the priority axioms imply their restricted priority counterparts. (Remember that a stronger requirement for cohesiveness results in weaker axioms.) When $q = 0$, and consequently $z = 0$, z-ℓ-cohesiveness reduces to ℓ-cohesiveness. Consequently, and analogously to the case for the restricted priority axioms, the priority adaptation of a rule that does not satisfy JR, PJR, or EJR, respectively, will violate p-JR, p-PJR, or p-EJR, respectively.

It is known that seq-CC satisfies JR (Aziz et al., 2017), and it turns out that all three of the priority adaptations of seq-CC satisfy p-JR.[5]

Proposition 5.1. *rf-seq-CC, rl-seq-CC and late-seq-CC all satisfy p-JR.*

Proof. We provide the proof for rl-seq-CC. The proofs for rf-seq-CC and late-seq-CC are entirely analogous.

Suppose, towards a contradiction, that rl-seq-CC does not satisfy p-JR. It follows that there are an election $E = (C, C^+, \boldsymbol{A}, k, q)$ and z-1-cohesive group $V \subseteq N$ such that no voter $i \in V$ has a representative in the winning committee $W = \text{rl-seq-CC}(E)$, i.e., $\left(\bigcup_{i \in V} A_i\right) \cap W = \emptyset$.

Now suppose $z = 0$, from which it follows that $|V| \geq n/k-q$, and let $c \in \bigcap_{i \in V} A_i$ be a jointly approved candidate. We know that c was eligible in each of the first $k - q$ rounds. Hence, since c was not elected, we know that in each of these rounds a candidate must have been selected representing at least $\lceil n/k-q \rceil$ voters who were not yet represented in that round. But then in the first $k - q$ rounds, at least n voters have a representative elected, which contradicts the fact that no voter in V has a representative.

If $z > 0$, i.e., $z \geq 1$, then $|V| \geq n/k$. Note that we cannot assume $|V| \geq n/k-q+z$ in this case due to condition (i) in the definition of z-ℓ-cohesiveness. Furthermore, we have $\left(\bigcap_{i \in V} A_i\right)^+ \neq \emptyset$, hence, let $c^+ \in \left(\bigcap_{i \in V} A_i\right)^+$. Then, since c^+ was eligible in each of the k rounds, at least $\lceil n/k \rceil$ previously unrepresented voters must have received a representative in each round. This contradicts the fact that the voters in V do not have a representative, concluding the proof. □

Corollary 5.1.1. *rf-seq-CC, rl-seq-CC and late-seq-CC all satisfy rp-JR.*

An alternative proof, which is more general, would make use of the fact that seq-CC is committee monotone and satisfies JR. Such a proof would then also work, e.g., for seq-Phragmén.

[5] Note that Aziz et al. (2017) refer to seq-CC as *Greedy Approval Voting.*

As we have seen above, p-JR is not able to distinguish between the different priority adaptations of seq-CC. This also holds, as we shall see, for the reserved-first and reserved-latest versions of seq-Phragmén. However, the two priority adaptations of seq-Phragmén come apart when we consider p-PJR: late-seq-Phragmén satisfies it, while rl-seq-Phragmén does not. For the proof of the former, the following lemma will be helpful.

Lemma 5.2. *Given some election $E = (C, \boldsymbol{A}, k)$, we say that a group of voters V is ℓ-cohesive for $k' \leq k$ (rounds) if V is ℓ-cohesive for $(C, \boldsymbol{A}, k')$. Now consider the execution of seq-Phragmén on E, where the provisional committee at the end of round r is denoted W_r. If a group V is ℓ-cohesive for $k' \leq k$ rounds, then $\left|\left(\bigcup_{i \in V} A_i\right) \cap W_{k'}\right| \geq \ell$, i.e., the group receives ℓ representatives in the first k' rounds.*

Proof. It is clearly the case that for any election $(C, \boldsymbol{A}, k)$, the k rounds of seq-Phragmén are going to be identical to the first k rounds for the election $(C, \boldsymbol{A}, k+1)$.[6] But then we have that for any election $(C, \boldsymbol{A}, k)$ and $k' \leq k$, the first k' rounds of seq-Phragmén on $(C, \boldsymbol{A}, k)$ will be identical to the k' rounds on $(C, \boldsymbol{A}, k')$. Now let V be a group of voters that is ℓ-cohesive for $k' \leq k$ and thus ℓ-cohesive for $(C, \boldsymbol{A}, k')$. We then have, since seq-Phragmén satisfies PJR, that ℓ candidates from $\bigcup_{i \in V} A_i$ must be elected in the k' rounds of $(C, \boldsymbol{A}, k')$ and thus also in the first k' rounds of $(C, \boldsymbol{A}, k)$. □

