LIBRARY
Te Tumu Herenga
THE UNIVERSITY OF AUCKLAND

RESEARCHSPACE@AUCKLAND

http://researchspace.auckland.ac.nz

ResearchSpace@Auckland

Copyright Statement

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

This thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of this thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from their thesis.

To request permissions please use the Feedback form on our webpage.
http://researchspace.auckland.ac.nz/feedback

General copyright and disclaimer

In addition to the above conditions, authors give their consent for the digital copy of their work to be used subject to the conditions specified on the Library Thesis Consent Form and Deposit Licence.

Note : Masters Theses

The digital copy of a masters thesis is as submitted for examination and contains no corrections. The print copy, usually available in the University Library, may contain corrections made by hand, which have been requested by the supervisor.

Strategic Manipulation in Voting Systems

Reyhaneh Reyhani

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Computer Science, The University of Auckland, 2013.

Abstract

In this thesis, we are going to study the strategic manipulation of voting rules, mostly scoring rules. In the first part, we focus on naive manipulation, where we have a coalition of manipulators and the other voters vote sincerely. In Section 1.4 we introduce a new measure of manipulability of voting rules, which reflects both the size and the prevalence of the manipulating coalitions and is adaptable to various concepts of manipulation. We place this measure in a framework of probabilistic measures that organizes many results in the recent literature. We discuss algorithmic aspects of computation of the measures and present a case study of exact numerical results in the case of 3 candidates for several common voting rules. In Section 1.5 we study manipulability measures as power indices in cooperative game theory. In Chapter 2, we study the asymptotic behaviour of a model of manipulation called safe manipulation for a given scoring rule under the uniform distribution on voting situations. The technique used is computation of volumes of convex polytopes. We present explicit numerical results in the 3 candidate case. In the second part of the thesis, we adopt a game-theoretic approach to study strategic manipulation. We try to explore more behavioural assumptions for our voters. In Chapter 3, we have an introduction to voting games and different factors such as the available amount of information, voters' strategies and

ability to communicate . In Chapter 4, we consider best-reply dynamics for voting games in which all players are strategic and no coalitions are formed. We study the class of scoring rules, show convergence of a suitably restricted version for the plurality and veto rules, and failure of convergence for other rules including k-approval and Borda. In Chapter 5, We discuss a new model for strategic voting in plurality elections under uncertainty. In particular, we introduce the concept of inertia to capture players' uncertainty about poll accuracy. We use a sequence of pre-election polls as a source of partial information. Under some behavioural assumptions, we show how this sequence can help agents to coordinate on an equilibrium outcome. We study the model analytically under some special distributions of inertia, and present some simulation results for more general distributions. Some special cases of our model yield a voting rule closely related to the instant-runoff voting rule and give insight into the political science principle known as Duverger's law. Our results show that the type of equilibrium and the speed of convergence to equilibrium depend strongly on the distribution of inertia and the preferences of agents.

This thesis is based on the results of the following papers [1], [2], [3], [4] and [5].

Acknowledgements

My deepest gratitude is to my main supervisor Dr. Mark C. Wilson for his continuous support and motivation. It was really a great experience working with him. His guidance helped me a lot in doing research and writing of this thesis. I would also like to thank my co-supervisor Dr. Geoffrey Pritchard for his valuable comments.

I would like to acknowledge the New Zealand Institute of Mathematics and its Applications for supporting me financially during my PhD.

I am so indebted to my parents Ahmad and Masoumeh for all they have done for me in my life. Their encouragement, love and support has been a constant source of motivation and energy and has enabled me to succeed in all stages of life.

My husband Javad was my real companion in the last six years. He patiently stood by me and supported me in all the challenges of life. I could not finish this thesis without his encouragement and valuable comments. I would like to express my gratitude for all his sacrifice.

I also would like to thank my 1.5 year old son Ryan for giving me the opportunity to feel the joy of being a mum. The day he was born is the most memorable day of my life. I dedicate this thesis to my parents, my husband and my son.

Contents

Part I

Strategic Manipulation

Chapter 1

The Manipulability of Voting Rules

1.1 An introduction to computational social choice

The novel field of computational social choice is an interdisciplinary field of study at the interface of social choice theory and computer science. It involves studying social choice mechanisms like voting systems and fair division systems with two approaches. In the first approach, computer science offers the computational complexity, algorithmic, mathematical and quantitative techniques for studying decision theory, social choice, welfare economics and game theory study of voting systems. In the second approach, social choice ideas are applied to computer science-related contexts, for example, in artificial intelligence, multiagent systems, social networks and agent optimisation. Chevaleyre et al. have provided a short introduction to this topic [6].

Some of the more studied problems in this topic are electing an alternative, allocating resources, reaching consensus, forming coalitions, aggregating judgements and beliefs.

In this thesis, we concentrate on the topic of strategic manipulation in voting systems. Voting as an aggregating method is widely used in collective decision making and network design. Voting rules can show some undesirable behaviour such as being vulnerable to strategic manipulation. We first explain our voting setup in Section 1.2 and then briefly explain strategic manipulation by an example in Section 1.3 . In Section 1.4, we study a new measure of manipulability of voting rules, and in Section 1.5 we study the relation between power measures and manipulability measures. Section 1.6 deals with relevant literature review and Section 1.7 discusses some future directions. In Chapter 2, we study the probability of safe manipulation. Chapter 3 discusses the integration of game theory in social choice where all voters are strategic. We study best reply dynamics and coordination via polling under uncertainty respectively in Chapters 4 and 5. Finally we have Chapter 6 which discusses a summary of the thesis and some future directions.

1.2 Basic terminology

Consider a set $V = \{v_1, \dots, v_n\}$ of agents (the **voters**) choosing from a given set $C = \{c_1, \dots, c_m\}$ of alternatives (the **candidates**). Each voter has an **opinion** or preference ranking (a complete strict linear ordering of the candidates). This convention is quite common in the field of voting theory.

The list of voters' preference orders $(R_1, \dots, R_n)$ forms the **sincere profile**. Each voter submits a linear ranking R'_i, which may or may not be the same as his sincere opinion, and this gives the **expressed profile**. For example, consider a set of 5 voters $V = \{v_1, \dots, v_5\}$ and a set of 3 candidates $C = \{a, b, c\}$ with sincere profile $(abc, abc, bac, cab, bac)$. The expressed profile $(abc, abc, bca, cab, bac)$ results from the voter v_3 not voting sincerely, while the other voters do.

A **social choice function** (also called a resolute voting rule) is a function that maps each profile to a single candidate, whereas a social choice correspondence (also called a voting rule) outputs a subset of the candidates. A key feature of most commonly used voting rules is **anonymity**: the function value is unchanged if voters are permuted, so the rule treats voters equally. In this case, the profile can be represented more succinctly as a **voting situation**, where we simply list the numbers of voters with each of the possible opinions. For example, for three candidates (a, b, c), with the standard order $abc, acb, bac, bca, cab, cba$ of opinions, the 6-tuple $\sigma = (n_1, \ldots, n_6)$ represents a voting situation with n_1 voters having preference order abc, etc. In the example above, this succinct input for the expressed profile would be $\sigma = (2, 0, 1, 1, 1, 0)$.

A voter v may try to manipulate the election result by submitting an expressed opinion that differs from his sincere opinion, so as to gain an outcome that v prefers to the sincere outcome. The fundamental result of Gibbard [7] and Satterthwaite [8] implies that for anonymous rules, provided that $m \geq 3$ and $n \geq 2$, some voter in some voting situation can succeed in such an attempt. This theorem shows that for a voting rule with more than 3 candidates, strategic manipulation can happen with an individual voter under some natural conditions such as non-imposition which conveys that each candidate can win under some conditions and non-dictatorial condition which means the result of the election is not based on the highest ranking alternative of one of voters.

A common class of anonymous voting rules consists of the (positional) **scoring rules**. For each m, a scoring rule is defined by a weight vector $(w_1, \ldots, w_m)$

with $w_1 \geq w_2 \geq \cdots \geq w_m$, and each voter gives score w_1 to his top-ranked candidate, w_2 to the next, etc. Uniqueness of the weight representation is obtained by imposing the restriction $w_1 = 1, w_m = 0$. The candidate with the highest total score wins. The most commonly used voting rules are listed below.

- Plurality rule, defined by the weight vector $(1, 0, 0, \ldots, 0)$;
- Borda's rule, defined by the weight vector $(m-1, m-2, \ldots, 1, 0)$;
- Antiplurality rule (veto rule), defined by the weight vector $(1, 1, \ldots, 1, 0)$.
- k-approval rule, defined by the weight vector $(1, 1, \ldots, 1, 0, \ldots, 0)$ (the number of 1's is exactly k).
- Instant-runoff rule, If no candidate receives a majority of the first choice, the candidate with the fewest number of votes is eliminated and the ballots cast for that candidate are redistributed to the continuing candidates according to the voters indicated preference. This process is repeated until one candidate obtains a majority.

In an election with approval rule, voters should decide whether they approve or disapprove a specific candidate. In k-approval, they should approve exactly k candidates. In fact, 1-approval rule is the same as plurality rule and $m-1$-approval rule is the same as antiplurality rule. For example, in an election with 4 candidates a, b, c and d and 2-approval voting rule, a voter with vote $abcd$, approves candidates a and b.

Another common class of anonymous rules consists of the Condorcet consistent rules based on the pairwise majority relation. We deal with the Copeland rule as a representative. For each pair of candidates a and b, the pairwise score $p(a, b)$ of a with respect to b equals the number of voters who rank a above b. The Copeland score of alternative a is given by $s(a) = \sum_{b \neq a} \operatorname{sign}(p(a, b) - p(b, a))$. The highest scoring candidate is the winner.

To ensure a unique winner in every situation, elections using scoring rules usually require an additional rule to deal with the possibility of tied scores for the first place. Different tie-breaking rules have been used in this context. For example, for deterministic tie-breaking rules, there is a fixed arbitrary order on candidates, and the winner is the first of the tied candidates with respect to this order. Random tie-breaking is more favourable for reasons of **neutrality** (symmetry between candidates) and tractability. However, random tie-breaking does not define a social choice function, because of nondeterminism, but rather a social choice correspondence.

For randomized tie-breaking, we choose one candidate uniformly at random. Scoring rules are neutral with this convention. When n is large, the probability of a tie occurring for a scoring rule under any of the most commonly studied preference distributions is asymptotically negligible, so tie-breaking conventions are not important. However, these assumptions can make major differences for small values of n. Copeland's rule must also deal with ties, and in two ways. First, the pairwise majority relation can (when n is even) result in a tie; the standard choices are to award both candidates involved 0, but other choices are possible. Second, the Copeland scores of candidates may be tied. In this case, we again can use random or deterministic tie-breaking as for scoring rules. Copeland's rule can have an asymptotically non-negligible fraction of ties under some preference distributions, and our tie-breaking assumptions definitely affect the values of the manipulability measures [9].

In some elections voters carry different weights that can represent their power in making decision, for example their amount of stock in share holders or constituency sizes. However, in most discussions of this thesis, voters have the same weight and we have unweighted voters.

1.3 Strategic manipulation

Strategic misrepresentation of a voter's true preferences, as a way of obtaining an outcome preferable to that which would be expected by voting sincerely, dates back thousands of years [10] and has generally been considered socially undesirable. This topic has been recently considered in many papers in computational voting theory. We will discuss some of them in Section 1.6.

Example 1.1. *Consider plurality rule and the following preference orders for 2000 US presidential election in Florida ($a \succ b$ means preferring a to b)*

49% *Bush* $\succ$ *Gore* $\succ$ *Nader,* 20% *Gore* $\succ$ *Nader* $\succ$ *Bush,* 20% *Gore* $\succ$*Bush* $\succ$ *Nader,* 11% *Nader* $\succ$*Gore* $\succ$ *Bush.*

If everyone votes sincerely, Bush will win this election. However, it would have been in the interest of the Nader supporters to misrepresent their preferences and vote for Gore. In that case, Gore will win provided others vote sincerely. This misrepresentation is called strategic manipulation and the Nader supporters form the coalition of manipulators.

Manipulability of social choice correspondences is a tricky subject and one has to have an order on subsets to define it. That is, one has to extend somehow preferences over alternatives to preferences over the sets of them. It can be done in numerous ways and there is a survey by Bossert, Barbera and Pattanaik in [11] about possible ways to do this. In this thesis, we use stochastic dominance improvement for defining the manipulation of sets of candidates.

Over the last few decades many papers (e.g. see [12, 13, 14] for a summary) have been published in the following framework: choose a set of social choice rules; choose a probability distribution over the set of preference orders; compute the probability P of a randomly chosen situation being manipulable with n voters; conclude which rules are asymptotically the best, that is, those for which P is least, for large n. The results depend strongly on

- the measure of manipulability which will be discussed in Section 1.4 and Chapter 2.
- assumptions on game-theoretic sophistication of the voters, and the information available to them which will be discussed in Chapters 4 and 5 .

For comparing different voting systems regarding the manipulability, different metrics have been used such as probability, complexity and social welfare and utility. Computational hardness of manipulation has been studied for voting rules and fair division mechanisms by some techniques as a barrier to susceptibility to manipulation. The methods which are used to study the computational voting rule consist of worst-case analysis, average-case, heuristic and approximate algorithms. When all voters behave strategically, game theory predicts the result of voting game by studying the outcome of interactions amongst multiple agents.

By reviewing the papers in this topic, it becomes clear that most papers consider strategic manipulation as an undesirable behaviour which should be minimized.

For example, Stensholt believes that the strategy most damaging to many preferential election methods is to give insincerely low rank to the main opponent of one's favourite candidate [15]. However, a small number of papers write in praise of manipulation [16] and believe that by strategic manipulation, the total social welfare in fact increases [17]. In this thesis, we consider strategic manipulation as an inevitable fact, and try to have a better understanding about this phenomenon.

1.4 A new measure of manipulability of voting rules

In almost all of the social choice literature, it is regarded as desirable to minimize the occurrence of manipulability of voting rules, that is, to design a social choice mechanism that incentivizes sincere expression of voter preferences as much as possible. Of course, the Gibbard-Satterthwaite theorem and related results [7, 8, 18] imply that completely avoiding manipulability has drastic consequences, and leads under very mild hypotheses to dictatorship. Thus many authors have tried to measure the manipulability of voting rules, typically by quantifying the probability of such an event, under various assumptions on the distribution of voter opinions (see Section 1.6 for detailed discussion of relevant literature and Section 1.4.1 for formal definitions). More recently the idea of using measures based on computational complexity has arisen (usually with a somewhat different definition of manipulability), leading to substantial activity in the "computational social choice" community.

Successful manipulation of an election, even in the case considered in the present article when the manipulators are opposed only by naive, sincere voters, requires considerable computational effort. To assemble a manipulating coalition, we must discover the preference rankings of voters, convince them to join the coalition, compute their strategy, and enforce their implementation of that strategy. Each of these becomes harder as the coalitions involved become larger.

However, measures based on the size of the manipulating coalition have been relatively little explored in the literature. By far the most commonly used measure is simply the probability that a random profile (chosen according to some standard distribution of voter preferences) allows some manipulation. The measures based on worst-case complexity mostly do not measure coalition size directly. Also, they are inherently crude, as they are defined only up to polynomial-time equivalence. This makes them less useful for comparing specific rules with respect to manipulability.

Furthermore, recent results using various models of manipulation show that at least for the most commonly studied distributions of preferences, there is a phase transition in the probability of manipulability as the coalition size grows relative to the total voting population, yet say little about how to compare rules near that threshold [19] .

Our contributions

We introduce (Section 1.4.1) a new general probabilistic measure of susceptibility to manipulation, describe its basic properties, and argue that it allows for more detailed comparisons of voting rules than existing measures. We investigate its values in detail in the 3-candidate case (Section 1.4.3) for several scoring rules and a representative Condorcet consistent rule, Copeland's rule. This is done for each of two standard probability models for voter preferences. We also investigate the relationship between the new measure and existing measures and put them in a common framework, thereby unifying much of the literature. We discuss the computation of these measures in detail and present several algorithms (Section 1.4.2).

1.4.1 Definition of the measures

We discuss three types of measures of manipulability of voting rules. All our measures are probabilistic and depend on a probability model for the distribution of opinions in the voter population. We consider in our numerical results two commonly studied distributions: the uniform distribution on profiles (known as the Impartial Culture hypothesis) and the uniform distribution on voting situations (known as the Impartial Anonymous Culture hypothesis). However the definitions make sense for any distribution.

The model of manipulation

Fix a voting rule. We define manipulability of a voting situation in stepwise fashion as in [20]. Our definition implies that, for example, a strategic vote by a voter with preference bac which changes the winner from a to c is not a valid manipulation. The result of the election must not only be changed, but changed in a way that incurs no loss to the manipulator. Other definitions are sometimes used in the literature. For example, the concept of **threshold manipulation** (where we promote b above a, ignoring the possibility that c might thereby overtake both of them) is studied in [21]. This is related to the idea of **destructive manipulation** used in many papers (we only care about defeating a, not who ends up winning).

However, the concept we define here (sometimes called **constructive manipulation**) is more standard.

A related concept, (unit cost) **bribery**, removes any constraint on the opinion of the manipulating voter about the new profile [22], [23], [22], [24] and [25]. In swap bribery, voters are willing to manipulate but not if that requires to depart too much from their sincere vote. In other words, their motivation for manipulation depends on the deviation from sincere ballots which is at most $m - 2$ for plurality. Campaign management is another type of strategic behaviour where the manager of campaign tries to make his desirable candidate win the election. He offers money to other voters for bringing that candidate forward. The amount of offered money depends on the number of changes the voter needs to apply. Campaign management for approval voting is considered in [26] .

Another type of strategic behaviour is **control** by adding or deleting voters or candidates [27]. In multi-mode control, we have 2 or more types of control actions at once. Agenda control, happens by adding small number of spoiler candidates. Teaming happens when adding more candidates actually helps the chances of any of them winning as can occur in Borda rule. In election control, only a small number of voters are added or deleted and the number of candidates is fixed.

Cloning is a type of control where the manipulator can replace each candidate by one or more new candidates. In this model, different voting rules show different reactions. For example, cloning for a fixed candidate can be useful in some voting rules such as antiplurality or be useless in some of them such as plurality. However, in some rules different cloning situations behave differently. For example, for Borda and Copeland's rule some cases are useful for that candidate and some cases are useless [28].

In some models of strategic manipulation the coalition of manipulators cast their votes after sincere voters. They choose their votes in a way that current winner changes. The sincere voters are always naive and just vote sincerely. This type of strategic behaviour has been used in computational social choice as strategic manipulation and is called **possible winners** problem [29, 30]. Zuckerman et al study coalitional constructive weighted manipulation of this model in [31]. They study whether there is a way for the manipulators T with weight vector W to choose an action profile which make alternative p win the election. In this model the sincere preference order of manipulators are not known, and ties are broken adversely to the manipulators [31] .

