Strategic Voting under Incomplete Information in Approval-Based Committee Elections

MSc Thesis *(Afstudeerscriptie)*

written by

Jason Tsiaxiras
(born March 28th, 1995 in Amsterdam, the Netherlands)

under the supervision of **Dr. Zoi Terzopoulou** and **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 5th, 2021

Members of the Thesis Committee:
Prof. Dr. Yde Venema (Chair)
Dr. Zoi Terzopoulou (Supervisor)
Prof. Dr. Ulle Endriss (Supervisor)
Prof. Dr. Aybüke Özgün
Dr. Arianna Novaro

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Contents

Abstract

In approval-based committee elections, voters vote by submitting an approval ballot, which indicates which candidates the voter approves of, with the purpose of electing a fixed size committee. Recent impossibility results have shown that approval-based committee voting rules that select committees that provide proportional representation to the voters are inherently susceptible to strategic manipulation. The impossibility results obtained in the literature rely on the assumption that voters have complete information about the preferences of their fellow voters. In this thesis, we extend the model of approval-based committee elections in order to account for strategic voting when voters have incomplete information about the preferences of their fellow voters. The purpose of this model is to be able to describe the strategic behavior of real voters more accurately and explore whether there are information restrictions under which voters do not have incentive to misrepresent their preferences.

In our analysis we employ formal methods to study voting rules that satisfy a set of normative properties. We also perform simulation experiments to study three well-established proportional voting rules. The results of our experiments show a high prevalence of incentive to manipulate when voters have complete information. Moreover, the experiments show that decreasing the information a voter has does not necessarily decrease the probability that this voter has incentive to manipulate. The negative results of our formal analysis show that, for two natural kinds of incomplete information voters may have, voting rules that satisfy a set of weak normative properties remain susceptible to manipulation. We also present two positive results for the Proportional Approval Voting rule which show that information restrictions exist under which this voting rule is (partially) strategy-proof.

Chapter 1

Introduction

This thesis is concerned with strategic voting in approval-based committee elections. In committee elections a group of voters vote on a set of candidates with the goal to elect a fixed size committee. In the approval-based model for committee elections, voters report their preferences by submitting an *approval ballot*, which indicates, for each candidate, whether the voter approves of this candidate or not. This approval-based model has in recent years received increased attention in the field of Computational Social Choice (Lackner and Skowron, 2020a). It stands in contrast to classical models in Social Choice Theory in which voters report their preferences as a linear ranking of the candidates from least preferred to most preferred.

The settings in which committee elections are held are ubiquitous. For example, national and local elections for representative bodies such as parliaments can be seen as committee elections. The election of a board in a company or a trade union are also committee elections. Moreover, business decisions such as deciding which group of products to advertise or which movies to offer on an airplane can be informed by voting on the alternatives and can therefore be regarded as committee elections (Skowron et al., 2016). Even voting on the locations of public facilities can be seen as a type of committee elections (Skowron et al., 2016). Within the large variety of types of committee elections there are different objectives. Consequently, different voting rules which aggregate the preferences of the voters in order to select a winning committee are appropriate. Elkind et al. (2017) identify three main objectives in committee elections. The first objective is to select the candidates which score best in terms of individual excellence. The second objective is to select a committee which provides proportional representation to voters. The third possible objective is to select a diverse subset of the candidates, which provides an approved option for as many voters as possible.

In this thesis we focus on the objective of proportional representation. We study approval-based committee voting rules that achieve outcomes which provide proportional representation to the voters. In particular, we are concerned with strategic voting for these proportional voting rules when voters have incomplete information about the preferences of the other voters.

1.1 Motivation for this Thesis

An argument in favor of committee elections based on approval ballots instead of linear rankings of the candidates is that voters may not always be able to provide a linear ranking of the candidates. This may be because voters do not hold such detailed opinions or because their opinions of the candidates cannot be reduced to a linear ranking. In such cases approval ballots may provide a more natural representation of the preferences, and additionally require less cognitive effort of the voters to construct compared to linear rankings.

In the field of Computation Social Choice, a central topic of study is that of strategic manipulation. This topic has received extensive attention ever since the influential work by Gibbard (1973) and Satterthwaite (1975). Gibbard and Satterthwaite showed that in classical Social Choice Theory, no voting rule can be simultaneously democratic, in the sense that every candidate has a chance to win and no voter acts as a dictator, and strategy-proof. These results are often referred to as *impossibility results*, because they show that defining a voting rule which satisfies a combination of multiple desirable properties is impossible. In the setting of approval-based committee elections analogous impossibility results have been obtained by Peters (2018), by Lackner and Skowron (2018) and by Kluiving et al. (2020). These studies provide formal proofs that the combination of strategy-proofness and proportionality cannot be attained. That is, there are no voting rules that are simultaneously proportional and strategy-proof.

Strategy-proof voting rules are voting rule for which a voter can never achieve a better outcome for herself by voting insincerely. The primary reason for the interest in strategy-proof voting rules is that strategic voting may lead election results away from the most optimal result. If many voters vote strategically, the reported preferences and the truthful preferences may differ to such an extend that the outcome produced by the voting rule does not optimally serve the electorate. Other attractive properties of the voting rule, such as fairness, proportionality or utilitarian welfare, are only guaranteed when voters vote truthfully. Therefore, the voting rule should elicit the truthful preferences by ensuring that each voter's best strategy is to report her truthful preferences.

A simple type of manipulation for proportional approval-based committee voting rules is called *free-riding*. Voters who manipulate by free-riding leave out certain popular candidates from their approval ballot in order to pretend to not be represented by these candidates and put more of their voting power behind their other candidates. Proportional voting rules are inherently susceptible to this type of manipulation because these voting rule attach more weight to voters who are not well represented. The impossibility results obtained in the literature provide formal confirmation of this fact.

The impossibility results, which establish that for any proportional voting rule there are cases in which voters have incentive to misrepresent their preferences, are based on the assumption that voters have complete information about the preferences of their fellow voters. In real-life elections, however, voters will not have complete information about the preferences of their fellow voters. The goal in this thesis is to account for incomplete information in order to more accurately model real voters and, hopefully, provide a way out of the impossibility results. We will investigate whether there are restrictions for the information voters can have under which proportional voting rules are strategy-proof.

Modelling incomplete information of voters in the context of strategic manipulation has been undertaken in different areas of Computational Social Choice. Reijngoud and Endriss (2012) and Endriss et al. (2016) study strategic manipulation under incomplete information for single-winner voting rules with linear preferences of voters. Terzopoulou (2017) studies strategic manipulation under incomplete information for judgement aggregation. The analysis in this thesis will make use of the model of incomplete information developed by Reijngoud and Endriss (2012) and apply it to approval-based committee voting.

1.2 About this Thesis

The following is a brief synopsis of the content of this thesis.

Chapter 2

The goal of this chapter is to establish the groundwork for our analysis. We begin by introducing the formal model of approval-based committee voting and the preference relation over election outcomes that we assign to the voters. We then present a set of normative properties, called axioms, which we would like a voting rule to satisfy. We discuss axioms that have been studied in the literature, and focus in particular on understanding what *proportionality* means in the context of approval-based committee voting. Moreover, we present a number of original axioms, the most notable of which is the *diminishing returns* axiom. In the last section of this chapter we introduce a number of voting rules that have been studied in the literature and which are considered to achieve proportional outcomes. We survey which of these voting rules satisfy which of our normative properties.

Chapter 3

In Section 3.1 we give an overview of the impossibility results that have been obtained in the literature. These results show that proportionality cannot be achieved in combination with strategy-proofness. The goal of this chapter is to understand why proportional rules are susceptible to strategic manipulation and how prevalent manipulations are for proportional rules. We present and analyse a number of examples of strategic manipulation and distinguish between different types of manipulation. In Section 3.3 we perform simulations to quantify the prevalence of these different types of manipulation for the most prominent voting rules. Finally, in Section 3.4 we present a special class of preference profiles on which proportionality and partial strategy-proofness may be achieved.

Chapter 4

In Section 4.1 we extend our model to account for strategic voting when voters have incomplete information about the preferences of the other voters. The main question addressed in this chapter is whether we can prevent manipulations by restricting the information voters have. In Section 4.2 we present a number of negative results. For some natural types of partial information, rules that satisfy our axioms are still susceptible to strategic manipulation. In Section 4.3 we perform simulations to quantify the prevalence of manipulable profiles for the most prominent rules when voters have partial information. In Section 4.4 we present two positive results that show (partial) information barriers can be achieved for one of the proportional rules. In Section 4.5 we discuss these results and reflect on the assumptions in our model.

Chapter 5

In this chapter we briefly summarise the main results in this thesis and we provide some directions for future research.

Chapter 2

Approval-Based Committee Voting

In this chapter, we begin by introducing the formal model of approval-based committee voting. We also introduce the preference relation over election outcomes that we assign to voters and define what we take to be incentive for strategic manipulation.

In Section 2.2 we present a number of normative properties, called *axioms*, which we would like a voting rule to satisfy. We discuss axioms that have been studied in the literature. Special attention here will be payed to understand what *proportionality* means in the context of approval-based committee voting. We also introduce a number of original axioms, the most notable of which is the *diminishing returns* axiom. The axioms we introduce can be seen as minimal requirements for our voting rules. They will play an important role in our analysis in Chapter 4.

In Section 2.4 we present a number of voting rules that have been studied in the literature and which are considered to achieve proportional outcomes. We survey which of these voting rules satisfy which of our normative properties.

2.1 The Model

In this section we lay the groundwork for our mathematical analysis by formally defining the concepts *approval set, election scenario* and *voting rule*, making explicit the assumptions about the voters' preferences and defining the notation we will use.

2.1.1 Preliminaries

We take $N = \{1, \dots, n\}$ to be the set of n voters, $C = \{c_1, \dots, c_m\}$ the set of m candidates and $k \in \mathbb{N}$ the desired committee size. Moreover, we denote by $\mathcal{P}(X)$ the powerset of X and by $\mathcal{P}_k(C)$ the set of all k-size subsets of C, which we call *committees*.

Each voter $i \in N$ has an *approval set* $A_i \subseteq C$ which represents her sincere dichotomous preference and we will represent the *approval profile* as a vector $\boldsymbol{A} = \big(A_1, \dots, A_n\big)$ of the approval sets of all voters in N. We denote the set of all approval profiles over (N, C) by $\mathcal{A}(N, C)$. Furthermore, we call a triple (N, C, k) an *election scenario* when the desired committee size k is equal to or less than the number of candidates m and we will call a pair $(\boldsymbol{A}, k)$ an *election instance* when $\boldsymbol{A}$ is an approval profile over (N, C) and (N, C, k) is an election scenario.

A *voting rule* is a function that takes as input election instances and returns non-empty sets of winning committees. That is, a voting rule is a function $F : \mathcal{A}(N, C) \rightarrow \mathcal{P}(\mathcal{P}_k(C)) \setminus \varnothing$. Strictly speaking, N, C and k are parameters and a voting rule F is defined for every choice of these parameters, as long as (N, C, k) is an election scenario. However, for ease of notation we omit this.

Throughout this thesis we will make use of the following additional notation. We will write $\boldsymbol{A}_{-i}$ to denote the partial profile resulting from removing voter i's approval set from $\boldsymbol{A}$. Similarly, we write $(\boldsymbol{A}_{-i}, B)$ to denote the profile $\boldsymbol{A}$ but with voter i's approval set replaced by some ballot $B \subseteq C$. In some cases, we will write $\boldsymbol{A}^{i+\{d\}}$ to denote the profile identical to $\boldsymbol{A}$ except that voter i adds candidate d to her approval ballot. These notations will be useful when we consider voter i's possibilities to manipulate the election by submitting a ballot B that is different from her approval set A_i. In this context we will refer to A_i as i's truthful approval set and to B as an insincere ballot.

2.1.2 Preference Relation

In order to study strategic behaviour in approval-based committee elections, assumptions have to be made about when a voter prefers one outcome over another. The two basic assumptions that we work with are the following. Firstly, we assume that the preferences of voters are accurately represented by dichotomous preferences over the set of candidates: voters approve of some of the candidates and disapprove of the remaining candidates. In particular, this means that voters do not have a preference order over the candidates in their approval set nor over the candidates outside of their approval set. There exist alternative models for committee voting in which each voter ranks the candidates from least to most preferred. We will not consider committee voting based on such linear preferences. For an overview of such models we refer the reader to a recent book chapter by Faliszewski et al. (2017). Our second assumption is that voters desire to maximise the representation they receive in the elected committee. A voter i prefers committee $W \in \mathcal{P}_k(C)$ over committee $W' \in \mathcal{P}_k(C)$ if and only if $|W \cap A_i| > |W' \cap A_i|$. That is, a voter prefers committee W over committee W' when the voter receives more representation in W than in W'.

The voting rules we consider are irresolute, which means the outcomes they return are sets of winning committees rather than single winning committees. Thus, in order to study strategic behaviour for these voting rules, assumptions have to be made about when a voter prefers one set of committees over another set of committees. To this end we use the concept of *stochastic dominance*.

Definition 1 (*Stochastic dominance*)**.** Let $A \subseteq C$ be a truthful approval set. A set of committees $X \subseteq \mathcal{P}_k(C)$ *stochastically dominates* a set of committees $Y \subseteq \mathcal{P}_k(C)$ relative to A when the following two conditions hold:

1. For each $\ell \in \mathbb{N}$ we have $\frac{|\{W \in X : |W \cap A| \geq \ell\}|}{|X|} \geq \frac{|\{W \in Y : |W \cap A| \geq \ell\}|}{|Y|}$.
2. There is an $\ell \in \mathbb{N}$ for which the above inequality is strict.

For elections in which a single committee is to be elected, the irresolute voting rules that we will consider will ultimately have to be extended with a tie-breaking procedure on the set of winning committees. The idea underlying the stochastic dominance preference relation is that this final tie-breaking is achieved by a random procedure employing the uniform probability distribution over the set of winning committees. With this in mind, the interpretation of the stochastic dominance relation is the following: a voter with approval set A prefers outcome $X \subseteq \mathcal{P}_k(C)$ over outcome $Y \subseteq \mathcal{P}_k(C)$ when the probability that the voter has ℓ representatives in the ultimately selected committee is at least as large on outcome X as on outcome Y for every natural number ℓ, and strictly larger for some natural number ℓ.

The stochastic dominance relation relative to an approval set $A \subseteq C$ gives a strict partial ordering of $\mathcal{P}_k(C)$. That is, the relation is *irreflexive, transitive* and *asymmetric*. In particular this means that some outcomes are incomparable.

Example 1. Consider the following scenario in which a voter i has approval set $A_i = \{a_1, a_2\}$, the set of candidates is $C = \{a_1, a_2, b, c\}$ and two possible outcomes are:

$$X = \{\{a_1, a_2\}, \{b, c\}\}$$
$$Y = \{\{a_1, b\}, \{a_1, c\}\}$$

If i is an optimistic voter, she might prefer outcome X over outcome Y because of the possibility to receive two representatives in outcome X. On the other hand, if i is pessimistic, she might prefer outcome Y over outcome X because in outcome Y she is guaranteed at least one representative. However, relative to A_i the two possible outcomes X and Y are incomparable according to the stochastic dominance definition: neither of the outcomes stochastically dominates the other. △

The example above shows that the stochastic dominance preference relation does not coincide with the preference order over $\mathcal{P}_k(C)$ that the the most optimistic voter may have, nor with the preference order over $\mathcal{P}_k(C)$ the most pessimistic voter may have. Rather, X stochastically dominates Y relative to A if and only if every utility-maximising voter with approval set A will prefer X over Y, no matter the voter's attitude towards risk. This observation can be formalised as is done in Theorem 1 below. The function u in the statement of the theorem models the degree of risk aversion of a voter. A convex function u corresponds to a risk-seeking voter whereas a concave function u corresponds to a risk-averse voter.

The stochastic dominance relation is, in some sense, a *minimal* ordering of $\mathcal{P}_k(C)$. In using this preference order we refrain from imposing any assumptions on the degree of risk aversion of the voters, but at the same time the ordering is strictly sparser than the preference order that a voter with a particular degree of risk aversion would have.

Theorem 1. (Levy, 2015) Let $u : \{0, \dots, k\} \to \mathbb{R}^+$ be a strictly increasing function. Let $X, Y \subseteq \mathcal{P}_k(C)$ and $S \subseteq C$. We have that X stochastically dominates Y subject to S if and only if $\mathbb{E}(X) > \mathbb{E}(Y)$ where $\mathbb{E}(X), \mathbb{E}(Y)$ are defined as follows:

$$\mathbb{E}(X) := \sum_{x \in \{0,\dots,k\}} \frac{|\{W \in X : |W \cap S| = x\}|}{|X|} u(x) \tag{2.1}$$

In this thesis we formalise the preferences of voters over sets of committees according to the stochastic dominance relation. There are, however, alternatives to this, and it is useful to state one alternative that has been studied in the literature in particular. This alternative is called the Kelly preference relation (Kelly, 1977). The Kelly preference relation is adopted, for instance, by Kluiving et al. (2020). It states that a voter prefers outcome X over outcome Y when she strictly prefers any committee in X over any committee in Y.

Definition 2 (*Kelly dominance*)**.** For truthful approval set $A \subseteq C$ and outcomes $X, Y \subseteq \mathcal{P}_k(C)$, X Kelly-dominates Y relative to A when $|W \cap A| > |W' \cap A|$ for all $W \in X$ and $W' \in Y$.

The Kelly preference relation provides a sparser ordering of $P_k(C)$ than the stochastic dominance relation because Kelly-dominance implies stochastic dominance. If a set of committees X Kelly-dominates a set of committees Y relative to a truthful approval set A, then a voter with this approval set strictly prefers any committee in X to any committee in Y. This means the voter will certainly receive strictly more representatives in the ultimately selected committee on outcome X than on outcome Y. In particular, for any number of representatives ℓ, the probability of receiving ℓ or more representatives on X is at least as large as on Y, and for the maximum number of representatives that the voter can receive on Y plus one, the probability of receiving (at least) this amount of representatives is strictly larger on X than on Y. Thus, Kelly dominance implies stochastic dominance.

Proposition 1. For $X, Y \subseteq \mathcal{P}_k(C)$ and $A \subseteq C$, if X Kelly dominates Y relative to A, then X stochastically dominates Y relative to A.

2.1.3 Strategic Manipulation

Having defined the preference relation with which we model voters' preferences over election outcomes, we can now define strategic manipulation with respect to this preference order. Manipulation in our setting refers to the submission of an insincere ballot. A voter has an incentive to manipulate when she can submit an insincere ballot and achieve a preferable outcome.

Definition 3. Let F be some voting rule and $(\boldsymbol{A}, k)$ some election instance. We say voter $i \in N$ has an *incentive to manipulate* on $(\boldsymbol{A}, k)$ when there is an insincere ballot $B \subseteq C$ such that $F((\boldsymbol{A}_{-i}, B), k)$ stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_i.

We say voting rule F is *susceptible to manipulation* when there is some election instance $(\boldsymbol{A}, k)$ for which some voter has an incentive to manipulate. We say F is *strategy-proof* when there is no election instance $(\boldsymbol{A}, k)$ for which some voter has incentive to manipulate.

Note that we treat the truthful approval set of a voter as the default ballot from which the voter will only deviate when there is a stochastically dominating outcome to be achieved. Moreover, the assumption underlying this definition of strategic manipulation is that voters have access to the full information about the truthful preferences of the other voters. Based on this information voters determine whether they can achieve a stochastically dominating outcome. In Chapter 4 we will revise the definition of strategic manipulation in order to account for voters having incomplete information about the preferences of the other voters.

Definition 3 gives the most general notion of strategy-proofness. That is, the full set of possible ballots is considered and only if, for all election instances and all voters, none of the possible insincere ballots achieve a stochastically dominant outcome, then the voting rule is strategy-proof. A number of the results presented in this thesis will apply to a more specific form of manipulation, namely manipulation by dropping one candidate from the truthful approval set.

Definition 4. Let F be some voting rule and $(\boldsymbol{A}, k)$ some election instance. We say voter $i \in N$ has an *incentive to manipulate* on $(\boldsymbol{A}, k)$ *by dropping one candidate* when there is some $d \in A_i$ such that $F((\boldsymbol{A}_{-i}, A_i \setminus \{d\}), k)$ stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_i.

We focus on this specific form of manipulation in some of the results because it greatly simplifies the technical analysis and because dropping one candidate is the most natural way to manipulate. As we shall see in Section 3.2, manipulation by submission of a subset of the truthful approval set is the most prevalent form of manipulation. Dropping one candidate, in turn, is the simplest form of manipulating by submission of a subset of the truthful approval set.

2.2 Axioms

In Computational Social Choice, a central method to compare the merits of different voting rules is to define formal normative properties called *axioms* and to analyse which voting rules satisfy these axioms. In this section we present a number of axioms that can be regarded as minimal requirements for our voting rules. These axioms feature in a number of manipulation results in Chapter 4. The strength of these results lies exactly in the fact that the axioms are minimal requirements. We also discuss a number of stronger axioms that have been studied in the literature with which particular voting rules can be distinguished. In the later chapters we will focus on these voting rules.

2.2.1 Basic Axioms

A basic fairness property of voting rules is that they do not give special treatment to any voter or candidate. This is captured by the two axioms *anonymity* and *neutrality* which, respectively, impose that all voters and candidates are treated equally.