Proposition 5.3. *late-seq-Phragmén satisfies p-PJR.*

Proof. Towards a contradiction, suppose the contrary. It follows that there are some election $E = (C, C^+, \boldsymbol{A}, k, q)$, winning committee $W =$ late-seq-Phragmén(E) and z-ℓ-cohesive group V such that W contains fewer than ℓ candidates that represent some member of V. For convenience, let $C_V = \bigcap_{i \in V} A_i$ denote the set of candidates unanimously approved by voters in V and let $C_V^{\text{rep}} = \left(\bigcup_{i \in V} A_i\right) \cap W$ denote the set of elected candidates approved by some member of V. We refer to any candidate $c \in C_V^{\text{rep}}$ as a *representative of* V. We thus have $|C_V^{\text{rep}}| \leq \ell - 1 < \ell$.

We first show that all priority candidates unanimously approved by voters in V must be elected, i.e., $C_V^+ \subseteq W$. Thereafter we show that, consequently, V must have at least ℓ representatives, which contradicts our assumption.

Suppose there is some unelected priority candidate c unanimously approved by the voters from V, i.e., $c \in C_V^+ \setminus W$. Since c is a priority candidate, they were eligible in the last round. Analogously to the proof that seq-Phragmén satisfies PJR (Brill et al., 2023), we can then show that the total load of the voters is less than k, which contradicts the fact that late-seq-Phragmén yields a valid load distribution. We provide the

[6] It follows, as is well-known, that seq-Phragmén is committee-monotone. However, we here need this stronger claim.

proof here for completeness.[7]

If $\left|C_V^+\right| < \ell$ then $|V| \geq \ell \cdot {}^{n}/_{k-q+z}$ and otherwise $|V| \geq \ell \cdot {}^{n}/_{k}$. Thus, in either case we have $|V| \geq \ell \cdot {}^{n}/_{k}$ (since $q \geq z$) and consequently $|V| > \ell \cdot {}^{n}/_{k+1}$. Now consider the last round k. Adding candidate c to the committee in this round would have resulted in the following maximum voter load:

$$\begin{aligned}
s_c^k &= \frac{1 + \sum_{i \in N_c} \bar{x}_i^{k-1}}{|N_c|} && \text{by Lemma 2.2} \\
&\leq \frac{1 + \sum_{i \in V} \bar{x}_i^{k-1}}{|V|} && \text{since } V \subseteq N_c \\
&\leq \frac{1 + (\ell - 1)}{|V|} && \text{since } |C_V^{\text{rep}}| \leq \ell - 1 \\
&< \frac{k+1}{n} && \text{since } |V| > \ell \cdot \frac{n}{k+1}
\end{aligned}$$

Now let c_k be the candidate elected in round k. It follows that the maximum load that resulted from electing c_k in the last round cannot be greater than what it would have been for c, i.e., $s_{c_k}^k \leq s_c^k$. But then, since $s_c^k < {}^{k+1}/_{n}$, we have that the maximum load after round k is $s_{c_k}^k < {}^{k+1}/_{n}$. That is, no voter has a load of ${}^{k+1}/_{n}$ or higher, i.e., $\bar{x}_i = \bar{x}_i^k < {}^{k+1}/_{n}$ for each $i \in N$. But then the following holds for the sum of all voter loads.

$$\begin{aligned}
\sum_{i \in N} \bar{x}_i^k &= \sum_{i \in V} \bar{x}_i^k + \sum_{i \in N \setminus V} \bar{x}_i^k \\
&\leq \ell - 1 + \sum_{i \in N \setminus V} \bar{x}_i^k && \text{since } |C_V^{\text{rep}}| \leq \ell - 1 \\
&\leq \ell - 1 + |N \setminus V| s_{c_k}^k && \text{since the maximum possible} \\
& && \text{load after round } k \text{ is } s_{c_k}^k \\
&< \ell - 1 + |N \setminus V| \frac{k+1}{n} && \text{since } s_{c_k}^k < {}^{k+1}/_{n} \\
&< \ell - 1 + \frac{n}{k+1}(k + 1 - \ell)\frac{k+1}{n} \\
&= k
\end{aligned}$$

To see that the last strict inequality holds, note that, since $|V| > \ell \cdot {}^{n}/_{k+1}$, we have

$$|N \setminus V| < n - \ell \cdot {}^{n}/_{k+1} = {}^{n}/_{k+1}(k + 1 - \ell).$$

Thus the sum of all the voter loads is strictly less than k, which contradicts the fact that (exactly) one unit of load is distributed over the voters

[7] Note that this part of the proof will also apply to rl-seq-Phragmén. That is, $|C_V^{\text{rep}}| < \ell$ implies $C_V^+ \subseteq W$. Accordingly, all of the cohesive group's unanimously approved priority candidates are elected in Example 5.1, which shows that late-seq-Phragmén violates p-PJR. This also agrees with the intuition discussed later on that rl-seq-Phragmén violates p-PJR since non-priority representatives that must be elected to meet the representation requirements for cohesive groups may not be elected due to the restriction to priority candidates.

for every elected candidate. We can thus conclude that all unanimously approved priority candidates are elected, i.e., $C_V^+ \subseteq W$.