Pattanaik discusses two types of manipulation: **counter-threat** and **reaction** [32]. Suppose the sincere outcome is a, then a voter with preference order bac, tries to manipulate in favour of b. In counter-threat model, the other voters try to punish him by making a worse result happening for him (making c win the election). Therefore, in this case a supporters just think about punishing the person that decides to manipulate not maximising their own utility. In reaction model, the other voters just decide to return the result to its sincere situation. Therefore, c supporters do not really care about making c win the election and just try to return the result to its sincere one (a becomes winner).

For single-winner outcomes with no ties, it is clear how to define the traditional definition of strategic manipulation.

Definition 1.2. *Fix a voting rule. Suppose that profiles π, π' each yield a unique winner, respectively c, c'. Then π' is* ***preferred*** *by voter v to profile π if and only if c' is no lower in v's preference order than c. If $c' \neq c$, so that c' is higher than c, then we say π' is* ***strongly preferred*** *to π.*

Remark 1.3. *Note that in our situation where indifference is not allowed (voters must break all ties between candidates before submitting their ordering), if c' is preferred to c, but not strongly preferred, then $c' = c$, so the concept "preferred" seems pointless at the first sight. However when we consider coalitions below, this distinction makes more sense, and we keep it in order to have consistency.*

If there is no unique winner, then deciding whether one outcome is preferred to another requires extra assumptions (essentially, we must extend the previous definition to preferences over sets of candidates). In our numerical results it is convenient to use uniform random tie-breaking and a particular such extension which we now describe. We again stress that the particular choice made here is not essential to the definitions of the new measures in Section 1.4.1.

Definition 1.4.

Let π be a profile. We say that π' is ***preferred*** *to π by voter v if and only if for each $k = 1 \ldots m$ the probability of electing one of v's most-favoured k candidates under π' is no less than under π. (If $\pi' \neq \pi$ the condition implies that this probability will be strictly greater for some k.)*

Remark 1.5. *Another way of stating this is to say that the probability distribution that describes the probability of each candidate winning under π is (first-order) stochastically dominated by the analogous distribution for π'. Equivalently, for every utility function that induces the preference order of v, the expected utility for v under π' is greater than the expected utility for v under π.*

Example 1.6. *(preferring one profile to another) Suppose that in profile π the outcome is that a and c tie as the winner, in profile π' the outcome is that b is the absolute winner, and in π'' the outcome is that a and b tie as the winner. The* *distribution of winning probability on* (a, b, c) *is* $(1/2, 0, 1/2)$ *for* π*,* $(0, 1, 0)$ *for* π' *and* $(1/2, 1/2, 0)$ *for* π''*. Thus, taking* $k = 1$ *in the definition, we see that* π' *is not preferred to* π *by a voter with sincere opinion* abc*. Also, taking* $k = 2$ *shows that* π *is not preferred to* π' *either. However* π'' *is preferred to both* π *and* π'*.*

We can now proceed to the remaining definitions.

Definition 1.7. *(i) A subset* $X \subset V$ *is a* ***manipulating coalition*** *at the profile* π *if and only if there is a profile* $\pi' \neq \pi$ *which agrees with* π *on* $V \setminus X$ *and is preferred to* π *by all members of* X*, and strongly preferred by some member. A manipulating coalition is* ***minimal*** *if it does not contain any proper subset that is also a manipulating coalition.*

(ii) A rule is ***manipulable at the profile*** π *if and only if there exists a manipulating coalition at this profile.*

(iii) An anonymous rule is ***manipulable at a voting situation*** σ *if and only if there exists a profile* π *giving rise to* σ*, at which the rule is manipulable.*

Example 1.8. *(manipulation) Consider the Borda rule, given by the weight vector* $(2, 1, 0)$*, and the voting situation with 2* abc*, 2* bac*, 2* bca*, 3* cab *voters. If one of the* cab *voters votes strategically as* acb*, then* a *and* b *tie. The new outcome is preferred by that voter because the winning probability distribution on the candidates has changed from* $(0, 1, 0)$ *to* $(1/2, 1/2, 0)$*.*

Example 1.9. *(manipulation in favour of bottom-ranked candidate) Consider the plurality rule, given by the weight vector* $(1, 0, 0)$*, and a voting situation having 4* abc*, 3* bca *and 2* cab *voters. The sincere winner is then* a*. There is no manipulating coalition in favour of* b *since the only voters preferring* b *to* a *are already* *contributing the maximum score to b and the minimum to a. However, if the bca voters all vote strategically as cba, then c wins.*

Example 1.10. *(manipulation possible in more than one way) Consider an election with 3 abc, 2 cba, 2 bca voters, and scoring rule plurality. The sincere scores are $(3, 2, 2)$. If both cba voters change their votes to bca in favour of b, we have a manipulating coalition with size 2 in favour of b. Also, we can consider a manipulating coalition with size 2 in favour of c. If both bca voters change their votes to cba, then the winner is c.*

A manipulating coalition may contain members whose manipulating strategy is to vote sincerely. The extreme case is as follows.

Example 1.11. *(the maximal coalition in favour of a candidate) For each candidate b other than the sincere winner a, the maximal coalition in favour of b consists of all voters having preference orders that rank b above a. Since voters in a manipulating coalition may vote sincerely, it follows that there exists some coalition that can manipulate in such a way as to make b the winner if and only if the same result can be achieved by the maximal coalition.*

Removing those members who vote sincerely still gives a manipulating coalition (which we might call an *active* coalition, although this term is not standard and will not be important here). A subset of voters contains a manipulating coalition if and only if it contains a minimal manipulating coalition.

Example 1.12. *(minimal coalitions) Consider the scoring rule with weights $(10, 9, 0)$ and a voting situation with 10 abc, 6 bac, 5 cab and 5 cba voters. The sincere result has the scores of a, b, c respectively being $(199, 195, 100)$. Consider manipulation* *in favour of* b. *If one of the* bac *voters changes to* bca, *the new scoreboard will be* $(190, 195, 109)$. *So it is a minimal manipulating coalition of size* 1, *and clearly also minimum. By contrast, if 4* cba *voters change their votes to* bca, *the new result will be* $(199, 199, 96)$. *This is also a minimal manipulating coalition but not a minimum one.*

Probabilistic measures of manipulability

We first fix a number n of voters and a probability distribution on the possible profiles (or voting situations). Let $\Sigma := \Sigma^n$ denote the set of of all voting situations equipped with a given probability measure and let S denote a sample from this distribution.

The first measure concerns the logical possibility of manipulation.

Definition 1.13.

$$P = \Pr\left[\textit{there is some coalition that can manipulate at } S\right].$$

This simple measure has been used extensively in the social choice literature. It is relatively simple to compute for standard rules and preference distributions, but fails to measure the computational effort required to assemble and coordinate a manipulating coalition. It is entirely possible *a priori* that two rules may have the same value of P, yet the manipulations for one require much effort (the recruitment of large coalitions of manipulating agents, perhaps with rare preference orders) while those for the other are relatively straightforward, in every voting situation.

Another measure takes account of the number of manipulators required.

Definition 1.14.

$$M = \min \{k \colon \textit{there is some coalition of size } k \textit{ that can manipulate at } S\}.$$

We would like to consider the expected value $E_S[M]$. However, this does not make sense because certain situations are not manipulable by any coalition, and so M is *defective*. We could therefore, consider $E_S[M \mid M < \infty]$. However it is *a priori* that this may be rather small for rules that are almost never manipulable, and larger for rules that are often manipulable. More information is obtained by considering the distribution function of M. For each k with $1 \leq k \leq n$, we consider $\Pr(M \leq k)$. This is precisely the probability that a randomly chosen voting situation can be manipulated by k or fewer agents.

The measure M allows us to consider the greater work required by larger coalitions. For example, the communication cost between coalition members may grow as M^2, if secret negotiations must be individually carried out. However, it is *a priori* possible that two rules may have the same value of M for every voting situation, yet one has very few manipulating coalitions of size M, while the other has many, in every voting situation. (See Example 1.18 below.)

A third measure, which is new as far as we know, is obtained by considering both the sizes and the prevalence of the manipulating coalitions. Both of these aspects are captured by the informational effort required to discover a manipulating coalition via the following procedure. We assume that although a potential instigator of manipulation knows the distribution of opinions (in other words, the sincere voting situation), he does not know which agents hold which opinions. We assume that such a person must simply interview agents one by one at random, until he has enough agents to form a manipulating coalition. This gives a random variable equal to the number of such queries.

Definition 1.15. *Let $V_1, \ldots, V_n$ be agents sampled without replacement from the set V of all agents, independently of S. Equivalently, $(V_1, \ldots, V_n)$ is a random ordering of V, with all possible orderings being equally likely, representing the order in which agents are queried. Let*

$$Q = \min\{k : \{V_1, \ldots, V_k\} \text{ contains a manipulating coalition at } S\}.$$

Note that Q is a random variable both because it depends on S and because it depends on $(V_1, \ldots, V_n)$. This random variable is in general defective, and is defined to be $+\infty$ if no manipulation is possible for S. In other words, we want Q to have a finite value. If k is the number of queries required to find a coalition, or determine that there isn't one, then $k = n$ wouldn't make sense, because we might find a coalition exactly after n queries. So $k = n + 1$ is the next value and we can use this to mean "not found".

Remark 1.16. *The dynamic query interpretation seems reasonable to us: it seems not unreasonable to assume that an instigator knows the voting situation (through polling) but not the exact profile. However, those readers who remain unconvinced will see in Section 1.4.2 that there is alternative, static interpretation of Q that does not depend on such a story.*

We illustrate these definitions using the following examples.

Example 1.17 (values of Q)**.** *Consider a setup with 2 agents, scoring rule Borda and 3 alternatives. There are 21 different voting situations, but by using symmetry,* *we need to only consider those with* $s(a) \geq s(b) \geq s(c)$ *(where for example* $s(a)$ *denotes the score of alternative* a*). It can be seen as in Section 1.4.2 that of these, only the voting situation* $(1, 0, 1, 0, 0, 0)$ *is manipulable. The sincere scoreboard is* $(3, 3, 0)$ *but if the* bac *agent changes his vote to* bca *then the result becomes* $(2, 3, 1)$*. Similarly, the* abc *can change to* acb*. Thus for this voting situation, we make* 1 *query with probability* 1 *so that* Q *is deterministic and equals* 1*. Now we weight this voting situation according to the culture. Under IAC the probability of such a voting situation will be* $3/21$*, so the values* $\Pr(Q \leq 1)$ *and* $\Pr(Q \leq 2)$ *are each* $1/7$*. Under IC, the situation corresponds to* 6 *profiles, so the weight is* $6/36$*, and the values of* $\Pr(Q \leq 1)$ *and* $\Pr(Q \leq 2)$ *are each* $1/6$*.*

Example 1.18. *(difference between* M *and* Q*) Consider the voting situation with* 6 cab *and* 2 bac *agents. Under the antiplurality voting rule, the sincere scoreboard is* $(8, 2, 6)$ *and the winner is* a*. There are no manipulating coalitions in favour of* b *but manipulation in favour of* c *is possible. The value of* M *is 2: if two of the* cab *agents vote* cba*, the new result is* $s' = (6, 4, 6)$*. If our first two queries discover* cab *agents, then* $Q = 2$ *otherwise,* $Q = 3$ *or* 4*. The expected value of* Q*, conditional on this voting situation, is* $18/7 \approx 2.57$*.*

Now consider the same situation under the $(3, 2, 0)$ *scoring rule. The sincere result is* $(16, 6, 18)$ *and the winner is* c*. Manipulation in favour of* b *is impossible, but we can manipulate in favour of* a *if the two* bac *agents change their votes to* abc*. Here again* $M = 2$*, and* Q *can have any value between* 2 *and* 8*. The expected value of* Q*, conditional on this voting situation, is* 6*.*

In this voting situation, both rules admit the logical possibility of manipulation by coalitions of two or more agents. However, such manipulating coalitions are much more prevalent under the antiplurality rule, because they involve a more *numerous type of agent. This difference between the rules is captured by Q.*

Analogues for other models of manipulation

The measures P, M, Q above depend only on the concept of manipulating coalition. If we change the model of manipulation, we obtain the obvious analogues of those measures. We discuss the case of bribery here, and leave other cases to the reader (for example, threshold manipulation). We denote the bribery-based analogues of the measures by P', M', Q'. The measure P' is rather uninteresting. Clearly, by bribing sufficiently many voters, we may make any given candidate win, provided the voting rule satisfies the nonimposition property (each candidate can win in some profile). Thus for most commonly used voting rules $P' = 1$ for each n and m. However, P' would be interesting if it is limited to a budget. The measure M' is more interesting, giving the minimum possible number of voters to bribe in order to change the result (and it is always finite, given nonimposition). For example, for plurality, M' equals the difference in scores between the first- and second-ranked candidates. The measure Q' gives the number of queries involved in determining a minimal set of agents who must be bribed.

Relations between the measures

We restrict to manipulation here; the analogous measures for different models satisfy the analogous relations.

We denote the distribution function of M by F_M, and analogously for Q, so that

$F_M(k) = \Pr(M \le k)$, etc. Note that

$$F_Q(k) \le F_M(k) \quad \text{for each } k;$$
$$F_Q(n) = F_M(n) = P.$$

Thus F_Q and F_M contain strictly more information than P. It does not appear that F_Q is strictly more informative than F_M, since $(\Pr(M \le k))_{k=1}^n$ cannot be recovered from $(\Pr(Q \le k))_{k=1}^n$. However, for a fixed voting situation S, $F_{M|S}(k)$ is either 0 or 1, and the smallest k for which it takes the value 1 is also the smallest k for which $F_{Q|S}(k) > 0$. Furthermore we have already seen that the value of M does not determine the entire distribution of Q on a given voting situation. We thus have strictly more information from Q than from M in this conditional sense.

We can unify the definitions of M and Q by considering a trivial query model for M. We assume that in this case we know all the voters' preferences, in other words the sincere profile, and our "query" consists of simply approaching a voter and inviting him to join a coalition (we assume that our invitations are never refused). We would make precisely M queries in order to minimize effort. Thus the values of $F_Q(k)$ and $F_M(k)$ correspond to the probability that we can form a coalition after k queries in the case of no extra information (only the voting situation) and full information (the complete profile), respectively. Analogues of these that consider various types of partial information could be considered, but they appear less compelling to us and we do not pursue them here.

We have already seen that M and Q can differ. It is easy to construct a rule for which M and Q differ enormously, if we allow non-anonymous rules.

Example 1.19. *Consider a rule ("oligarchy") that fixes 3 voters and lets them decide the result using plurality, no matter what the other voters do. In this case M is at most 1 for each manipulable situation, but Q will with high probability be of order n.*

The relation between M and Q is further clarified by considering minimal manipulating coalitions. In a minimal manipulating coalition, no member votes sincerely and all of them must act together in order to manipulate. Clearly every minimum size coalition is minimal, but the converse is not true in general.

The definition of Q implies that when the query sequence terminates, we have for the first time in that sequence found a set of voters that contains a minimal manipulating coalition. Let μ be the smallest size of such a coalition; like Q, μ is a random variable. Also $Q \geq \mu \geq M$. Thus the excess of Q over M measures not only how many wasted queries we make, but also the difference between μ and M. If σ is a voting situation in which even one minimal coalition of size larger than M exists, then $E[Q \mid S = \sigma] \geq E[\mu \mid S = \sigma] > E[M \mid S = \sigma]$.

Example 1.20. *In Example 1.12, the minimum coalition size is 1 but there exists a minimal manipulating coalition of size 4. Thus $E[Q] > M$, conditional on this voting situation.*

We now show that there are anonymous rules for which $M|S$ and $Q|S$ can be very different.

Example 1.21. *Consider the plurality rule and denote the sincere winner by a, and let b be some other candidate. Let x denote the number of voters who rank a first, y the number who rank b first, and z the number who do not rank b first, but* *who rank b over a (thus there are $n-x-y-z$ voters who rank neither a nor b first, but rank a over b). Manipulation in favour of b is only possible if the voters of the last type express a preference that ranks b first, and this can succeed if and only if $y + z \geq x$. The minimal coalition size in this case is $x - y$. The query sequence ends when we have found $x - y$ elements from the set of z elements above.*

This has the flavour of a coupon-collector problem. If we assume that n is very large compared to z, then the expected length of the query sequence is well approximated by $(x - y)n/z$. The ratio of $E[Q]$ to M is then not bounded by any constant factor, even for a fixed number of candidates.

Remark 1.22. *Based on the analysis of scoring rules in [33], we believe, via a heuristic argument, that for the IC preference distribution, the ratio $E[Q]/M$ will be bounded by a constant depending only on m and the weight vector, outside a set of exponentially (in n) small probability.*

1.4.2 Computation of the measures

Algorithms

All the measures discussed so far can be computed for anonymous rules in time that is polynomially bounded in n, provided that m remains bounded. The rest of this section elaborates on this theme. Not surprisingly, it seems that P is easier to compute than M, which is easier than Q. We give several algorithms.

We consider here only algorithms that first compute the value of the measure conditional on each voting situation, and aggregate this according to the chosen culture. There may exist other algorithms that are more efficient and act by considering several voting situations at once, but we have not yet found any. The number of possible voting situations is the number of solutions in nonnegative integers to the equation $n_1 + n_2 + \cdots + n_{m!} = n$. This equals $\binom{n+m!-1}{n} = \frac{n(n+1)\ldots(n+m!-1)}{(m!-1)!}$. Such objects are represented as vectors of length $m!$ if we fix an order of the types, and we call these (as usual) *compositions* of n with $m!$ parts.

Before proceeding we note an alternative characterization of Q that will be useful.

Definition 1.23. *For each $k \geq 1$ consider the set V^k of all k-subsets of V equipped with the uniform measure and consider the product space $\Sigma \times V^k$. Let E denote the event*

$$E := \{(S, A) \in \Sigma \times V^K \mid A \textit{ contains a manipulating coalition at } S\}.$$

Lemma 1.24. *Let $k \geq 1$ and let A denote a sample from V^k and S a sample from Σ. Then*

$$\Pr(Q \leq k \mid S) = \Pr(E|S)$$

and hence

$$F_Q(k) = \Pr(E) = \Pr(A \textit{ contains a manipulating coalition}).$$

In other words, the probability that we require at most k queries to find a manipulating coalition equals the probability that a randomly chosen k-subset contains a manipulating coalition.