Axiom 1 (*Anonymity*)**.** F satisfies *anonymity* when for all election instances $(\boldsymbol{A}, k)$ and every bijection $\pi : N \to N$ we have $F(\boldsymbol{A}, k) = F(\pi(\boldsymbol{A}), k)$, where $\pi(\boldsymbol{A})$ denotes the profile obtained after permuting the voters according to π.

Axiom 2 (*Neutrality*)**.** F satisfies *neutrality* when for all election instances $(\boldsymbol{A}, k)$ and every bijection $\pi : C \to C$ we have $\pi(F(\boldsymbol{A}, k)) = F(\pi(\boldsymbol{A}), k)$, where $\pi(\boldsymbol{A})$ denotes the profile obtained after permuting the candidates according to π.

For resolute voting rules (which always return a single winning committee) anonymity and neutrality are often given up in order to break ties between committees. However, the voting rules we consider are all irresolute and do not make use of any tie-breaking orders. All voting rules we consider in the remainder of this thesis will satisfy anonymity and neutrality. The *candidate monotonicity* axiom imposes that additional support for a candidate does not hurt the candidate.

Axiom 3 (*Candidate monotonicity*)**.** F satisfies *candidate monotonicity* when for any election instance $(\boldsymbol{A}, k)$, any voter $i \in N$ and any candidate $d \in C$ it holds that:

1. if $d \in W$ for all $W \in F(\boldsymbol{A}, k)$ then $d \in W'$ for all $W' \in F(\boldsymbol{A}^{i+\{d\}}, k)$, and
2. if $d \in W$ for some $W \in F(\boldsymbol{A}, k)$ then $d \in W'$ for some $W' \in F(\boldsymbol{A}^{i+\{d\}}, k)$.

A fundamental utilitarian concept in Social Choice Theory is Pareto efficiency. The idea of Pareto efficiency is that some outcomes are dominated in an uncontroversial sense. That is, for some outcomes there exist other outcomes that give every voter at least the same utility and some voters greater utility. Pareto efficiency requires that dominated outcomes should be avoided, as they are sub-optimal.

Axiom 4 (*Pareto efficiency*)**.** We say committee W_1 *Pareto dominates* committee W_2 on profile $\boldsymbol{A} \in \mathcal{A}(N, C)$ when:

1. every voter receives at least as much representation in W_1 as in W_2 (for all $i \in N$ we have $|A_i \cap W_1| \geq |A_i \cap W_2|$), and
2. some voter receives more representation in W_1 than in W_2 (for some $i \in N$ we have $|A_i \cap W_1| > |A_i \cap W_2|$.

F satisfies *Pareto efficiency* when, for every election instance $(\boldsymbol{A}, k)$, $F(\boldsymbol{A}, k)$ does not contain Pareto dominated committees.

Even though Pareto efficiency is a natural axiom, it is surprisingly demanding in the context of approval based committee voting. Many important voting rules do not satisfy Pareto efficiency (Lackner and Skowron, 2020b). For this reason we work with the following much weaker efficiency requirement which we call *minimal efficiency*.

Axiom 5 (*Minimal efficiency*)**.** For election instance $(\boldsymbol{A}, k)$ let $S_U = \bigcap_{i \in N} A_i$ be the set of all unanimously approved candidates and $S_A = \bigcup_{i \in N} A_i$ be the set of all approved candidates. F is *minimally efficient* when for any election instance $(\boldsymbol{A}, k)$ the following two conditions are met:

1. $S_U \subseteq W$ for all $W \in F(\boldsymbol{A}, k)$ if $|S_U| \leq k$ and
   $W \subseteq S_U$ for all $W \in F(\boldsymbol{A}, k)$ if $|S_U| > k$.
2. $W \subseteq S_A$ for all $W \in F(\boldsymbol{A}, k)$ if $|S_A| \geq k$ and
   $S_A \subseteq W$ for all $W \in F(\boldsymbol{A}, k)$ if $|S_A| < k$.

The idea behind the minimal efficiency axiom is that voting rules should (i) exhaust the set of unanimously approved candidates before electing candidates that are not unanimously approved and (ii) exhaust the set of candidates that are approved at least once before electing candidates that are not approved at all. Note that Pareto efficiency implies minimal efficiency.

The last basic axiom we introduce is the *diminishing returns* axiom.

Axiom 6 (*Diminishing returns*)**.** F satisfies *diminishing returns* when for any election instance $(\boldsymbol{A}, k)$, voter $i \in N$, candidate $d \in C \setminus A_i$ and committees $W_1, W_2 \in \mathcal{P}_k(C)$ such that $d \in W_1$, $d \in W_2$ and $|W_1 \cap A_i| > |W_2 \cap A_i|$ the following holds:

If $W_2 \in F(\boldsymbol{A}, k)$ then $W_1 \notin F(\boldsymbol{A}^{i+\{d\}}, k)$.

The rationale behind this axiom is that the value of giving a voter an additional representative is greater when the voter has fewer representatives. This is reflected in the condition the axiom imposes on the voting rule as follows. When committee W_2 is judged by the voting rule to have at least as much merit as committee W_1 and a voter has fewer representatives in W_2 than in W_1, then, if this voter would be represented by an additional candidate in both W_2 and W_1, the voting rule should judge committee W_2 to be strictly better than committee W_1. Besides the normative appeal of this axiom, we will see that it is an important axiom because it is satisfied by a broad class of voting rules.

2.2.2 Proportionality

For general approval profiles, understanding what proportionality means and defining a formal normative property that corresponds to proportionality is not a straightforward task. This is because profiles may consist of arbitrarily overlapping approval sets and for such profiles the idea of representing minority groups of voters is lost.

In order to study proportionality for approval based committee voting, several directions have been undertaken in the literature. One approach is to study the behaviour of voting rules on a class of well-structured profiles for which the notion of proportionality can be more easily captured. The argument is that if a voting rule meets our proportionality requirements on these well-structured profiles, we may expect the voting rule to also perform well for general profiles. This approach has been undertaken by Brill et al. (2018) and Lackner and Skowron (2021). An other approach is to define proportionality properties more generally using the notion of *cohesive groups*. This approach has been taken by Aziz et al. (2017). A notable third approach has been undertaken by Skowron (2018), who has defined a quantitative measure of proportionality called *proportionality degree* and has studied the guaranteed proportionality degree for various voting rules. We will briefly review the first and second approach, as they are most in line with the normative approach taken in this thesis. After this brief review we present the proportionality axiom that we will use in the remainder of this thesis.

The class of approval profiles on which the meaning of proportionality is most easily understood is the class of party-list profiles.

Definition 5. A profile $\boldsymbol{A} \in \mathcal{A}(N, C)$ is called a *party-list profile* when, for each pair of voters $i, j \in N$ we have $A_i = A_j$ or $A_i \cap A_j = \emptyset$.

For a party-list profile $\boldsymbol{A}$ the set of voters N and the set of candidates C can be divided into p disjoint parties. In this way, party-list profiles resemble parliamentary elections with political parties. Typically parliamentary elections employ a system in which voters pick and vote for a political party rather than for individual candidates within the party. The problem of assigning seats in a parliament to political parties based on the number of votes each party received is known as an *apportionment problem*. Proportional representation in apportionment problems has been extensively studied and is well-understood. Numerous apportionment methods for assigning seats in parliamentary elections are used around the world. For a comprehensive treatment of apportionment methods and proportional representation in parliamentary democracies we refer the reader to a book by Pukelsheim (2017).

A straightforward proportionality requirement for apportionment methods is that the method assigns to any party i that receives v_i votes at least $\lfloor k \times \frac{v_i}{n} \rfloor$ of the k seats. This requirement is called *lower quota*. After distributing seats to satisfy lower quota, there are typically a number of fractional seats left to distribute. Many apportionment methods satisfy lower quota but differ in the way these fractional seats are distributed. The analyses of Brill et al. (2018) and Lackner and Skowron (2021) is centered around the *D'Hondt* apportionment method. This apportionment method is considered to deliver proportional outcomes. It satisfies the lower quota property and is widely used in parliamentary elections around the world (Pukelsheim, 2017).

In the framework of approval based committee voting, party-list profiles are a very specific type of profiles. They are profiles in which the voters that belong to the same group are in full agreement with each other. We can apply the idea of representing minority groups of voters more broadly than only to groups of voters that are in complete agreement. To this end we can use the notion of *cohesive groups*, which are large enough groups of voters who are, though not necessarily in complete agreement, in agreement about approving some common candidates.

Definition 6. A group $V \subseteq N$ is ℓ-cohesive if: (i) $|V| \geq \ell \cdot \frac{n}{k}$ and (ii) $|\bigcap_{i \in V} A_i| \geq \ell$.

An ℓ-cohesive group of voters approves at least ℓ common candidates and also has a size of at least ℓ times a kth of the electorate. The idea is that such a group of voters should receive some amount of representation. In fact, we would like to impose the requirement that an ℓ-cohesive group of voters can fill ℓ times a kth of the committee spots (which is ℓ committee spots) since the group has a size of ℓ time a kth of the electorate. The voters in such a group also agree on at least ℓ candidates, so they could fill the ℓ committee spots with these candidates. Thus, we would like to require that every voter in an ℓ-cohesive group of voters receives ℓ representatives in the elected committee. However, as the example below shows, this requirement is not attainable for all election instances.

Example 2. (Lackner and Skowron, 2020a) Let N be a set of 12 voters, $C = \{a, b, c, d\}$ the set of four candidates and $k = 3$ the desired committee size. Let $\boldsymbol{A}$ be the following profile over (N, C):

$$
\begin{array}{llll}
1 \times \{a, d\} & 1 \times \{a, b\} & 1 \times \{b, c\} & 1 \times \{c, d\} \\
2 \times \{a\} & 2 \times \{b\} & 2 \times \{c\} & 2 \times \{d\}
\end{array}
$$

Since $\frac{n}{k} = \frac{12}{3} = 4$, a 1-cohesive group has a size of at least 4 voters. There are four 1-cohesive groups for the election instance $(\boldsymbol{A}, k)$, namely: $\{1, 2, 3, 4\}$, who agree on candidate a, $\{4, 5, 6, 7\}$, who agree on candidate b, $\{7, 8, 9, 10\}$, who agree on candidate c and $\{10, 11, 12, 1\}$, who agree on candidate d. In order to give each voter who belongs to a 1-cohesive group one representative, we would have to elect all four candidates. However, the desired committee size is 3 so this is not possible. △

Example 2 shows that there can be no voting rule which always outputs committees for which every voter in an ℓ-cohesive group receives ℓ representatives. However, as is shown by Aziz et al. (2017) and Sánchez-Fernández et al. (2017), there are weaker requirements for the representation of ℓ-cohesive groups that can be attained. They propose two axioms called *Extended Justified Representation* (Aziz et al., 2017) *Propotional Justified Representation* (Sánchez-Fernández et al., 2017) and show that there are voting rules that satisfy these axioms.

Axiom 7 (*Extended Justified Representation*)**.** F satisfies extended justified representation (EJR) if, for any election instance $(\boldsymbol{A}, k)$, each $W \in F(\boldsymbol{A}, k)$ and each ℓ-cohesive group of voters V there is a voter $i \in V$ such that i receives at least ℓ representatives in W. That is, $|A_i \cap W| \geq \ell$.

Axiom 8 (*Proportional Justified Representation*)**.** F satisfies proportional justified representation (PJR) if for any election instance $(\boldsymbol{A}, k)$, each $W \in F(\boldsymbol{A}, k)$ and each ℓ-cohesive group of voters V we have $|\bigcup_{i \in V} A_i \cap W| \geq \ell$.

The EJR axiom requires that for any ℓ-cohesive group there is always some voter who receives at least ℓ representatives in the outcome. The axiom does not require that every voter in an ℓ-cohesive group is represented. Strictly speaking, as long as one of the voters in the ℓ-cohesive group receives ℓ representatives, EJR is satisfied even if all other voters in the ℓ-cohesive group receive no representation at all. However, as the voters in ℓ-cohesive groups have similar preferences, the expectation is that providing ℓ representatives for one of the voters in the group will inevitably provide some representation for the other voters in the group as well. The PJR axiom was proposed by Sánchez-Fernández et al. (2017) as a weaker property than EJR because EJR is very demanding and also incompatible with a desirable property called *perfect representation*. Perfect representation requires that if for an election instances there is a way to divide the set of voters into k disjoint groups of equal size and assign to each group one candidate that represents this group, then the voting rule should output a set of candidates that achieves this. The PJR axiom also intuitively comes closer to the spirit of proportionality than EJR, because it imposes that the union of voters in an ℓ-cohesive group is represented by ℓ candidates in the outcome, rather than one voter. Note that EJR implies PJR and that for party-list instances, both EJR and PJR imply the lower quota property.

To define the proportionality axiom that we will use in the remainder of this thesis, we will turn our attention back to groups of voters who are in complete agreement, as is the case for the groups in party-list profiles.

Axiom 9 (*Minimal proportionality*)**.** F is *minimally proportional* when for any election instance $(\boldsymbol{A}, k)$ and set of candidates $S \subseteq C$ such that $|\{i \in N \mid A_i = S\}| \geq \frac{n}{k} * |S|$ it is the case that $S \subseteq W$ for all $W \in F(\boldsymbol{A}, k)$.

The minimal proportionality axiom requires that any group of voters who are in complete agreement, have a size of at least l times a kth of the electorate and who, moreover, approve of l or fewer candidates, should be represented by all the candidates they approve of. The minimal proportionality axiom is a weak axiom in the sense that it only imposes a requirement on the outcome for specific election instance $(\boldsymbol{A}, k)$, namely those which contain a group of voters as was just described. For election instance $(\boldsymbol{A}, k)$ that do not contain such groups of voters, the axiom does not impose any conditions on the outcome. We work with this particular formalisation of proportionality because it is a straightforward application of the idea of proportionality and because it simplifies the technical analysis. Furthermore, as we shall see in the subsequent chapters, even this weak form of proportionality in combination with the axioms *minimal efficiency, candidate monotonicity* and *diminishing returns* poses problems in the form of strategic manipulation.

Proposition 2. Any voting rule that satisfies Proportional Justified Representation (Axiom 8) satisfies minimal proportionality (Axiom 9).

Proof. Let $(\boldsymbol{A}, k)$ be some election instance and $S \subseteq C$ some set of candidates such that $|\{i \in N | A_i = S\}| \geq \frac{n}{k} * |S|$. This means that $V := \{i \in N | A_i = S\}$ is an $|S|$-cohesive group. Let $F(\boldsymbol{A}, k)$ denote the outcome of election $(\boldsymbol{A}, k)$ under some rule F that satisfies PJR. By PJR we have $\bigcup_{i \in V} |A_i \cap W| \geq |S|$ for all $W \in F(\boldsymbol{A}, k)$. Since $A_i = S$ for all $i \in V$ we have $S \subseteq W$ for all $W \in F(\boldsymbol{A}, k)$. $\square$

2.3 Approval-Based Committee Voting Rules

Now that we have defined our formal model and the normative properties we are interested in, we can turn our attention to the main object of our study:

the voting rules. In this section we have two goals. The first goal is to introduce a number of important proportional voting rules and discuss some of their properties. The second goals is to examine whether there are voting rules that satisfy the axioms *candidate monotonicity, minimal efficiency, diminishing returns* and *minimal proportionality.* These axioms will feature in a number of results presented in Chapter 4, and it is important to verify that there exist voting rules that satisfy them.

A broad class of approval based committee voting rules was introduced by the Danish mathematician Thorvald N. Thiele in the late 19th century (Thiele, 1895). The voting rules introduced by Thiele are based on the idea of maximizing the total satisfaction of the voters, where the satisfaction of a voter for a given committee is a function of the number of representatives the voter has in the committee.

Definition 7 (*w-Thiele methods*)**.** A w-Thiele method is parametrized by a non-decreasing function $w : \mathbb{N} \to \mathbb{R}$ with $w(0) = 0$. The score of committee W given a profile $\boldsymbol{A}$ is given by:

$$sc_w(\boldsymbol{A}, W) = \sum_{i \in N} w(|W \cap A_i|)$$

For election instance $(\boldsymbol{A}, k)$, the w-Thiele method returns the committees of size k with maximum score.

The function $w : \mathbb{N} \to \mathbb{R}$ that parametrizes the w-Thiele method represents the satisfaction of a voter as a function of the number of representatives the voter has. For any number of representatives $r \in \mathbb{N}$, the quantity $\delta(r) := w(r+1) - w(r)$ can be seen as the increase in satisfaction a voter experiences when going from r to $r+1$ representatives. If $\delta(r)$ is a decreasing function in r, then the Thiele method may, for example, favour giving a small group of voters a first representative over giving a larger group of voters a second representative. w-Thiele methods may exhibit very different properties depending on w. The following three voting rules are all w-Thiele methods.

Rule 1 (*Approval Voting* (AV))**.** AV is the w-Thiele method parametrized by $w(n) = n$

Rule 2 (*Proportional Approval Voting* (PAV))**.** Let $h(n) := \sum_{i=1}^{n} \frac{1}{i}$ be the harmonic function. *PAV* is the w-Thiele method parametrized by the function $h(n)$.

Rule 3 (*Chamberlin–Courant* (CC))**.** CC is the w-Thiele method parametrized by $w(n) = min(1, n)$.

We can also define a w-Thiele method in terms of its weight sequence $u = (\delta(0), \delta(1), \dots)$. For the three voting rules AV, PAV, and CC, the respective weight sequences are $u_{AV} = (1, 1, 1, \dots)$, $u_{PAV} = (1, \frac{1}{2}, \frac{1}{3}, \dots)$ and $u_{CC} = (1, 0, 0, \dots)$. These sequences indicate, for every number of current representatives, what the increase of satisfaction of a voter would be when receiving an additional representative. For AV, this increase is independent of the number of representatives the voter already has while for CC only the first representative adds value. For PAV the added value decreases for each additional representative, following the harmonic series. As a result, the three voting rules AV, PAV and CC, capture, respectively, the three different objectives of *individual excellence, proportional representation* and *diversity*.

Example 3. Let N be a set of 10 voters, $C = \{a, b, c, d, e\}$ the set of five candidates and $k = 3$ the desired committee size. Let $\boldsymbol{A}$ be the following profile over (N, C):

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

The candidates a, b, c, d, e are approved $6, 5, 2, 3, 4$ times, respectively. The AV rule picks the candidates that are approved most often, so the outcome under AV is $F_{AV}(\boldsymbol{A}, k) = \{a, b, e\}$. Note that this outcome leaves four voters unrepresented while giving some voters three representatives.

The outcome under CC is $F_{CC} = \{\{a, c, d\}, \{b, c, d\}\}$ because these two committees are the only committees that maximize the number of voters that receives at least one representative. Namely, by give every voter at least one representative. Note that the outcome includes both c and d which are the least approved candidates.

The outcome under PAV is $F_{PAV} = \{\{a, b, d\}\}$. This committee gives six voters two representatives and two voters one representative. The score of the committee $\{a, b, d\}$ is $sc_{PAV}(\{a, b, d\}, \boldsymbol{A}) = 6 * 1.5 + 2 * 1 = 11$ which is optimal. Note that in contrast with the outcome under AV, PAV elects candidate d instead of candidate e. The result of replacing e for d in committee $\{a, b, e\}$ is that three voters go from having three representatives to two and two voters go from having no representatives to having one representative. △

The w-Thiele methods can be seen as global optimization problems in the sense that they return sets of candidates (committees) that optimize the score

$sc_w(\boldsymbol{A}, W)$. For each w-Thiele method there also exists a local optimization problem which selects candidates in k rounds, at each round picking the candidate that maximizes the increase in the score $sc_w(\boldsymbol{A}, W)$. These local optimization variants of w-Thiele methods are usually called *sequential w-Thiele methods*. We are interested in one sequential w-Thiele method in particular; sequential Proportional Approval Voting. Aziz et al. (2015) show that computing the winning committees under the PAV is NP-hard. The sequential version of PAV, however, can be seen as a polynomial time computable approximation of PAV (Skowron et al., 2016).

Rule 4 (*Sequential Proportional Approval Voting* (seq-PAV))**.** Seq-PAV elects k candidates in k rounds. We start with the empty set $W_0 = \emptyset$ and for each round $r \in \{1, \ldots, k\}$ we construct $W_r = W_{r-1} \cup \{c\}$ where c is the candidate that maximizes the score $sc_{PAV}(W_r, \boldsymbol{A}) = \sum_{i \in V} h(|W_r \cap A_i|)$. If at any round there is a tie between candidates, some tie-breaking order over the candidates is used. The irresolute rule returns all committees that win for some tie-breaking order.

An interesting approval based committee voting rule that is not a Thiele method was introduced by the Swedish mathematician Lars Edvard Phragmén in the late 19th century (Phragmén, 1894). The voting rules introduced by Phragmén are based on the sharing of costs among voters. The idea is that including a candidate in the winning committee has a cost which is shared among the voters who approve the candidate. For a treatment of different versions of Phragmén's methods as well as historical information on both Phragmén's and Thiele's methods in English we refer the reader to a paper by Janson (2016).

Rule 5 (*Phragmén's rule*)**.** For election instance $(\boldsymbol{A}, k)$ Phragmén's rule selects k candidates one by one with the following procedure.

Each voter has a budget which starts at 0 and continuously increases as time goes on. When a group of voters with a common approved candidate reach a combined budget of 1, the candidate is added to the committee and the budget of each voter in the group is set to 0. The budget of the voters who did not approve of this candidate remains the same. This procedure continues until k candidates are elected. If at any point two candidates could be selected at the same time, some tie-breaking order over the candidates is used. The irresolute version of this rule returns all committees that are winning for some tie-breaking order.