We now show that, given $C_V^+ \subseteq W$, V must have at least ℓ representatives in W, which contradicts the assumption to the contrary. If $\left|C_V^+\right| \geq \ell$ then we are done, so suppose $\left|C_V^+\right| < \ell$, from which it follows that $z < \ell$. Now, let $x = \left|C_V^+\right|$ be the number of priority candidates unanimously approved by voters in V (note that it is possible that $x > z$). Since $C_V^+ \subseteq W$, this is also the number of unanimously approved priority candidates elected, i.e., $x = |C_V^+ \cap W|$. We then have $\ell - 1 \geq x \geq z$. Furthermore, let y be the number of these priority candidates that are elected before the restriction to priority candidates comes into force. It follows that $x-y$ priority candidates from C_V^+ are elected in the restricted rounds.

If the restriction never came into force, then the outcome would be that of seq-Phragmén; hence as seq-Phragmén satisfies PJR and V is ℓ-cohesive, we would be done. So suppose the restriction does come into force, from which it follows that $y < q$. Then at least the first $k - q + y$ rounds were unrestricted. If $y \geq z$ then V is ℓ-cohesive for the first $k - q + y$ rounds, since then $|V| \geq \ell \cdot n/_{k-q+z} \geq \ell \cdot n/_{k-q+y}$. Consequently, by Lemma 5.2, there must be at least ℓ representatives for V elected in the first $k - q + y$ rounds. If, on the other hand, $y < z$, then, we claim, V is $[\ell - (z - y)]$-cohesive for the first $k - q + y$ rounds.[8] This is indeed the case, since

$$
\begin{aligned}
|V| &\geq \ell \cdot \frac{n}{k-q+z} && \text{since } V \text{ is } z\text{-}\ell\text{-cohesive and } \left|C_V^+\right| < \ell \\
&\geq [\ell - (z-y)] \cdot \frac{n}{k-q+z-(z-y)} && \text{since } \ell > z - y > 0 \\
&= [\ell - (z-y)] \cdot \frac{n}{k-q+y}
\end{aligned}
$$

But then, by Lemma 5.2, V must receive at least $\ell-(z-y)$ representatives in the first $k-q+y$ rounds. Hence, as we know that another $x-y \geq z-y$ representatives are elected in the restricted rounds, V receives at least $\ell - (z-y) + (z-y) = \ell$ representatives in total.

Thus, in both cases V receives at least ℓ representatives, which contradicts our assumption to the contrary. We can thus conclude, by *reductio*, that late-seq-Phragmén satisfies p-PJR. □

Corollary 5.3.1. *late-seq-Phragmén satisfies rp-PJR, p-JR and rp-JR.*

In contrast with late-seq-Phragmén, rl-seq-Phragmén does not satisfy p-PJR, as the following example shows.

[8] Note that since $\ell > z > y$, we have $\ell - (z - y) \geq 1$.

Example 5.1. Let $C = \{\overline{a}, \overline{b}, \overline{c}, \overline{d}, e\}$ $k = 4$, $q = 2$ and consider the profile

$$1 \times \{\overline{a}, \overline{b}, \overline{c}, \overline{d}\} \qquad\qquad 3 \times \{\overline{a}, \overline{b}, e\}.$$

In the first two rounds $\overline{a}$ and $\overline{b}$ will be elected (the order will depend on tiebreaking between them). In the last two rounds the candidate pool is restricted to priority candidates, which means $\overline{c}$ and $\overline{d}$ will be elected (again, the order depends on tiebreaking). Now note that the three voters approving of $\{\overline{a}, \overline{b}, e\}$ are a 2-3-cohesive group since they jointly approve of 3 candidates, among which 2 priority candidates and since $3 \geq 3 \cdot {}^{4}/_{4-2+2}$. Thus, p-PJR requires that $\overline{a}$, $\overline{b}$ and e all be elected, which is not the case.