Proof. Given S, the event $Q \leq k$ means precisely that the initial subset A_Q formed by the first k queries contains a manipulating coalition at S. Each subset of V of size k occurs with equal probability $\binom{n}{k}^{-1}$ as an initial subset of queries of the query sequence, so that A_Q is distributed as A. This gives the first equality and the second set of equalities follows from standard probability computations. $\square$

Remark 1.25. *The distribution function of Q can thus be computed by simply counting the number of subsets of a fixed size that contain a manipulating coalition.*

Note that it is also true that for each fixed A,

$$\Pr(A \textit{ contains a manipulating coalition}) = F_Q(k).$$

This is because of the symmetry between voters. Without the symmetry, we know only that the expectation over A of $P(E|A)$ equals $F_Q(k)$.

General algorithms

We now discuss the computation in more detail. Throughout, we assume the existence of a subroutine C that, given a voting situation and a subset X of V, determines whether there is some subset of X that is a manipulating coalition. For scoring rules, we describe such a C in Section 1.8.1.

A direct computation of $P \mid S$, $M \mid S$ and $Q \mid S$ proceeds by enumerating subsets X and using C to test each one. For P, we need only do this for $X = V$. For M and Q we should enumerate all size 1 subsets, then all size 2, etc. Once we find a manipulating coalition, this finds a minimal manipulating coalition of minimum size, and thus determines M. To determine Q, we must continue to generate all subsets of all possible sizes.

An obvious improvement is to generate compositions subsets of size 1, 2, $\ldots$,

n in turn, adding new minimal manipulating coalitions to a table. Each newly generated subset is checked to see whether it contains any elements of the table. If yes, we can update $\Pr(M \leq i)$ and $\Pr(Q \leq i)$ accordingly, for all $i \geq k$. Otherwise, check the subset to see whether it is itself a minimal manipulating coalition by invoking C (we add it to the table if so, and update probabilities accordingly). Since checking containment is simply a coordinatewise operation on the compositions and is faster than C itself, this gives a clear speedup especially for large k. Also, we only invoke C to check whether a given subset *is* a minimal manipulating coalition, not whether it contains one. This allows for simplification of C in some circumstances.

Of course, non-anonymous rules require the entire profile. A general rule may require generation of $\binom{n}{k}$ subsets for each k, and hence 2^n in the worst case, when the situation is not manipulable. We restrict to anonymous rules from now on. In this case we can generate instead all types of subsets (compositions of k into t parts), the number of which for each k is $\binom{k+t-1}{k}$, where t is the number of possible types of voters to consider in a coalition (in other words we generate the compositions of k into t parts). For M we can take $t = (m-1)!$, since we need only consider subsets consisting of voters not ranking the sincere winner highest, but for Q we take $t = m!$ since all types may be found by our query process.

Finally, as noted above, to compute P, M, Q we know no better method in general than to aggregate the conditional probabilities. There is one idea which seems very promising at the first sight. To compute Q, we can simply fix A and iterate over all voting situations, instead of looping over all voting situations and then over all A. But this overlooks the fact that in the first method, we must consider each profile represented by the voting situation separately. This requires $m!^n$ invocations of C

overall and will be uncompetitive almost always with the more direct method.

Specialized algorithms for scoring rules

For scoring rules, there exist substantial improvements to the above procedure. The key point is that manipulability by a coalition may be described by systems of linear (in)equalities, and some steps above can be combined. We give a brief description below and refer to [20] for more details. We note that Copeland's rule lends itself to completely analogous computations which is presented in appendix.

For each candidate b different from the sincere winner a, and each subset X consisting of voters who prefer b to a, we have a system S_b of linear (in)equalities describing manipulations which result in b winning. The subroutine C_b simply checks whether this system is feasible, and C simply combines the results of these subroutines with a logical "OR".

To describe S_b, we begin with the variables. There is one nonnegative integer variable x_i for each sincere preference order occurring in X, and one nonnegative variable y_j for each strategic vote that can occur. It appears that in the worst case the number of x's is $m!/2$, the number of types that rank b above a. The number of y's could be even larger for a general rule. However it is readily seen that for scoring rules, only strategic votes that rank b first should be considered (other strategies are dominated by strategies of this type), so the number of y's required is at most $(m-1)!$. Furthermore the number of x's can be reduced. For those types who sincerely rank a last and b first, voting sincerely is a dominant strategy and hence these voters can be removed from any coalition, so that we need consider only $(m-2)!(m+1)(m-2)/2$ types. The total number of x's and y's to consider then equals $m!(m^2 + m - 4)/(2m^2 - 2m)$. For some special scoring rules this can be reduced even further. For example, for plurality and antiplurality, we need only consider 1 possible strategy (rank b first (respectively a last) and the others in any fixed order), so the number of y's is only 1. And in this case, those voters sincerely ranking b first (respectively a last) cannot do better than by using the sincere strategy, so the number of x's is $m!(m-2)/(2m)$.

We now describe the constraints in S_b. We first have the constraints that $s(a) \geq s(b) \geq s(c) \geq \dots$. Our random tie-breaking assumption allows this and gives a speedup by a factor close to $m!$, because we do not need to generate all voting situations. The scores after manipulating satisfy $s(b') \geq s(c')$ for all candidates c (there may be some strict inequalities depending on tie-breaking cases, but we keep them all non-strict for simplicity). There is also an equality constraint that the sum of x's equals the sum of y's. The total number of constraints is $2m - 1$. Note that although the scores when expanded in terms of the weights and numbers of voters of each type will involve more variables, the constraints listed only involve the stated variables, because of cancellation.

In order to compute $P|S$ we can simply invoke the subroutine C with $X = V$, as mentioned above. In fact we can go even further for some special distributions. For example, for the IAC distribution, the linear systems described above allow direct computation of the aggregate measure P as follows. We need to count the number of lattice points in the polytope determined by the system S_b. This is accomplished by algorithms to compute Ehrhart polynomials as described in [34, 35]. Inclusion-exclusion then allows the computation of P.

To compute $M|S$ we consider the integer linear programming problem associated to S_b, with objective function equal to $\sum_i x_i$. The minimum over b of the optimal values of these optimization problems is precisely $M|S$.

We now move on to discuss the new measure Q. We know that it contains more information than P and M, so it should be expected to be harder to compute. Although we have no theoretical justification for such an assertion, our efforts to find algorithms have convinced us that it is true.

To compute $Q \mid S$, the now-obvious method is to generate all (types of) subsets of size k, for each k, and check them in turn (using $m-1$ iterations of the improved algorithm C involving the linear system above, one for each losing candidate). We also use the lookup table approach above. We must still generate all possible types of subsets.

Alternative algorithm for Q There is an alternative method that avoids generating all types of subsets, which works particularly well for $m = 3$. The idea is to first enumerate systematically all equivalence classes of minimal manipulating coalitions, and then use inclusion-exclusion to compute the number of subsets of each size k that contain at least one of these minimal coalitions.

Definition 1.26. *Equivalence classes of minimal manipulating coalitions consist of coalitions which have the same size and the same distribution of types.*

Let N denote the number of such equivalence classes. Under plurality and antiplurality, $N \leq m-1$, because the minimal coalitions that can elect a fixed losing candidate b simply consist of all subsets of a certain fixed size M from the set of voters having one of the $(m-2)!(m+1)(m-2)/2$ x-types as discussed above. These can be represented as compositions of M with (at most) $(m-2)!(m+1)(m-2)/2$ parts in the usual way. Although there are many different types, they are all equivalent and we do not need to distinguish between types. For more general scoring rules, mixed coalitions where we must keep track of types are possible, and N can be larger than M, where M as usual is the minimum coalition size. To find them, we can first find the minimal "pure" coalitions consisting of elements of the same type using t calls to C, and then determine the mixed ones systematically by search, which may invoke C of the order of N times.

Consider the uniform distribution on the set of all subsets of V of size k, and let E_i be the event that a size k subset contains a minimal manipulating coalition of type i. We seek to compute $\Pr(Q \leq k) = \Pr(\cup_i E_i)$. By the inclusion-exclusion formula, we have

$$\Pr(Q \leq k) = \Pr(\cup_i E_i) = \sum \Pr(E_i) - \sum \Pr(E_i \cap E_j) + \cdots .$$

The number of terms in the inclusion-exclusion formula is $2^N - 1$. Also, the intersection of p terms requires the computation of the union of p types of coalitions, which takes time of order $pm!$ using the obvious algorithm of taking the coordinatewise maximum of the compositions. This gives a running time of order $N^2 2^N m!$. When N is sufficiently small, the inclusion-exclusion method is superior to the method described above. However as we have seen N can grow rapidly with M and t. Thus it seems that, for a general scoring rule, this method will only be competitive with the other method above when t and N are fairly small (however, for (anti)plurality it appears to be much superior). We do not have a clear description of exactly when each method is best.

The case $m = 3$ In this case much simplification is possible, as we now describe. The system S_b can be reduced to a linear integer program with 3 variables and 4 inequality constraints. The solution of the feasibility and optimality problems for this linear system can be simplified. As shown in [20], the system can be replaced, using Fourier-Motzkin elimination, by a real linear system in the x's only, that gives necessary conditions for manipulability that are sometimes sufficient. For example, when $m = 3$, this latter procedure works exactly for antiplurality and all rules definable by weight vectors $(1, \lambda, 0)$ with $\lambda \leq 1/2$ — the so-called "easy rules"), but only gives bounds for the other values of λ ("hard rules").

For the purposes of computing $P|S, M|S$ and $Q|S$, we may simplify the linear system when dealing with minimal coalitions. The number of types of minimal coalitions is even lower than the general bound given above. This is because when $m = 3$, the minimal coalitions that can manipulate in favour of b are disjoint from those that can manipulate in favour of c. The minimal coalitions consist only of cba and bac voters, or of bca voters. For certain rules there are even fewer: minimal coalitions under plurality consist only of cba voters or only of bca voters, while minimal coalitions under antiplurality consist only of bac voters (it is never possible to manipulate in favour of c in antiplurality, because cab voters can only increase the advantage of b over c).

In addition, the number of strategies to check is very small, since a minimal coalition containing bac and cba voters will all vote as bca, without loss of generality, while a bca coalition requires consideration only of cba. Thus when testing whether a coalition is minimal, it suffices to check whether switching all cba and bac to bca is a valid manipulation, and whether switching all bca to cba is a valid manipulation.

If we use the alternative method with inclusion-exclusion, the inclusion-exclusion formula has N of order $2M$ and $m! = 6$ for a general rule, whereas the direct method requires $2\binom{n+6}{n}$ calls to C. The first method will be better for (anti)plurality and also for other rules provided n is small enough.

Statistics

We can readily compute the conditional expectations $E[M \mid M < \infty]$, etc, from the distribution functions as follows. We have

$$\sum_{k=0}^{n} \Pr(M > k) = \sum_{k=1}^{n} k \Pr(M = k) + (n+1) \Pr(M = \infty).$$

Reorganizing this equation yields

$$\begin{aligned} E[M \mid M < \infty] &= \frac{\sum_{k=1}^{n} k \Pr(M = k)}{\Pr(M < \infty)} \\ &= \frac{\sum_{k=0}^{n}[1 - \Pr(M \leq k)] - (n+1)\Pr(M = \infty)}{\Pr(M < \infty)} \\ &= n + 1 - \frac{\sum_{k=1}^{n} \Pr(M \leq k)}{\Pr(M \leq n)} \\ &= n - \frac{\sum_{k=1}^{n-1} \Pr(M \leq k)}{\Pr(M \leq n)}. \end{aligned}$$

1.4.3 Basic numerical results

To get a feeling for the behaviour of Q and to allow for comparison with other measures such as P and M, we have carried out detailed computations of Q for $m = 3$ and for the same scoring rules and preference distributions used in [20].

Details of the implementation are found in Section 1.8.1. We present a few representative results here, with discussion. In addition, we present some analogous results for a completely different type of rule, namely the Copeland rule.

In the Appendix, we give details of the algorithm implementation for $m = 3$, and more detailed numerical results which we feel would obscure the overall picture if presented in the present section.

We first consider small parameter values (note that these results will be more affected by our tie-breaking assumptions). Table 1.1 gives values of $\Pr(Q \leq k)$ for IC when $1 \leq k \leq n \leq 6$. Table 1.2 presents the analogous data for IAC.

Table 1.1: Values of $\Pr(Q \leq k)$ under IC

n	k	plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	antiplurality	Copeland
2	1	0.0000	0.3333	0.1667	0.1667	0.1667	0.3333	0.1667
2	2	0.0000	0.5000	0.1667	0.1667	0.1667	0.3333	0.1667
3	1	0.0000	0.0000	0.1111	0.1667	0.1389	0.1111	0.0000
3	2	0.0000	0.0000	0.1944	0.2222	0.2222	0.1111	0.0000
3	3	0.0000	0.0000	0.2500	0.2500	0.2500	0.1111	0.0000
4	1	0.1111	0.2083	0.1528	0.1759	0.1852	0.1481	0.1111
4	2	0.2037	0.3519	0.2176	0.2917	0.3009	0.2222	0.1991
4	3	0.2778	0.4583	0.2639	0.3657	0.3704	0.2685	0.2639
4	4	0.3333	0.5417	0.2917	0.4028	0.4028	0.2963	0.2917
5	1	0.0741	0.1173	0.1296	0.1620	0.1119	0.2099	0.0000
5	2	0.1481	0.2099	0.2122	0.2662	0.1767	0.3148	0.0000
5	3	0.2222	0.2901	0.2793	0.3465	0.2191	0.3580	0.0000
5	4	0.2963	0.3611	0.3472	0.4120	0.2531	0.3688	0.0000
5	5	0.3704	0.4167	0.4167	0.4630	0.2816	0.3750	0.0000
6	1	0.0412	0.1260	0.1376	0.1472	0.1229	0.1070	0.0823
6	2	0.0905	0.2168	0.2155	0.2423	0.2252	0.1523	0.1556
6	3	0.1451	0.2946	0.2760	0.3230	0.3169	0.1677	0.2189
6	4	0.2016	0.3629	0.3283	0.3969	0.3956	0.1718	0.2706
6	5	0.2572	0.4218	0.3774	0.4623	0.4594	0.1741	0.3099
6	6	0.3086	0.4707	0.4237	0.5163	0.5071	0.1754	0.3369

We then choose $n = 32$ as a moderate number of voters, even and not divisible by

Table 1.2: Values of $\Pr(Q \leq k)$ under IAC

n	k	plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	antiplurality	Copeland
2	1	0.0000	0.2857	0.1429	0.1429	0.1429	0.4286	0.1429
2	2	0.0000	0.4286	0.1429	0.1429	0.1429	0.4286	0.1429
3	1	0.0000	0.0000	0.1429	0.2143	0.1786	0.2143	0.0000
3	2	0.0000	0.0000	0.2500	0.2857	0.2857	0.2143	0.0000
3	3	0.0000	0.0000	0.3214	0.3214	0.3214	0.2143	0.0000
4	1	0.0714	0.1548	0.1190	0.1667	0.1905	0.2619	0.0714
4	2	0.1310	0.2540	0.1548	0.2619	0.2857	0.3413	0.1310
4	3	0.1786	0.3333	0.1786	0.3214	0.3333	0.3810	0.1786
4	4	0.2143	0.4048	0.1905	0.3571	0.3571	0.4048	0.1905
5	1	0.0429	0.0905	0.1381	0.1524	0.1286	0.2810	0.0000
5	2	0.0857	0.1595	0.2214	0.2476	0.1952	0.3857	0.0000
5	3	0.1286	0.2190	0.2810	0.3190	0.2333	0.4286	0.0000
5	4	0.1714	0.2762	0.3333	0.3714	0.2619	0.4429	0.0000
5	5	0.2143	0.3333	0.3810	0.4048	0.2857	0.4524	0.0000
6	1	0.0260	0.0931	0.1126	0.1580	0.1385	0.1970	0.0433
6	2	0.0589	0.1537	0.1684	0.2411	0.2433	0.2619	0.0844
6	3	0.0961	0.2045	0.2156	0.3032	0.3286	0.2857	0.1234
6	4	0.1351	0.2506	0.2602	0.3563	0.3935	0.2948	0.1576
6	5	0.1732	0.2944	0.3052	0.4026	0.4394	0.3009	0.1840
6	6	0.2078	0.3377	0.3506	0.4416	0.4740	0.3052	0.1948

3 to reduce the number of ties and therefore the effect of our specific tie-breaking assumptions (it turns out that for odd n and $m = 3$, Copeland's rules is never manipulable under the randomized tie-breaking assumption). In Figures 1.1, 1.2 and 1.3 we plot F_Q under IC and IAC. For small values of k it is hard to distinguish the different scoring rules, so we provide more detail in Tables 1.5 and 1.8.2 in the Appendix.

In Table 1.3 we display the expected values of M and Q, conditional on the voting situation being manipulable.

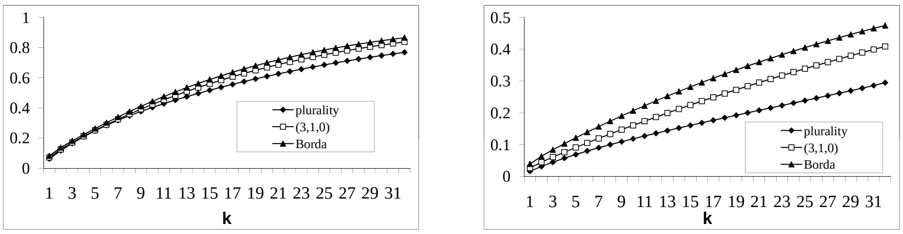

Figure 1.1: Values of $\Pr(Q \leq k)$ when $n = 32$, under IC and IAC.

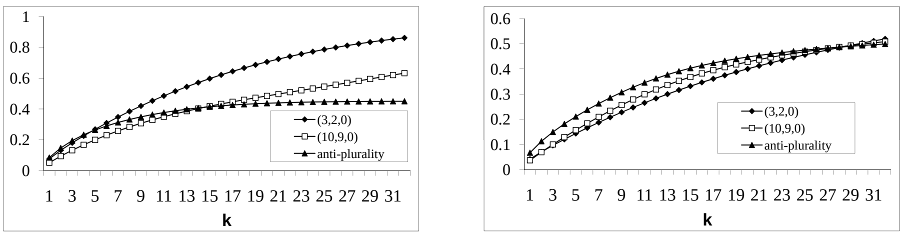

Figure 1.2: Values of $\Pr(Q \leq k)$ when $n = 32$, under IC and IAC.