Example 4. Let N be a set of 10 voters, $C = \{a, b, c, d, e\}$ the set of five candidates and $k = 3$ the desired committee size. Let $\boldsymbol{A}$ be the following profile over (N, C):

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

The first time a group of voters with a common candidate reach a combined budget of 1 is at $t_1 = \frac{1}{6}$. At this moment, the six voters who approve a have sufficient budget to elect a. After the selection of a, the budgets towards the remaining candidates as a function of t are given by: $b_b = 5t - 5\frac{1}{6}$, $b_c = 2t$, $b_d = 3t - 1\frac{1}{6}$, $b_e = 4t - 4\frac{1}{6}$. The first time t_2 for which one of these budgets equals 1 is $t_2 = \frac{11}{30}$, when the voters approving b have a combined budget of 1. After this, the budgets towards the remaining candidates as a function of t are: $b_c = 2t$, $b_d = 3t - 1\frac{1}{6}$, $b_e = 4t - 3\frac{11}{30} - \frac{1}{6}$. The first time t_3 for which one of the budgets equals 1 is $t_3 = \frac{7}{18}$, when the voters who approve of d have a combined budget of 1. Thus, the outcome is $F_{Phragmén}(\boldsymbol{A}, k) = \{\{a, b, d\}\}$. △

The voting rules *PAV*, *seq-PAV* and *Phragmén's rule* all exhibit attractive proportionality properties. For example, all three rules extend the D'Hondt apportionment method. That is, on party-list profiles, their outcomes correspond to the outcome of the D'Hondt apportionment method. However, among the three rules, PAV stands out in particular because it is the only rule that satisfies EJR. In fact, within the class of Thiele methods, EJR can be used to characterize PAV (Aziz et al., 2017). Even though sequential-PAV can be seen as an approximation algorithm for PAV, it does not share these proportionality properties with PAV: it neither satisfies EJR nor PJR (Aziz et al., 2017). Phragmén's rule, on the other hand, does satisfy PJR but fails to satisfy EJR (Brill et al., 2017). Out of the three rules, PAV is also the only rule that satisfies Pareto efficiency (Lackner and Skowron, 2020b). On the other hand, the computational complexity properties speak in favour of Phragmén's rule and seq-PAV, especially for applications to large elections, as computing winning committees under PAV is NP-hard. For a more detailed comparison of the proportionality properties of the three rules we refer the reader to a paper by Skowron (2018).

	EJR	PJR	D'Hondt	Pareto	Comp. complexity
Sequential-PAV	✗	✗	✓	✗	P
Phragmén's rule	✗	✓	✓	✗	P
PAV	✓	✓	✓	✓	NP-hard

Let us now turn our attention to the axioms *candidate monotonicity, minimal efficiency, diminishing returns* and *minimal proportionality.* The candidate monotonicity axiom is a special case of a more general axiom studied by Sánchez-Fernández and Fisteus (2019) which they call *support monotonicity* *without population increase*. The paper by Fernández and Fisteus provides an extensive treatment of monotonicity axioms for approval based committee voting. Among other results, they show that seq-PAV, PAV and Phragmén's rule satisfy candidate monotonicity. As was shown in Proposition 2, the minimal proportionality axiom is satisfied by all voting rules that satisfy Proportional Justified Representation. This means, in particular, that PAV and Phragmén's rule satisfy minimal proportionality.

It is easy to see that minimal efficiency is satisfied by PAV, Phragmén and seq-PAV. For Phragmén's rule, candidates that are not approved by any voter will never have sufficient budget towards their election whereas candidates that are approved at least once will at some time t have sufficient budget towards their election. Similarly, candidates that are approved unanimously will always have a sufficient budget towards their election before candidates that are not unanimously approved, independently of which candidates are already selected in the procedure, because the voters contributing to a non-unanimously approved candidate are a strict subset of the voters contributing to a unanimously approved candidate. Similarly straightforward arguments show that PAV and seq-PAV satisfy minimal efficiency.

Lastly, Diminishing Returns is satisfied by a broad subclass of w-Thiele methods, namely the w-Thiele methods for which $w : \mathbb{N} \to \mathbb{R}$ is a strictly concave function.

Definition 8. We call a function $f : \mathbb{N} \to \mathbb{R}$ *strictly concave* when for any $\ell_1, \ell_2 \in \mathbb{N}$ such that $\ell_1 > \ell_2$ we have $f(\ell_1 + 1) - f(\ell_1) < f(\ell_2 + 1) - f(\ell_2)$.

The weight sequences $u = (\delta(0), \delta(1), \dots)$ of strictly concave w-Thiele methods is such that $\delta(r)$ is a strictly decreasing function of r. This means that the satisfaction of the voters is modelled such that voters care more about their first representatives than about receiving more representation when they are already well represented. Note that PAV is a strictly concave w-Thiele method.

Theorem 2. Every strictly concave w-Thiele method satisfies Diminishing Returns (Axiom 6).

Proof. Let $w : \mathbb{N} \to \mathbb{R}$ be a non-decreasing and strictly concave function with $w(0) = 0$. Let $F(\boldsymbol{A}, k)$ denote the outcome of election instance $(\boldsymbol{A}, k)$ under the w-Thiele method parametrized by w.

Take any election instance $(\boldsymbol{A}, k)$, voter $i \in N$, candidate $d \in C \setminus A_i$ and committees $W_1, W_2 \in \mathcal{P}_k(C)$ such that $d \in W_1$, $d \in W_2$ and $|W_1 \cap A_i| > |W_2 \cap A_i|$. Suppose that $W_1 \in F(\boldsymbol{A}^{i+\{d\}}, k)$. We will show that $W_2 \notin F(\boldsymbol{A}, k)$, thereby showing the contrapositive of the implication in axiom 6. We have:

$$
\begin{aligned}
sc_w(W_1, \boldsymbol{A}) &= sc_w(W_1, \boldsymbol{A}^{i+\{d\}}) - \big(w(|W_1 \cap A_i| + 1) - w(|W_1 \cap A_i|)\big) \\
sc_w(W_2, \boldsymbol{A}) &= sc_w(W_2, \boldsymbol{A}^{i+\{d\}}) - \big(w(|W_2 \cap A_i| + 1) - w(|W_2 \cap A_i|)\big)
\end{aligned}
$$

Since w is strictly concave and $|W_1 \cap A_i| > |W_2 \cap A_i|$ we have:

$$
\big(w(|W_1 \cap A_i| + 1) - w(|W_1 \cap A_i|)\big) < \big(w(|W_2 \cap A_i| + 1) - w(|W_2 \cap A_i|)\big)
$$

Moreover, since $W_1 \in F(\boldsymbol{A}^{i+\{d\}}, k)$ we have:

$$
sc_w(W_1, \boldsymbol{A}^{i+\{d\}}) \geq sc_w(W_2, \boldsymbol{A}^{i+\{d\}})
$$

From the above equations we get:

$$
\begin{aligned}
sc_w(W_1, \boldsymbol{A}) &= sc_w(W_1, \boldsymbol{A}^{i+\{d\}}) - \big(w(|W_1 \cap A_i| + 1) - w(|W_1 \cap A_i|)\big) \\
&\geq sc_w(W_2, \boldsymbol{A}^{i+\{d\}}) - \big(w(|W_1 \cap A_i| + 1) - w(|W_1 \cap A_i|)\big) \\
&> sc_w(W_2, \boldsymbol{A}^{i+\{d\}}) - \big(w(|W_2 \cap A_i| + 1) - w(|W_2 \cap A_i|)\big) \\
&= sc_w(W_2, \boldsymbol{A})
\end{aligned}
$$

Thus, $sc_w(W_1, \boldsymbol{A}) > sc_w(W_2, \boldsymbol{A})$. We conclude that $W_2 \notin F(\boldsymbol{A}, k)$. □

The table below gives an overview of the satisfaction of the axioms. The row corresponding to the strictly concave w-Thiele methods indicates whether the axiom are satisfied for *all* strictly concave w-Thiele methods. In Appendix A.1 proofs are given for some of the entries in the table that are not discussed in this section. The question marks indicate open questions.

	C. Monotonicity	Min. Efficiency	D. Returns	Min. Prop.
S.C. w-Thiele	✓	✓	✓	✗
PAV	✓	✓	✓	✓
Seq-PAV	✓	✓	?	?
Phragmén	✓	✓	?	✓
AV	✓	✓	✗	✗

Chapter 3

Strategic Voting

In this chapter we lay out the main problem that this thesis is concerned with, namely that proportional voting rules are susceptible to strategic manipulation when voters have complete information about the preferences of the other voters. We also present a positive result for PAV, namely that on party-list profiles, PAV is not susceptible to manipulations which involve dropping one candidate.

In Section 3.1 we give an overview of the impossibility results that have been obtained in the literature. These results show that proportionality cannot be achieved in combination with strategy-proofness. In this section we also provide a number of examples of strategic manipulation for our proportional voting rules. We find that there are several types of manipulation that a voter can engage in. In Section 3.2 we present two simulation experiments in which we quantify the prevalence of these different types of manipulation for the most prominent voting rules. In Section 3.3 we present the proof of the positive result for PAV on party-list profiles.

3.1 Susceptibility to Strategic Voting

In this section we discuss the impossibility results obtained in the literature and provide a number of examples of strategic manipulation of proportional voting rules.

3.1.1 Impossibility Results

Ever since the influential work by Gibbard (1973) and Satterthwaite (1975), Social Choice Theory has been concerned with impossibility results. Gibbard and Satterthwaite showed that in classical Social Choice Theory, no voting rule can be simultaneously democratic, in the sense that every candidate has a chance to win and no voter acts as a dictator, and strategy-proof. In the setting of approval based committee elections analogous impossibility results have been obtained by Peters (2018), by Kluiving et al. (2020) and by Lackner and Skowron (2018). These studies provide formal proofs that the combination of strategy-proofness and proportionality cannot be attained. That is, that there is no voting rules that is proportional and strategy-proof.

Lackner and Skowron (2018) define strategy-proofness equivalently to the current framework (using the notion of stochastic dominance) and show that for the class of counting rules, AV is the only non-trivial rule that is strategy-proof. The class of counting rules is a broad class of voting rules which include the w-Thiele methods (Lackner and Skowron, 2021). Since AV is not a proportional voting rule, the result entails that there are no proportional counting rules that are strategy-proof.

Peters (2018) shows that for a weak form of proportionality and a weak form of strategy-proofness, there is no resolute voting rule which meets both of these properties. A notable aspect of this work is that the result relies on a computer-generated proof. Kluiving et al. (2020) use a similar proof method and show that there is no irresolute voting rule which satisfies a weak proportionality property and a weak strategyproofness property in combination with Pareto effiency. The proportionality axiom used by Kluiving et al. (2020) is weaker than minimal proportionality. Moreover, they define strategy-proofness using the notion of Kelly-dominance, which, as Proposition 1 states, implies stochastic dominance. In other words, if a voter can submit an insincere ballot to achieve a Kelly-dominant outcome, then the voter can achieve a stochastically dominant outcome. Therefore, if a voting rule is susceptible to manipulation for the Kelly preference relation, it is also susceptible to manipulation for the stochastic dominance preference relation. For these reasons, the result obtained by Kluiving et al. (2020) applies directly to our current framework. Their result shows us that there is no voting rule that is strategy-proof (in our definition of strategy-proofness), Pareto efficient and minimally proportional.

These impossibility results show us that proportional voting rules are inherently susceptible to manipulation. Strictly speaking, the results do not apply to Phargmén's rule and seq-PAV, as they are neither Pareto efficient nor counting rules. However, it is not difficult to find election instances on which seq-PAV and Phragmén's rule are also susceptible to manipulation.

3.1.2 Examples of Strategic Voting

The impossibility results tell us that proportional voting rules are generally susceptible to manipulation. It is worthwhile to look more closely at how voters can manipulate. Here we give a number of examples of manipulation for the voting rules PAV, seq-PAV and Phragmén's rule.

Example 5. Let N be a set of 20 voters, $C = \{a, b, c\}$ the set of three candidates and $k = 2$ the desired committee size. Let F be any of the voting rules PAV, seq-PAV or Phragmén's rule and let $\boldsymbol{A}$ be the following profile over (N, C):

$$10 \times \{a, b\} \quad 10 \times \{a, c\}$$

The outcome of this election for voting rule F is: $F(\boldsymbol{A}, k) = \{\{a, b\}, \{a, c\}\}$. This can be easily seen from the fact that each of these rules satisfies minimal efficiency, anonymity and neutrality. Consider now the profile $\boldsymbol{A'}$ in which the first voter drops candidate a from her approval ballot:

$$1 \times \{b\} \quad 9 \times \{a, b\} \quad 10 \times \{a, c\}$$

The outcome of election instance $(\boldsymbol{A'}, 2)$ under each of these rules now is $F(\boldsymbol{A'}, k) = \{\{a, b\}\}$. This means that the first voter achieves (for each of the rules) a stochastically dominating outcome by dropping candidate a from her approval set. △

Example 6. Let N be a set of 15 voters, $C = \{a, b, c, d\}$ the set of four candidates and $k = 3$ the desired committee size. Let F be any of the voting rules PAV, seq-PAV or Phragmén's rule and let $\boldsymbol{A}$ be the following profile over (N, C):

$$4 \times \{a, c\} \quad 4 \times \{b, c\} \quad 5 \times \{d\} \quad 1 \times \{b\} \quad 1 \times \{a, d\}$$

The outcome of the election is $F(\boldsymbol{A}, 3) = \{\{b, c, d\}\}$. Note that for five of the voters, candidate d is the only candidate that represents these voters. Consider the profile $\boldsymbol{A'}$ when the last voter drops candidate d from her approval ballot:

$$4 \times \{a, c\} \quad 4 \times \{b, c\} \quad 5 \times \{d\} \quad 1 \times \{b\} \quad 1 \times \{a\}$$

The outcome of election instance $(\boldsymbol{A'}, 3)$ under each of the rules is $F(\boldsymbol{A'}, 3) = \{\{a, c, d\}, \{b, c, d\}\}$. The first committee in this outcome gives voter 15 two representatives. This means voter 15 can achieve a stochastically dominating outcome by dropping candidate d from her approval ballot. △

The previous two examples make a connection between proportionality and strategic maniplation clear. Certain candidates are assured a spot in the winning committee, either because they are very popular or because they provide the only way to represent a minority group of voters. Voters who approve of such a candidate but also approve of other candidates which are not assured a spot in the winning committee may use this fact to put more of their voting power behind other candidates by dropping the candidate that will be elected anyway from their approval ballot. This works because proportional voting rules tend to give more weight to voters who are not well represented. This greater weight comes in the form of a larger budget for Phragmén's rule or in the form of a greater contribution to the score of a committee or candidate for PAV and seq-PAV. Essentially, for these voting rules, voters can pretend that they are not represented by certain candidates and demand to be be represented by other candidates. This type of manipulation is called *free-riding*, as the voter who manipulates in this way is represented by a candidate without having contributed to the election of this candidate. In this sense, the manipulator receives this representation for free. We give the following precise definition of free-riding for our voting rules.

Definition 9. We say a voter $i \in N$ can manipulate election $(\boldsymbol{A}, k)$ by *free-riding* if there is some $B \subsetneq A_i$ such that:

1. $F((\boldsymbol{A}_{-i}, B), k)$ stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_i and
2. for all $a \in A_i \setminus B$ and $W \in F((\boldsymbol{A}_{-i}, B), k)$ we have $a \in W$.

The second condition in the definition above expresses that all candidates that the voter drops are candidates that will certainly be elected, even when the voter drops these candidates. With this definition we deviate slightly from the way the term free-riding is generally used in the literature, as many authors use the term to refer to any manipulation by submission of a subset of the truthful approval set (for example Botan, 2021). We use the term free-riding for this specific kind of manipulation because, as the following example shows, there also exist subset manipulations for our voting rules that are conceptually different from free-riding.

Example 7. Let N be a set of 6 voters, $C = \{a, b, c, d\}$ the set of four candidates and $k = 2$ the desired committee size. Let $\boldsymbol{A}$ be the following profile over (N, C):

$$1 \times \{a, b, c\} \quad 1 \times \{a, b, d\} \quad 1 \times \{a, d\} \quad 1 \times \{a, b\} \quad 1 \times \{b, c\} \quad 1 \times \{c, d\}$$

The outcome of the election instance under PAV is $F(\boldsymbol{A}, 2) = \{ab, ac, bd\}$. Note that every candidate appears in some winning committee, but not every pair of candidates forms a winning committee. This is a result of the correlations between the candidates in the profile. For example, even though a and b are both approved four times, bd is a winning committee but ad is not. This is because out of the four voters who approve of a more voters are also represented by d than out of the four voters who approve of b. Because the rule gives more weight to voters who have less representation, this mean that, when electing d, a greater score can be achieved by adding b than by adding a.

Notice that for this outcome, voter 1 would prefer to remove the committees that contain candidate b. That is, the voter would prefer to have outcome $\{ac\}$, on which she certainly receives two representatives, over $\{ab, ac, bd\}$, on which she might only receive one representative. Since voter 1 approves of candidate b, she has a way to achieve this, namely by dropping candidate b from her approval set. The resulting profile $\boldsymbol{A}'$ is the following.

$$1 \times \{a, c\} \quad 1 \times \{a, b, d\} \quad 1 \times \{a, d\} \quad 1 \times \{a, b\} \quad 1 \times \{b, c\} \quad 1 \times \{c, d\}$$

The outcome for election instance $(\boldsymbol{A}', 2)$ under the PAV rule indeed is $F(\boldsymbol{A}', 2) = \{ac\}$. Thus, voter 1 can achieve a stochastically dominating outcome by dropping candidate b from her approval ballot. $\triangle$

The manipulation in Example 7 has a different character than the manipulation in Examples 5 and 6, even though in both cases the voter drops a candidate to achieve a stochastically dominating outcome. In Example 7 the voter drops a candidate with the goal to *prevent* it's election. This is fruitful because the candidate is associated with committees the voter dislikes. In Examples 5 and 6 on the other hand, the voter drops a candidate because it will be elected anyway. That is, the voter manipulated by free-riding. Examples such as Example 7 can also be found for our two sequential rules. Moreover, as the following example shows, a similar mechanism can be used to manipulate by adding a candidate to the approval ballot, as oppossed to dropping a candidate.

Example 8. Let N be a set of 11 voters, $C = \{a, c, d, e\}$ the set of four candidates and $k = 2$ the desired committee size. Let $\boldsymbol{A}$ be the following profile over (N, C):

$$2 \times \{a, d\} \quad 2 \times \{a, e\} \quad 2 \times \{c, d\} \quad 3 \times \{c, e\} \quad 1 \times \{a\} \quad 1 \times \{d\}$$

For each of the rules PAV, seq-PAV and Phragmén, the outcome of $(\boldsymbol{A}, k)$ is the same, namely, $F(\boldsymbol{A}, k) = \{ac, de\}$. Note that in the profile every candidate is approved five times. Moreover, there are no voters who approve of both a and c and there are no voters who approve of both d and e, while for any other pair of candidates there are voters who approve of both candidates. The outcome of the election is a result of this structure.

Notice that the 10th voter prefers the first committee in the outcome over the second committee in the outcome. This gives the voter an opportunity to manipulate by adding candidate c to her approval ballot. The resulting profile $\boldsymbol{A'}$ is:

$$2 \times \{a, d\} \quad 2 \times \{a, e\} \quad 2 \times \{c, d\} \quad 3 \times \{c, e\} \quad 1 \times \{a, c\} \quad 1 \times \{d\}$$

Indeed, the outcome of election instance $(\boldsymbol{A'}, 2)$ under each of the rules is $F(\boldsymbol{A'}, 2) = \{ac\}$. Thus, the 10th voter can achieve a stochastically dominating outcome by adding candidate c to her ballot. △

3.2 The Prevalence of Manipulability

The examples and impossibility results discussed in Section 3.1 reveal that PAV, seq-PAV and Phragmén's rule are susceptible to manipulation for at least some election instances. In order to judge the severity of the situation it is useful to quantify the prevalence of manipulability for such voting rules. In the case that only a few rare election instances are the cause of the impossibility results, we might find the risk of manipulation acceptable, whereas in the case that most election instances are manipulable we might find this risk unacceptable. Beyond the prevalence of manipulable profiles it is interesting to quantify the prevalence of the different kinds of manipulations that have been identified in Section 3.1.

In this section we present two simulation experiments in which the prevalence of different types of manipulation is measured for the voting rules PAV, seq-PAV and Phragmén's rule. In the first experiment we randomly generate election instance $(\boldsymbol{A}, k)$ for an election scenario (N, C, k) and compute, for five categories of manipulation, the relative number of election instances which are susceptible to this type of manipulation. In the second experiment we use the same categories of manipulation and randomly generate profile-voter pairs $(\boldsymbol{A}, i)$ to measure, for each category of manipulation, the relative number of profile-voters pairs for which the voter has opportunity to manipulate with this type of manipulation.

In both experiments, the profiles are generated by randomly sampling approval sets for each voter using the uniform probability distribution over the full set of possible approval sets $\mathcal{P}(C) \setminus \{\varnothing\}$. This sampling distribution is known in Social Choice Theory as the *impartial culture* assumption. Even-though this probability distribution over preferences does not accurately model real-world preferences of voters, it is commonly used in simulation studies. For example, Lackner and Skowron (2018) use a similar sampling distribution to study strategic voting in approval-based committee voting.