Though rl-seq-Phragmén does not satisfy p-PJR, it does satisfy rp-PJR, and consequently rp-JR. To see this, note that an rp-ℓ-cohesive group V is ℓ cohesive for the first $k - q$ rounds, from which the result follows directly in light of Lemma 5.2. In the above counterexample, since $3 \geq 1 \cdot {}^{4}/_{4-2}$ and $3 < 2 \cdot {}^{4}/_{4-2}$, rp-PJR requires only that one of the candidates from $\{\overline{a}, \overline{b}, e\}$ is elected, which is the case. This also illustrates the difference in strength of the two axioms. In this case, with committee size $k = 4$, p-PJR, in keeping with PJR, requires three candidates for the 3-cohesive group, while rp-PJR requires only one.

Intuitively, rl-seq-Phragmén fails p-PJR because a cohesive group's (potential) non-priority representatives can only be assigned unrestricted seats. Hence, if too many of a cohesive group's priority representatives are elected in unrestricted rounds, the non-priority candidates that must be elected in order to meet the representation requirements prescribed by p-PJR may not be able to be elected. This intuition also makes clear why late-seq-Phragmén does satisfy p-PJR: if a priority representative of a group is elected in an unrestricted seat, the restriction comes into force later, meaning the necessary non-priority candidates can still be elected.

We have seen that rl-seq-Phragmén does satisfy rp-PJR and rp-JR. It also satisfies p-JR.

Proposition 5.4. *rl-seq-Phragmén satisfies p-JR.*

Proof. Let $E = (C, C^+, \boldsymbol{A}, k, q)$ be an election and V a z-1-cohesive group and let $C_V = \bigcap_{i \in V} A_i$. It follows that $|C_V| \geq 1$. Furthermore, if $\left|C_V^+\right| \geq 1$ we have $|V| \geq {}^{n}/_{k}$, and otherwise, if $C_V^+ = \emptyset$, we have $z = 0$ and $|V| \geq {}^{n}/_{k-q}$. Let $W = \text{rl-seq-Phragmén}(E)$. We show that there is some voter $i \in V$ with $A_i \cap W \neq \emptyset$, considering the two cases separately.

Suppose $C_V^+ = \emptyset$, from which it follows that $z = 0$. We then have that V is 1-cohesive for the first $k - q$ rounds, since then $|V| \geq {}^{n}/_{k-q}$. As the first $k - q$ rounds are unrestricted, it follows by Lemma 5.2, that at least one representative $c \in \bigcup_{i \in V} A_i$ is elected in the first $k - q$ rounds.

Now suppose $\left|C_V^+\right| \geq 1$. It follows that there is some unelected priority candidate $c^+ \in C_V^+ \setminus W$. Now let $C' = C^+ \cup W$, i.e., C' is the set of priority candidates minus those non-priority candidates that were not elected, and consider the election $E' = (C', \boldsymbol{A}, k)$. In light of Lemma 2.5, we know that $W = \text{seq-Phragmén}(E')$. Now, since $|V| \geq {}^{n}\!/_{k}$ and $c^+ \in C'$, we have that V is 1-cohesive for E'. But then, as seq-Phragmen satisfies JR, it must be the case that there is some voter $i \in V$ with a representative in W, i.e., $A_i \cap W \neq \emptyset$.

We can thus conclude that $\left(\bigcup_{i \in V} A_i\right) \cap W \neq \emptyset$. That is, rl-seq-Phragmén satisfies p-JR. $\square$

This brings us to the end of the chapter, which presents a first step in the adaptation of well-studied proportionality axioms to the priority setting. We saw that late-seq-Phragmén fares particularly well. It may be expected that the priority adaptations of MES would do so as well, though this is yet to be confirmed.

Chapter 6

Conclusion

In this chapter we first summarise the key findings and contributions made in this thesis. Following the summary, we will briefly discuss the results, paying particular attention to potential directions for future research.

Motivated by a wide variety of multiwinner voting instances in which certain kinds of candidates are to be afforded preferential treatment, this thesis presents a first attempt at a formal study of such settings. We introduced a priority model for multiwinner voting that enriches the standard model by distinguishing between priority and non-priority candidates and by introducing a quota q that specifies the minimum number of priority candidates that are to be elected. Several priority rules were defined that respect the quota, i.e., that guarantee the election of at least q priority candidates. These rules were all based on the idea that q of the k available seats are to be reserved for priority candidates, while the other $k-q$ seats are open to all candidates. We considered three ways of reserving seats: reserving the first seats, the last seats, or, for sequential rules, reserving the remaining seats at the latest possible moment this has to be done to guarantee the quota is respected. For candidate-wise scoring rules, it turned out, there is no difference between the reserved-first and the reserved-latest approach, while for seq-Phragmén and MES the reserved-first approach is either ill-defined or obviously flawed.