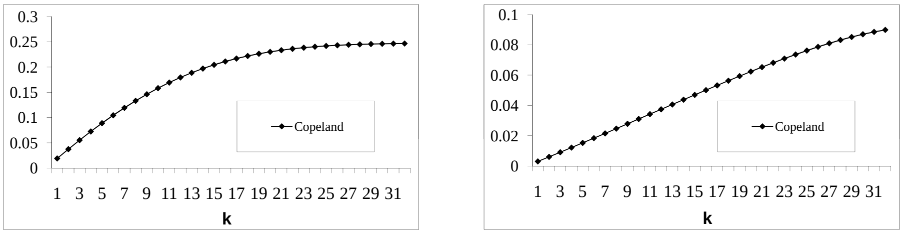

Figure 1.3: Values of $\Pr(Q \leq k)$ when $n = 32$, under IC and IAC.

Table 1.3: Expected values under IC/IAC for $n = 32$.

	IC			IAC		
voting rule	$E(Q\|Q < \infty)$	$E(M\|M < \infty)$	P	$E(Q\|Q < \infty)$	$E(M\|M < \infty)$	P
plurality	11.9323	2.20268	0.768301	14.8069	3.06844	0.294847
(3,1,0)	12.0874	3.49139	0.837584	14.6438	4.85753	0.408826
Borda	11.8601	3.90602	0.865632	13.6529	4.95002	0.474621
(3,2,0)	11.5922	3.54276	0.86109	12.7878	4.3114	0.519419
(10,9,0)	12.3713	3.07908	0.63231	11.2213	3.63562	0.508407
antiplurality	6.18951	1.61894	0.45002	9.10156	3.00334	0.499054
Copeland	9.13436	2.11101	0.246735	15.1917	3.85937	0.089856

Comments on results

The results obtained shed light on the differences between the measures P, M, Q and show that they can rank rules in very different ways. We give a few details below.

The most obvious feature of the results is the different shape of the graphs for M and Q (the former can be found in the Appendix). Indeed, the graphs shown exhibit (slightly) fewer crossings with Q than with M, indicating more robustness to coalition size for Q. For example, when $n = 32$ there is a clear ordering of the rules plurality, (3,1,0), Borda with respect to susceptibility to manipulation under IAC. In particular, there is a single dominant rule with respect to Q (which minimizes the measure for each $k = 1 \dots n$). The difference between Q and M relates to the distribution of each type and the power of each type in changing the result based on the voting rule.

The measure P gives a simple way to linearly order rules for a given n, by their susceptibility to manipulation. However, comparing distribution functions of the form F_M and F_Q is harder. A natural choice is the partial order in which rule r with associated distribution function F_r dominates rule s with associated distribution function F_s if and only if $F_r(k) \leq F_s(k)$ for each k, $1 \leq k \leq n$.

Looking deeper, we see that this dominance ordering among all our scoring rules holds fairly often for small n under IAC for Q, but rarely under IC, whereas the opposite is true for M. A specific example: when $n = 5$, the plurality rule is dominant over all our other scoring rules under IAC when measured by Q, but this is not the case when M is used as the measure, while antiplurality is dominant under IC with respect to M and not Q. Of course, when k is large compared to n, the graphs of M and Q are the same, as they all report the value P for the given rule.

Restricting to conditional values computed at those situations which are manipulable for the given rule, we find that the ordering of rules based on the data in Table 1.3 is different for M and Q. These induced orderings also differ from that induced by P. In fact the 6 combinations of measures $P/M/Q$ and cultures IC/IAC all yield different orderings of the 7 rules!

The results in Table 1.3 suggest that although antiplurality rather often cannot be manipulated at all under IC, it generally requires smaller coalitions in the situations where it is manipulable. This can be observed by checking the values of P and M. However, a similar result is true for plurality, yet the value of Q shows that finding the coalitions under antiplurality is considerably easier. Thus Q adds valuable extra information even in this case. The small values of M and Q for antiplurality presumably occur because for this rule, any voter not ranking the sincere winner first or last has the same power to manipulate, and can only do so by ranking the sincere winner last, which makes the maximum possible change in scores. For other rules, there are more constraints on the coalition members and they have less power to change scores, so larger coalitions are required and they are less numerous.

We describe some conjectures about the behaviour of these measures for large n in Section 1.7.

One might expect that as the weight vector approaches the vector $(1, 1 \ldots, 1, 0)$ that defines the antiplurality rule, the behaviour of the measures P, M and Q smoothly approaches that for antiplurality. However this is not always the case for large k, as can be guessed from our results here, and also from known facts. In fact under IC the asymptotic value of $\Pr(M \leq n)$ (in other words, the value of P) tends to 1 for all rules except plurality but some value less than 1 for antiplurality (when $m = 3$, this latter value equals $1/2$, and the value converges to 1 as $m \to \infty$ [33]). However, this is only a limit result for $k = n$ and $n \to \infty$. For each fixed k and n, convergence does occur as expected. Also, for other distributions such as IAC, this phenomenon does not occur.

The results show that the Copeland rule is considerably less manipulable than scoring rules under all (unconditional) measures, and indeed dominates our chosen scoring rules in most cases presented. The low value of P for this rule is not surprising (of course, under our tie-breaking assumptions, this value is exact 0 for odd n). The rule is defined in terms of the pairwise majority relation and has quite different properties from those of scoring rules. In [36] it was shown that for $m = 3$ under IAC and using lexicographic tie-breaking, Copeland's rule is considerably less manipulable when measured by P than Borda's rule. This resistance to manipulability appears mostly to be a result of the indecisiveness of Copeland's rule: for $m = 3$ the probability of a 3-way tie under either IC or IAC does not approach zero as $n \to \infty$ (unlike scoring rules), yet under our random tie-breaking assumption, it turns out that such situations are never manipulable.

However, conditional information shows that finding a manipulating coalition for Copeland when one exists is sometimes relatively easy, and the rule is not more resistant to manipulation than our scoring rules in the conditional sense.

1.5 Power measures and manipulability measures

In this section we try to study the connection of the theory of manipulability measures with the better-known, but still controversial, theory of power indices. Strategic manipulation can be interpreted in terms of a simple game. We study measure Q as a power index in cooperative game theory which measures the importance of an individual in forming a manipulating coalition. Collective and individual power measures in simple games can be modelled using a sequential model for the discovery of winning coalitions. This link allows for the use of manipulability measures that are specializations of well understood and axiomatically described measures for simple games, and also suggests new general power measures generalizing previously used measures of manipulability.

1.5.1 The simple game associated with a profile

Definition 1.27. *A **simple game** on a finite set X is a subset of 2^X, whose elements are called **winning coalitions**.*

For each profile π and social choice correspondence R we define a simple game $G(R, \pi)$ as follows: declare a subset of X to be winning if and only if it contains a manipulating coalition for R at π. Note that the game may be empty, and this occurs if and only if the rule is not manipulable at π. We call it the *manipulation game* determined by R and π.

Other definitions of "manipulation"

Clearly, any model of coalitional manipulation of a voting rule leads to an associated simple game. We simply define a winning coalition to be one that contains a manipulating coalition (assuming always that the complement of the coalition consists of naive, sincere voters). There are arbitrarily many restrictions one could make on coalition formation (for example, only players adjacent in some fixed network can belong to a coalition). Those that we have observed in the study of manipulation are listed below.

Example 1.28. *(Unit cost bribery) A subset of V is winning if it contains a set of voters who can change the result of the election by changing their votes (not necessarily in accordance with their preferences). Winning coalitions always exist unless the rule is* imposed *(i.e., the voters' preferences are irrelevant). A restricted budget variant exists: let B be a positive integer, and define a coalition to be winning if it is winning as above and the number of members who do not prefer the new winner to the old one is at most B.*

Example 1.29. *(Threshold manipulation) Declare a subset of V to be winning if it can ensure that the sincere winner no longer wins.*

Example 1.30. *(Truth-biased manipulation) Declare a subset of V to be winning if it can manipulate as in our standard definition, and each member has ranked the sincere winner last in its preference order. Voters in this model are very risk-averse and only try to manipulate if there is nothing to lose.*

Example 1.31. *(Bloc voting) Declare a subset of V to be winning if it can manipulate as in our standard definition, via a manipulation where all coalition members of the same type vote the same way.*

Noncooperative games

For each noncooperative game given in normal form, and a distinguished action profile $\mathbf{a}$, we may define a simple game in the following way. A coalition is winning if and only if it contains a subset of players who can jointly deviate from $\mathbf{a}$ (assuming all other players stick with $\mathbf{a}$) in such a way that they each have higher payoff from the resulting action profile. There is a winning coalition if and only if $\mathbf{a}$ is not a strong Nash equilibrium of the game.

Transferable utility games

A simple game is a special case of a TU-game where the characteristic function takes only the values 0 and 1. The more general TU-game assigns a value by its **characteristic function** $v : 2^X \to \mathbb{R}$, such that $v(\emptyset) = 0$. We denote the class of all TU-games on X by $\mathcal{G}(X)$.

1.5.2 The random query process

Consider the following stochastic process. We choose elements of X sequentially without repetition, at each step choosing uniformly from the set of elements not yet chosen, until the set of elements seen so far first becomes a winning coalition. This is the same process considered by Shapley and Shubik [37] in defining their power index. We first consider the random variable equal to the number of queries required.

Definition 1.32. *Let $V_1, \ldots, V_n$ be elements sampled without replacement from X, where $n = |X|$. Equivalently, $\pi := (V_1, \ldots, V_n)$ is a uniformly random permutation of X, representing the order in which elements are to be chosen. Let*

$$Q_\pi = \min\{k : \{V_1, \ldots, V_k\} \text{ contains a winning coalition}\}.$$

Remark 1.33. *If the game is empty we will not find a winning coalition. In this case we define Q_π to have the value $n+1$. If the game is monotone, in Definition 1.32 the word "contains" can be replaced by "is".*

Definition 1.34. *The quantity $\overline{Q}$ is defined to be the expectation of Q with respect to the uniform distribution on permutations of X. In symbols, $\overline{Q} = E_\pi[Q_\pi]$.*

Non-sequential interpretation

The sequential nature of the process is only apparent, once we have averaged over all possible orders. Thus we ought to be able to find a representation of $\overline{Q}$ that does not mention order of players. In order to do this, we assume from now on that the game is monotone.

Definition 1.35. *For each natural number k, define the probability measure m_k to be the uniform measure on the set of all subsets of X of size k. Thus each subset of X of size k is equally likely to be chosen, with probability $\binom{n}{k}^{-1}$.*

For each natural number k, we let W_k (respectively, L_k) denote the set of all winning (respectively, losing) coalitions of size k.

Lemma 1.36. *For each k with $0 \leq k \leq n$,*

$$\Pr(Q \leq k) = \Pr(W_k)$$

where the latter probability is with respect to m_k.

In other words, the probability that we require at most k queries to find a winning coalition equals the probability that a randomly chosen k-subset is a winning coalition.

Remark 1.37. *The cumulative distribution function F_Q of Q can thus be computed by simply counting the number of winning coalitions of each fixed size.*

We can now derive a simple explicit formula for $\overline{Q}$.

Lemma 1.38.

$$\overline{Q} = n + 1 - \sum_{k=0}^{n} \frac{|W_k|}{\binom{n}{k}}.$$

The quantity $\overline{Q}$ intuitively seems to be a measure of inertia or **resistance** (as discussed in [38]): its value is large if winning coalitions are scarce, and small if they are plentiful. The rescaled quantity $1 - \overline{Q}/(n+1)$ looks like an index of what has been called **complaisance** [38, 39] and **decisiveness** [40].

We can consider more general changes of variable F than the affine rescaling above, and show that measure Q is analogues of the Shapley value under a non-standard, but natural, definition of a simple game. For the Shapley-Shubik index, the measures extend naturally to measures for TU-games. In particular, the individual measures include all weighted semivalues. The details of these computations and some other results can be found in [2]. [2] is not included in this thesis as it is a work in progress.

1.6 Comparison with existing literature

Here we discuss work of other authors, viewed through the framework of the present chapter. The papers in question mostly do not use this terminology, and we aim to unify past work.

1.6.1 Results concerning P and M

After the initial news that manipulability is essentially inevitable [7, 8] much work has been done to minimize the likelihood of manipulation without restricting the expressed preferences of voters, with a smaller literature dealing with IAC, and very little with other distributions.

Early research on manipulability focuses on computing $F_M(1)$, the probability that an individual can manipulate. The measure $P = F_M(n)$ has been studied in many papers. [41, 42, 43] have considered coalitional manipulation. The idea of studying M is introduced in [44]. It is investigated in detail for scoring rules in [20] (see also [21] and [45]). The measure $F_M(k)$ is used implicitly in [46], where it is shown that for IAC, and for all faithful scoring rules (those where all weights in the weight vector defining the rule are different) and runoff rules based on them, there is a constant C, depending on m and the rule, so that $F_M(k) \leq Ck/n$. Similar result for IC but with the upper bound $Ck/\sqrt{n}$ is obtained in [47]. Precise asymptotics (when m is fixed, as $n \to \infty$) for $F_M(k)$ under IC are given for scoring rules in [33]. Xia and Conitzer [48] prove that M must be of order at least $\sqrt{n}$ for a wide class of rules under rather general assumptions on preference distributions, with a different definition of manipulation.

In the case of IC, our results and also the results in [33] clarify the conventional wisdom on the relative manipulability of scoring rules, and Borda's in particular. For example, Saari [49] claims that (under IC when $m = 3$) the Borda rule is the scoring rule that is least susceptible to "micro manipulations" (only individual voters or small coalitions) but is quite poor among the scoring rules for macro manipulations. His definition of "micro manipulation" deals with the case $k = o(\sqrt{n})$, where there are few manipulating coalitions for scoring rules as we have seen. In [33] the authors have proved Saari's assertion in more generality. Also, it appears likely from our numerical results that under IC, Borda is the most manipulable scoring rule when k is of order $\sqrt{n}$ or greater, by all our measures. Peleg has also studied the probability that some k voters can manipulate the election and has proved that every scoring rule under the IC conjecture is asymptotically coalitionally non-manipulable by coalitions of size $k \leq o(\sqrt{n})$ [50].

1.6.2 Results concerning Q

The quantity Q has not appeared explicitly before in the literature to our knowledge. Several authors [51, 52, 53] have used probabilistic arguments to give lower bounds on the probability of individual manipulation via a random change in the preference order, again under IC. These results yield (weak) lower bounds on $F_Q(1)$ that decay polynomially in n and exponentially in m; they are more closely related to the bribery analogue of Q.

1.6.3 Complexity measures

Another way to measure hardness of manipulation is by computational complexity, in particular NP-completeness, and this has led to active research in recent years such as [54, 55, 56, 57, 58, 59, 48, 60]. The computational difficulty of manipulation is studied firstly in [61]. Faliszewski et al. have given a survey in computational aspect of strategic manipulation [62].

One feature of this stream of research is that it deals not with manipulation as normally defined, but with a weaker problem, namely that of *winner determination*, or *terminating preference elicitation* (hardness results for this model automatically imply hardness results for manipulation, but not vice versa) . A set S of sincere voters with preferences is given, as is a set T of potential manipulators, who have no preferences. The question is whether the result of the election with electorate S can be changed by the addition of the voters from T (who may vote in any way they choose). Note that if it is possible for T to influence the result in some way, then starting from this new result and then abstaining, T can also change the result. For scoring rules, this means that the members of T can also change the result by changing their votes. Hence T is a coalition "worthy of bribery". Thus if the subset T is not given, then finding it is equivalent to the *unit bribery* problem [22].

When voters are weighted, several manipulation and bribery problems become NP-hard even for fixed m [57]. A more serious issue with complexity results is that several recent papers such as [31, 63] have given evidence that although manipulation problems may be NP-hard in the worst case, these problems are polynomial-time to solve in practice. It is true that these results do not directly yield strong results on Q, being based on a different model of manipulation, but they do suggest several conjectures (as well as proving results for the analogue $\tilde{Q}$ of Q for bribery).

In the standard unweighted case considered in the present article, almost all manipulation problems are solvable in polynomial time unless m is unbounded. Thus unweighted complexity results have little relevance to traditional applications of social choice theory to politics and economics, although they may well be important in newer areas such as search engine aggregation.

Menton and Singh have a survey of voting rules which are NP-hard for unweighted coalitional manipulation for a constant number of manipulators and polynomial in winner determinations [64].

When coalition is small, it cannot change the result but when it is large enough, there is a phase transition based on the fixed number of candidates. The phase transition of manipulation is discussed by several authors independently [65, 66, 47, 46], [31], [19]. Recently, Mossel et al. have studied the phase transition of the coalitional manipulation problem for generalized scoring rules [67].

[68] has considered a survey on the complexity of manipulation and control. The vulnerability of elections to control by adding or deleting votes is studied in [69]. The complexity of winner determination and control problems is discussed in [70]. Computational complexity of control is also studied in [71]. Margin of victory is discussed in [72]. Faliszewski et al have studied control and bribery for Copeland in [9, 73].

1.7 Extensions and future work

There are several obvious directions in which to extend the work of the present chapter. These all relate to asymptotic results, which are most easily obtained under the IC hypothesis, and we restrict to that case here. For scoring rules, the probability of manipulability approaches 1 for all rules other than antiplurality as $n \to \infty$, for fixed m.

A heuristic argument is as follows. The query process for large n is closely modelled by a random walk in $m!$ dimensions, starting at the origin. Each step corresponds to a new query and the transition probabilities are equal for each direction. The walk terminates when it hits the polytope defining manipulability, and this should happen with high probability in order $\sqrt{n}$ steps. This leads to the conjecture: there is a constant C depending on m and the rule such that $Q \leq CM$ asymptotically almost surely as $n \to \infty$. Thus for scoring rules manipulating coalitions of size close to the possible minimum should be fairly common, and how common they are is measured by Q. It may also be true more generally for other voting rules, although clearly there are rules for which it is false (for example, consider a rule that fixes 5 voters and lets them decide the result using plurality, no matter what the other voters do - in this case M is at most 2 but Q should be of order n). A weaker conjecture would be that $Q \leq C\mu$ where μ is the size of the smallest minimal coalition found by the query process, as described in Section 1.4.1.

A related question is: when scoring rules are compared asymptotically on the basis of Q, are their relative merits the same as when compared on the basis of M? We already know that the relative order induced by M and Q can differ for various small n and k.

Our numerical results here and the results in [33] allow us to make some further conjectures.

- For fixed $v > 0$, $\Pr(Q \leq v\sqrt{n})$ tends to a limit $g(v)$ as $n \to \infty$;
- g is a strictly increasing function with $g(0) = 0$ and (for all scoring rules except antiplurality) $g(\infty) = 1$;
- When $m \in \{3, 4\}$, the minimum value of $g'(0)$ is attained by the Borda rule ("Borda is the most resistant to micro-manipulation", and otherwise, this position is held by the $\lceil m/2 \rceil$-approval rule).