In these experiments we are interested in quantifying the prevalence of the different types of manipulations identified in Section 3.1. We want to know how prevalent free-riding manipulations are compared to manipulations that do not rely on free-riding. We are also interested in the prevalence of manipulations that rely on adding candidates to the approval ballot. For voting rule F, election instance $(\boldsymbol{A}, k)$ over (N, C, k) and voter $i \in N$ we define the following five categories in which a ballot $B \subseteq C$ can fall.

- Ballot B provides a manipulation for voter i on $(\boldsymbol{A}, k)$.
- Ballot B provides a *subset* manipulation for voter i on $(\boldsymbol{A}, k)$.
- Ballot B provides a manipulation for voter i on $(\boldsymbol{A}, k)$ which is *not* based on a subset manipulation.
- Ballot B provides a *free-riding* manipulation for voter i on $(\boldsymbol{A}, k)$
- Ballot B provides a manipulation for voter i on $(\boldsymbol{A}, k)$ which is *not* based on free-riding.

Category one contains all ballots that give the voter a stochastically dominating outcome compared to the truthful outcome. The other four categories contain more specific manipulations. A ballot is a subset manipulation when it is a subset of the voter's approval set and provides a stochastically dominating outcome. Free-riding manipulations are defined in Definition 9. The two remaining categories require an explanation about when we take a manipulation to be based on subset manipulation or based on free-riding. Consider the following example, which is a modified version of Example 5.

Example 9. Let N be a set of 22 voters, $C = \{a, b, c, e, f, g\}$ the set of six candidates and $k = 2$ the desired committee size. Let F be any of the voting rules PAV, seq-PAV or Phragmén's rule and let $\boldsymbol{A}$ be the following profile:

$$1 \times \{a, b, e\} \quad 1 \times \{a, c, f\} \quad 10 \times \{a, b\} \quad 10 \times \{a, c\}$$

Note that candidates e and f are only approved once and candidate g is not approved at all. The outcome of this election is $F(\boldsymbol{A}, 2) = \{ab, ac\}$. The first voter can manipulate by dropping candidate a from her approval ballot. The resulting outcome is $F((\boldsymbol{A}_{-1}, \{b, e\}), k) = \{ab\}$. Note that this manipulation is a subset manipulation and, more specifically, a manipulation by free-riding.

Now, the first voter can achieve the same insincere outcome by submitting ballot $B_1 = \{b, e, g\}$ since candidate g is so unpopular that, even if the voter adds candidate g to her ballot, g will not be part of any winning committee. The same insincere outcome can also be achieved by submitting ballot $B_2 = \{b\}$ which drops candidate a and e from the sincere approval set because candidate e was certainly not elected anyway.

The manipulations with ballots B_1 and B_2 are based on the initial free-riding manipulation by dropping candidate a. Additionally dropping candidate e or adding candidate g is, in some sense, irrelevant. △

Example 9 shows that there are cases in which a voter can manipulate by free-riding and as a result can also manipulate by submitting a ballot that is not a subset or does not fit the definition of free-riding. However, in the simulation experiments, we are interested in quantifying the prevalence of manipulations which rely on adding candidates or in which candidates are dropped in order to prevent their election, as in Examples 7 and 8 of the previous section. We are not interested in including ballots such as B_1 and B_2 in Example 9 in these figures. Therefore, we define what it means for a manipulation to be based on a subset manipulation or on a free-riding manipulation as follows.

Definition 10. For voting rule F, election instance $(\boldsymbol{A}, k)$ and voter $i \in N$ we say ballot B provides a manipulation for voter i which is *based on subset manipulation* when:

1. $F((\boldsymbol{A}_{-i}, B), k)$ stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_i and
2. for all $a \in B \setminus A_i$ and all $W \in F((\boldsymbol{A}_{-i}, B), k)$ we have $a \notin W$.

Definition 11. For voting rule F, election instance $(\boldsymbol{A}, k)$ and voter $i \in N$ we say ballot B provides a manipulation for voter i which is *based on free-riding* when:

1. $F((\boldsymbol{A}_{-i}, B), k)$ stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_i and
2. for all $a \in B \setminus A_i$ and $W \in F((\boldsymbol{A}_{-i}, B), k)$ we have $a \notin W$ and
3. for all $a \in A_i \setminus B$ we have:
   (i) $a \in W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$ or

(ii) $a \notin W$ for all $W \in F(\boldsymbol{A}, k)$

Below the results of both simulation experiments are presented for the PAV rule for election instances with five candidates, a desired committee size of three and several sizes of the electorate. Analogous data for sequential PAV and for Phragmén's rule are presented in Appendix B.1.

	Susceptible	Subset	NON Subset	Free-riding	NON Free-riding
n=5	0.56	0.56	0.06	0.38	0.4
n=10	0.68	0.68	0.09	0.58	0.33
n=15	0.63	0.63	0.09	0.55	0.3
n=20	0.59	0.59	0.08	0.55	0.24

Figure 3.1: For every type of manipulation the relative number of election instances $(\boldsymbol{A}, k)$ for which some voter can manipulate PAV for $m = 5$, $k = 3$ and different values of n based on 1000 instances. The columns 'NON Subset' and 'NON Free-riding' contain the figures for manipulations that are not based on subset manipulation and not based on free-riding, respectively.

Figure 3.1 shows, for each category of manipulation, the relative number of election instances $(\boldsymbol{A}, k)$ for which some voter can manipulate the PAV rule. The first column in the table shows the relative number of election instances for which some voter can manipulate with some insincere ballot B. We observe that the prevalence of manipulable election instance for this election scenario is significant. Moreover, election instances that are susceptible to subset manipulations are significantly more prevalent than election instances that are susceptible to manipulations that are not based on subset manipulations. Similarly, election instances susceptible to free-riding are significantly more prevalent than the ones susceptible to non-free-riding based manipulations. Figure 3.2 below shows the result of the second experiment for PAV on the same election scenario.

	Opportunity	Subset	NON Subset	Free-riding	NON Free-riding
n=5	0.2	0.2	0.01	0.15	0.1
n=10	0.18	0.18	0.01	0.15	0.06
n=15	0.14	0.14	0.01	0.12	0.03
n=20	0.12	0.12	0.01	0.11	0.02
n=25	0.11	0.11	0.0	0.1	0.02
n=30	0.11	0.11	0.0	0.1	0.02
n=35	0.11	0.11	0.0	0.1	0.01
n=40	0.1	0.1	0.01	0.09	0.01

Figure 3.2: The relative number of instances $(\boldsymbol{A}, i)$ for which i has incentive to manipulate, for PAV with $m = 5$, $k = 3$ and different values of n based on 2000 instances. The columns 'NON Subset' and 'NON Free-riding' contain the figures for manipulations that are not based on subset manipulation and not based on free-riding, respectively.

We observe that the disparity between the prevalence of subset and non-subset based manipulations and between free-riding and non-free-riding based manipulations is even greater for the second experiment. In experiment 1, free-riding manipulation are around twice as prevalent as the non-free-riding based manipulations. For experiment 2, the prevalence of free-riding manipulations is around ten times that of the non-free-riding based manipulations. This greater disparity in the second experiment can be explained by the fact that when an election instance is manipulable by free-riding there are generally many voters who can manipulate. When an election instance is manipulable by non-free-riding, there may generally be only a few voters with specific approval sets who can manipulate. The experimental results for sequential PAV and Phragmén's rule show similar behavior. However, for these two sequential rules, the prevalence of manipulations that are not based on subset manipulations or not based on free-riding are considerably higher than for PAV. Further analysis would be required to understand why the sequential rules provide more opportunity for these two types of manipulation.

3.3 Party-list Profiles

In this section we show that PAV is immune to manipulations by dropping one candidate on the class of party-list profiles. Related results have been obtained by Botan (2021), who shows that Thiele methods are immune to three types of manipulations on party-list profiles. The preferences relation used in that study, however, is different from the stochastic dominance relation.

The examples of manipulation given in Section 3.1 show two mechanisms by which voters can manipulate. The first is by free-riding. That is, by dropping candidates that will certainly be elected anyway in order to give more support to other candidates. The second mechanism relies on specific correlations between candidates in the profile that result in correlations between candidates in the outcome. For such outcomes a voter may be able to pick a subset of the outcome she prefers to the actual outcome by increasing or decreasing the support for certain candidates. Neither of these two types of manipulation seem possible on party-list profiles for the following reasons. Recall that party-list profiles are well-structured profiles on which the candidates can be divided into parties and each voter supports one of those parties. That is, a voter approves of all candidates in a party and does not approve of any candidate outside of this party. The voting rules PAV, seq-PAV and Phragmén's rule will treat all candidates in a party equally. Therefore, if there is a candidate who is certainly elected then all candidates in the party are certainly elected. Any voter who approves of a candidate that is certainly elected will therefore be completely satisfied with the outcome and will not be able to achieve a better outcome. For this reason, manipulation by free-riding cannot be an option on party-list profiles. Moreover, even though candidates in a party-list profile are highly correlated - two candidates are either always approved together or never approved together - there are no voters who can exploit correlations in the outcome as is done in the examples in Section 3.1. This is because for any voter and any winning committee that gives this voter, say, l representatives, replacing these l representatives with different candidates from the same party will also result in a winning committee. As a result of this symmetry in the outcome, no voter will be able to pick a preferable subset of the outcome by dropping a candidate. Theorem 3 confirms these intuitions for the PAV rule. The theorem states that on party-list profiles, PAV is immune to manipulations that involve dropping one candidate. The proof of the theorem makes use of the following technical lemma.

Lemma 1. Let $r : \mathbb{N} \to [0, 1]$ be some decreasing function in n, $k \in \mathbb{N}$ some natural number and $N : \{0, \dots, k\} \to \mathbb{N}$ an arbitrary function. Then for any $l \in \mathbb{N}$ such that $l \leq k$ it is the case that:

$$\frac{\sum_{j=0}^{l} N(j)r(j)}{\sum_{j=0}^{k} N(j)r(j)} \geq \frac{\sum_{j=0}^{l} N(j)}{\sum_{j=0}^{k} N(j)}$$

Proof. For a proof of this lemma see Appendix A.2. □

Theorem 3. Let (N, C, k) be some election scenario and $\boldsymbol{A}$ a party-list profile over (N, C). For PAV there is no voter who can manipulate by dropping one candidate from her approval ballot.

Proof. Let (N, C, k) be some election scenario and $\boldsymbol{A}$ a party-list profile over (N, C). Let $\boldsymbol{A'} = (\boldsymbol{A}_{-i}, B)$ with $B = A_i \setminus \{c\}$ for some $c \in A_i$. That is, $\boldsymbol{A'}$ is the approval profile in which voter i has dropped candidate c from her ballot.

Let $F(\boldsymbol{A}, k)$ denote the set of k-size committees returned by the PAV rule for the election instance $(\boldsymbol{A}, k)$. Let $F(\boldsymbol{A'}, k)$ denote the set of k-size committees returned by the PAV rule for the election instance $(\boldsymbol{A'}, k)$. We will show that for all $l \in \mathbb{N}$ such that $l \leq k$:

$$\frac{|\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| \geq l\}|}{|F(\boldsymbol{A}, k)|} \geq \frac{|\{W \in F(\boldsymbol{A'}, k) : |W \cap A_i| \geq l\}|}{|F(\boldsymbol{A'}, k)|} \tag{3.1}$$

Note firstly that the following holds:

$$\text{For all } W \in \mathcal{P}_k(C) \text{ we have } sc_{PAV}(\boldsymbol{A}, W) \geq sc_{PAV}(\boldsymbol{A'}, W) \tag{3.2}$$

More specifically:

$$\text{For all } W \in \mathcal{P}_k(C) \text{ with } c \in W \text{ we have } sc_{PAV}(\boldsymbol{A}, W) > sc_{PAV}(\boldsymbol{A'}, W) \tag{3.3}$$

$$\text{For all } W \in \mathcal{P}_k(C) \text{ with } c \notin W \text{ we have } sc_{PAV}(\boldsymbol{A}, W) = sc_{PAV}(\boldsymbol{A'}, W) \tag{3.4}$$

Moreover, we have the following:

$$\begin{gathered}\text{For all } W \in \mathcal{P}_k(C), \text{ all } v \in N \text{ and all } S \subseteq A_v \text{ such that } |S| = |W \cap A_v| \\ \text{we have } sc_{PAV}(\boldsymbol{A}, W) = sc_{PAV}(\boldsymbol{A}, (W \setminus A_v) \cup S)\end{gathered} \tag{3.5}$$

That is, since $\boldsymbol{A}$ is a party-list profile, for any committee $W \in \mathcal{P}_k(C)$ with some candidates from, say, party A_v, replacing those candidates with different candidates from the same party A_v does *not* change the PAV-score of the committee. Since PAV selects the committees with maximum PAV-score we have as a result:

$$\begin{gathered}\text{For all } l \in \mathbb{N} \text{ and all } W \in F(\boldsymbol{A}, k) \text{ such that } |W \cap A_i| = l \text{ we have:} \\ (W \setminus A_i) \cup S \in F(\boldsymbol{A}, k) \text{ for all } S \in \mathcal{P}_l(A_i)\end{gathered} \tag{3.6}$$

We can now distinguish two cases: (i) for all $W \in F(\boldsymbol{A}, k)$ we have $A_i \subseteq W$ or (ii) there exists a $W \in F(\boldsymbol{A}, k)$ such that $|W \cap A_i| < |A_i|$. In case (i) there can clearly be no set of committees that stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_i. Condition (3.1) holds for all $X \subseteq \mathcal{P}_k(C)$, and thus also for $F(\boldsymbol{A'}, k)$. Next, let us look at case (*ii*). Note that this case implies that there is some $W \in F(\boldsymbol{A}, k)$ such that $c \notin W$. This fact follows from (3.6) above. Now, given that the PAV-score of no committee is higher with respect to $\boldsymbol{A'}$ than with respect to $\boldsymbol{A}$ (this is inequality (3.2) above), this means that the maximum PAV-score with respect to $\boldsymbol{A'}$ is the same as the maximum PAV-score with respect to $\boldsymbol{A}$. As a result of this and, respectively, (3.3) and (3.4) we have:

For all $W \in F(\boldsymbol{A}, k)$ such that $c \in W$ we have $W \notin F(\boldsymbol{A'}, k)$ and
for all $W \in F(\boldsymbol{A}, k)$ such that $c \notin W$ we have $W \in F(\boldsymbol{A'}, k)$

Moreover, by (3.2) and the fact that the maximum PAV-score with respect to $\boldsymbol{A'}$ is the same as with respect to $\boldsymbol{A}$ we have:

For all $W \in \mathcal{P}_k(C)$ such that $W \notin F(\boldsymbol{A}, k)$ we have $W \notin F(\boldsymbol{A'}, k)$

In other words, $F(\boldsymbol{A'}, k)$ is exactly $F(\boldsymbol{A}, k)$ but with all committees containing c dropped. It remains to be shown that this fact together with (3.6) implies (3.1). To this end, define $d(l)$ as follows:

$$d(l) := \frac{|\{S \subseteq A_i : |S| = l \text{ and } c \in S\}|}{|\{S \subseteq A_i : |S| = l\}|} \tag{3.7}$$

That is, $d(l)$ is the ratio of l-size subsets of A_i that contain candidate c. We established that, for any committee $W \in F(\boldsymbol{A}, k)$ with overlap size l with A_i, every $(W \setminus A_i) \cup S$ is in $F(\boldsymbol{A}, k)$ for $S \subseteq A_i$ such that $|S| = l$. This means that $d(l)$ is the ratio of committees that will be dropped going from $F(\boldsymbol{A}, k)$ to $F(\boldsymbol{A'}, k)$ for every overlap size l. Thus, we have the following:

$$\begin{aligned} |F(\boldsymbol{A}, k)| &= \sum_{j=0}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}| \\ |F(\boldsymbol{A'}, k)| &= \sum_{j=0}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}| \times (1 - d(j)) \end{aligned} \tag{3.8}$$

Also, for any $l \in \mathbb{N}$:

$$\begin{aligned} &|\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| \geq l\}| = \sum_{j=l}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}| \\ &|\{W \in F(\boldsymbol{A'}, k) : |W \cap A_i| \geq l\}| = \\ &\qquad \sum_{j=l}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}| \times (1 - d(j)) \end{aligned} \tag{3.9}$$

Thus, showing (3.1) reduces to showing that for all $l \in \mathbb{N}$ such that $l \leq k$:

$$\frac{\sum_{j=l}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|}{\sum_{j=0}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|} \geq \frac{\sum_{j=l}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|(1 - d(j))}{\sum_{j=0}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|(1 - d(j))}$$

Or, equivalently, for all $l \in \mathbb{N}$ such that $l \leq k$:

$$\frac{\sum_{j=0}^{l} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|}{\sum_{j=0}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|} \leq \frac{\sum_{j=0}^{l} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|(1 - d(j))}{\sum_{j=0}^{k} |\{W \in F(\boldsymbol{A}, k) : |W \cap A_i| = j\}|(1 - d(j))} \tag{3.10}$$

Now, crucially, $d(l)$ is an increasing function in l. This follows by the following argument:

$$d(l) := \frac{|\{S \subseteq A_i : |S| = l \text{ and } c \in S\}|}{|\{S \subseteq A_i : |S| = l\}|} = \frac{\binom{|A_i|-1}{l-1}}{\binom{|A_i|}{l}} = \frac{l}{|A_i|}$$

Thus, $r(l) := (1 - d(l))$ is a *decreasing* function in l. By Lemma 1 it follows that (3.10) holds. We conclude that $F(\boldsymbol{A'}, k)$ does not stochastically dominate $F(\boldsymbol{A}, k)$ relative to A_i. □

3.4 Discussion

In this chapter we gave a detailed description of the problem that this thesis is concerned with. This problem is the inherent susceptibility to strategic manipulation of proportional voting rules and, moreover, the significant prevalence of opportunity to manipulate.

In Section 3.1 we reviewed the impossibility results obtained in the literature. These results show us that, given some additional constraints such as Pareto efficiency or being a counting rule, there are no voting rules that are both proportional and strategy-proof. Strictly speaking, the impossibility result in the literature do not apply to sequential PAV or Phragmén's rule. However, we gave a number of example of strategic manipulation that confirm that sequential PAV and Phragmén's rule are also susceptible to strategic manipulation. The examples presented in the second part of Section 3.1 show that the approval-based model in combination with proportional voting rules gives rise to a variety of ways to strategically misrepresent ones preferences. We showed that, for the three proportional voting rules we study, there are instance for which a voter can manipulate by *adding* candidates to her approval set and there are also instances for which the a voter can manipulate by *dropping* candidates from her approval set. We have given a precise definition of manipulation by free-riding for our proportional voting rules. We have also seen that there are manipulations that are conceptually different from free-riding and rely on correlations between candidates in the profile.

In Section 3.2 we performed simulation experiments to measure the prevalence of the different types of manipulation. The results show us that free-riding is, for all three of the voting rules considered, a significant type of manipulation in terms of the prevalence. Moreover, manipulations that rely on adding candidates are relatively rare for the PAV rule. For sequential PAV and Phragmén's rule the experiments show that manipulations that are based on adding candidates are more prevalent than for PAV. However, for these rules as well, on the election scenarios studied, the prevalence of subset manipulations is significantly higher than for non- subset based manipulations. The results of the experiments therefore justify our focus of attention to subset-manipulations. In particular, in a number of formal results in Chapter 4, we will focus on manipulation by dropping *one* candidate, as this is the simplest form of subset manipulation.

In Section 3.3 we presented a formal proof that shows that on party-list profiles, PAV is not susceptible to manipulations which involve dropping one candidate from the approval ballot. By showing that, on this class of well-structured profiles, strategic manipulation by dropping one candidate is not possible, this result helps us understand better the mechanism behind strategic manipulations in our framework. However, the domain of party-list profiles is very restrictive. The value of approval-based committee voting lies exactly in the fact that it gives voters more freedom than voters have in conventional parliamentary elections with political parties. That is, the value of approval-based committee voting is that voters can approve any candidate they want and that candidates do not need to be organised in parties.

Chapter 4

Strategic Voting under Incomplete Information

In this chapter we extend the model of approval-based committee voting to account for strategic manipulation when voters have incomplete information. We follow the model developed by Reijngoud and Endriss (2012), which is built on the notion of *information functions*. We define two information functions that capture natural types of information that voters may have about the preferences of their fellow voters.

In Section 4.2 we present two negative results which show that voting rules that satisfy our minimal requirements *candidate monotonicity*, *minimal efficiency*, *diminhising returns* and *minimal proportionality* remain susceptible to manipulation when voters have limited information about the preferences of other voters. In Section 4.3 we present the results of a simulation experiment in which the prevalence of manipulation under incomplete information is measured. In Section 4.4 we present two positive results that show (partial) information barriers can be achieved for the PAV rule. We conclude this chapter with a discussion of these results and the assumptions in our model.

4.1 Information Functions

In actual voting scenarios, voters may have different kinds of information about the preferences of the other voters. For example, a voter may know who the most approved candidate is or know which candidates some of the other voters approve of. We model the different kinds of information a voter may have formally through the notion of *information functions*. An information function $\pi : N \times \mathcal{A}(N, C) \to \mathcal{I}$ maps voters and approval profiles to pieces of information in the set $\mathcal{I}$. An information function represents the available information to each voter given a truthful approval profile. For the purpose of readability we will write $\pi_i(\boldsymbol{A})$ for the information voter i has given profile $\boldsymbol{A}$. In this thesis we will work with the information functions listed below.