We first considered the defined rules with regards to anonymity and priority neutrality, a weakening of neutrality that takes into account the differential treatment of priority and non-priority candidates. Whereas all rules considered are anonymous, we saw that rules satisfy priority neutrality only when the underlying non-priority rule satisfies neutrality.

Subsequently, we introduced two axioms, *possible win* and *guaranteed win*, to capture the requirement that candidates ought never to be worse off under a priority rule as priority candidates than as non-priority candidates. We showed that all considered priority rules meet these minimal requirements. We also considered analogous axioms concerning entire committees, namely *committee possible win* and *committee guaranteed win*, which proved too strong to be considered minimal requirements for priority rules. Instead, we argued, these can be interpreted as encoding a form of priority treatment. Accordingly, we found that the reserved-last priority rules satisfy these requirements, while the reserved-first and reserved-latest adaptations do not.

In the rest of the thesis we restricted our attention to two different kinds of elections: excellence-based elections and elections concerning proportional representation. For excellence-based elections, we saw that the dual election rules based on candidate-wise scoring functions were most suitable. All such rules satisfy *priority merit* and *non-priority merit*, implying that a candidate c can only ever be preferred over a higher-scoring candidate c' if the former is a priority candidate and the latter is not. The candidate-wise reserved-first dual rules are then characterised by *respect of quota* and *limit to quota*. Together with priority merit and non-priority merit, these axioms enforce the minimal deviation from score maximisation required to elect q priority candidates. We also introduced *active quota*, which demands a stronger form of preferential treatment, requiring, roughly, that q *protected* priority candidates are elected that would otherwise not have been elected. The candidate-wise reserved-first dual rules are characterised by this stronger form of preferential treatment and *limit to active quota*.

Consistent with the finding that the candidate-wise reserved-last dual rules considered provide a greater degree of preferential treatment, we showed that these rules fail *quota priority efficiency*, though they do satisfy the weaker requirement of *priority efficiency*. That is, though the outcomes of these rules cannot be Pareto improved upon without decreasing the number of priority candidates elected, such improvements may be possible without violating respect of quota. In contrast, we found that the candidate-wise reserved-first dual rules considered do satisfy quota priority efficiency, which accords with the finding that these rules minimally deviate from the goal of maximising the score.

In the context of excellence-based elections, opting for a stronger form of priority treatment, as exemplified by the candidate-wise reserved-last rules, comes at the expense of having a committee with lower-scoring candidates. If committee quality and voter satisfaction, particularly regarding outcome efficiency, are the primary concern, then the candidate-wise reserved-first rules are most suitable. In contrast, the candidate-wise reserved-last rules are more suitable if the aim is to use reserved seats to aid priority candidates who would otherwise not have been elected.

Finally, we considered proportional representation in the context of the priority model. We adapted the notion of ℓ-cohesiveness from the standard setting to take into account, firstly, the worst-case scenario in which no reserved seats can be used to provide representatives for a cohesive group, and secondly, the fact that the more priority candidates a cohesive group jointly approves of, the easier it is to grant them representatives. This yielded a hierarchy of three axioms (from weakest to strongest): *priority justified representation* (p-JR), *priority proportional justified representation* (p-PJR), and *priority extended justified represen-*

tation (p-EJR). We saw that all priority adaptations of seq-CC and seq-Phragmén satisfy p-JR, while only late-seq-Phragmén satisfies p-PJR. Accordingly, the latter rule is a suitable choice when the main concern is proportional representation.

Having summarised the key developments and findings, the subsequent and final section of this thesis reflects on the accomplished work and presents potential directions for future research.

Evaluation and future research

We set out to develop a simple multiwinner voting framework in which the preferential treatment of certain kinds of candidates can be modelled. This priority model proved applicable to various kinds of elections and valuable in the analysis of essential properties of voting rules. It allowed us to investigate the preferential treatment of priority candidates, as well as previously studied properties that remain relevant in the priority setting. Moreover, we successfully identified rules that exhibit the desired properties for the different settings studied.

A notable omission from our results concerns the priority adaptations of MES. Given that MES satisfies strong proportionality axioms such as EJR, we may expect its priority adaptations to perform well in this context too. Thus far, however, it is unclear whether, e.g., late-MES satisfies p-EJR.

Necessarily, there are many other relevant aspects that could not be covered in this thesis. In that light we give some suggestions for possible future research, which fall, roughly, into the following categories: expanding the set of axioms considered in the priority setting, developing new rules, generalising the studied model, complexity analysis, and empirical analysis.

A great many axioms have been studied in multiwinner voting, and many more of these are applicable to the priority setting than we had room to cover here. Most obvious, since we consider proportionality, are axioms concerning the fair distribution of voting power. However, other axioms such as *committee monotonicity*, *consistency* and *strategyproofness* are also relevant in the priority setting.