We also conjecture that for IAC, the plurality rule dominates all other scoring rules when measured by Q, for all n (at least for $m = 3$).

Our numerical results were only for the case $m = 3$. In [33] it was shown that this case is rather special for the asymptotic behaviour of M, and that "steady-state" behaviour sets in when $m \geq 6$. It would be interesting to investigate whether the same remains true for Q, and it is also interesting to study the computational complexity of measure Q.

1.8 Appendix: details from our studies with $m = 3$

1.8.1 Details of the algorithm implementation

The computer code used to generate the numerical results in this thesis, is available on request.

When $m = 4$ and $n = 100$, the number of voting situations exceeds 10^{24}, and so exhaustive enumeration of voting situations as above is practically impossible for large n. In this thesis we present computational results only for $m = 3$, so as not to have to resort to stochastic simulation. Even when $m = 3$, some care is required. For example when $n = 100$, the number of possible voting situations is nearly 10^8. Also for a fixed voting situation, the computation of Q using enumeration of all types of coalitions for each k can take time of order n^6. Hence small speedups can make the difference between practical and impractical computation. We now discuss some of these.

First, as mentioned above, we need only perform computation for those voting situations for which $s(a) \geq s(b) \geq s(c) \ldots$, because of our tie-breaking convention. This means that each such voting situation is weighted by the size of its orbit under the symmetric group of permutations of the candidates. This size divides $m!$. The probability of a given voting situation under IAC is $\binom{n+5}{5}^{-1}$, while probability under IC of the voting situation $(n_1, \ldots, n_6)$ is $\frac{n!}{6^n n_1! n_2! \cdots n_6!}$. We use this also to weight the voting situations above appropriately.

Scoring rules

Our algorithms described above work particularly well for $m = 3$. We give details for one special case, other cases being very similar (see Table 1.4 and [20] for more details). Suppose that the three candidates a, b, c have sincere scores $s(a) > s(b) \geq s(c)$. The variables x_1, x_2 correspond respectively to the number of voters of type bac and cba, while the variables y_1, y_2 to bac and bca. We also have the equality constraint $x_1 + x_2 = y_1 + y_2$. The sincere scores are expressed in terms of linear combinations of 6 variables that give the parts of the composition that is the voting situation. The restrictions on the sincere scores above yield two inequalities between these scores.

As described in Section 1.4.2 we can omit the linear system entirely, since by the use of the lookup table of minimal manipulating coalitions we only test whether a subset is a minimal manipulating coalition or not. We know that such coalitions must consist only of cba and bac voters, all of whom vote insincerely as bca, or only of bca voters, all of whom switch to cba. Thus we need only make the relevant switch in votes and compute the new election result.

The scores after manipulation are expressed as:

$$
\begin{aligned}
s(a') &= s(a) + (y_1 - x_3)w_2 + (y_2 - x_4)w_3 - x_6 w_3 \\
s(b') &= s(b) + (y_1 - x_3)w_1 + (y_2 - x_4)w_1 - x_6 w_2 \\
s(c') &= s(c) - (y_1 - x_3)w_3 + (y_2 - x_4)w_2 - x_6 w_1
\end{aligned}
$$

and two more inequalities result from these. If we eliminate one of y_1, y_2 using the equality constraint, we obtain an integer linear programming problem with 4 variables and 4 constraints.

Table 1.4: Different types of manipulation: scoring rules, $m = 3$.

Sincere outcome	Manipulated outcome	Possible?	Coalition member types
$s(a) > s(b) \geq s(c)$	b wins	Yes	bac, cba
	a, b tie	Yes	bac, cba
	c wins	Yes	cab, bca
	a, c tie	Yes	cab, bca
	b, c tie	No	
	3-way tie	No	
$s(a) = s(b) \geq s(c)$	a wins	Yes	abc, cab
	b wins	Yes	bac, cba
	c wins	No	
$s(a) = s(b) = s(c)$		No	

Copeland's rule

The details above for scoring rules carry over almost completely to Copeland's rule (we have omitted details, but they are routine to verify). The types of manipulations shown in Table 1.4 are the same. The difference is that a coalition of bac and cba voters has a dominant manipulating strategy, namely for them all to switch to cba. The linear system interpretation also holds, provided we use the Copeland score instead of the positional score.

One simplification we can make is that when n is odd, under our random tie-breaking assumption, Copeland's rule is never manipulable. This is easily verified as follows. In each pairwise contest, there cannot be a tie. So without loss of generality the Copeland scores are $a : 2, b : 1, c : 0$ or $a : 1, b : 1, c : 1$. In the second case no voter has incentive and power to manipulate according to our definition. In the first case any manipulating coalition must increase the score of either b or c (or both) relative to a. The bac voters have power by voting bca, but this only helps c, so is not preferred. The cab voters have power by voting cba, but this only helps b so is not preferred. The cba and bca also cannot succeed as they cannot change the winner.

1.8.2 Additional numerical results

We collect here some basic values of Q to enable comparison of the graphs in Section 1.4.3. We also graph M when $n = 32$.

Table 1.5: $\Pr(Q \leq k)$ for $n = 32$ under IC

n	k	plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	antiplurality	Copeland
32	1	0.0622452	0.0709853	0.0818004	0.0762216	0.0516847	0.084842	0.0189515
32	2	0.116312	0.124421	0.136257	0.132179	0.0943801	0.147295	0.0373838
32	3	0.164568	0.169929	0.181409	0.180549	0.132563	0.195052	0.0552198
32	4	0.208212	0.211586	0.223018	0.225521	0.16763	0.232952	0.0723872
32	5	0.247915	0.250851	0.262945	0.268256	0.20007	0.26403	0.0888203
32	6	0.284158	0.288264	0.301659	0.308993	0.230108	0.290186	0.104461
32	7	0.317369	0.324062	0.339131	0.347779	0.257907	0.312615	0.119258

Table 1.6: $\Pr(Q \leq k)$ for $n = 32$ under IAC

n	k	plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	antiplurality	Copeland
32	1	0.0163955	0.0263586	0.039171	0.0415325	0.0371333	0.0669051	0.0029835
32	2	0.0311806	0.0448672	0.0633464	0.0704502	0.0693989	0.112552	0.00599814
32	3	0.0448065	0.0610597	0.0838489	0.0960968	0.0998481	0.149296	0.00904392
32	4	0.0574015	0.0763073	0.102983	0.120294	0.129006	0.18138	0.0121186
32	5	0.0690667	0.0910389	0.121394	0.143487	0.156969	0.210501	0.0152196
32	6	0.079941	0.105418	0.139293	0.1658	0.183747	0.237398	0.0183441
32	7	0.0901798	0.119519	0.156744	0.187277	0.209326	0.262402	0.0214889

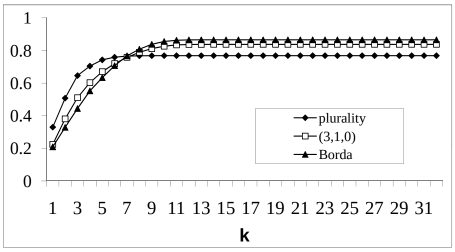

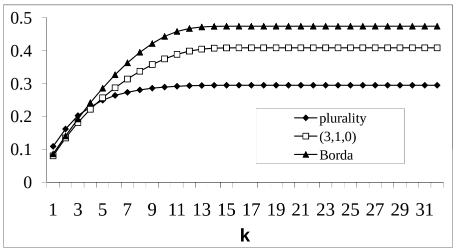

Figure 1.4: Values of $\Pr(M \leq k)$ when $n = 32$, under IC and IAC.

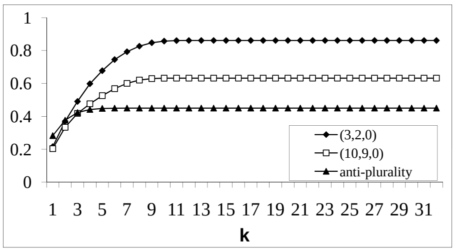

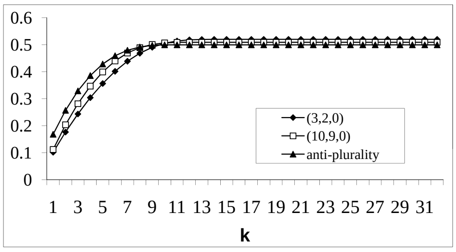

Figure 1.5: Values of $\Pr(M \leq k)$ when $n = 32$, under IC and IAC.

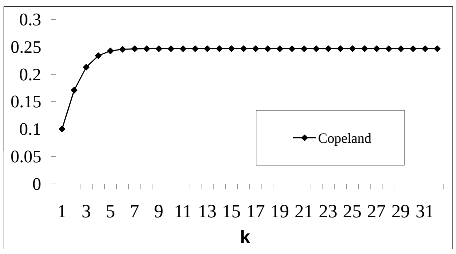

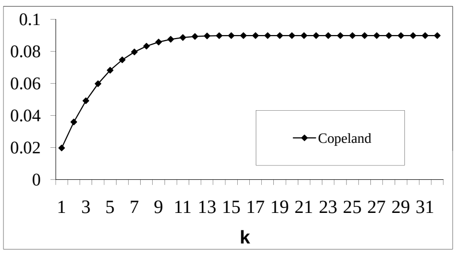

Figure 1.6: Values of $\Pr(M \leq k)$ when $n = 32$, under IC and IAC.

Chapter 2

The Probability of Safe Manipulation

2.1 Introduction

For common social choice functions, the probability that a single individual can succeed in changing the election result under commonly used preference models converges to zero as n, the number of voters, tends to ∞. Thus the question of coalitional manipulation is more interesting.

Coalitions must be of fairly large size in order to manipulate effectively. For example, under the IC hypothesis (uniform distribution on profiles) the manipulating coalitions are typically of order $\sqrt{n}$, while they can be considerably larger under other preference distributions [46, 47]. Thus the question of coalition formation becomes important, because there are substantial coordination difficulties to be overcome in order to manipulate successfully.

In the previous chapter, we studied the query sequence model for forming a manipulating coalition. In this chapter we study another model for coalition formation which was proposed by Slinko and White [74]. In this model a "leader" publicizes a strategic vote and voters sharing the leader's preference order decide whether to follow this strategy or vote sincerely. As a topic for further research, [74] lists the study of the probability that such an attempt succeeds sometimes and the coalition members never fare worse than with the sincere outcome. In this chapter we study this topic for a well-known preference distribution, namely the Impartial Anonymous Culture.

2.2 Definitions and basic properties

Let $m \geq 1$ be an integer and let C be a set of size m, the set of *alternatives* (or *candidates*). Let $n \geq 1$ be an integer and let V be a set of size n, the set of *agents* (or *voters*). Each agent is assumed to have a total order of the alternatives, the agent's *preference order*. An agent a *strongly prefers* alternative i to alternative j if and only if i is strictly above j in a's preference order; if we also allow the possibility $i = j$ then we just use the term *prefers*. There are $M := m!$ possible such preference orders, which we call *types*. We denote the set of all types by T and the set of all agents of type t by V_t. A multiset from T with total weight n is a *voting situation*, whereas a function taking V to T is a *profile*. Each voting situation corresponds naturally to several profiles, corresponding to the different permutations of the multiset.

Let F be a social choice function, a map that associates an element of C to each profile. If this map depends only on the voting situation, then the rule is called *anonymous*.

In the following definitions it is assumed that agents not mentioned continue to vote sincerely.

Definition 2.1. *A voting situation is* manipulable *if there is some subset X of voters such that, if all members of X vote insincerely, the result is strongly preferred by all members of X to the sincere outcome. Such a set X is called a* manipulating coalition.

*A voting situation is **safe** for voters of type t if there is some type t' such that for all x with $0 \leq x \leq n_t$, whenever x agents of type t change their vote to t', these agents weakly prefer the resulting outcome to the sincere outcome.*

*A voting situation is **safely manipulable** by voters of type t if it is safe for them, and there is some value of x for which the agents concerned strongly prefer the resulting outcome to the sincere outcome.*

There are three main points in the definition of safe manipulation:

- the manipulating coalition consists only of voters of a single type;
- the manipulating strategy is the same for all coalition members;
- the size of this coalition is unknown and there must be no risk of obtaining a worse outcome than the sincere one (through "undershooting" or "overshooting").

Overshooting occurs when the following situation holds. If some number x change from t to t', the result is strongly preferred to the sincere one, but if some number $y > x$ change, the sincere result is strongly preferred to the latter outcome. **Undershooting** is the same, but with $y > x$ replaced by $y < x$. Examples in [74] show that both phenomena can occur. In fact they can both occur in the same voting situation as shown by the following example.

Example 2.2. *Let $m = 5$ and consider the voting situation with 3 voters having each of the possible preference orders, except the order 12345 which has 4 voters. The scoring rule (see Section 2.3 for definitions if necessary) with weights $(55, 39, 33, 21, 0)$ yields scores that induce an overall ordering 12345 (meaning candidate 1 wins, candidate 2 is second, etc). Consider voters of type 53124 and the strategy of voting 35241. If 1 voter switches to this strategy, the new winner is candidate 2; if 2 voters switch, then the new winner is candidate 3; if 3 voters switch, the new winner is candidate 4. This shows that undershooting and overshooting can be possible for the same type and choice of insincere strategy in the same voting situation.*

Remark 2.3. *We can consider a game in which the set T of types of voters is partitioned into two subsets, T', T''. The set T'' consists of all types of voters whose only action is to vote sincerely, while voters corresponding to types in T' have all possible votes open to them (we do not allow abstention). In the case where $T'' = \emptyset$ and this is common knowledge, we have a fully strategic game. A situation is manipulable if and only if it is not a strong Nash equilibrium of this game.*

When $T' = T_i$ for some fixed type T_i, there is a different game that is easier to analyse. A situation is safe for members of T' if and only if there exists a pure strategy that weakly dominates the sincere strategy, and safely manipulable if and only if there exists a pure strategy that dominates the sincere strategy.

Remark 2.4. *Note that for each type of voter that ranks the sincere winner lowest, every situation is safe (in fact a stronger statement is true: such voters have nothing to lose by strategic voting, no matter what T' and T'' are). On the other hand, types that rank the sincere winner highest can never manipulate.*

2.3 Algorithms and polytopes

We restrict to scoring rules. However the method described works more generally (for some rules, much more care may be needed when considering ties).

Scoring rules

Definition 2.5. *Let $w = (w_1, \ldots, w_m)$ be such that all $w_i \geq 0$, $w_1 \geq w_2 \ldots w_m$ and $w_1 > w_m$. The* scoring rule defined by w *gives the following score to each candidate c:*

$$|c| = \sum_{t \in T} n_t w_{r(c,t)}$$

where $r(c,t)$ denotes the rank of c according to type t. The candidates with largest score are the winners. The scores give a ***social ordering*** *of candidates (the value of the associated social welfare function).*

Remark 2.6. *If a tie occurs for largest score, then a separate tie-breaking procedure is needed in order to obtain a social choice function. This can be a difficult issue, but fortunately when considering asymptotic results under IAC as in this thesis, we do not need to consider it further. This is because the set of tied situations has measure zero in the limit as $n \to \infty$.*

We now impose an order on the candidates, and write $C = \{c_1, c_2, \ldots, c_m\}$. The types are then identified with permutations of $\{1, \ldots, m\}$ and can be written in the usual way. We describe the scores by the **scoreboard**, the tuple $s = (|c_1|, \ldots, |c_m|)$ of scores. The group of types acts on the scoreboard w via permuting candidates and we denote the action of t on w by w^t. In terms of our current notation, we have

$$s = \sum_{t \in T} n_t w^{t^{-1}}.$$

Example 2.7. *Let $m = 3$ and consider the voting situation in which 6 agents have preference order 312 and 2 agents have order 213. Under the plurality rule given by $w = (1, 0, 0)$, the scoreboard is $(0, 2, 6)$ and c_3 wins, the social ordering being 321. Under the Borda rule given by $(2, 1, 0)$, the scoreboard is $(8, 4, 12)$ and the order of second and third place is reversed, the social ordering being 312. Under the antiplurality rule given by $w = (1, 1, 0)$, the scoreboard is $(8, 2, 6)$ and social ordering is 132. There is no weight vector for which c_2 can win, as c_3 always has a higher score.*

Without loss of generality we assume from now on that the sincere social ordering is $123 \ldots m$.

2.3.1 When t and t' are specified

Fix types t and t' until further notice. We now describe the set S of safely manipulable voting situations. S is the union $\bigcup_{t \in T} S_t$, where S_t is the set of situations that are safely manipulable by voters of type t. This can be further refined to $S = \bigcup_{t \neq t'} S_{t,t'}$ where $S_{t,t'}$ is the set of situations that are safely manipulable by voters of type t using strategy t'.

To describe $S_{t,t'}$, we use the following basic observations.

Let x denote the number of members in a coalition of type t who vote insincerely and suppose they vote t'. Then the new and old scoreboards are related by

$$s' - s = x \left(w^{(t')^{-1}} - w^{t^{-1}} \right).$$

For brevity we refer to those candidates ranked above candidate 1 by agents of type t as **good**, and those ranked below 1 as **bad**. For example, when $m = 3$ and the social ordering is 123, then according to an agent of type 213, c_2 is good and c_3 is bad. The new outcome is preferred by type t agents if and only if no bad candidate is the new winner. It is strongly preferred if and only if some good candidate is the new winner.

Proposition 2.8. *When $m = 3$, undershooting can never occur, and overshooting occurs if and only if some bad candidate wins when $x = n_t$.*

Proof. First note that as a function of x, the differences in scores of each alternative between the sincere and strategic voting situation are (linearly) either increasing or decreasing. Thus if candidate i is above candidate j for some x but below for some larger value of x, it will remain below for all even larger values of x. For types 123 and 132, no better result can be achieved by strategic voting; for types 231 and 321, no worse result. The only other cases are types 213 and 312. In each case there is only one good and one bad candidate: once one overtakes the other and the sincere winner, it stays ahead and cannot be subsequently beaten by another candidate of the opposite type. □

Proposition 2.9. *The following algorithm solves the decision problem for safe* *manipulation for scoring rules, and runs in polynomial time provided the tie-breaking procedure does.*

Let $|c|_x$ denote the score of candidate c when x agents have switched from t to t', and let L be the set of points of intersection of the graphs of the functions $x \mapsto |c|_x$ for $0 \leq x \leq n_t$. Sort the elements of L. For each interval formed by successive elements, compute the maximum score B of all bad candidates, and the maximum score G of all good candidates. If $B > G$ for any interval (or $B = G$ and the tie-breaking procedure says that a valid manipulation in favour of a bad candidate has occurred) then safe manipulation is not possible; otherwise it is possible.