1. *Profile.* The Profile information function returns the complete approval profile for each voter: $\pi_i(\boldsymbol{A}) = \boldsymbol{A}$ for all $i \in N$ and $\boldsymbol{A} \in \mathcal{A}(N, C)$.

2. *Zero.* The Zero information function returns no information for each voter: $\pi_i(\boldsymbol{A}) = 0$ for all $i \in N$ and $\boldsymbol{A} \in \mathcal{A}(N, C)$.

3. *Approval Score.* The Approval Score information function returns the approval score $as(c)$ of each candidate c:

   $\pi_i(\boldsymbol{A}) = \big(as(c_1, \boldsymbol{A}), \ldots, as(c_m, \boldsymbol{A})\big)$ for all $i \in N$ and $\boldsymbol{A} \in \mathcal{A}(N, C)$.

   The approval score $as(c, \boldsymbol{A})$ of candidate $c \in C$ is defined as the number of voters who approve of c on profile $\boldsymbol{A}$: $as(c, \boldsymbol{A}) = |\{j \in N | c \in A_j\}|$.

4. *t-Profile* returns the approval sets of the first t voters of the approval profile, skipping the voter who receives the information:

$$\pi_i(\boldsymbol{A}) = \begin{cases} \big(A_1, \ldots, A_{i-1}, A_{i+1}, \ldots, A_{t+1}\big) \text{ for all } \boldsymbol{A} \in \mathcal{A}(N, C) \text{ and } i \in N \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{such that } i \leq t \\ \big(A_1, \ldots, A_t\big) \text{ for all } \boldsymbol{A} \in \mathcal{A}(N, C) \text{ and } i \in N \text{ such that } i > t \end{cases}$$

The Approval Score IF and the t-Profile IF represent two natural kinds of information voters may have. The Approval Score IF may correspond to situations in which voters receive information about the popularity of the candidates through polls, whereas the t-Profile information function models, for example, situations in which voters know the complete preferences of their friends. Strictly speaking, the t-Profile IF defined above is an information function schema rather than an information function. For each $t \in \{1, \ldots, n-1\}$ the definition above gives an information function. Note that the Profile IF and Zero IF are particular cases of the t-Profile IF schema, lying on the two extremes.

Now that we have defined information functions, we can define the second central concept, which is the concept of *information sets*. An information set contains the partial profiles that a voter considers possible when the voter has information $\pi_i(\boldsymbol{A})$. We denote by $\mathcal{W}_i^{\pi(\boldsymbol{A})}$ the information set of voter i on profile $\boldsymbol{A}$ induced by π. It is defined as follows:

$$\mathcal{W}_i^{\pi(\boldsymbol{A})} := \{\boldsymbol{A'}_{-i} \in \mathcal{A}(N \setminus \{i\}, C) \mid \pi_i(\boldsymbol{A'}_{-i}, A_i) = \pi_i(\boldsymbol{A})\}$$

We can think of profile $\boldsymbol{A}$ as the actual world and of $\mathcal{W}_i^{\pi(\boldsymbol{A})}$ as the set of worlds i considers possible. Voter i considers the profiles in $\mathcal{W}_i^{\pi(\boldsymbol{A})}$ possible upon learning information $\pi_i(\boldsymbol{A})$ because these profiles are consistent with the information she has. Note that implicit in this model is that every voter knows how many voters will participate in the election. As Reijngoud and Endriss (2012) remark, $\mathcal{W}$ satisfies the basic properties of a knowledge operator: reflexivity, symmetry and transitivity. For all information functions π, all voters $i \in N$ and any approval set $A_i \subseteq C$ we have the following:

- (REF) $\boldsymbol{A}_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A}_{-i},A_i)}$ for any $\boldsymbol{A}_{-i} \in \mathcal{A}(N \setminus \{i\}, C)$.
- (SYM) $\boldsymbol{A}_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{B}_{-i},A_i)}$ implies $\boldsymbol{B}_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A}_{-i},A_i)}$ for any $\boldsymbol{A}_{-i}, \boldsymbol{B}_{-i} \in \mathcal{A}(N \setminus \{i\}, C)$.
- (TRA) $\boldsymbol{A}_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{B}_{-i},A_i)}$ and $\boldsymbol{B}_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{C}_{-i},A_i)}$ implies $\boldsymbol{A}_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{C}_{-i},A_i)}$ for any $\boldsymbol{A}_{-i}, \boldsymbol{B}_{-i}, \boldsymbol{C}_{-i} \in \mathcal{A}(N \setminus \{i\}, C)$.

Reflexivity entails that a voter always considers the actual profile $\boldsymbol{A}_{-i}$ possible. Symmetry and transitivity together entail that a voter would have the same information set were the actual profile any of the profiles she considers possible.

Now that we have defined the central concepts *information function* and *information sets*, we can give our definition of manipulation under partial information.

Definition 12 (*π-manipulation*)**.** Let F be some voting rule, $(\boldsymbol{A}, k)$ some election instance, $i \in N$ some voter and π some information function. We say i has *incentive to π-manipulate* on $(\boldsymbol{A}, k)$ when there is an insincere ballot $B \subseteq C$ such that:

1. There is a profile $\boldsymbol{A}^*_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A})}$ for which $F((\boldsymbol{A}^*_{-i}, B), k)$ stochastically dominates $F((\boldsymbol{A}^*_{-i}, A_i), k)$ relative to A_i.
2. There is no profile $\boldsymbol{A}'_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A})}$ for which $F((\boldsymbol{A}'_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}'_{-i}, B), k)$ relative to A_i.

We say voting rule F is *susceptible to π-manipulation* when there is some election instance $(\boldsymbol{A}, k)$ on which some voter has incentive to π-manipulate. We say voting rule F is *π-strategy-proof* when F is not susceptible to π-manipulation.

This definition reflects our assumption that voters are *cautious* when it comes to manipulating under incomplete information. Note that, for the Profile information function, the definition of π-manipulation reduces to the definition of manipulation under complete information given in Chapter 2.

Definition 13. We say information function π is *at least as informative as* information function σ when $\mathcal{W}_i^{\pi(\boldsymbol{A})} \subseteq \mathcal{W}_i^{\sigma(\boldsymbol{A})}$ for all voters $i \in N$ and all profiles $\boldsymbol{A} \in \mathcal{A}(N, C)$.

Lemma 2. Let π and σ be information functions such that π is at least as informative as σ. If rule F is susceptible to σ-manipulation, then F is susceptible to π-manipulation.

Proof. Suppose rule F is susceptible to σ-manipulation. This means there are $\boldsymbol{A} \in \mathcal{A}(N, C)$, $k \in \mathbb{N}$, $i \in N$ and $B \subseteq C$ such that: (i) there is $\boldsymbol{A}^*_{-i} \in \mathcal{W}_i^{\sigma(\boldsymbol{A})}$ such that $F((\boldsymbol{A}^*_{-i}, B), k)$ stochastically dominates $F(\boldsymbol{A}^*_{-i}, A_i), k)$ relative to A_i and (ii) there is no $\boldsymbol{A}'_{-i} \in \mathcal{W}_i^{\sigma(\boldsymbol{A})}$ such that $F((\boldsymbol{A}'_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}'_{-i}, B), k)$ relative to A_i.

Consider $\mathcal{W}_i^{\pi(\boldsymbol{A}^*_{-i}, A_i)}$. Since the information sets satisfy (REF) we have $\boldsymbol{A}^*_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A}^*_{-i}, A_i)}$. Since π is at least as informative as σ (by assumption), we have $\mathcal{W}_i^{\pi(\boldsymbol{A}^*_{-i}, A_i)} \subseteq \mathcal{W}_i^{\sigma(\boldsymbol{A}^*_{-i}, A_i)}$. Moreover, by the properties (SYM) and (TRANS) of information sets and the fact that $\boldsymbol{A}^*_{-i} \in \mathcal{W}_i^{\sigma(\boldsymbol{A})}$ we have $\mathcal{W}_i^{\sigma(\boldsymbol{A}^*_{-i}, A_i)} = \mathcal{W}_i^{\sigma(\boldsymbol{A})}$. Thus, we have that $\boldsymbol{A}^*_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A}^*_{-i}, A_i)} \subseteq \mathcal{W}_i^{\sigma(\boldsymbol{A})}$. This means that for profile $(\boldsymbol{A}^*_{-i}, A_i)$, $k \in \mathbb{N}$, $i \in N$ and $B \subseteq C$ we have (i) $\boldsymbol{A}^*_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A}^*_{-i}, A_i)}$ such that $F((\boldsymbol{A}^*_{-i}, B), k)$ stochastically dominates $F(\boldsymbol{A}^*_{-i}, A_i), k)$ relative to A_i and (ii) there is no $\boldsymbol{A}'_{-i} \in \mathcal{W}_i^{\pi(\boldsymbol{A}^*_{-i}, A_i)}$ such that $F((\boldsymbol{A}'_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}'_{-i}, B), k)$ relative to A_i. This, in turn, means that F is susceptible to π-manipulation. □

The significance of Lemma 2 is that voting rules that are strategy-proof when voters have some amount of information about the preferences of the other voters are also strategy-proof when voters receive less information. Contrapositively, when a voting rule is π-manipulable when voters have some amount of information, then the voting rule is also π-manipulable when voters have more information.

4.2 Manipulability under Incomplete Information

In this section we present two π-manipulability results for voting rules that satisfy the axioms *minimal efficiency, miminal proportionality*, *candidate monotonicity* and *diminishing returns.*

First, we prove the following lemma which establishes that, for voting rules that satisfy *candidate monotonicity* and *diminhising returns*, a voter can safely drop a candidate from her approval set when, after doing so, this candidate is still guaranteed a spot in the elected committee.

Lemma 3. Let F be an ABC rule that satisfies *candidate monotonicity* and *diminishing returns*. For any election instance $(\boldsymbol{A}, k)$, any voter $i \in N$ and insincere ballot $B = A_i \setminus \{d\}$ for some $d \in A_i$ we have that if $d \in W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$ then $F(\boldsymbol{A}, k)$ does *not* stochastically dominate $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i.

Proof. Let F be some ABC rule that satisfies *candidate monotonicy* and *diminishing returns*. Take any election instance $(\boldsymbol{A}, k)$, voter $i \in N$ and $B = A_i \setminus \{d\}$ for some $d \in A_i$. Suppose that $d \in W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$.

By candidate monotonicity we have $d \in W'$ for all $W' \in F(\boldsymbol{A}, k)$. Let l_{min} denote the minimum representation that voter i can receive on the insincere outcome $F((\boldsymbol{A}_{-i}, B), k)$ with non-zero probability. That is, let $l_{min} = |W \cap A_i|$ for $W \in F((\boldsymbol{A}_{-i}, B), k)$ that minimizes $|W \cap A_i|$. It holds that voter i receives at most l_{min} representation on the sincere outcome $F(\boldsymbol{A}, k)$ by the following argument.

Take any $W_2 \in F((\boldsymbol{A}_{-i}, B), k)$ such that $|W_2 \cap A_i| = l_{min}$ and take any $W_1 \in \mathcal{P}_k(C)$ such that $|W_1 \cap A_i| > l_{min}$. If $d \notin W_1$ then $W_1 \notin F(\boldsymbol{A}, k)$ because, by candidate monotonicity, we had $d \in W'$ for all $W' \in F(\boldsymbol{A}, k)$. If $d \in W_1$ we have (1) $d \in W_1$ and $d \in W_2$, (2) $|W_1 \cap A_i| > |W_2 \cap A_i|$ and (3) $W_2 \in F((\boldsymbol{A}_{-i}, B), k)$ and thus, by diminishing returns, $W_1 \notin F(\boldsymbol{A}, k)$.

This means we have $|W \cap A_i| \leq l_{min}$ for all $W \in F(\boldsymbol{A}, k)$. We conclude that $F(\boldsymbol{A}, k)$ does *not* stochastically dominate $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i. $\square$

The definition of π-manipulation we follow contains two conditions that have to be met in order for a voter to have incentive to manipulate. The first condition is that there is some profile the voter considers possible on which submitting the insincere ballot achieves a stochastically dominant outcome. The second condition is that the voter does not consider any profile possible for which submitting the insincere ballot leaves her worse of than the truthful ballot. The significance of Lemma 3 in relation to the second condition is that, if a voter were to identify that one of her approved candidates will certainly be elected for any profile she consider possible, then dropping this candidate will never leave her worse off. We can therefore analyse the susceptibility to π-manipulations that involve dropping one candidate for different information functions by analysing when the information about a profile is sufficient to determine that one of the candidates will certainly be elected. For election instances for which a voter can manipulate by free-riding by dropping one candidate, the question about which information functions allow the voter to π-manipulate by dropping one candidate reduces to the question about when the voter has sufficient information to determine that the candidate will certainly be elected.

Analysing the t-Profile IF and the Approval Score IF with this approach produces the π-manipulability results in Theorem 4 and Theorem 5. Both theorems make use of the following lemma in which, for every election scenario (N, C, k) such that $n - k > \frac{n}{k}$, a profile is given on which voter n can manipulate by dropping a candidate.

Lemma 4. Let F be a rule that satifies the axioms *anonymity, neutrality, minimal efficiency, minimal proportionality* and *diminishing returns*. Let (N, C, k) be an election scenario such that $n - k > \frac{n}{k}$. For the election instance $(\boldsymbol{A}, k)$ with $\boldsymbol{A}$ defined below, voter n can manipulate by dropping candidate d.

$$\boldsymbol{A}\text{:}\quad n-k \times \{d\} \quad 1 \times \{d, c_1\} \quad 1 \times \{d, c_2\} \quad \ldots \quad 1 \times \{d, c_k\}$$

That is, $d, c_1, \ldots, c_k$ denote $k+1$ candidates and in profile $\boldsymbol{A}$, every voter approves of candidate d and for each candidate c_j in $\{c_1, \ldots, c_k\}$ there is exactly one voter who in addition to d approves of c_j and every voter approves of at most one candidate in $\{c_1, \ldots, c_k\}$.

Proof. Let F and $(\boldsymbol{A}, k)$ be as described in the statement of the lemma. Voter n can manipulate by dropping candidate d from her approval ballot.

Note that by minimal efficiency candidate d is part of any committee in the outcome. Moreover, by minimal efficiency all winning committees have to be subsets of $\{d, c_1, \ldots, c_k\}$. By anonymity and neutrality *all* subsets of $\{d, c_1, \ldots, c_k\}$ that contain d are winning committees. Thus, we have the following outcome:

$$F(\boldsymbol{A}, k) = \{W \in \mathcal{P}_k(\{d, c_1, \ldots c_k\}) | d \in W\}$$

Let $\boldsymbol{A'}$ be the profile when voter n drops candidate d from her approval ballot:

$$\boldsymbol{A'}\text{:}\quad n-k \times \{d\} \quad 1 \times \{d, c_1\} \quad 1 \times \{d, c_2\} \quad \ldots \quad 1 \times \{c_k\}$$

Since F is minimally proportional and $n - k \geq \frac{n}{k}$, d is part of any winning committee in $F(\boldsymbol{A'}, k)$. Moreover, by minimal efficiency, again, all committees in the outcome are subsets of $\{d, c_1, \ldots, c_k\}$. We have:

$$F((\boldsymbol{A}_{-1}, \{c_1\}), k) \subseteq \{W \in \mathcal{P}_k(\{d, c_1, \ldots c_k\}) | d \in W\}$$

Now, because F satisfies diminishing returns and there are committees in $F(\boldsymbol{A}, k)$ which contain c_k (in addition to d) we have that $F(\boldsymbol{A'}, k)$ contains only the committees that contain both d and c_k. Take some $W_1 \in \{W \in \mathcal{P}_k(\{d, c_1, \ldots c_k\}) | d \in W\}$ such that $c_k \in W_1$ and any $W_2 \in \{W \in \mathcal{P}_k(\{d, c_1, \ldots c_k\}) | d \in W\}$ such that $c_k \notin W_2$. We have $|W_1 \cap A_n| > |W_2 \cap A_n|$ and $W_1 \in F(\boldsymbol{A}, k)$. By diminishing returns, then, we have $W_2 \notin F((\boldsymbol{A}_{-n}, \{c_k\}), k)$.

This means that $F(\boldsymbol{A'}, k)$ consists of committees that contain both d and c_k. Moreover, by anonymity and neutrality, it consists of *all* committees in $\mathcal{P}_k(\{d, c_1, \ldots c_k\})$ that contain d and c_k. That is:

$$F(\boldsymbol{A'}, k) = \{W \in \mathcal{P}_k(\{d, c_1, \ldots, c_k\}) | d \in W \text{ and } c_k \in W\}$$

$F((\boldsymbol{A}_{-n}, \{c_k\}), k)$ stochastically dominates $F(\boldsymbol{A}, k)$ relative to A_n. We conclude that voter n can manipulate $(\boldsymbol{A}, k)$ under F by dropping candidate d. □

Theorem 4. Let F be a rule that satisfies the axioms *anonymity, neutrality, minimal efficiency, candidate monotonicity, minimal proportionality* and *diminishing returns*. Let (N, C, k) be an election scenario such that $n - k > \frac{n}{k}$ and let π denote the Approval Score information function. Then F is susceptible to π-manipulation.

Proof. Let (N, C, k) and F be as described in the statement of the theorem. We will take profile $\boldsymbol{A}$ to be the profile described in Lemma 4:

$$\boldsymbol{A}: \quad n - k \times \{d\} \quad 1 \times \{d, c_1\} \quad 1 \times \{d, c_2\} \quad \ldots \quad 1 \times \{d, c_k\}$$

For this profile, voter n has incentive to π-manipulate by dropping candidate d from her approval ballot. The approval score information for this profile is the following.

$$\begin{aligned} as(d, \boldsymbol{A}) &= n \\ as(c_1, \boldsymbol{A}) &= 1 \\ as(c_2, \boldsymbol{A}) &= 1 \\ &\vdots \\ as(c_k, \boldsymbol{A}) &= 1 \end{aligned}$$

For any partial profile $\boldsymbol{A'}_{-n}$ consistent with the voter n's information, there must be at least $n - k$ voters with approval set $\{d\}$. Since we assumed that $n - k > \frac{n}{k}$ and F satisfies minimal proportionality, we know that, for any $\boldsymbol{A'}_{-n} \in \mathcal{W}_n^{\pi(\boldsymbol{A})}$ that voter n considers possible, we have $d \in W$ for all W

in $F((\boldsymbol{A'}_{-n}, \{c_k\}), k)$. By Lemma 3 and the assumption that F satisfies diminishing returns and candidate monotonicity, we find that $F((\boldsymbol{A'}_{-n}, A_n), k)$ does not stochastically dominates $F((\boldsymbol{A'}_{-n}, \{c_k\}), k)$ relative to A_n, for every $\boldsymbol{A'}_{-n} \in \mathcal{W}_n^{\pi(\boldsymbol{A})}$. In other words, voter n is never worse off when she drops candidate d from her approval ballot. Moreover, by Lemma 4 voter n can manipulate on the actual profile $\boldsymbol{A}$ by dropping candidate d.

Thus, for election instance $(\boldsymbol{A}, k)$ both conditions for π-manipulation in Definition 12 are met: there is no profile consistent with the voter's information for which dropping candidate d leaves her worse off, and there is at least one consistent profile (namely the actual profile $\boldsymbol{A}$), on which dropping d is beneficial. In conclusion, F is π-manipulable on election scenarios that satisfy $n - k > \frac{n}{k}$. □

Theorem 5. Let F be a rule that satisfies the axioms *anonymity, neutrality, minimal efficiency, candidate monotonicity, minimal proportionality* and *diminishing returns*. Let (N, C, k) be an election scenario such that $n - k > \frac{n}{k}$ and let π be the t-Profile information function with $t = \lceil \frac{n}{k} \rceil$. Then F is susceptible to π-manipulation.

Proof. Let F, (N, C, k) and π be as described in the statement of the theorem. We will take profile $\boldsymbol{A}$ to be the profile described in Lemma 4:

$$\boldsymbol{A}\colon \quad n - k \times \{d\} \quad 1 \times \{d, c_1\} \quad 1 \times \{d, c_2\} \quad \dots \quad 1 \times \{d, c_k\}$$

For this profile, voter n has incentive to π-manipulate by dropping candidate d from her approval ballot.

Firstly, by Lemma 4, voter n can manipulate election instance $(\boldsymbol{A}, k)$ by dropping candidate d from her approval ballot. Secondly, there is no partial profile consistent with the information voter n has for which the voter is worse off when dropping candidate d. Take any $\boldsymbol{A'}_{-n} \in \mathcal{W}_n^{\pi(\boldsymbol{A})}$. Since π is the t-Profile IF with $t = \lceil \frac{n}{k} \rceil$ and $n - k > \frac{n}{k}$, voter n knows that the first $\lceil \frac{n}{k} \rceil$ voters have approval set $\{d\}$. This means that for any $\boldsymbol{A'}_{-n} \in \mathcal{W}_n^{\pi(\boldsymbol{A})}$ the first $\lceil \frac{n}{k} \rceil$ will also have approval set $\{d\}$. By the minimal proportionality of F we then have $d \in W$ for all $W \in F((\boldsymbol{A'}_{-n}, B), k)$. By Lemma 3, we have that $F((\boldsymbol{A'}_{-n}, A_n), k)$ does *not* stochastically dominate $F((\boldsymbol{A}_{-n}, \{c_k\}), k)$ relative to A_n. Since $\boldsymbol{A'}_{-n}$ was arbitrary, it follows that there is no $\boldsymbol{A'}_{-n} \in \mathcal{W}_n^{\pi(\boldsymbol{A})}$ for which $F((\boldsymbol{A'}_{-n}, A_n), k)$ stochastically dominates $F((\boldsymbol{A'}_{-n}, \{c_k\}), k)$ relative to A_n.