In this thesis we have been concerned exclusively with rules that meet the quota, which can, in this context, be seen as the minimal requirement regarding the preferential treatment of priority candidates. Regarding excellence-based elections, we were able to explore different forms of priority treatment, since candidate-wise scores allow for the direct comparison of candidates. However, in the context of proportionality, where a candidate's suitability depends fundamentally on the other candidates chosen, this is more challenging. The axioms *committee possible win* and *committee guaranteed win* offer more universal measures of preferential treatment, but their underlying motivation is somewhat unclear. It would be valuable to develop a better understanding of different forms of preferential treatment in general, or at least for other settings, such as in the context of proportionality.

The method of reserving q out of k seats for priority candidates yielded a number of interesting rules. However, we encountered challenges when applying the reserved-first approach to seq-Phragmén and MES. Addressing these issues, though beyond the scope of this thesis, would yield rules comparable to the reserved-last and reserved-latest adaptations. Additionally, besides the non-priority rules that we have considered, there are others, such as *Greedy Monroe* or *Minimax Approval Voting*, that may be adapted using the same method of reserving seats.

Besides reserving q seats for priority candidates, there is another natural approach to adapting rules so that they respect the quota. For any scoring function we may define a rule that returns the committees respecting the quota that maximise this score. (Note that for candidate-wise scores this is equivalent to the reserved-first adaptation.) It may be expected that, e.g., such an adaptation of PAV would fare well with regards to proportionality. Although we chose to focus on computationally tractable rules, an analysis of these alternative rules would be valuable, offering fruitful comparisons with the rules that we did consider.

While our focus has been on rules that respect the quota, there are situations in which preferential treatment is to be given to certain candidates without having an explicit requirement regarding the number of such candidates to be elected. In such cases a different class of rules seems natural. For example, rules that require voters to pay a price or shoulder a load for the election of a candidate may be adapted by making priority candidates relatively cheaper. Similarly, for scoring rules, greater weight could be assigned to the contribution of (voters who approve of) priority candidates to the overall score. Given that the axioms that we have considered all assume respect of quota, these types of settings without a quota would necessitate significantly different axioms.

The most obvious shortcoming of the model studied is its simplicity, assuming a bipartite set of candidates. There are often more types of candidates with differing requirements for preferential treatment. For instance, the Rwandan parliament reserves a different numbers of seats for women, youth members and disabled members. A naive extension of the priority model we consider would treat types as mutually exclusive, i.e., as partitioning the set of candidates. Though this will be applicable in certain situations, a more comprehensive model that allows candidates to have multiple types, such as being both a woman and disabled, is necessary for many scenarios.

We did not consider computational complexity in this thesis. Since the priority rules here considered do not fundamentally alter the non-priority rules on which they are based, but simply restrict the set of eligible candidates in certain steps, these priority rules will generally inherit the computational tractability of their underlying non-priority rules.

Nevertheless, there are other interesting questions concerning computational complexity in the priority model. For example, we may consider what the complexity is of determining whether a certain committee satisfies p-JR, p-PJR or p-EJR.

Finally, it would be interesting to study the distinctions and implications of the studied rules from a more empirical perspective. By analysing actual voting data, we can get insight into the practical ramifications of these rules and may identify any deviations from the theoretical expectations yielded by the axiomatic analysis.

References

Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, *93*(3), 729–747. https://doi.org/10.1257/000282803322157061

Arrow, K. J., Sen, A., & Suzumura, K. (2002). *Handbook of Social Choice and Welfare* (Vol. 1). Elsevier.

Aziz, H., Brill, M., Conitzer, V., Elkind, E., Freeman, R., & Walsh, T. (2017). Justified representation in approval-based committee voting. *Social Choice and Welfare*, *48*(2), 461–485. https://doi.org/10.1007/s00355-016-1019-3

Aziz, H., Elkind, E., Huang, S., Lackner, M., Sanchez-Fernandez, L., & Skowron, P. (2018). On the complexity of extended and proportional justified representation. *Proceedings of the 32nd Conference on Artificial Intelligence (AAAI-18)*, 902–909. https://doi.org/10.1609/aaai.v32i1.11478

Aziz, H., & Monnot, J. (2020). Computing and testing Pareto optimal committees. *Autonomous Agents and Multi-Agent Systems*, *34*(1), 24. https://doi.org/10.1007/s10458-020-09445-y

Balinksi, M. L., & Young, H. P. (2001). *Fair Representation: Meeting the Ideal of One Man, One Vote*. Brookings Institution Press. https://doi.org/10.2307/2130978