Proof. The winner is constant on each interval, so we need only check one point in each interval, plus endpoints to deal with ties. There are at most $m(m-1)/2$ intersections of the lines which are the graphs of the functions $x \mapsto |c|_x$ for $0 \leq x \leq n_t$. The condition on maximum good and bad scores can be checked for each interval in time proportional to m. □

Corollary 2.10. *When $m = 3$, we need only calculate which candidate wins when $x = n_t$, and safe manipulation is possible if and only if the winner is good.*

2.3.2 The general case

When at least one of t and t' is not specified, there are obviously more possibilities, and a brute force approach that simply tries each pair (t, t') in turn will work. However, we can clearly do better than this.

There are some values of t for which S_t is empty. This means that no matter what the situation and the differences in the sincere scores, safe manipulation is impossible by type t. For example, every t for which the sincere winner 1 is ranked first has no incentive to manipulate. Other types have incentive but as we see in Example 2.11, S_t may still be empty.

For those t for which S_t is nonempty, we can still remove strategies t' for which $S_{t,t'}$ is empty. Similarly, we can express the union defining S_t with as few terms as possible. This is done by discarding dominated strategies (in any particular voting situation, even more strategies may be dominated, but we consider here those that are never worth including for any situation). For example, any type that ranks a bad candidate ahead of a good one is dominated by the type that differs only by transposing those two candidates. Thus all good candidates should be ranked ahead of all bad ones. The sincere winner should not be ranked ahead of any good candidate for the same reason. Furthermore, each strategy that does not allow some good candidate to catch the sincere winner should be rejected, as should each strategy that further advantages a bad candidate higher in the social ordering over all good candidates.

Example 2.11 ($m = 3$)**.** *Consider type* 312. *The only possibly undominated strategy that we need to consider, according to the above discussion, is* 321. *However* 321 *cannot lead to successful manipulation, as it increases the score of* 2 *and not of* 3. *Thus type* 312 *cannot manipulate at all, let alone safely. Types* $231, 213$ *and* 321 *have respectively the strategies* $321, 231, 231$ *available.*

Example 2.12. *When* $m = 4$, *the strategies that are worth considering in some situation are as follows. For any type starting with* 1, *only the sincere strategy. For any type ending with* 1, *any strategy that keeps* 1 *at the bottom. For types starting* 41, *only the sincere strategy; for types starting* 31, *any strategy that lowers* 1 *while keeping* 3 *at the top and not promoting* 2; *for types starting* 21, *any strategy* *that lowers* 1*, keeping* 2 *first. For types ranking* 1 *third, transpose the two good candidates.*

When there are very few distinct entries in w, there are many fewer strategies to consider. The extreme cases are plurality ($w = (1, 0, \ldots, 0)$) and antiplurality $(1, 1, \ldots, 1, 0)$). For plurality (respectively antiplurality), safe manipulation is possible by a type t voter if and only if it is possible by ranking some good candidate first (respectively some bad candidate last). The player is indifferent between the different strategies satisfying this criterion (if the good candidate is fixed) and the analysis does not distinguish between them, so we can assume that any such voter uses a standard strategy that makes a chosen good candidate the favoured one and orders the others by increasing value of index. Thus, for example, for $m = 3$ under plurality we consider 213 and 312 as possible values for t'.

We have so far expressed S_t in terms of a union of $S_{t,t'}$ which is as small as possible. However the terms in the union may not be disjoint. For example, with $m \geq 4$ a voter of type ranking c_1 last may use any of the $(m - 1)! - 1$ insincere strategies that leave c_1 at the bottom (when $m = 3$ there is only one such strategy).

To compute the final probability of safe manipulation, we need to compute the volume of the union of all S_t. This union is in general not disjoint even for $m = 3$, as the following example shows.

Example 2.13. *Let* $m = 3$ *and consider the voting situation with* 3 *agents having preference* 123*,* 2 *having preference* 231 *and* 2 *having preference* 321*. Under the plurality rule, the last two types can each manipulate safely.*

We use inclusion-exclusion to compute the volume of the union. The number of terms in the inclusion-exclusion formula is $2^p - 1$ where p is the number of types involved.

2.4 Numerical results

We restrict to $m = 3$ and some selected scoring rules including the commonly studied plurality, Borda ($w = (2, 1, 0)$), and antiplurality.

For a situation in which the sincere result is 123, types $123, 132$ and 312 cannot manipulate safely. We need to deal with only the remaining types, each of which has only one insincere strategy to consider. The linear systems in question are as follows. We denote $w_i - w_j$ by w_{ij}.

The fact that 123 is the sincere result is expressed as $|c_1| \geq |c_2| \geq |c_3|$. This translates to

$$\begin{aligned}
0 &\leq n_1 w_{12} + n_2 w_{13} + n_3 w_{21} + n_4 w_{31} + n_5 w_{23} + n_6 w_{32} \\
0 &\leq n_1 w_{23} + n_2 w_{32} + n_3 w_{13} + n_4 w_{12} + n_5 w_{31} + n_6 w_{21} \\
n_i &\geq 0 \text{ for all } i \\
n &= n_1 + \cdots + n_6.
\end{aligned}$$

For type 213, safe manipulation is possible if and only the following additional conditions are satisfied.

$$|c_2| \geq |c_1| - n_3 w_{23}$$

$$|c_2| \geq |c_3| + n_3 w_{23}$$

which simplifies to the following system.

$$0 \geq n_1 w_{12} + n_2 w_{13} + n_3 w_{31} + n_4 w_{31} + n_5 w_{23} + n_6 w_{32}$$

$$0 \leq n_1 w_{23} + n_2 w_{32} + n_3 w_{12} + n_4 w_{12} + n_5 w_{31} + n_6 w_{21}$$

Every voting situation can be represented in this way up to a permutation of alternatives.

Thus the asymptotic probability under IAC that type 213 can safely manipulate is given by the ratio of the volume of the "strategic" polytope to that of the "sincere" polytope. A completely analogous method works for other types. The volumes can be computed using standard software as described in [34, 35].

The results for several voting rules are shown in Table 2.1. The column labelled "P(manip)" gives the asymptotics probability of a voting situation begin manipulable (possibly by a coalition of more than one type) and was computed using the methods in [20] (note that the results for plurality, antiplurality and Borda have been computed exactly elsewhere [34]). Note that the ordering of rules according to their susceptibility to manipulation and the corresponding order for safe manipulation differ. Also the entries in the last column, giving conditional probabilities, are decreasing. This last fact is not surprising in hindsight and probably not dependent on the culture IAC. For example, plurality allows only one type of member in a minimal manipulating coalition, and such members have nothing to lose, so manipulation is possible if and only if it is safely possible. At the other extreme, only one type of voter can manipulate under antiplurality, but whether this is safe or not depends strongly on the voting situation.

The Borda rule is often criticized for its susceptibility to manipulation. While it is

Table 2.1: Asymptotic probability under IAC of a situation being (safely) manipulable.

scoring rule	P(manip)	P(safely)	P (safely | manip)
plurality	0.292	0.292	1.00
(3,1,0)	0.422	0.322	0.76
Borda	0.502	0.347	0.69
(3,2,0)	0.535	0.330	0.62
(10,9,0)	0.533	0.264	0.49
antiplurality	0.525	0.222	0.42

still the most manipulable here by both measures, it is clear that many manipulable situations under Borda require unsafe manipulations. The plurality rule seems the least manipulable when complicated coalitions are used, but its advantage disappears when safety is considered. These results, which of course depend on the particular distribution IAC, nevertheless indicate that when communication is restricted, traditional ratings of voting rules may need to be revised.

Table 2.2 shows the probability that a given rule is safely manipulable by all of the individual types listed. We see for example that type 213 has the most manipulating power under the $(3, 2, 0)$ rule, whereas 231 and 321 are strongest under plurality. Note that, for example, there is an appreciable probability that both types 213 and 321 can manipulate safely. If each proceeds, ignoring the other, the result may no longer be safe. On the other hand, if both 231 and 321 try simultaneously to manipulate safely, the cancellation effect means that they are less likely to be disappointed.

Table 2.2: Asymptotic probability under IAC of safe manipulation by various types

scoring rule	213	231	321	213, 231	213, 321	231, 321	213, 231, 321
plurality	0.0000000	0.156250	0.246528	0.000000	0.000000	0.111111	0.000000
(3,1,0)	0.178369	0.086670	0.216913	0.000080	0.104229	0.053084	0.000067
Borda	0.225000	0.047950	0.196759	0.000033	0.093542	0.027400	0.000024
(3,2,0)	0.239297	0.020019	0.152812	0.000007	0.070438	0.010926	0.000005
(10,9,0)	0.234375	0.001687	0.051107	0.000000	0.022681	0.000866	0.000000
antiplurality	0.2222222	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000

2.5 Further discussion

The uniform distribution on profiles (the Impartial Culture hypothesis) has been used in many analyses in voting theory, because of its analytical tractability. However, for the asymptotic study of safe manipulation it seems less useful. This is because under IC for scoring rules, much weight is placed on situations that are nearly tied: a typical situation has almost equal numbers of each type, and the differences between the scores are of order $\sqrt{n}$. Thus as $n \to \infty$ the probability that, for example, a voter of type 321 can safely manipulate will approach 1 rapidly, while the probability that a type 213 can do so will approach 0 rapidly.

The inclusion-exclusion procedure used is probably exponential in m, since the number p of types used seems to grow linearly in m (we have not formally proved this). Thus a better algorithm is needed for large m.

The argument of Section 2.3.2 involves a monotonicity property that should be satisfied by more than just the scoring rules, but we have not pursued such a generalization here, leaving it for possible future work.

The literature on safe manipulation is very small still - our literature search turned up only one preprint of unknown publication status, dealing with complexity issues (though a similar idea was apparently used in [75] without explicit mention). However the basic model is attractive and some obvious generalizations should be investigated. For example, we can use a probability distribution to model the number of followers, instead of considering the worst case outcome, and thereby consider whether strategic voting even with lack of coordination can lead to better outcomes in the sense of expected utility.

Part II

Voting Games

Chapter 3

An Introduction to Voting Games

3.1 Introduction

In the previous chapters we discussed naive manipulation in which just a coalition of manipulators vote strategically and the others vote sincerely. However, a natural question that arises here is that what will happen if all voters behave strategically and all of them know that too. In this case the strategy of each agent depends on the strategy of other agents. Game theory is a useful tool in interactive decision theory. The cooperation and conflict of self-interested and autonomous agents can be modelled mathematically in the game-theoretic study of voting systems. Self-interested agents do not necessarily want to hurt each other, or even that they care only about themselves. Instead, it means that each agent has his own description of which states of the world he likes [76]. This can include good or bad things happening to the other agents and he tries to bring about these states of the world. This analysis can be considered for two different purposes: as a mechanism designer for analyzing the behaviour of people in the system and using those information for designing the system, or as a player for finding the best response against the action of the other players. From one perspective, games are either simultaneous or sequential. We are going to discuss **simultaneous** voting games in which voters act simultaneously. If voters do not play simultaneously, we have a **sequential** game. Another way of categorizing games is to divide them to cooperative and non-cooperative. In **non-cooperative** game theory the basic modelling unit is the individual (including his beliefs, preferences, and possible actions). In **cooperative** voting games, players are able to form binding commitments. We are going to study the non-cooperative models of voting games.

3.2 Voting game

Definition 3.1. *A **normal or strategic form game** consists of:*

- N*, a finite set of **players**,*
- *For each player* $i \in N$*, a finite set of **pure strategies*** S_i*,*
- *For each player* $i \in N$*, a **payoff function*** u_i *that specifies a utility value for each profile of pure strategies* $(s_1, \ldots, s_i, \ldots, s_n)$*. The range of this function is normally the set of real numbers, where the number represents a cardinal utility. However, in our model, the payoff function is the ordinal utilities given by the voters' preferences of that profile.*

In voting games, the set of players are voters, states, special interest groups, or politicians. Each player's action is his vote or decision, and his strategy determines his action. One kind of strategy is to select a single action and play it. Such a strategy is called pure strategy. Players could also follow another (less obvious) type of strategy which is called mixed strategy: randomizing over the set of possible strategies according to some probability distribution. Note that pure strategy is a special case of mixed strategy.

A **dominated** strategy is a strategy for which there is some other strategy that is always better whatever the other players are doing . A strategy is **strictly dominant** when no matter how the other players may play, it is the best strategy. For example, in elections with 3 candidates and approval voting, the dominant strategy is to vote for the most desirable element and voting for the least desirable candidate is dominated [77]. A game is **dominance solvable** if iterated removal of dominated strategies ends in a unique equilibrium which is a reasonable guess for what will happen if we have rational players and complete information [78].

Definition 3.2 (**Nash equilibrium**). *A Nash equilibrium is a list (profile) of strategies of all players, from which no player is willing to deviate unilaterally. In other words, the profile* $(s_1^*, ..., s_N^*)$ *is a Nash equilibrium if*

$\forall i \in N$*and* $\forall s_i$:

$$u_i(s_1^*, ..., s_i^*, ..., s_N^*) \geq u_i(s_1^*, ..., s_i, ..., s_N^*)$$

Theorem 1 (Nash Theorem)**.** *At least one (mixed strategy) Nash equilibrium exists in a non-cooperative game with a finite set of actions.*

Nash equilibrium is a stable situation when there is no voter who has motivation to deviate unilaterally. If a strictly dominant strategy exists for one player , that player will play that strategy in each of the game's Nash equilibria.

A voting rule is **strategy-proof** if for all possible profiles of preferences, "every one votes sincerely" is a Nash equilibrium. In other words, if all voters behave sincerely, no voter benefits by being insincere. However, as we discussed in Chapter 1, the Gibbard-Satterthwaite theorem [7, 8] shows that for each nondictatorial social choice function allowing unrestricted preferences of voters over m alternatives ($m \geq 3$) and such that each alternative can win in some profile, there always exists a profile which is unstable. In other words, in the voting game with ordinal utilities given by the voter preferences of that profile, the strategy where all voters express their sincere preferences may not be a Nash equilibrium. Therefore, there exists at least one strategy profile where one voter has incentive to deviate unilaterally by expressing an insincere preference.

Strategic voting is clearly a question of game theory. However, it has still been little studied from this viewpoint, perhaps since its main questions go beyond the Nash equilibrium concept (which applies only to individual manipulation). So far most of the studies in the area of computational voting game are dealing with the cooperative models of coalitional voting games or the complexity analysis of relevant solution concepts (e.g. Nash equilibrium [79]) such as, exploring the voting power of coalitions in weighted voting games (e.g. weighted threshold games) [80, 81, 82], the compact representation of such games or studying the complexity of the core and Banzhaf and Shapley values [83].

A start has been made in filling the gaps between that case and the classical game theory situation of full common knowledge [75]. The game-theoretic study of strategic manipulation is discussed in [84, 85, 86, 87, 88, 89]. The aims and objectives of these papers are to define a proper model for a voting game and to find the equilibrium outcome. Beside Nash equilibrium, other solution concepts are studied in this context such as correlated equilibrium [90] and regret minimization [91].

Correlated equilibrium as a signalling device tells each player about the strategy that he should choose (it gives a joint probability distribution over the set of outcomes). However, they are not informed about the outcome of the experiment, and they may choose to follow or not according to their utility function. Regret minimization is a relatively new solution concept. In this solution concept, each voter does not know about the other players' actions and he just tries to choose a strategy that ensures that he has done reasonably well compared to the best possible action without paying attention to the other players' actions. In fact the regret of each action represents the utility difference of the best possible outcome and that action. The quality of solution concepts has been measured in some papers by the price of stability and the price of anarchy e.g. in [92, 93]. The price of stability and anarchy are the best and the worst possible ratio between the cost of an outcome at Nash equilibrium and that of an optimal one.

Predicting the result of the game is challenging, as voting games can have many equilibria. Therefore, we are interested in studying how we can omit some of the possible equilibria and reach a unique equilibrium. We concentrate on best reply voting games and study the convergence of dynamic process in Chapter 4. We study the factors that influence this convergence.

One important factor which has significant effect on the strategies of players in voting games is the available amount of information. Complete information models where the preferences are common knowledge among the voters are more common in this area. However, recently there are more papers studying partial information models such as [94, 95, 96, 97]. Poisson model for population uncertainty is discussed in [98].

We study the effect of information by introducing a new model of voting games in Chapter 5. In this model voters achieve partial information via a series of pre-election polls, and also each voter has some uncertainty about the announced result of polls. We study the different distributions of uncertainty for plurality voting games.

Chapter 4

Best Reply Dynamics for Scoring Rules

4.1 Introduction

The strategic misrepresentation of a voter's true preferences, as a way of obtaining an outcome preferable to that which would be expected by voting sincerely, dates back thousands of years. The amount of information available to voters and their ability to communicate influence voter's behaviour greatly. Here we consider the case in which all players behave strategically, but coalitions are not formed. The natural setting then is that of a normal form game with ordinal preferences, or more generally a game form.

The voting games of this type have enormously many Nash equilibria and are not necessarily dominance solvable [87]. Eliminating dominated strategies is not also helpful because typically far too many equilibria remain for the Nash equilibrium to be a credible prediction. Other refinements such as strong and coalition-proof Nash equilibria may not always exist [99]. One natural direction of enquiry is to consider best-reply dynamics, where players take turns in moving myopically in response to previous moves by other players (these moves are pure strategies of the associated game). For many games this process leads to convergence (necessarily at a pure Nash equilibrium). It can also be interpreted in the voting context as a method of reaching consensus, and is in fact used in this way in some applications such as Doodle (for scheduling). According to Fudenberg and Levine [100], in some cases, most learning models do not converge to any equilibrium and just coincide with the notion of rationalizability, but if best-reply dynamics converges, it necessarily finds a NE. Therefore, the question that arises here is in which cases these best-reply dynamics converge for voting games. To our knowledge, in the voting context the first paper to discuss best-reply dynamics is [101], which concentrated on the plurality rule. The authors considered the effect of initial state, tie-breaking rule, the players' strategy and weights on convergence. The results show that this definition of best reply, even for such a rule which restricts voter expression severely, is too general to guarantee convergence. Sequential and simultaneous voting games for plurality with abstention have been discussed in [88]. For the sequential case, they provide a complete analysis of the setting with two candidates, and show that for three or more candidates the equilibria of sequential voting may behave in a counterintuitive manner. The strategy of each voter depends strongly on the information he has about the other players' preference orders.