For election instance $(\boldsymbol{A}, k)$ both conditions for π-manipulation in Definition 12 are met: there is no profile voter n considers possible on which dropping candidate d will leave her worse off, and there is some profile she considers possible on which dropping d will achieve a stochastically dominating outcome.

Thus election instance $(\boldsymbol{A}, k)$ is π-manipulable by voter n. We conclude that F is π-manipulable on election scenarios that satisfy $n - k > \frac{n}{k}$. □

Theorem 4 and Theorem 5 show that no voting rule F that satisfies the axioms *anonymity, neutrality, minimal efficiency, candidate monotonicity, minimal proportionality* and *diminishing returns* is π-strategy-proof for the information functions $\lceil \frac{n}{k} \rceil$-Profile and Approval Score. Strictly speaking, a particular instance of $(\boldsymbol{A}, k)$ constructed in Lemma 4 would provide this result. Theorem 4 and Theorem 5 are more general, in the sense that they show that for any election scenario for which $n - k > \frac{n}{k}$ there is at least one election instance that is π-manipulable for the $\lceil \frac{n}{k} \rceil$-Profile and Approval Score information functions.

The condition $n - k > \frac{n}{k}$ can be satisfied for any desired committee size $k \in \mathbb{N}$. In particular, for any desired committee size $k \in \mathbb{N}$ there is a minimum number of voters $n \in \mathbb{N}$ above which the condition is always satisfied, and this minimum number of voters lies close to k.

Since the t-Profile information function for any $t \in \mathbb{N}$ such that $n \geq t > \frac{n}{k}$ is at least as informative as the $\lceil \frac{n}{k} \rceil$-Profile information function, we have the following corollary of Theorem 5 by Lemma 2.

Corollary 1. Let F be a rule that satisfies the axioms *anonymity, neutrality, minimal efficiency, candidate monotonicity, minimal proportionality* and *diminishing returns.* Let (N, C, k) be an election scenario such that $n - k > \frac{n}{k}$ and let π be the t-Profile information function with $t > \frac{n}{k}$. Then F is susceptible to π-manipulation for (N, C, k).

The results presented in this section show that restricting the information voters have to the Approval Score information or the t-Profile information for any $t > \frac{n}{k}$ does not provide a sufficient information barrier to manipulation for the voting rules we are interested in.

4.3 The Prevalence of π-Manipulability

Given the π-manipulability results presented in the previous section, it is natural to continue our investigation by analysing the prevalence of manipulability under incomplete information. Even-though the $\lceil \frac{n}{k} \rceil$-Profile IF and Approval Score IF do not provide full information barriers to π-manipulation for the voting rules that we are interested in, we may expect that at least the prevalence of manipulability decreases when voters have incomplete information. In this section we present experimental results which quantify the prevalence of manipulability under incomplete information. Surprisingly, it is not generally the case that less information results in a decreased prevalence of manipulability.

The simulation experiment which is presented in this section is performed for small elections and for information functions that provide close to complete information. The reason for this is that the computational problem of deciding whether an election instance is π-manipulable becomes intractable very quickly. For example, for any t-Profile information function, the number of possible partial profiles to consider for each voter grows linearly in t but exponentially in the number of candidates m. This is because the approval sets which are unknown can be any out of the $2^m - 1$ possible approval sets. Similar intractability problems apply for the Approval Score information function.

We present a simulation experiment in which the prevalence of π-manipulable election instances is quantified for the $(n-1)$-Profile and $(n-2)$-Profile information functions. These information functions correspond, respectively, to the situation in which voters know the preferences of all but one or all but two of the other voters. We randomly generate election instances and, for each election instance, compute how many voters can π-manipulate. As before, we generate election instances by sampling approval sets for every voter using the uniform probability distribution over the set of possible approval sets $\mathcal{P}(C) \setminus \varnothing$.

In Figure 4.1 we observe that the proportion of election instances that are *not* susceptible to manipulation is significantly higher when voters have complete information than when voters know the approval sets of all but one of their fellow voters. The proportion of election instances for which there is *one* manipulator when voters have complete information is equal to the proportion of election instance for which there is *one* manipulator when voters know the approval sets of all but one of their fellow voters. However, for each number of manipulators higher than one, the proportion of election instances for which there are such a number of manipulators is larger when voters lack information about one of the approval sets than when voters have complete information. We can conclude from this that the probability that a voter has incentive to manipulate increases when the voter lacks information about one of the approval sets, relative to when the voter has complete information. On the other hand, we also observe in this figure that, when voters lack information about the approval set of two of their fellow voters, the proportion of election instances which are *not* susceptible to manipulation increase again, and is even higher than when voters have complete information.

To understand the data it is useful to look at what effect receiving additional information can have on the incentive a voter has to manipulate in our model.

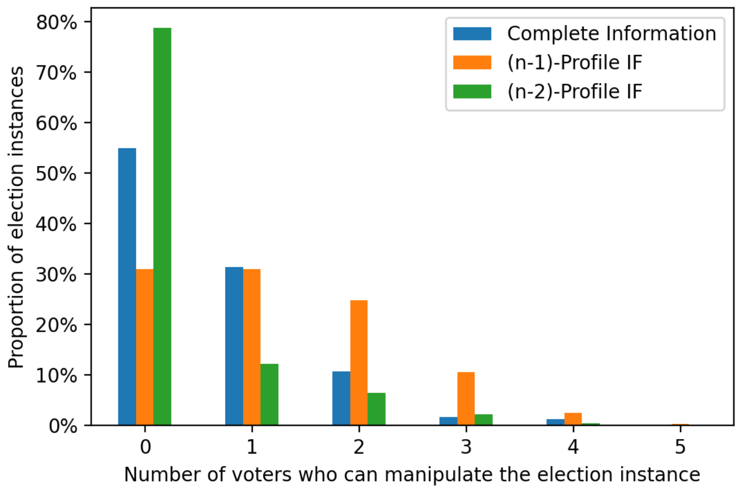

Figure 4.1: This figure shows π-manipulability data for sequential-PAV for the election scenario parametrized by $n = 5, m = 4$ and $k = 3$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

For a particular election instance $(\boldsymbol{A}, k)$, increasing the information available to a voter refines the information set of this voter. That is, it shrinks the set of profiles the voter considers possible. This may have two opposing effects on the voter's opportunity to manipulate. The first is that the voter may learn that a profile which previously prevented a manipulation is no longer consistent with the information, thereby creating an opportunity to manipulate. The second is that the voter may learn that a profile which previously provided the incentive to submit an insincere ballot is now no longer consistent with the information, thereby eliminating the opportunity to manipulate. In the first case the election instance may go from non-manipulable to manipulable whereas in the second case the election instance may go from manipulable to non-manipulable.

Figure 4.2 and Figure 4.3 both show simulation results for PAV for an election scenario with three candidates and a desired committee size of two. The difference between the two figures is that in Figure 4.2 the number of voters is four, whereas in Figure 4.3 the number of voters is ten. In both cases the prevalence of π-manipulable election instances is greater for the $(n-1)$-Profile IF than under complete information. For the $(n-2)$-Profile IF the two election scenarios show different results: with four voters, the $(n-2)$-Profile IF seems to prevent any incentive to π-manipulate, whereas for ten voters the prevalence

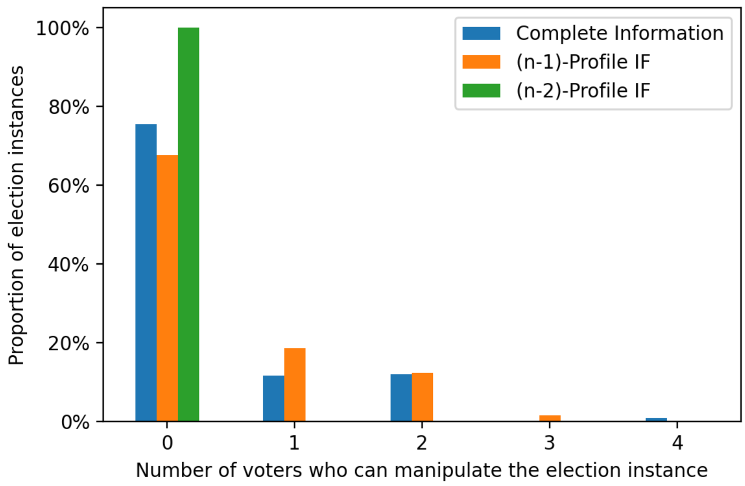

Figure 4.2: This figure shows π-manipulability data for PAV for the election scenario parametrized by $n = 4, m = 3$ and $k = 2$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

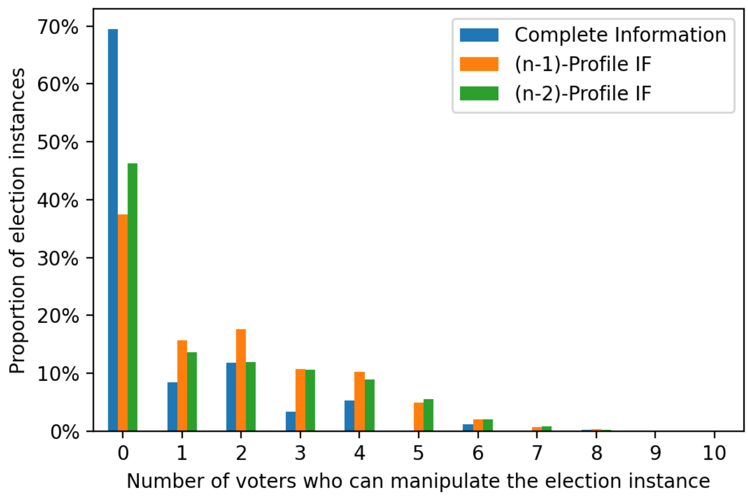

Figure 4.3: This figure shows π-manipulability data for PAV for the election scenario parametrized by $n = 10, m = 3$ and $k = 2$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

is increased and similar to the $(n-1)$-Profile IF. Note that for four voters, the $(n-2)$-Profile IF provides voters with very little information, namely only the approval set of one other voter is known. For ten voters, on the other hand, the $(n-2)$-Profile IF provides information about a majority of the voters, namely the approval sets of seven of the fellow voters are known. Additional data for each of the three voting rules and for an additional election scenario is presented in Appendix B.3.

The results of the experiment draw our attention to the fact that decreasing the amount of information voters have does not always decrease the prevalence of manipulability. In fact, for most of the election scenarios studied in the experiment, withholding information about one or two of the approval sets increases the probability that an arbitrary voter has incentive to manipulate. On the other hand, the results of the experiment for election scenarios with a very small number of voters - for which the $(n-2)$-Profile IF comes close to providing zero information - show that withholding information about a significantly large part of the electorate does decrease the prevalence of π-manipulation relative to complete information.

These observations both motivate the need for complete information barriers to manipulation and suggest that such information barriers exist when the number of approval sets that voters have knowledge of approaches zero.

4.4 Information Barriers

In this section we present two results that give information barriers to manipulation for the PAV rule. The first result is that PAV is Zero-strategyproof for all election scenarios (N, C, k) with $n \geq k$. That is, when voters have no information about the preferences of their fellow voters, they do not have an incentive to manipulate. This result reassures us that there indeed exist information barriers to manipulation for PAV. The second result is that voters can be allowed to know some of the approval sets in the profile without gaining an incentive to manipulate by dropping one candidate. That is, voters can know the approval set of a certain number of their fellow voters (this number depends on the total number of voters and the desired committee size) without having incentive to manipulate by dropping a candidate for any election instance.

Theorem 6. PAV is Zero-strategyproof for all election scenarios (N, C, k) with $n \geq k$.

Proof. Let (N, C, k) be an arbitrary election scenario for which $n \geq k$ and $i \in N$ an arbitrary voter. $F(\boldsymbol{A}, k)$ denotes the outcome of PAV on the election instance

$(\boldsymbol{A}, k)$. Note, firstly, that if $A_i = C$ or if $n = 1$ voter i will not have incentive to π-manipulate because i will be completely satisfied with the outcome for any election instance, if she submits the truthful ballot A_i. Completely satisfied here means that $|W \cap A_i| = min(k, |A_i|)$ for all $W \in F(\boldsymbol{A}, k)$. Thus, we assume that $n \geq 2$ and $A_i \subsetneq C$.

Under these assumptions, we will show that for every election scenario (N, C, k) such that $n \geq k$, sincere approval set A_i and insincere ballot $B \subseteq C$ that i may submit, there is a partial profile $\boldsymbol{A}_{-i} \in \mathcal{A}(N \setminus \{i\}, C)$ such that $F((\boldsymbol{A}_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i. That is, for every election scenario (N, C, k) such that $n \geq k$, sincere approval set A_i and insincere ballot $B \subseteq C$, there is some partial profile that voter i considers possible for which submitting ballot B results in an outcome that is stochastically dominated by the truthful outcome. This fact implies that for any election instance $(\boldsymbol{A}, k)$, voter $i \in N$ and ballot B, voter i does not have incentive to Zero-manipulate.

We will distinguish two cases: (i) $B \subseteq C$ is such that for some $d \in A_i$ we have $d \notin B$ and (ii) $B \subseteq C$ is such that $A_i \subsetneq B$.

(i) We have $B \subseteq C$ and $d \notin B$ for some $d \in A_i$. We consider two subcases: $(i.a)$ $|A_i| < k$ and $(i.b)$ $|A_i| \geq k$.

$(i.a)$ We have $B \subseteq C$ such that $d \notin B$ for some $d \in A_i$ and $|A_i| < k$. Let $\boldsymbol{A}_{-i}$ be the partial profile in which all voters have approval set $A_j = C$. We have $F((\boldsymbol{A}_{-i}, A_i), k) = \{W \in \mathcal{P}_k(C) | A_i \subseteq W\}$ so voter i is completely satisfied with the outcome when she submits her truthful approval ballot. When $|B| < k$ we have $F((\boldsymbol{A}_{-i}, B), k) = \{W \in \mathcal{P}_k(C) | B \subseteq W\}$ and when $|B| \geq k$ we have $F((\boldsymbol{A}_{-i}, B), k) = \mathcal{P}_k(B)$. In both cases there is some $W \in F((\boldsymbol{A}_{-i}, B), k)$ such that $d \notin W$. For such winning committee W we have $|A_i \cap W| < |A_i|$ so voter i is not completely satisfied with the (insincere) outcome. This means $F((\boldsymbol{A}_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i.

$(i.b)$ We have $B \subseteq C$ such that $d \notin B$ for some $d \in A_i$ and $|A_i| \geq k$.

We will construct a partial profile $\boldsymbol{A}_{-i}$ in which a candidate c who is not approved by voter i and the candidate d are competing for a spot in the winning committee(s). By dropping candidate d from her approval ballot, voter i ensures that candidate c is elected, leaving i worse off.

Let $\boldsymbol{A}_{-i} \in \mathcal{A}(N \setminus i, C)$ be such that, for all $j \in N \setminus \{i\}$, $A_j = S_1 \cup \{c\}$ where S_1 is some k-size subset of A_i that includes $d \notin B$, and c is a candidate which is not in A_i. Now, $F((\boldsymbol{A}_{-i}, A_i), k)$ is such that i is fully satisfied because there are exactly k candidates that are approved unanimously (and therefore there is exactly one winning committee consisting of these candidates) and these candidates are all approved by i.

Next, we consider $F((\boldsymbol{A}_{-i}, B), k)$. There must be some $W \in F((\boldsymbol{A}_{-i}, B), k)$ with $c \in W$ by the following argument. The set $B \cap A_j$ contains exactly the candidates that are approved unanimously. Since A_j contains k candidates that may also be contained in B and one candidate that is not in B (namely d), there are at most k candidates approved unanimously. This means that $A_j \cap B \subseteq W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$. If $c \in B$ then $c \in W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$ and we would be done; voter i would not be fully satisfied with the outcome because $|W \cap A_i| < k$ for all winning committees W. Suppose that $c \notin A_j \cap B$. This means that $|A_j \cap B| < k$ and therefore there is at least one more candidate needed outside of $A_j \cap B$ to form a (winning) committee. Since $A_j \cup B$ is the set of all candidates that are approved at least once (and $A_j \cup B \geq k$) the remaining candidate(s) to fill the committee must come from $A_j \setminus B$ or from $B \setminus A_j$. It cannot be that only candidates in $B \setminus A_j$ are chosen to fill the remaining spot(s) in the winning committee because candidates in $B \setminus A_j$ are only approved once (by voter i) and candidates in $A_j \setminus B$ are approved at least once (because we assume $n \geq 2$) and, crucially, all voters, including i, are represented by the candidates in $A_j \cap B$ which means voter i cannot outweigh the other $(n-1)$ voters. There must thus be at least one winning committee W that contains a candidate $e \in A_j \setminus B$. Moreover, if a candidate $e \in A_j \setminus B$ is included in some winning committee $W \in F((\boldsymbol{A}_{-i}, B), k)$ then there must also be a winning committee W' that contains candidate c because e and c are approved by exactly the same candidates. In conclusion, there is a winning committee $W' \in F((\boldsymbol{A}_{-i}, B), k)$ such that $c \in W'$. Since $c \notin A_i$ and $|A_i| \geq k$ this means that $F((\boldsymbol{A}_{-i}, A_i B), k)$ stochastically dominates $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i.

(ii) We have $B \subseteq C$ such that $A_i \subsetneq B$. We distinguish two cases: *(ii.a)* in which $|B| \geq k$ and *(ii.b)* in which $|B| < k$.

(ii.a) We consider the potential manipulation where i submits a superset of her approval ballot. That is, submitting ballot B for which $A_i \subsetneq B$. Additionally we assume $|B| \geq k$. Let $\boldsymbol{A}_{-i}$ be the partial profile where all voters other than i have the approval set $A_j = C$. On the profile $(\boldsymbol{A}_{-i}, A_i)$ all candidates in A_i are approved unanimously and all candidates outside of A_i are not approved unanimously. If $|A_i| \leq k$ we have $A_i \subseteq W$ for all $W \in F((\boldsymbol{A}_{-i}, A_i), k)$. If $|A_i| > k$ we have $F((\boldsymbol{A}_{-i}, A_i), k) = \mathcal{P}_k(A_i)$ and $|A_i \cap W| = k$ for all $W \in F((\boldsymbol{A}_{-i}, A_i), k)$. So, if the voter submits her truthful ballot, she is guaranteed the highest number of representatives possible in case the partial profile is $\boldsymbol{A}_{-i}$. Now consider $F((\boldsymbol{A}_{-i}, B), k)$. Since $|B| > k$ by assumption, and B contains

exactly the candidates approved unanimously, we have $F((\boldsymbol{A}_{-i}, B), k) = \mathcal{P}_k(B)$. This means there is a winning committee $W \in F((\boldsymbol{A}_{-i}, B), k)$ such that $|W \cap A_i| < min(k, |A_i|)$. So, if i submits ballot B she is not guaranteed the highest number of representative possible. This means $F((\boldsymbol{A}_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i.

$(ii.b)$ We consider the potential manipulation where voter i submits a superset B of her approval ballot A_i and $|B| < k$. Let $\delta := k - |B|$ be the difference between the size of the ballot B and the target committee size k. We will construct the profile $(\boldsymbol{A}_{-i}, A_i)$ such that $\delta + 1$ candidates outside of B will certainly have a spot in the winning committee. The remaining $k - (\delta + 1)$ committee spots will be filled with candidates from B.

Let $S \subseteq C \setminus B$ be a set of size $\delta + 1$ and let $\lceil n \times \frac{\delta+1}{k} \rceil$ voters have approval set S. Since PAV satisfies minimal proportionality, we have $S \subseteq W$ for all $W \in F((\boldsymbol{A}_{-i}, A_i), k)$ and $S \subseteq W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$. Let the remaining voters have approval set B. Note that $\lceil n \times \frac{\delta+1}{k} \rceil + 1 \leq n$ is implied by $n \geq k$ (see Appendix A.3 for a short proof of this) so we can indeed construct such a profile for (N, C, k).

Consider the outcome $F((\boldsymbol{A}_{-i}, A_i), k)$ of the election instance $((\boldsymbol{A}_{-i}, A_i), k)$. Since the only candidates that are approved at least once are the candidates in $S \cup B$ (and $|S \cup B| > k$), we will have $W \subseteq S \cup B$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$. Moreover, for each pair of candidates $a \in A_i$ and $b \in B \setminus A_i$, the set of voters approving b is a strict subset of the set of voters approving a. As a result, candidates in A_i are prioritised to fill the remaining $k - (\delta + 1)$ spots of the committee. Additionally, there are enough remaining spots in the committee to elect all candidates in A_i. That is, we have $|A_i| \leq k - (\delta + 1)$ because $|A_i| < |B| = k - \delta$. Thus, we will have $A_i \subseteq W$ for all $W \in F((\boldsymbol{A}_{-i}, A_i), k)$.