Barberà, S., Bossert, W., & Pattanaik, P. K. (2004). Ranking sets of objects. In S. Barberà, P. J. Hammond, & C. Seidl (Eds.), *Handbook of Utility Theory* (pp. 893–977). Springer. https://doi.org/10.1007/978-1-4020-7964-1_4

Barberà, S., & Coelho, D. (2008). How to choose a non-controversial list with k names. *Social Choice and Welfare*, *31*(1), 79–96. https://doi.org/10.1007/s00355-007-0268-6

Brams, S. J., & Kilgour, D. M. (2015). Satisfaction approval voting. In R. Melnik (Ed.), *Mathematical and Computational Modeling* (pp. 273–298). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781118853887.ch11

Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (Eds.). (2016). *Handbook of Computational Social Choice*. Cambridge University Press. https://doi.org/10.1017/CBO9781107446984

Bredereck, R., Faliszewski, P., Igarashi, A., Lackner, M., & Skowron, P. (2018). Multiwinner elections with diversity constraints. *Proceedings of the 32nd Conference on Artificial Intelligence (AAAI-18)*, 933–940. https://doi.org/10.1609/aaai.v32i1.11457

Brill, M., Freeman, R., Janson, S., & Lackner, M. (2023). Phragmén's voting methods and justified representation. *Mathematical Programming*. https://doi.org/10.1007/s10107-023-01926-8

Brill, M., Laslier, J.-F., & Skowron, P. (2018). Multiwinner approval rules as apportionment methods. *Journal of Theoretical Politics*, *30*(3), 358–382. https://doi.org/10.1177/0951629818775518

Celis, L. E., Huang, L., & Vishnoi, N. K. (2018). Multiwinner voting with fairness constraints. *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18)*, 144–151. https://doi.org/10.24963/ijcai.2018/20

Cevallos, A., & Stewart, A. (2021). A verifiably secure and proportional committee election rule. *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, 29–42. https://doi.org/10.1145/3479722.3480988

Chamberlin, J. R., & Courant, P. N. (1983). Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, *77*(3), 718–733. https://doi.org/10.2307/1957270

Doğan, B. (2016). Responsive affirmative action in school choice. *Journal of Economic Theory*, *165*, 69–105. https://doi.org/10.1016/j.jet.2016.04.007

Doğan, B., Doğan, S., & Yıldız, K. (2021). Lexicographic choice under variable capacity constraints. *Journal of Public Economic Theory*, *23*(1), 172–196. https://doi.org/10.1111/jpet.12482

Dur, U., Kominers, S. D., Pathak, P. A., & Sönmez, T. (2018). Reserve design: Unintended consequences and the demise of Boston's walk zones. *Journal of Political Economy*, *126*(6), 2457–2479. https://doi.org/10.1086/699974

Dur, U., Pathak, P. A., & Sönmez, T. (2020). Explicit vs. statistical targeting in affirmative action: Theory and evidence from Chicago's exam schools. *Journal of Economic Theory*, *187*, 104996. https://doi.org/10.1016/j.jet.2020.104996

Echenique, F., & Yenmez, M. B. (2015). How to control controlled school choice. *American Economic Review*, *105*(8), 2679–2694. https://doi.org/10.1257/aer.20130929

Elkind, E., Faliszewski, P., Skowron, P., & Slinko, A. (2017). Properties of multiwinner voting rules. *Social Choice and Welfare*, *48*(3), 599–632. https://doi.org/10.1007/s00355-017-1026-z

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

Faliszewski, P., Sawicki, J., Schaefer, R., & Smolka, M. (2017a). Multiwinner voting in genetic algorithms. *IEEE Intelligent Systems*, *32*(1), 40–48. https://doi.org/10.1109/MIS.2017.5

Faliszewski, P., Skowron, P., Slinko, A., & Talmon, N. (2017b). Multiwinner rules on paths from k-Borda to Chamberlin–Courant. *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*, 192–198. https://doi.org/10.24963/ijcai.2017/28

Faliszewski, P., Skowron, P., Slinko, A., & Talmon, N. (2017c). Multiwinner voting: A new challenge for social choice theory. In U. Endriss (Ed.), *Trends in Computational Social Choice* (pp. 27–47). AI Access.