4.1.1 Our contribution

A natural extension of [101] is to consider general positional scoring rules, which we do. We find that non-convergence occurs much more often in this case, as might be expected because of the much larger strategy spaces involved. For the antiplurality (veto) rule, which restricts strategy spaces as much as plurality, we give a complete analysis and show convergence under rather general conditions. We also give unified simple proofs for plurality and antiplurality and give more details on the boundary between convergence and nonconvergence when tie-breaking methods are considered. We study cycles in the scoring rules between plurality and antiplurality. For a general scoring rule, the order in which players respond in the best reply dynamics influences the convergence considerably. Our results show that some tightening of the definition of best reply is indeed required for convergence for plurality and antiplurality. However, a natural extension of this tighter definition to general scoring rules fails to guarantee convergence.

4.2 Problem description

4.2.1 Voting setup

There is a set C of alternatives (candidates) and a set V of players (voters), with $m := |C|$, $n := |V|$. Each voter has a strict total order on candidates, the preference order of that voter, denoted σ_v. This defines the set $\mathcal{T}$ of types of voters, and $|\mathcal{T}| = m!$. The function mapping $v \mapsto \sigma_v$ is the profile. A voting rule (or social choice correspondence) that maps each profile to a nonempty subset of C (the winner set).

For a voting rule R, we study the game $G(V, C, R)$ where each voter v submits a permutation π_v of the candidates as an action. The set of pure strategies available to voter i, S_i, consists of the $m!$ possible types. In other words, a voter can report a preference order, which may not be his sincere one. We denote the sincere profile and the profile at time t respectively by p_0 and p_t. We order the types lexicographically, based on a fixed order of candidates.

A **voting situation** is a multi-set from T with total weight n. For anonymous rules (those invariant under permutations of the voters), the voting situation gives a more compact description than the full profile, with no loss of information. For example, if we have 3 candidates a, b and c, and 4 voters with preference orders abc, bca, cab and bca, the voting situation coinciding with that profile is $(1, 0, 0, 2, 1, 0)$.

A **voting rule** (or social choice correspondence) is a mapping taking each profile to a nonempty subset of C (the **winners**). A voting rule is **resolute** (or a social choice function) if the set of winners always has size 1.

The **scoring rule** determined by a weight vector w with

$$1 = w_1 \ \geq \ w_2 \geq \cdots \geq \ \ w_{m-1} \ \geq \ w_m = \ 0$$

assigns the score

$$s(c) := \sum_{t \in \mathcal{T}} |\{v \in V | \pi_v = t\}| \, w_{\pi_v{}^{-1}}(c) \tag{4.1}$$

to each candidate. For example, several well-known scoring rules are:

- Plurality: $w = (1, 0, \ldots, 0, 0)$ in which each voter in effect votes for one candidate.

- Antiplurality (veto): $w = (1, 1, \ldots, 1, 0)$ in which each voter in effect votes against one candidate.
- Borda: $w = (m - 1, m - 2, \ldots, 1, 0)$.

The winners are the candidates with the highest score. These rules allow ties in scores and to make them resolute, we choose to use a deterministic tie-breaking rule. However, for neutrality (symmetry between candidates) we need to consider randomized tie-breaking.

4.2.2 Improvement step

Let p be a profile. Suppose that voter v changes his vote. We say this is an **improvement step** if p' (the new profile) is **preferred** to p by voter v. The fundamental results on strategic manipulation initiated by Gibbard [7] and Satterthwaite [8] imply that, provided the voting rule is resolute, under very mild additional conditions (such as not being dictatorial), and provided that $m \geq 3$ and $n \geq 2$, some agent in some sincere voting situation has an improvement step.

In order to describe improvement steps in more detail, we need to discuss outcomes and payoffs (at least ordinal, if not cardinal). The obvious way to do this in the case of resolute voting rules is to declare that the outcome in which the winner is a is preferred by voter v to the outcome in which the winner is b if and only if a is higher than b in v's sincere preference order.

Example 4.1. *(alphabetical tie-breaking) Consider the Borda rule, given by the weight vector* $(2, 1, 0)$*, and the voting situation with 2* abc*, 2* bac*, 2* bca*, 3* cab

voters. The current winner is b. If one of the cab voters changes as acb, then a wins. The new outcome is preferred by that voter because he prefers a to b.

Stochastic dominance

In the case of multiple winners (or randomized tie-breaking), more assumptions are needed. We unify the two cases by using the idea of stochastic dominance as in [20]. This corresponds to a rather risk-averse model of manipulation, as we now describe. It can be described in probabilistic language as follows. For each winner set constructed by the voting rule, we have a uniform distribution on the candidates in that set, and other candidates have probability zero associated with them. Voter v prefers an outcome with winner set W to an outcome with winner set W' if and only if the following condition holds. List the candidates in decreasing order of preference for voter v, and consider the probability distributions as described above. We say that W is preferred to W' if and only if for each $k = 1 \cdots m$ the probability of electing one of the first k candidates given outcome W should be no less than given W'. (If $W' \neq W$ the condition implies that this probability will be strictly greater for some k).

Our definition of improvement step implies that, for example, a vote by a voter with preference bac which changes the winner set from a to $\{b, c\}$ is not an improvement. Of course, if we assigned cardinal utilities to outcomes, there might be some voters for which such a move increases expected utility. In fact, it is easily shown that our definition above says that the probability distribution associated with W first order stochastically dominates the distribution associated with W'. It is well known [102] that this is equivalent to requiring that W is preferred to W'

in terms of expected utility, for all cardinal utilities consistent with the preference order of the voter.

Example 4.2. *(random tie-breaking) Suppose that in profile p the outcome is that a and c tie as the winner, in profile p' the outcome is that b is the absolute winner, and in p'' the outcome is that a and b tie as the winner. The probability distribution of winning on (a, b, c) is $(1/2, 0, 1/2)$ for p, $(0, 1, 0)$ for p' and $(1/2, 1/2, 0)$ for p''. Thus, taking $k = 1$ in the definition, we see that p' is not preferred to p by a voter with sincere opinion abc. Also, taking $k = 2$ shows that p is not preferred to p' either. However, p'' is preferred to both p and p'.*

Other possibilities For example, [101] has considered the case where voters have fixed but arbitrary cardinal utilities. This allows for situations in which more moves are considered to be improvement steps than in our stochastic dominance model above.

4.3 Best reply dynamics

We make the following assumptions in our analysis of best reply dynamics for scoring rules.

- No fixed order for players' turns: in fact, whichever voter has an improvement step can move next.
- Myopic moves: Voters act as though each move is their only chance for improving the result, regardless of considering any chance of changing in the future.

- Costly voting: if there would be no change in the winner set, no move is made.
- Restricted best reply (RBR): we may have several improvement steps which give the same outcome, in which case we choose the one that maximizes the winning score margin of the new winner.
- Stochastic dominance-based improvement step for non-resolute rules.

All the assumptions except the last one are consistent with those in [101]. The fourth applies only for scoring rules, but the others make sense for all voting rules.

Example 4.3. *Consider the antiplurality rule with 2 voters $V = \{1, 2\}$ and 4 candidates $C = \{a, b, c, d\}$, alphabetically tie-breaking. The sincere profile is $p_0 = (acbd, bacd)$. Vetoing candidate c is represented by $-c$ in the strategy profile of voters. The number above the arrow represents the player who moves, and the candidate in braces shows the winner. If voters start from sincere state, we have:*

$(-d, -d)\{a\} \xrightarrow{2} (-d, -a)\{b\} \xrightarrow{1} (-b, -a)\{c\} \xrightarrow{2} (-b, -c)\{a\}$

As you can see in the example, best reply is not unique, for example, the last move by the second player can instead be $-d$. However, $-c$ (vetoing the current winner) is what we call RBR for antiplurality .

4.4 Antiplurality

In this section we show convergence of best reply dynamics under rather general conditions, for a very special scoring rule, namely the antiplurality rule.

For the game $G(V, C, A)$, since $w = (1, 1, \ldots, 1, 0)$, we can without loss of generality assume that $S_i = \{-c|c \in C\}$ (because subtracting the vector $(1, 1, \ldots, 1)$ from the weight vector makes no difference to the outcome of the game or to the differences in scores). In fact, there are $(m-1)!$ possible orders that give the same score. Thus, each improvement step can be written $-a \rightarrow -b$ where $b \neq a$.

Remark 4.4. *We define o_t as the winner set after the move of player i at time t. For alphabetical tie-breaking this set is a singleton.*

Analogous to the case for plurality [101], there are 3 types of improvement steps.

Type 1: $a \notin o_t$ and $b \in o_{t-1}$

Type 2: $a \in o_t$ and $b \notin o_{t-1}$

Type 3: $a \in o_t$ and $b \in o_{t-1}$

Remark 4.5. *It can easily be shown that if $a \notin o_t$ and $b \notin o_{t-1}$, this move does not change the winner set. Therefore, it is not an improvement step.*

Example 4.6. *Suppose we have 2 voters and 3 candidates using antiplurality rule with alphabetical tie-breaking. The sincere profile is $p_0 = (abc, bac)$. If voters start from the sincere state, the current winner is a. If the second player changes his vote from $-c$ to $-a$, the winner switches to b. According to our definition, it is a type 1 move.*

Some notations We define some notations that we use through the rest of the paper.

- We write $c \rhd c'$ if c has a lower index (higher priority) than c' in tie-breaking.
- We write $s(c') \preccurlyeq s(c)$ if either $s(c') < s(c)$ or $s(c) = s(c')$ and $c \rhd c'$(note that it is not a logical notation, and we just use it for simplicity).
- We use the symbol $a \succ_i b$ when voter i prefers candidate a to b.
- We denote the score of candidate a after the improvement step at time t by $s_t(a)$.
- We use the notation $x \xrightarrow{i} y$ when voter i changes his vote from x to y.

Theorem 4.7. *Suppose that $-a \rightarrow -c$ is a type 2 improvement step at time t, and let $b \in o_{t-1}$. Then $-a \rightarrow -b$ is a type 3 improvement step leading to the same set o_t. Furthermore, in this case the margin of victory of the new winner will be more than in the original case.*

Proof.

After the improvement step $-a \rightarrow -c$ at time t, we have

$$s_t(a) = s_{t-1}(a) + 1$$
$$s_t(c) = s_{t-1}(c) - 1.$$

Since $a \in o_t$ (according to the definition of type 2) and $b \in o_{t-1}$ and $s_{t-1}(b) = s_t(b)$, in alphabetical tie-breaking, we have

$$s_t(a) \succcurlyeq \; s_t(b) \succcurlyeq \; s_t(c) \quad \text{and} \quad s_t(a) \succcurlyeq s_t(y) \; y \in C \setminus \{a, b\} \tag{4.2}$$

If we had the improvement step $-a \to -b$ at time t instead, (we denote the score in this case with s'_t)

$$s'_t(a) = s_t(a) \quad \text{and} \quad s'_t(b) = s_t(b) - 1;$$

$$s'_t(c) = s_t(c) + 1 \quad \text{and} \quad s'_t(y) = s_t(y).$$

By substituting in Equation (4.2), we have $s'_t(a) \succcurlyeq s'_t(y)$ for each $y \in C$. Therefore, a is the new winner. For randomized tie-breaking, we can substitute $\succcurlyeq$ by $\geq$. Also, the margin of victory with a type 3 improvement step would be $s'_t(a) - s'_t(b) = s_t(a) - s_t(b) + 1$ which is more than the original margin $s_t(a) - s_t(b)$. $\square$

We now make a key definition of the allowed moves. Allowing type 2 moves can lead to a cycle. An example for plurality has been presented in [101] (Proposition 4). We present a similar example for antiplurality with 7 candidates and 10 voters below. Suppose the sincere preference is

$P0 = (3251764, 4653721, 1245673, 4275631, 2541637, 6351472, 3765214, 7345261, 4561723, 6725134)$

and voters start by voting sincerely. We present the first several iterations (symbol $\Diamond$ shows the stage from which the cycle becomes apparent):

$(-4,-1,-3,-1,-7,-2,-4,-1,-3,-4)\{5\} \xrightarrow{2} (-4,-5,-3,-1,-7,-2,-4,-1,-3,-4)\{6\} \xrightarrow{3}$
$(-4,-5,-6,-1,-7,-2,-4,-1,-3,-4)\{2\} \xrightarrow{7} (-4,-5,-6,-1,-7,-2,-2,-1,-3,-4)\{3\} \xrightarrow{2}$
$(-4,-1,-6,-1,-7,-2,-2,-1,-3,-4)\{5\} \xrightarrow{8} (-4,-1,-6,-1,-7,-2,-2,-5,-3,-4)\{3\} \xrightarrow{3}$
$(-4,-1,-3,-1,-7,-2,-2,-5,-3,-4)\{6\} \xrightarrow{8} (-4,-1,-3,-1,-7,-2,-2,-1,-3,-4)\{5\} \xrightarrow{10}$
$(-4,-1,-3,-1,-7,-2,-2,-1,-3,-5)\{6\} \xrightarrow{2} (-4,-6,-3,-1,-7,-2,-2,-1,-3,-5)\{4\} \xrightarrow{7}$
$(-4,-6,-3,-1,-7,-2,-4,-1,-3,-5)\{2\} \xrightarrow{2} (-4,-1,-3,-1,-7,-2,-4,-1,-3,-5)\{6\} \xrightarrow{3}$
$(-4,-1,-6,-1,-7,-2,-4,-1,-3,-5)\{2\} \xrightarrow{7} (-4,-1,-6,-1,-7,-2,-2,-1,-3,-5)\{3\} \ldots$
$(-4,-1,-3,-1,-7,-2,-4,-1,-3,-4)\{5\} \quad \Diamond$.

Definition 4.8. *(**RBR**) A **restricted best reply** is any improvement step of type 1 or type 3, in which the player making the step vetoes his least preferred member of o_{t-1}, denoted β_{t-1}.*

From now on, we consider only improvement steps using restricted best replies. It is also clear from the definition that no two consecutive improvement steps can be made by the same voter.

Example 4.9. *When voters start from the sincere initial state, and the sincere scoreboard is a tie among all candidates, all improvement steps would be type 3 ones. Therefore, no improvement step can occur, as voters have already voted against their least desirable candidate, and any change will allow that candidate to win.*

Definition 4.10. *(**set of potential winners**) The set of potential winners at time t, W_t consists of those candidates who have a chance of winning at the next step (time $t+1$), depending on the different RBR of voters.*

Remark 4.11. *If candidate c can win by type 1, it can also win by type 3 because when a candidate can win without increasing its score, it is obviously still a winner when its score is increased by 1. Therefore,*

$$W_t = \{c \mid \textit{if some player moves } -c \to -b \textit{ at time } t+1\textit{, then } \; c \in o_{t+1}\} \quad (4.3)$$

4.4.1 Alphabetical tie-breaking

Lemma 4.12. *If $t < t'$ then $W_t \subseteq W_t'$.*

Proof. Consider an improvement step $-a \rightarrow -b$ at time t. According to Definition 4.8, $o_{t-1} = b$. Let $c \in W_{t-1}$ and $y \in C \setminus \{a, b\}$. Then, by considering that the scores of c and y, $\forall y \in C;\ y \neq a, b$ don't change at time t, we have:

$$s_t(c) + 1 = s_{t-1}(c) + 1 \succcurlyeq s_{t-1}(b) - 1 = s_t(b) \tag{4.4}$$

$$s_t(c) + 1 = s_{t-1}(c) + 1 \succcurlyeq s_{t-1}(y) = s_t(y) \tag{4.5}$$

If the improvement step is of type 3, then best reply $-c \rightarrow -b$ at time t gives the same scores as the best reply $-a \rightarrow -b$ followed by $-c \rightarrow -a$ at time $t + 1$. Therefore, $c \in W_t$.

If the improvement step is of type 1, let $b' = o_t$. Note that $b' \notin \{a, b\}$.

According to equation (4.5), for $y = b'$,

$$s_t(c) + 1 \succcurlyeq s_t(b') > s_t(b') - 1 \tag{4.6}$$

According to the definition of winner,

$$s_t(b') \succcurlyeq s_t(y); \forall y \in C \tag{4.7}$$

In particular for $y = a$,

$$s_t(c) + 1 \succcurlyeq s_t(b') \succcurlyeq s_t(a) \tag{4.8}$$

Thus, by transitivity of $\succcurlyeq$ (which follows from the underlying transitive lexicographic order on C), $c \in W_t$. $\square$

A counter-example for an arbitrary deterministic tie-breaking rule Consider a situation with candidates a, b, c and x under the antiplurality rule. Suppose the set of candidates with the highest score after round $t-1$ is $\{b, x\}$ and $s_{t-1}(a) = s_{t-1}(c) = s_{t-1}(b) - 1$. Suppose further that the order of candidates in tie-breaking is as follows: $b \rhd x$ and $c \rhd x$ and $x \rhd a$ and $a \rhd c$. Based on Definition 4.8, $c \in W_{t-1}$. Consider a best reply $-a \rightarrow\ -b$ at time t. If it is a type 3 move then $o_t = a$ and c is still in W_t, as $-c \rightarrow -a$ makes c winner. Suppose the move is of type 1 and $o_t = x$. According to the tie-breaking rule, $b \rhd x$ and $c \rhd x \rhd a$ but, $a \rhd c$. Thus, c is not in W_t because $-c \rightarrow -x$ does not make c win.

Lemma 4.13. *There is at most one type 1 move and each voter has at most $m-1$ moves of type 3.*

Proof. Suppose a step $-a \longrightarrow -b$ is a type 1 move at time t. We claim this improvement step is the first improvement step. If it is not the first improvement step, according to Definition 4.8, a has been a winner before. Therefore, a has been in the winner set in the past. In other words, $\exists t' : t' < t\ \ a = o_{t'}$ and therefore, $a \in W_{t'}$. According to Lemma 4.12, $a \in W_{t-1}$ which means after improvement step $-a \rightarrow -b$ at time t, a is a winner. However, this has contradiction with the assumption of improvement step of type 1. Therefore, there is at most one type 1 move. According to the definition of improvement step, at every step $-a \xrightarrow{i} -b$ of type 3, it must hold that $a \succ_i b$. Therefore, each voter has at most $m-1$ steps of type 3.

$\square$

Theorem 4.14. *Restricted Best Reply Dynamics (RBRD) for $G(V, C, A)$ with alphabetical tie-breaking will converge to a NE from any state in at most $1 + (m - 1)n$ steps.*

Proof. If we have n voters, Lemma 4.13 implies that each voter makes at most $m - 1$ moves of type 3 and there is at most one type 1 move. □

4.4.2 Randomized tie-breaking

Lemma 4.15. *If $t < t'$ then $W_t \subseteq W_t'$.*

Proof. The proof is very similar to the alphabetical case (Lemma 4.12). Except, we do not need to deal with tie-breaking. Therefore, we can substitute the notation $\succcurlyeq$ by $\geq$. For the second part of the proof where we consider a type 1 improvement step, we can always find such a b'. To see this, note that according to the definition of improvement step, the winner set should be changed and the score of b decreases. Therefore, b cannot be the unique winner at time t as it results in b being the unique winner at time $t - 1$, contradicting the definition of improvement step.