Next, consider the outcome of the election on profile $(\boldsymbol{A}_{-i}, B)$. As before, we have $S \subseteq W$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$ and $W \subseteq S \cup B$ for all $W \in F((\boldsymbol{A}_{-i}, B), k)$. However, in contrast to the outcome of the election on profile $(\boldsymbol{A}_{-i}, A_i)$, we now have that all candidates in B are approved by the same set of voters. As a result, we have:

$$F((\boldsymbol{A}_{-i}, B), k) = \{W \in \mathcal{P}_k(S \cup B) | S \subseteq W\}$$

Since $|B|$ is larger than the number of spots remaining in the committee after electing the candidates in S, there is some $W \in F((\boldsymbol{A}_{-i}, B), k)$ such that $a \notin W$ for some $a \in A_i$. We conclude that $F((\boldsymbol{A}_{-i}, A_i), k)$ stochastically dominates $F((\boldsymbol{A}_{-i}, B), k)$ relative to A_i. □

The proof of the following theorem makes use of the concept of *marginal* *contribution* for the PAV rule.

Definition 14. For profile $\boldsymbol{A} \in \mathcal{A}(N, C)$ and committee $W \in \mathcal{P}_k(C)$ we define the *marginal contribution* of candidate $a \in W$ to committee W on profile $\boldsymbol{A}$ as the difference between the PAV-score of W and $W \setminus \{a\}$ on profile $\boldsymbol{A}$.

Theorem 7. Let (N, C, k) be some election scenario and π a t-Profile information function such that the following is satisfied:

$$\lceil n \times \frac{k-1}{k} \rceil + 3(t+1) \leq n \tag{4.1}$$

Let $(\boldsymbol{A}, k)$ be some election instance over (N, C, k). Under PAV, no voter can π-manipulate by dropping a candidate from her approval ballot.

Proof. Let $(\boldsymbol{A}, k)$ and π be as described above and let $i \in N$ be some voter. Suppose that i has information $\pi(\boldsymbol{A})$. Let $d \in A_i$ be some candidate that i considers dropping (i considers submitting ballot $B = A_i \setminus \{d\}$). Let $\boldsymbol{A}_I$ denote the partial profile consisting of the approval sets of the t voters of which i knows the approval set, plus the approval set of i herself. We will extend $\boldsymbol{A}_I$ to a profile $\boldsymbol{A'}$ over (N, C) on which i is worse off when she drops candidate d. The partial profile $\boldsymbol{A'}_{-i}$ will be considered possible by voter i and therefore will prevent her from having an incentive to manipulate by dropping d.

Take some $S \subseteq C \setminus \{d\}$ of size $k-1$. We will construct the profile $\boldsymbol{A'}$ such that the $k-1$ candidates in S will certainly be elected and candidate d and some candidate a that voter i does not approve of compete for the remaining spot in the winning committee.

Let F_X denote the set of committees $W \in \mathcal{P}_k(C)$ such that $S \subseteq W$ with the highest PAV-score under $\boldsymbol{A}_I$. That is, F_X contains the committees $S \cup \{c\}$ for which c has maximal marginal contribution to $S \cup \{c\}$ under profile $\boldsymbol{A}_I$.

We can distinguish three cases:

1. $S \cup \{d\} \in F_X$.
2. $S \cup \{d\} \notin F_X$ but $S \cup \{a\} \in F_X$ for some $a \notin A_i$.
3. For all $S \cup \{c\} \in F_X$ we have $c \in A_i \setminus \{d\}$.

Case 1. Let a be some candidate that is not approved by i. Let $\boldsymbol{A}_I^1$ be the partial profile obtained by (1) removing all candidates outside of $S \cup \{a\} \cup \{d\}$ from the approval sets in $\boldsymbol{A}_I$ and (2) swapping candidates a and d in each approval set in $\boldsymbol{A}_I$. Construct $\boldsymbol{A'}$ as follows[1]:

[1]Strictly speaking, the first t voters in $\boldsymbol{A'}$ should have the approval sets according to $\pi(\boldsymbol{A})$ and the ith voter should have approval set A_i, but this can be achieved by rearranging the order of the voters

1. Let $\lceil n \times \frac{k-1}{k} \rceil$ voters have approval set S.
2. Let $t+1$ voters have the approval sets in $\boldsymbol{A}_I$.
3. Let $t+1$ voters have the approval sets in $\boldsymbol{A}_I^1$.
4. Let the remaining voters have approval set C.

Now consider $F(\boldsymbol{A'}, k)$. Since PAV satisfies minimal proportionality, we have $S \subseteq W$ for all $W \in F(\boldsymbol{A'}, k)$. We will have $S \cup \{c\} \in F(\boldsymbol{A}, k)$ for candidates $c \in C \setminus S$ with maximum marginal contribution to $S \cup \{c\}$ on profile $\boldsymbol{A'}$. To see for which candidates c this marginal contribution is maximum, it is sufficient to consider for which candidates the marginal contribution is maximum on the partial profile $(\boldsymbol{A}_I, \boldsymbol{A}_I^1)$. Note that a and d are the only candidates $c \in C \setminus S$ for which the marginal contribution to $S \cup \{c\}$ is (possibly) larger under profile $(\boldsymbol{A}_I, \boldsymbol{A}_I^1)$ than under profile $\boldsymbol{A}_I$. This is because the other candidates are not contained in any approval set in $\boldsymbol{A}_I^1$. Moreover, the marginal contribution of a to $S \cup \{a\}$ under $(\boldsymbol{A}_I, \boldsymbol{A}_I^1)$ and the marginal contribution of d to $S \cup \{d\}$ under $(\boldsymbol{A}_I, \boldsymbol{A}_I^1)$ are equal because of the symmetry between $\boldsymbol{A}_I$ and $\boldsymbol{A}_I^1$ with respect to a and d. More precisely, for any number $l \in \mathbb{N}$, there are as many approval sets containing exactly l candidates in S *and* candidate a as there are approval sets containing exactly l candidates in S *and* d. Since $S \cup \{d\} \in F_X$ we will thus have $S \cup \{d\} \in F((\boldsymbol{A}_I, \boldsymbol{A}_I^1)), k)$ and $S \cup \{a\} \in F((\boldsymbol{A}_I, \boldsymbol{A}_I^1)), k)$. This means we will have $S \cup \{d\} \in F(\boldsymbol{A'}, k)$ and $S \cup \{a\} \in F(\boldsymbol{A'}, k)$.

Next, consider the outcome of the election when i drops candidate d. We will have $F((\boldsymbol{A'}_{-i}, B), k) = F(\boldsymbol{A'}, k) \setminus S \cup \{d\}$ because the effect of dropping candidate d is that only committees containing candidate d will have a strictly lower PAV-score. Since $S \cup \{a\} \in F(\boldsymbol{A'}, k)$ and $a \notin A_i$ this means that $F(\boldsymbol{A'}, k)$ stochastically dominates $F((\boldsymbol{A'}_{-i}, B), k)$ relative to A_i.

Case 2. Let a be some candidate for which $S \cup \{a\} \in F_X$ and $a \notin A_i$. The construction of $\boldsymbol{A}_I^1$ and $\boldsymbol{A'}$ is equivalent to the construction in case 1. We again find that $S \cup \{a\} \in F(\boldsymbol{A'}, k)$, $S \cup \{d\} \in F(\boldsymbol{A'}, k)$ and $F((\boldsymbol{A'}_{-i}, B), k) = F(\boldsymbol{A'}, k) \setminus S \cup \{d\}$, which means that $F(\boldsymbol{A'}, k)$ stochastically dominates $F((\boldsymbol{A'}_{-i}, B), k)$ relative to A_i.

Case 3. Take some $a \notin A_i$. Let $\boldsymbol{A}_I^1$ be the partial profile obtained by (1) removing all candidates outside of $S \cup \{a\} \cup \{d\}$ from $\boldsymbol{A}_I$ and (2) swapping candidates a and d in each approval set in $\boldsymbol{A}_I$. Additionally, let $\boldsymbol{A}_I^2$ be the partial profile consisting of $t-1$ times the approval set $\{a, d\}$. Construct $\boldsymbol{A'}$ as follows:

1. Let $\lceil n \times \frac{k-1}{k} \rceil$ voters have approval set S.

2. Let $t+1$ voters have the approval sets in $\boldsymbol{A}_I$.

3. Let $t+1$ voters have the approval sets in $\boldsymbol{A}_I^1$.

4. Let $t+1$ votes have the approval sets in $\boldsymbol{A}_I^2$.

5. Let the remaining n_r voters have approval set C.

By minimal proportionality, we have $S \subseteq W$ for all $W \in F(\boldsymbol{A'}, k)$.

Note that the marginal contribution of candidate a to $S \cup \{a\}$ and the marginal contribution of candidate d to $S \cup \{d\}$ are equal due to the symmetry in $\boldsymbol{A'}$ with respect to a and d. More precisely, for any number $l \in \mathbb{N}$, there are as many approval sets containing exactly l candidates in S *and* candidate a as there are approval sets containing exactly l candidates in S *and* d. As a result we have $S \cup \{d\} \in F(\boldsymbol{A'}, k)$ if and only if $S \cup \{a\} \in F(\boldsymbol{A'}, k)$. Furthermore, the marginal contribution of d to $S \cup \{d\}$ under $\boldsymbol{A'}$ is at least $(t+1) + \frac{n_r}{|S|+1}$ because of partial profile $\boldsymbol{A}_I^2$. And, the marginal contribution to $S \cup \{c\}$ of any $c \notin S \cup \{d\} \cup \{a\}$ is at most $(t+1) + \frac{n_r}{|S|+1}$ because the only approval sets that may contain candidates outside of $S \cup \{d\} \cup \{a\}$ are the $t+1$ approval sets in $\boldsymbol{A}_I$ and the n_r approval sets of the remaining voters who have approval set C. Thus, we will have $S \cup \{d\} \in F(\boldsymbol{A'}, k)$ and $S \cup \{a\} \in F(\boldsymbol{A'}, k)$.

Next, consider $F((\boldsymbol{A'}_i, B), k)$. Since only the score of the committee $S \cup \{d\}$ decreases with respect to the truthful profile $\boldsymbol{A}$, we will have $F((\boldsymbol{A'}_i, B), k) = F(\boldsymbol{A'}, k) \setminus S \cup \{d\}$. We conclude that outcome $F(\boldsymbol{A'}, k)$ stochastically dominates outcome $F((\boldsymbol{A'}_i, B), k)$ relative to A_i. □

Theorem 7 provides an information barrier to manipulation by dropping one candidate for the PAV rule. It states that any voter who has t-Profile information of the truthful approval profile does not have incentive to π-manipulate by dropping one candidate when t satisfies Inequality 4.1. To evaluate the significance of this information barrier, Inequality 4.1 should be analysed. For election scenarios for which the number of voters is a multiple of the desired committee size, Inequality 4.1 becomes easier to analyse because we can drop the ceiling operators appearing in the expression. In this case we can rewrite 4.1 to:

$$t < \frac{1}{3}\frac{n}{k} \tag{4.2}$$

When t satisfies 4.2 we are guaranteed that a potential manipulator cannot identify that she has incentive to manipulate by dropping one candidate. That is, for any t that satisfies 4.2, a voter may know t approval sets for any election instance without having incentive to manipulate by dropping one candidate. In this sense the theorem provides an upper bound to the information voters can safely have. However, it may be that the best upper bound is higher than the upper bound provided by Inequality 4.2. That is, Theorem 7 does not entail that whenever t does not satisfy the inequality, there are instances for which a voter with t-Profile information has incentive to manipulate.

Inequality 4.2 shows a linear relation between the upper bound of t and the value of $\frac{n}{k}$. Figure 4.4 shows a plot of this linear relation. For low values of $\frac{n}{k}$ the number of safely known approval sets t is close to zero, but for large values of $\frac{n}{k}$ the upper bound may become significant. For example, when $n = 240$ voters elect a committee of size $k = 3$, voters may know up to 26 approval sets in the truthful approval profile without having incentive to manipulate by dropping one candidate.

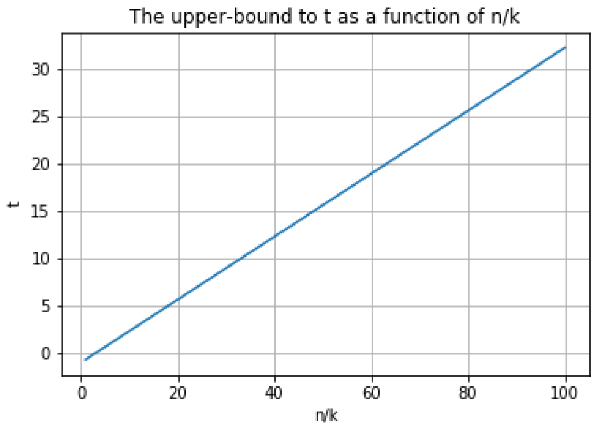

Figure 4.4: The upper-bound to t as a function of $\frac{n}{k}$. The vertical axis shows the number of approval sets a voter may known for any election instance without having incentive to manipulate by dropping one candidate.

4.5 Discussion

In this chapter we set out to extend the model of approval-based committee voting to account for strategic manipulation when voters have incomplete information. We explored whether there are information barriers to manipulation for two natural types of information. In this section we reflect on the results obtained and on the assumptions that we made about the strategic behavior of voters.

In the current model we make two important assumptions about when a voter has incentive to manipulate under incomplete information. The first assumption is that the preferences of voters over election outcomes are modelled according to the stochastic dominance relation. As has been pointed out in Chapter 2, the stochastic dominance preference relation is sparser than the preference relation that a utility maximizing voter with a particular attitude towards risk would have. We use the stochastic dominance relation exactly because we do not want to make assumptions about the attitude voters have towards risk. However, in the complete information setting, a utility maximizing voter with a particular attitude towards risk will more frequently have an incentive to manipulate than the voter we model. For strategic manipulation under incomplete information, we make a second important assumption about when a voter has incentive to manipulate. Namely, we assume that voters treat their uncertainty about the actual profile with caution. That is, two conditions have to be met for a voter to have incentive to manipulate under incomplete information. The first condition is that the considered insincere ballot gives the voter an outcome which she prefers to the truthful outcome for at least one possible election instance. The second conditions is a safety condition: the considered insincere ballot should never result in an outcome that is dominated by the truthful outcome.

Both of these assumptions impose strong conditions that have to be met in order for a voter to have incentive to manipulate. Consequently, we may judge that susceptibility results based on these definitions are strong results that will continue to hold when voters have a complete or less sparse preference order over possible election outcomes. However, one fact about the combination of the two assumptions should be noted in order to properly reflect on any results based on this model. Consider a utility maximizing voter with a complete preference order over possible election outcomes. This voter will, in the complete information case, more frequently have incentive to manipulate. If we stick to the same two conditions for manipulation under incomplete information, but instead take the underlying preference relation to be the preference relation of this voter rather than the stochastic dominance relation, then we do not necessarily find that the voter will more frequently have incentive to manipulate under incomplete information. There will be more cases in which a possible profile invites manipulation, but there will also be more cases in which a possible profile prevents manipulation. That is, the first condition in the definition will be satisfied more frequently but the second condition will be satisfied less frequently. It is therefore not immediately clear whether the susceptibility results will remain applicable when we consider utility maximizing voters with a complete preference order over election outcomes.

The crucial observation on which the susceptibility results in Section 4.2 rely is that limited information is required to identify that a candidate will certainly be elected, and that when a candidate is certainly elected, a voter can safely drop this candidate for any voting rule that satisfies candidate monotonicity and diminishing returns. This fact, in combination with the prevalence of manipulations by free-riding makes the voting rules we are interested in susceptible to manipulation, even when voters have limited information about the preferences of their fellow voters.

Lemma 3 formally states the observation described above. The proof of this lemma, in fact, shows something stronger than is given in the statement of the lemma. It shows that, whenever a candidate is certainly elected, dropping this candidate will result in an outcome which is not only not stochastically dominated by the truthful outcome, but even result in an outcome on which the minimum representation the voter receives is at least as high as the maximum representation the voter would have received on the truthful outcome. For this reason, this result may be applied more broadly than our current framework. That is, for rules that satisfy candidate monotonicity and diminishing returns, dropping a candidate which is certainly elected is safe for any utility maximizing voter with a particular attitude towards risk.

One important aspect of strategic manipulation which we have not investigated in this thesis is that of the computational complexity of identifying possible ways to manipulate under incomplete information. We have indicated in Section 4.3 that the brute force algorithms for deciding whether a voter has incentive to π-manipulate for the voting rules PAV, sequential PAV and Phragmén's rule for both the t-Profile information functions and the Approval Score information functions are intractable. The reason for this is the exponentially increasing number of profiles the voter considers possible when increasing the size of the election scenario. It remains an open question whether there exists, for these information functions, polynomial time algorithms to decide whether a voter can π-manipulate a given election instance. Two related computational problems are deciding, based on the Approval Score information or based on the t-Profile information, whether one of the candidates will certainly be elected. In case there exist polynomial time algorithms for the latter decision problems, then this will have implications for the accuracy of our model. In particular, it might not be reasonable to expect that voters who can identify that one of their approved candidates will be elected even without their support, which by Lemma 3 implies that they can safely drop this candidate, will also spend computational resources to determine that the first condition in Definition 12 is satisfied.

The previous considerations and the susceptibility results in Section 4.2 demonstrate that information restrictions to prevent manipulation have to be sought below the level of information required to identify that a candidate will certainly be elected. The two positive result presented in Section 4.4 show that (partial) information barriers for PAV can be found below this level.

Chapter 5

Conclusion

The objective in this thesis has been to extend the model of approval-based committee voting to account for strategic manipulation when voters have incomplete information. The purpose of our model is to be able to describe more accurately the knowledge and behavior of real voters who will not always have access to the complete information about the preferences of their fellow voters. The ultimate purpose of our model was to investigate whether there are information restrictions under which voters do not have incentive to misrepresent their preferences.

The context of this thesis consisted of the extensive literature on proportionality for approval-based committee voting rules and the impossibility results that show that strategic manipulation is an inherent problem for proportional voting rules. In this thesis we have given further evidence that strategic manipulation is indeed a critical problem for proportional voting rules, by performing simulation experiments to measure the prevalence of manipulation for three proportional voting rules. Moreover, we have seen that proportional voting rules are susceptible to different kinds of manipulation. The simulation experiments showed that among the types of manipulation we identified, manipulation by free-riding is a significant type of manipulation in terms of prevalence. The third simulation experiment showed that voters who do not have information about the preferences of one or two of their fellow voters may have an incentive to manipulate with increased probability, rather than decreased probability, relative to voters with complete information. The experimental results presented in this thesis further motivate the need for information barriers to strategic manipulation for proportional voting rules.

In Chapter 4 we studied whether there are information restrictions under which voters never have incentive to misrepresent their preferences. The analysis in this chapter has yielded some positive and some negative results. The negative results apply to voting rules that satisfy the axioms *diminishing returns*, *candidate monotonicity*, *minimal proportionality* and *minimal efficiency*. When voters possess the approval score information of the candidates, strategic manipulations in the form of dropping one candidate cannot be completely ruled out for voting rules that satisfy the axioms. Similarly, when voters know the approval set of at least $\frac{n}{k}$ of the voters, incentive to misrepresent preferences is not ruled out. The two positive results in Chapter 4 apply to the PAV voting rule. We showed that PAV is $Zero$-strategy-proof. That is, PAV is not susceptible to manipulation when voters have no information about the preferences of other voters. We also showed that when t is below a certain upper bound, PAV is not susceptible to manipulations that involve dropping one candidate when voters have t-Profile information.

There are several directions of research to continue the current work. One direction is to examine whether sequential PAV and Phragmén's rule satisfy the diminishing returns axiom. In this thesis we have shown that a broad class of w-Thiele methods, which includes PAV, satisfy the diminishing returns axiom. However, it remains an open question whether the other proportional voting rules satisfy this axiom. There is some informal reason to suspect that there is a connection between proportionality and the diminishing returns axiom. We expect that proportional voting rules achieve proportional outcomes by applying the idea of diminishing returns. That is, a voting rule achieves proportional outcomes by giving more weight to voters who are not well represented. Essentially, giving a voter an additional representative is attached greater value when the voter is not well represented. This is also the idea that is captured in the diminishing returns axiom. However, as far as the current work is concerned, this observation is merely informal. It would be interesting to determine whether there are any formal relations between the diminishing returns axiom and proportionality axioms.

Another direction for future research would be to examine the computational complexity of the problems of identifying strategic manipulations under incomplete information for the t-Profile and Approval Score information functions. As has been noted in the discussion of Chapter 4, results about the computational complexity of these problems, as well as the computational complexity of determining whether a candidate will certainly be elected based on the Approval Score information or the t-Profile information, may have implications for the behavior we can expect of real voters. To be able to model the behavior of real voters more accurately, it would be fruitful to study the algorithmic side of strategic voting under incomplete information.

Appendix A

Additional Proofs

A.1 Strictly Concave w-Thiele Methods

In this section we show that every strictly concave w-Thiele method satisfies *candidate monotonicity* and *minimal efficiency.*

Lemma 5. Every strictly concave non decreasing function $w : \mathbb{N} \to \mathbb{R}$ is strictly increasing.