Faliszewski, P., Slinko, A., & Talmon, N. (2020). Multiwinner rules with variable number of winners. *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-20)*, 67–74. https://doi.org/10.3233/FAIA200077

Faliszewski, P., & Talmon, N. (2018). Between proportionality and diversity: Balancing district sizes under the Chamberlin-Courant rule. *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-18)*, 14–22. https://dl.acm.org/doi/10.5555/3237383.3237931

Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, *69*(1), 9–15. https://doi.org/10.2307/2312726

Hafalir, I. E., Yenmez, M. B., & Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, *8*(2), 325–363. https://doi.org/10.3982/TE1135

International Institute for Democracy and Electoral Assistance. (n.d.). *Global Database on Gender Quotas*. Retrieved May 6, 2023, from https://www.idea.int/data-tools/data/gender-quotas

Janson, S. (2018). Phragmén's and Thiele's election methods. arXiv:1611.08826 [math.HO]. https://doi.org/10.48550/arXiv.1611.08826

Jaworski, M., & Skowron, P. (2022). Phragmén rules for degressive and regressive proportionality. *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-22)*, 328–334. https://doi.org/10.24963/ijcai.2022/47

Kilgour, D. M. (2010). Approval balloting for multi-winner elections. In J.-F. Laslier & M. R. Sanver (Eds.), *Handbook on Approval Voting* (pp. 105–124). Springer. https://doi.org/10.1007/978-3-642-02839-7_6

Kilgour, D. M. (2016). Approval elections with a variable number of winners. *Theory and Decision*, *81*(2), 199–211. https://doi.org/10.1007/s11238-016-9535-2

Kilgour, D. M., & Marshall, E. (2012). Approval balloting for fixed-size committees. In D. S. Felsenthal & M. Machover (Eds.), *Electoral Systems: Paradoxes, Assumptions, and Procedures* (pp. 305–326). Springer. https://doi.org/10.1007/978-3-642-20441-8_12

Kojima, F. (2012). School choice: Impossibilities for affirmative action. *Games and Economic Behavior*, *75*(2), 685–693. https://doi.org/10.1016/j.geb.2012.03.003

Lackner, M., & Skowron, P. (2018). Consistent approval-based multi-winner rules. *Proceedings of the 2018 ACM Conference on Economics and Computation (EC-18)*, 47–48. https://doi.org/10.1145/3219166.3219170

Lackner, M., & Skowron, P. (2020). Utilitarian welfare and representation guarantees of approval-based multiwinner rules. *Artificial Intelligence*, *288*, 103366. https://doi.org/10.1016/j.artint.2020.103366

Lackner, M., & Skowron, P. (2023). *Multi-Winner Voting with Approval Preferences*. Springer Nature. https://doi.org/10.1007/978-3-031-09016-5

Lang, J., & Skowron, P. (2018). Multi-attribute proportional representation. *Artificial Intelligence*, *263*, 74–106. https://doi.org/10.1016/j.artint.2018.07.005

Laslier, J.-F., & Sanver, M. R. (Eds.). (2010). *Handbook on Approval Voting*. Springer. https://doi.org/10.1007/978-3-642-02839-7

Norris, P. (1997). Choosing electoral systems: Proportional, majoritarian and mixed systems. *International Political Science Review*, *18*(3), 297–312. https://doi.org/10.1177/019251297018003005

Peters, D. (2018). Proportionality and strategyproofness in multiwinner elections. *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-18)*, 1549–1557. https://doi.org/10.5555/3237383.3237931

Peters, D., & Skowron, P. (2020). Proportionality and the limits of welfarism. *Proceedings of the 21st ACM Conference on Economics and Computation (EC-20)*, 793–794. https://doi.org/10.1145/3391403.3399465

Phragmén, E. (1896). Sur la théorie des élections multiples. *Öfversigt af Kongliga Vetenskaps-Akademiens Förhandlingar*, *53*, 181–191.

Phragmén, E. (1899). Till frågan om en proportionell valmetod. *Statsvetenskaplig Tidskrift*, *2*(2), 297–305.

Procaccia, A. D., Rosenschein, J. S., & Zohar, A. (2008). On the complexity of achieving proportional representation. *Social Choice and Welfare*, *30*(3), 353–362. https://doi.org/10.1007/s00355-007-0235-2

Pukelsheim, F. (2017). *Proportional Representation: Apportionment Methods and Their Applications* (2nd ed.). Springer. https://doi.org/10.1007/978-3-319-64707-4

Sánchez-Fernández, L., Elkind, E., Lackner, M., Fernández, N., Fisteus, J., Basanta Val, P., & Skowron, P. (2017). Proportional justified representation. *Proceedings of the 31st Conference on Artificial Intelligence (AAAI-17)*. https://doi.org/10.1609/aaai.v31i1.10611

Sayedahmed, D. (2021). *Refugee settlement and other matching problems with priority classes and reserves: A market design perspective* (Doctoral dissertation). Concordia University.

Skowron, P., Faliszewski, P., & Lang, J. (2016). Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, *241*, 191–216. https://doi.org/https://doi.org/10.1016/j.artint.2016.09.003

Thiele, T. N. (1895). Om flerfoldsvalg. *Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, *1895*, 415–441.