Proof. Let $w : \mathbb{N} \to \mathbb{R}$ be a strictly concave and non decreasing function. Suppose w were not strictly increasing. That is, for some $l' \in \mathbb{N}$ we have $w(l'+1) - w(l') \leq 0$. Then, because w is strictly concave, we would have $w(l'+2) - w(l'+1) < w(l'+1) - w(l') \leq 0$. This would contradict the assumption that w is non decreasing. So, we have $w(l+1) - w(l) > 0$ for all $l \in \mathbb{N}$. Equivalently, $w(l+1) > w(l)$ for all $l \in \mathbb{N}$. $\square$

Proposition 3. Every strictly concave w-Thiele methods satisfies *candidate monotonicity*

Proof. Let $w : \mathbb{N} \to \mathbb{R}$ be a strictly concave and non decreasing function. This means, by Lemma 5, w is a strictly increasing function. Consider any election instance $(\boldsymbol{A}, k)$, voter $i \in N$ and candidate $d \in C$. We have, for all $W \in \mathcal{P}_k(C)$ such that $d \in W$:

$$\begin{aligned} sc_w(\boldsymbol{A}^{i+\{d\}}, W) &= \sum_{i \in N} w(|W \cap A_i|) \\ &= \sum_{i \in N \setminus \{i\}} w(|W \cap A_i|) + w(|W \cap (A_i \cup \{d\})|) > sc_w(\boldsymbol{A}, W) \end{aligned}$$

and for all $W \in \mathcal{P}_k(C)$ such that $d \notin W$:

$$sc_w(\boldsymbol{A}^{i+\{d\}}, W) = sc_w(\boldsymbol{A}, W)$$

Since the w-Thiele method selects the committees with the highest score, we find that, whenever $d \in W$ for all $W \in F(\boldsymbol{A}, k)$ we have $d \in W'$ for all $W' \in F(\boldsymbol{A}^{i+\{d\}}, k)$ and whenever $d \in W$ for some $W \in F(\boldsymbol{A}, k)$ we have $d \in W'$ for some $W' \in F(\boldsymbol{A}^{i+\{d\}}, k)$ □

Proposition 4. Every strictly concave w-Thiele methods satisfies *minimal efficiency.*

Proof. Let $w : \mathbb{N} \to \mathbb{R}$ be some strictly concave non decreasing function. By Lemma 5 we have that w is strictly increasing.

Let $(\boldsymbol{A}, k)$ be some election instance. Define the marginal contribution of a candidate $a \in W$ to committee W as the difference between the w-score of W and the w-score of $W \setminus \{a\}$ on profile $\boldsymbol{A}$. That is, $mc(a, W, \boldsymbol{A}) := sc_w(\boldsymbol{A}, W) - sc_w(\boldsymbol{A}, W \setminus \{a\})$.

Since w is strictly increasing, the marginal contribution of a candidate who is unanimously approved is strictly greater than the marginal contribution of a candidate who is not unanimously approved, for any committee W. As a result, replacing a non-unanimously approved candidate with a unanimously approved candidate will always strictly increase the w-score of the committee. Similarly, the marginal contribution of an approved candidate is strictly greater than the marginal contribution of a candidate who is not approved, for any committee. Replacing an unapproved candidate with an approved candidate will therefore always increase the score of the committee.

Since the w-Thiele method selects the committees with highest score, there will never be winning committees that contain unapproved candidates while there are still approved candidate available, or winning committees that contain non-unanimously approved candidates when there are still unanimously approved candidates available. □

A.2 Proof of lemma 1

Lemma 1. Let $r : \mathbb{N} \to [0, 1]$ be some decreasing function in n, $k \in \mathbb{N}$ some natural number and $N : \{0, \ldots, k\} \to \mathbb{N}$ an arbitrary function. Then for any $l \in \mathbb{N}$ such that $l \leq k$ it is the case that:

$$\frac{\sum_{j=0}^{l} N(j) r(j)}{\sum_{j=0}^{k} N(j) r(j)} \geq \frac{\sum_{j=0}^{l} N(j)}{\sum_{j=0}^{k} N(j)}$$

Proof. Take $\overline{r}_l$ to be defined as follows:

$$\overline{r}_l := r(0)\frac{N(0)}{\sum_{j=0}^{l} N(j)} + \cdots + r(l)\frac{N(l)}{\sum_{j=0}^{l} N(j)}$$

Note that

$$\overline{r}_l \geq r(l) \geq r(l+1) \geq \cdots \geq r(k) \tag{A.1}$$

This is because $\overline{r}_l$ is a weighted average of $r(0)\ldots r(l)$ and we have $r(j) \geq r(l)$ for all $j \leq l$. Moreover:

$$\sum_{j=0}^{l} N(j)r(j) = \Big(\sum_{j=0}^{l} N(j)\Big)\overline{r}_l \tag{A.2}$$

From (A.1) it follows that:

$$\sum_{j=l+1}^{k} N(j)\overline{r}_l \geq \sum_{j=l+1}^{k} N(j)r(j) \tag{A.3}$$

From (A.2) and (A.3) it follows that:

$$\begin{aligned}
\frac{\sum_{j=0}^{l} N(j)r(j)}{\sum_{j=0}^{k} N(j)r(j)} &= \frac{\sum_{j=0}^{l} N(j)r(j)}{\sum_{j=0}^{l} N(j)r(j) + \sum_{j=l+1}^{k} N(j)r(j)} \\
&= \frac{\big(\sum_{j=0}^{l} N(j)\big)\overline{r}_l}{\big(\sum_{j=0}^{l} N(j)\big)\overline{r}_l + \sum_{j=l+1}^{k} N(j)r(j)} \\
&\geq \frac{\big(\sum_{j=0}^{l} N(j)\big)\overline{r}_l}{\big(\sum_{j=0}^{l} N(j)\big)\overline{r}_l + \big(\sum_{j=l+1}^{k} N(j)\big)\overline{r}_l} = \frac{\sum_{j=0}^{l} N(j)}{\sum_{j=0}^{k} N(j)}
\end{aligned}$$

□

A.3 The Condition in the Proof of Theorem 6

In this section we give a short derivation to show that the condition that is used in case (*ii.b*) in the proof of Theorem 6 is satisfied.

Proposition 5. For $n, k \in \mathbb{N}$ such that $n \geq k$, $|B|$ and $|A_i|$ such that $1 \leq |A_i| < |B| < k$ and $\delta := k - |B|$ we have:

$$\lceil n \times \frac{\delta + 1}{k} \rceil + 1 \leq n$$

Proof. Let δ, n and k be as described in the statement of the proposition. Each of the following relations are equivalent:

$$k \leq n$$
$$-2n \leq -2k$$
$$k \times n - 2n \leq k \times n - 2k$$
$$n \times (k-2) \leq k \times (n-2)$$
$$n \times \frac{k-2}{k} \leq n - 2$$
$$n \times \frac{k-2}{k} + 2 \leq n$$

Since $2 \leq |B|$ we have $\delta + 1 \leq k - 2$. So, $k \leq n$ implies:

$$n \times \frac{\delta+1}{k} + 2 \leq n$$

Since $\lceil n \times \frac{\delta+1}{k} \rceil + 1 \leq n \times \frac{\delta+1}{k} + 2$ we have:

$$\lceil n \times \frac{\delta+1}{k} \rceil + 1 \leq n$$

We conclude that the profile constructed in case $(ii.b)$ of the proof of Theorem 6 does not exceed the number of voters available. □

Appendix B

Simulation Experiment Results

B.1 Experiment 1

In this section we present the results of our first experiment. In this experiment we measured, for five different types of manipulation, the proportion of election instances that are susceptible to this type of manipulation when voters have complete information.

	Susceptible	Subset	NON Subset	Free-riding	NON Free-riding
n=5	0.59	0.59	0.09	0.38	0.42
n=10	0.73	0.73	0.18	0.6	0.45
n=15	0.68	0.67	0.2	0.58	0.39
n=20	0.62	0.62	0.2	0.55	0.36

Figure B.1: For every type of manipulation the relative number of election instance $(\boldsymbol{A}, k)$ for which some voter can manipulate sequential PAV for $m = 5$, $k = 3$ and different values of n based on 1000 instances.

	Susceptible	Subset	NON Subset	Free-riding	NON Free-riding
n=5	0.56	0.55	0.16	0.32	0.42
n=10	0.72	0.71	0.33	0.58	0.51
n=15	0.69	0.68	0.36	0.62	0.47
n=20	0.63	0.62	0.34	0.58	0.42

Figure B.2: For every type of manipulation the relative number of election instance $(\boldsymbol{A}, k)$ for which some voter can manipulate Phragmén's rule for $m = 5$, $k = 3$ and different values of n based on 1000 instances.

B.2 Experiment 2

In this section we present the results of our first experiment. In this experiment we measured, for five different types of manipulation, the proportion of profile-voter pairs for which the voter has incentive to manipulate with this type of manipulation, under the assumption that voters have complete information.

	Opportunity	Subset	NON Subset	Free-riding	NON Free-riding
n=5	0.2	0.2	0.02	0.12	0.12
n=10	0.18	0.18	0.03	0.14	0.09
n=15	0.15	0.15	0.03	0.12	0.07
n=20	0.17	0.16	0.03	0.14	0.06
n=25	0.14	0.13	0.03	0.12	0.06
n=30	0.13	0.12	0.02	0.1	0.05
n=35	0.11	0.11	0.02	0.1	0.04
n=40	0.13	0.12	0.03	0.11	0.04

Figure B.3: The relative number of instance $(\boldsymbol{A}, i)$ for which i has incentive to manipulate, for sequential PAV with $m = 5$, $k = 3$ and different values of n based on 2000 instances.

	Opportunity	Subset	NON Subset	Free-riding	NON Free-riding
n=5	0.16	0.15	0.05	0.09	0.11
n=10	0.17	0.16	0.06	0.12	0.1
n=15	0.16	0.15	0.05	0.13	0.08
n=20	0.14	0.13	0.05	0.11	0.07
n=25	0.11	0.11	0.04	0.09	0.05
n=30	0.12	0.11	0.05	0.1	0.06
n=35	0.1	0.09	0.04	0.08	0.05
n=40	0.12	0.11	0.03	0.1	0.04

Figure B.4: The relative number of instance $(\boldsymbol{A}, i)$ for which i has incentive to manipulate, for Phragmén's rule with $m = 5$, $k = 3$ and different values of n based on 2000 instances.

B.3 Experiment 3

In this section, we present the results of the third experiment, in which we measured the proportion of election instances that are susceptible to π-manipulation for the $(n-1)$-Profile IF, the $(n-2)$-Profile IF and complete information.

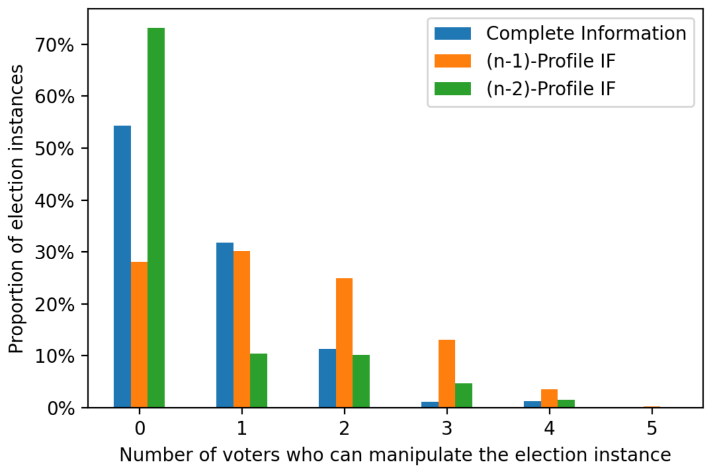

Figure B.5: π-manipulability data for PAV for the election scenario parametrized by $n = 5, m = 4$ and $k = 3$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

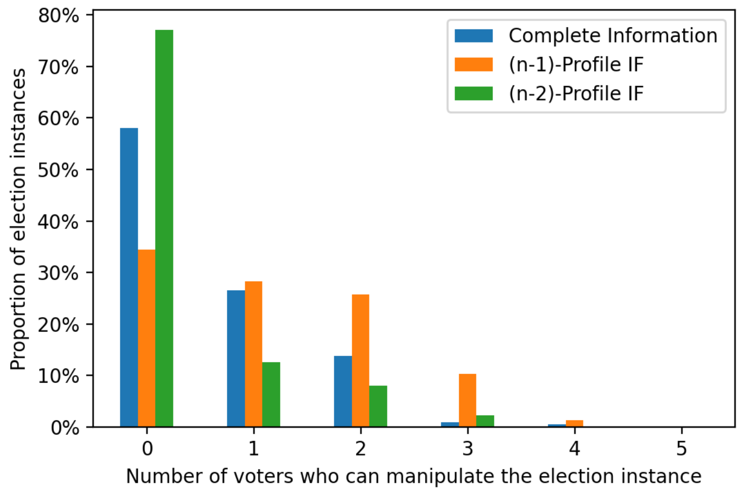

Figure B.6: π-manipulability data for Phragmén's rule for the election scenario parametrized by $n = 5, m = 4$ and $k = 3$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

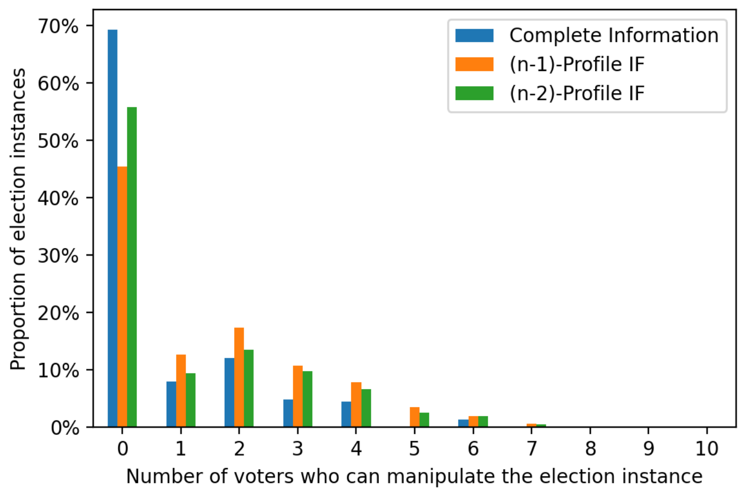

Figure B.7: π-manipulability data for sequential PAV for the election scenario parametrized by $n = 10, m = 3$ and $k = 2$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

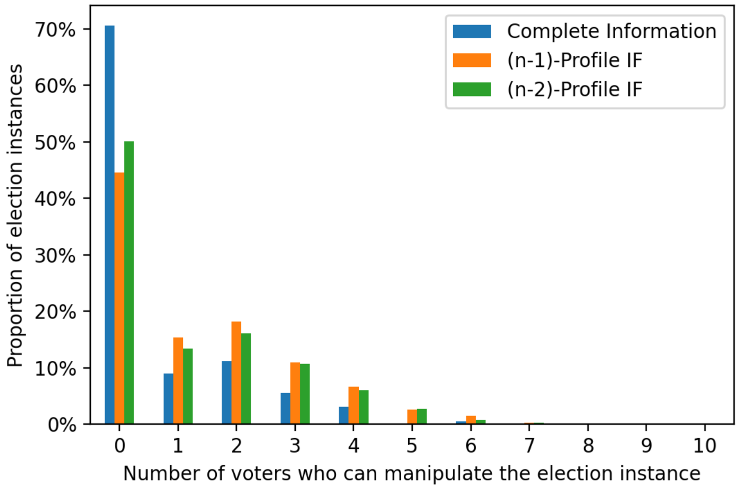

Figure B.8: π-manipulability data for Phragmén's rule for the election scenario parametrized by $n = 10, m = 3$ and $k = 2$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

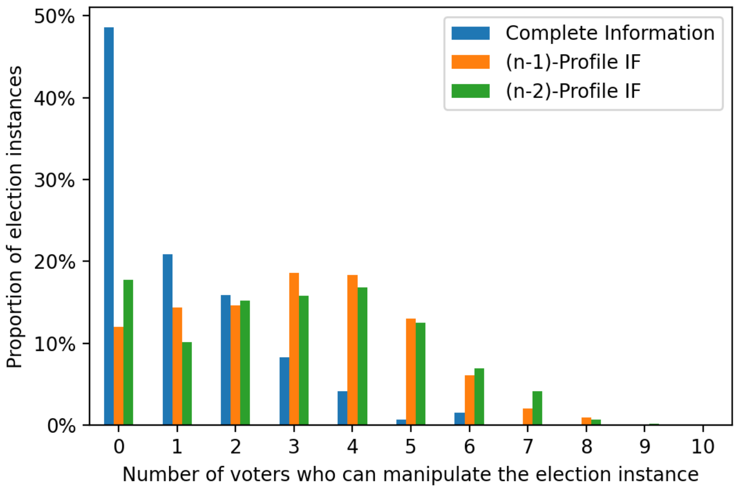

Figure B.9: π-manipulability data for PAV for the election scenario parametrized by $n = 10, m = 4$ and $k = 3$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

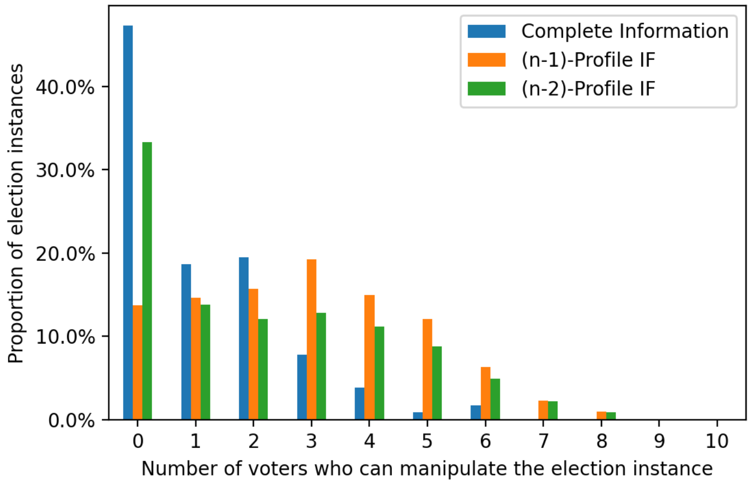

Figure B.10: π-manipulability data for Sequential PAV for the election scenario parametrized by $n = 10, m = 4$ and $k = 3$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

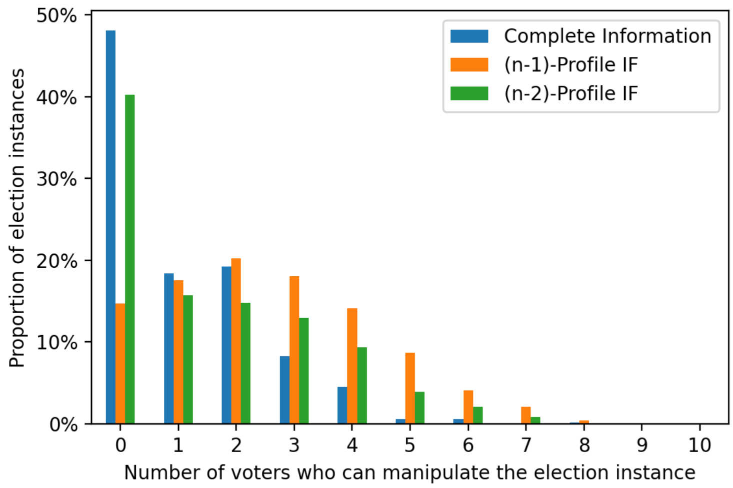

Figure B.11: π-manipulability data for Phragmén's rule for the election scenario parametrized by $n = 10, m = 4$ and $k = 3$. Per number of manipulators the plot shows the proportion of election instances that are π-manipulable by this number of voters. The data is based on 1000 election instances.

Bibliography

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

Aziz, H., Gaspers, S., Gudmundsson, J., Mackenzie, S., Mattei, N., and Walsh, T. (2015). Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems*, pages 107–115.

Botan, S. (2021). Manipulability of Thiele methods on party-list profiles. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 223–231.

Brill, M., Freeman, R., Janson, S., and Lackner, M. (2017). Phragmén's voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 406–413.

Brill, M., Laslier, J.-F., and Skowron, P. (2018). Multiwinner approval rules as apportionment methods. *Journal of Theoretical Politics*, 30(3):358–382.

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

Endriss, U., Obraztsova, S., Polukarov, M., and Rosenschein, J. S. (2016). Strategic voting with incomplete information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 236–242.

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

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

Janson, S. (2016). Phragmén's and Thiele's election methods. *arXiv preprint arXiv:1611.08826*.

Kelly, J. S. (1977). Strategy-proofness and social choice functions without singlevaluedness. *Econometrica: Journal of the Econometric Society*, 45(2):439–446.

Kluiving, B., de Vries, A., Vrijbergen, P., Boixel, A., and Endriss, U. (2020). Analysing irresolute multiwinner voting rules with approval ballots via SAT solving. In *In Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-2020)*.

Lackner, M. and Skowron, P. (2018). Approval-based multi-winner rules and strategic voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18)*, pages 340–346.

Lackner, M. and Skowron, P. (2020a). Approval-based committee voting: Axioms, algorithms, and applications. *arXiv preprint arXiv:2007.01795*.

Lackner, M. and Skowron, P. (2020b). Utilitarian welfare and representation guarantees of approval-based multiwinner rules. *Artificial Intelligence*, 288:103366.

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

Levy, H. (2015). *Stochastic dominance: Investment decision making under uncertainty*. Springer.

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

Phragmén, E. (1894). *Sur une méthode nouvelle pour réaliser, dans les élections, la représentation proportionnelle des partis.*

Pukelsheim, F. (2017). *Proportional representation.* Springer.

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

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

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

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

Skowron, P. (2018). Proportionality degree of multiwinner rules. *arXiv preprint arXiv:1810.08799*.

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

Terzopoulou, Z. (2017). Manipulating the manipulators: Richer models of strategic behavior in judgment aggregation. Master's thesis, ILLC, University of Amsterdam.

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