□

Lemma 4.16. *There is at most one type 1 move and each voter has at most $m - 1$ moves of type 3.*

Proof. The first part can be proved in a similar way to Lemma 4.13. For the second part, similarly, we show that $a \succ_i b$ if voter i makes the type 3 improvement step $-a \rightarrow -b$. According to the definition of type 3 improvement step, $b \in o_{t-1}$ and $a \in o_t$. We define $p(a)$ as the probability of winning of a. Two cases can occur.

Case 1: $a \in o_{t-1}$

$p(a)$ increases to 1 and $p(b)$ decreases to 0. The probability of winning of candidates in the set o_{t-1} decreases and for other candidates stay 0.

In this case, a becomes the unique winner at time t. Therefore, according to the definition of stochastic dominance improvement step, a should be preferred to all other elements of o_{t-1}.

Case 2: $a \notin o_{t-1}$

$i)$ $b = o_{t-1}$ In this case, $p(a)$ and $p(c)$ increases to $\frac{1}{k+2}$ and $p(b)$ decreases from 1 to $\frac{1}{k+2}$ (assuming the number of candidates (c) whose score is 1 point behind b is k) and for other candidates it remains the same.

$ii)$ $b \in o_{t-1}$ therefore, $p(a)$ increases and $p(b)$ decreases and $p(c)$ stays the same. Therefore $a \succ_i b$, otherwise, it is not an improvement step. □

The analogue of Theorem 4.14 now follows.

Theorem 4.17. *RBRD for $G(V, C, A)$ with randomized tie-breaking, will converge to a NE from any state in at most $(m-1)n+1$ steps.*

Remark 4.18. *The only part in the proof for randomized tie-breaking, where we used stochastic dominance assumption of improvement step is for the bound on type 3 moves. An example of cycle is already shown in [101] for a fixed utility case.*

4.4.3 Who can win?

In this part, we describe W_t in more detail.

$$W_t = W_t^0 \cup W_t^1 \cup W_t^2 \tag{4.9}$$

where W^0 is the level of winner set which includes the candidates who are tied with the winner, W^1 contains the candidates who can win by a type 1 move and W^2 those who can win by a type 3 move and not a type 1 move. Let $M_t = s_t(o_t)$ and $d_t(c) = M_t - s_t(c)$. In fact $d_t(c)$ represents the score difference of candidate c and the winner after move t. Therefore, $W^0 = \{c \mid d(c) = 0\}$. The description of the other two subsets is straightforward.

Proposition 4.19. *For alphabetical tie-breaking,*

$$W_t^1 = \{c \mid d(c) = 1, c \rhd c'; \forall c' \in W_t^0\} \tag{4.10}$$

$$W_t^2 = \{c \mid d(c) = 2 \ \textit{ and unique winner and } \ c \rhd c'; \forall c' \in W_t^1 \cup W_t^0\}. \tag{4.11}$$

For the case of randomized tie-breaking,

$$W_t = \{c \mid d_t(c) \leq 1 \textit{ or } d_t(c) = 2 \ \textit{ and there is a unique winner } \}. \tag{4.12}$$

To obtain a better idea about who is really winning in practice at equilibrium, we ran several simulation experiments with different initial profiles (sincere, random). The numerical results suggest that in the cases with sincere initial state, the winner set of equilibrium is contained in W_0. However, this is not true when we start from an arbitrary state.

4.5 Plurality

The results in this section are completely analogous to those in subection 4.4, and are quite similar to [101] but with easier proofs. We remove some details of proofs as they are similar to previous section.

Definition 4.20. *(**RBR**) For plurality rule, a **restricted best reply** is any improvement step of type 1 or type 3, in which*

Type 1: $a \notin o_{t-1}$ *and* $b \in o_t$

Type 3: $a \in o_{t-1}$ *and* $b \in o_t$

The restricted best replies defined above are similar to the best replies in [101], where the phrase "better reply" is used for non-restricted best replies.

Remark 4.21. *(**set of potential winners**) For plurality also, we just consider the candidates who can win by type 3 moves because of the same argument as antiplurality. Therefore, the set of potential winners is*

$$W_t = \{c \mid \textit{if some player moves } a \to c \textit{ and } a \in o_t \textit{ then } c \in o_{t+1}\} \tag{4.13}$$

4.5.1 Alphabetical tie-breaking

Lemma 4.22. *If* $t < t'$ *then* $W_t' \subseteq W_t$.

Proof. Consider an improvement step $a \to b$ at time t. By the definition of best reply in Definition 4.20, $b = o_t$. Let $c \in W_t$. Considering the new scores of b, c

and y, $\forall y \in C;\; y \neq \; a, b$ we have:

$$s_{t-1}(c) + 1 = s_t(c) + 1 \succcurlyeq s_t(b) - 1 = s_{t-1}(b) \tag{4.14}$$

$$s_{t-1}(c) + 1 = s_t(c) + 1 \succcurlyeq s_t(y) = s_{t-1}(y) \tag{4.15}$$

If the improvement step $a \rightarrow b$ is of type 3, then best reply $a \rightarrow b$ followed by $b \rightarrow c$ at time $t+1$ give the same scores as best reply $a \rightarrow c$ at time t. Therefore, $c \in W_{t-1}$.

If the improvement step is of type 1, let $a' = o_{t-1};$ Note that $a' \notin \{a, b\}$.

According to Equation (4.15), for $y = a'$,

$$s_{t-1}(c) + 1 \succcurlyeq \; s_{t-1}(a') \tag{4.16}$$

According to the definition of winner,

$$s_{t-1}(a') \succcurlyeq s_{t-1}(y);\; \forall y \in C \tag{4.17}$$

In particular for $y = a$,

$$s_{t-1}(c) + 1 \succcurlyeq s_{t-1}(a') \succcurlyeq s_{t-1}(a) \tag{4.18}$$

Thus, by transitivity of $\succcurlyeq$ (which follows from the underlying transitive lexicographic order on C), $c \in W_{t-1}$. $\square$

Lemma 4.23. *The number of type 1 moves is at most m and each voter has at most $m-1$ moves of type 3.*

Proof. Suppose a step $a \rightarrow b$ is a type 1 move at time t. We claim $a \notin W_t$. If $a \in W_t$ then $b \rightarrow a$ makes a winner but we know $b \rightarrow a$ makes a' win (the two consecutive moves have cancelled out each other). Therefore, $a \notin W_t$. According to Lemma 4.22, $a \notin W_{t'}$; $\forall t' > t$. Therefore, the number of type 1 moves is limited and equals the maximal set of potential winners which at most can have m elements. Also, as at every step $a \xrightarrow{i} b$ of type 3, it must hold that $b \succ_i a$ because of the definition of improvement step, each voter has at most $m-1$ moves of type 3. □

Theorem 4.24. *RBRD for $G(V, C, P)$ with alphabetical tie-breaking will converge to a NE from any state in at most $m + (m-1)n$ steps.*

Proof. If we have n voters, Lemma 4.23 implies that convergence must occur with at most $m + (m-1)n$ steps. □

4.5.2 Randomized tie-breaking

Lemma 4.25. *If $t < t'$ then $W'_t \subseteq W_t$.*

Proof. The proof is very similar to the alphabetical case (Lemma 4.22). Except, we do not need to deal with tie-breaking. Therefore, we can substitute the notation $\succcurlyeq$ by $\geq$. For the second part of the proof where we consider a type 1 improvement step, we can always find such a a' by similar reasoning as in proof of Lemma 4.15.

□

Lemma 4.26. *The number of type 1 moves is at most m and each voter has at most $m-1$ moves of type 3.*

Proof. The proof is very similar to Lemma 4.16 by considering the differences of Lemma 4.23 and 4.13.

□

Theorem 4.27. *RBRD for $G(V, C, P)$ with randomized tie-breaking will converge to a NE from any state in at most $m + (m-1)n$ steps.*

Proof. If we have n voters, Lemma 4.26 implies that convergence must occur with at most $m + (m-1)n$ steps.

□

Remark 4.28. *The only part in the proof for randomized tie-breaking where we used the assumption of stochastic dominance is for the bound on type 3 moves. Note that an example is given in [101] showing that if we use fixed utility function, and improvement is defined by expected utility increase, a cycle can occur. The stronger definition of improvement step using stochastic dominance allows us a general convergence result.*

4.6 Counterexamples and interesting phenomena

Best reply dynamics for scoring rules other than plurality and antiplurality does not necessarily converge. Each of the examples in this section starts from the sincere initial state.

Example 4.29. *(Cycle for Borda) Consider the sincere profile $p_0 = (abc, bca)$ and voting rule Borda and alphabetical tie-breaking.*

$(abc, bca)\{b\} \xrightarrow{1} (acb, bca)\{a\} \xrightarrow{2} (acb, cba)\{c\} \xrightarrow{1} (abc, cba)\{a\} \xrightarrow{2} (abc, bca)\{b\}$ ◊.

Remark 4.30. *The allowed moves in the previous example are reasonable for restricted best replies with 3 candidates. Putting the desirable candidate (the new winner) at the top and the current winner at the bottom maximizes the winning score margin of the new winner.*

Cycle for scoring rules "close to Plurality":

- Suppose we have 3 candidates a, b and c and $p_0 = (abc, bca)$. The scoring rule is $w = (1, \alpha, 0)$; $0{<}\alpha \leq \frac{1}{2}$ and we use alphabetical tie-breaking.

  $(abc, bca)\{b\} \xrightarrow{1} (acb, bca)\{a\} \xrightarrow{2} (acb, cba)\{c\} \xrightarrow{1} (abc, cba)\{a\} \xrightarrow{2}$ $(abc, bca)\{b\}$ $\diamondsuit$

- general m and $n = 2$

  $(ab\cdots c, bc\cdots a)\{b\} \xrightarrow{1} (a\cdots cb, bc\cdots a)\{a\} \xrightarrow{2} (a\cdots cb, cb\cdots a)\{c\} \xrightarrow{1}$ $(ab\cdots c, cb\cdots a)\{a\} \xrightarrow{2} (ab\cdots c, bc\cdots a)\{b\}$ $\diamondsuit$

Cycle for scoring rules "close to antiplurality":$m = 3, n = 4$ Suppose we have 3 candidates a, b and c. The sincere profile is $p_0 = (abc, bac, cab, bca)$. Our scoring rule is $(1, \alpha, 0)$; $\frac{1}{2} \leq \alpha < 1$ with alphabetical tie-breaking.

$(abc, bac, cab, bca)\{b\} \xrightarrow{1} (acb, bac, cab, bca)\{a\} \xrightarrow{4} (acb, bac, cab, cba)\{c\} \xrightarrow{1}$ $(abc, bac, cab, cba)\{a\} \xrightarrow{4} (abc, bac, cab, bca)\{b\}$ $\diamondsuit$

Example 4.31. *(Order of players matters) To understand the impact of the order of players on the dynamics, we consider Borda rule with 4 voters and 3 candidates. Suppose $p_0 = (acb, acb, cab, cba)$ and players start from the sincere state.*

The winner is c. The first player is not satisfied with the result and changes his vote to abc to make a the sole winner. For simplicity, we show the moves of players as below:

$(acb, acb, cab, cba)\{c\} \xrightarrow{1} (abc, acb, cab, cba)\{a\} \xrightarrow{3} (abc, acb, cba, cba)\{c\} \xrightarrow{2} (abc, abc, cba, cba)\{a\} \xrightarrow{4} (abc, abc, cba, bca)\{b\} \xrightarrow{1} (acb, abc, cba, bca)\{1\} \xrightarrow{4} (acb, abc, cba, cba)\{c\} \xrightarrow{1} (abc, abc, cba, bca)\{b\}$ $\diamondsuit$

Note $p_4 = p_7$ and we have a cycle.

Now let's consider another order for the players. We start with another profile coinciding with $V = (0, 2, 0, 0, 1, 1)$.

$(acb, acb, cba, cab)\{c\} \xrightarrow{1} (abc, acb, cba, cab)\{a\} \xrightarrow{4} (abc, acb, cba, cba)\{c\} \xrightarrow{2} (abc, abc, cba, cba)\{a\} \xrightarrow{3} (abc, abc, bca, cba)\{b\} \xrightarrow{4} (abc, abc, bca, cab)\{a\}$ *(equilibrium)*

Thus, in contrast with previous order, we reach an equilibrium with this order of players. 8 of 12 profiles coinciding with this voting situation do not converge.

Example 4.32 (an example of cycle for 2-approval voting)**.** *Consider 4 candidates $C = \{a, b, c, d\}$ and 2 voters with $p_0 = \{acdb, dbca\}$ under 2-approval voting rule with weight vector $w = (1, 1, 0, 0)$. Players start from the sincere state and we use alphabetical tie-breaking. Therefore, the sincere winner is c. As voters need to approve two candidates we show the dynamic process as below:*

$(ac, db)\{a\} \xrightarrow{2} (ac, dc)\{c\} \xrightarrow{1} (ab, dc)\{a\} \xrightarrow{2} (ab, db)\{b\} \xrightarrow{1} (ac, db)\{a\}$ $\diamondsuit$

4.7 Conclusion and future directions

A summary of results:

- The upper bound of convergence for plurality in our paper is $m+(m-1)n$. However, it is m^2n^2 in paper [101]. Our upper bound for antiplurality is mn.
- The possibility of winning of a candidate depends on the type of improvement step and also the candidate's priority in tie-breaking.
- The number of type 2 moves is not bounded, so we need to use RBR for convergence.
- We need to use stochastic dominance RBR for randomized tie-breaking for plurality and antiplurality. Without this assumption we can have cycles, as shown in [101] and [103].
- Convergence fails for some deterministic tie-breaking rules.
- The order of players influences convergence, the equilibrium result and also the speed of convergence.
- We have examples of cycling for 2-approval.

During the writing of this paper, we noticed that Lev and Rosenschein have also considered similar questions and have obtained quite similar results [103]. However, our paper is completely independent from their work and has a different approach. We now give a brief discussion of the similarities and differences between these papers.

Both papers give convergence results for antiplurality under alphabetical tie-breaking: our Theorem 4.14 corresponds to [103, Theorem 13]. Both show nonconvergence for k-approval (Example 4.32 vs Theorem 19) and Borda (Example 4.29 vs Theorem 11). The counterexample for Borda in [103] works for any tie-breaking rule, and for $m \geq 4$, whereas ours works for $m \geq 3$ but uses a specific tie-breaking rule. In addition, [103] gives a counterexample for the maximin rule with a non-lexicographic deterministic tie-breaking rule, while we consider only scoring rules.

[103] deals only with deterministic tie-breaking, while we discuss randomized tie-breaking in some detail and show that stochastic dominance is the sufficient condition for ensuring convergence. Furthermore, we consider plurality and show how the proofs for antiplurality and plurality are essentially dual to each other. Our convergence proofs are shorter and, in our view, simpler. The upper bound in [103, Lemma 17] for antiplurality is $(m-2)n$ which can be contradicted by considering $p_0 = (bac, cab)$. If voters start from $(-b, -c)\{a\} \xrightarrow{1} (-a, -c)\{b\} \xrightarrow{2} (-a, -b)\{c\} \xrightarrow{1} (-c, -b)\{a\ \}\ \ \Diamond$ where for $m = 3$, first voter has 2 moves. Therefore, first voter has $m - 1$ improvement steps.

As far as future directions go, an important issue in extending to other voting rules is to properly define a notion of restricted best reply which is general enough to encompass all "reasonable" moves by rational agents seeking to maximize their payoff at each step, yet doesn't allow cycles. Already Example 4.29 shows that this will be difficult for Borda. Our proof skeleton for plurality and antiplurality could be adopted provided this difficulty is overcome. However for this approach to work easily, we would need the composition of two improvement steps to yield the same situation as a single improvement step (as in the discussion of type 3 moves in the proof of Lemma 4.12). One possible way of overcoming this problem would be to impose a domain restriction (do not allow all possible preference profiles to occur). Conceivably this might even allow type 2 moves as defined above to be reinstated as allowable improvement steps, while still maintaining convergence.

Chapter 5

Coordination via Polling in Plurality Voting Games under Inertia

5.1 Introduction

Voting as a preference aggregation method is widely used in human society and artificially designed systems of software agents. A large amount of recent research has considered the situation where a single individual or a small coalition attempts to manipulate an election result in its favour, assuming the remaining agents are naive (that is, always vote sincerely). Such an assumption on agent behaviour can be justified if the goal is to prove computational hardness results. However, if we wish to understand how voting rules function under fully strategic behaviour, we need to study a game-theoretic model of strategic manipulation.

The plurality rule is the most widely used voting rule, despite substantial criticism from social choice theorists. One point in its favour is its simplicity and spaceefficiency: an agent needs only report a single alternative instead of submitting a full preference order, a list of utilities, or a binary approval vector, as is the case with most other rules. However, even such a simple rule can become complicated when strategic voting behaviour is considered. In this paper, we study plurality voting under the assumption that all agents act strategically, as a starting point for a study of further classes of rules.

Voting games notoriously have many equilibria, and agents often cannot coordinate on a particular equilibrium outcome. Hence, voting games are hard to understand. The lack of publicly known information can exacerbate the lack of coordination of agents. A commonly used device that addresses the coordination issue, especially for plurality elections, is to use publicly announced pre-election polls. Such polls, which amount to an approximate simulation of an election with the same agents and alternatives, increase the commonly known information among agents and may influence their strategic behaviour. However, the beliefs of agents regarding the accuracy of these results can be different. This is a key point in the present paper, and we introduce the concept of inertia to describe these differences in beliefs.

Several authors from the political science and economics disciplines have discussed the influence of pre-election polls in plurality elections, both empirically and theoretically. The key topic of interest is what is called "Duverger's law", a general political science principle stating that plurality voting tends to lead to two-party competition [104]. More recently some papers have appeared that study equilibria in plurality voting games from a more algorithmic viewpoint (e.g. [101, 88]). Most of the models that have been used, with a few exceptions (e.g. [85, 101]), concern static equilibria, classifying them as "duvergerian" or "non-duvergerian", and do not attempt to discuss the dynamic process of converging to equilibria via the use of polls. There are several important differences between our work and existing literature. One of the differences is related to the different amount of information and strategic behaviour of agents. The other extra feature considered in the present paper is agent-dependent beliefs about the reliability of this information.

5.1.1 Our contribution

We present a model for plurality elections that allows for heterogeneous agents. We introduce the concept of an agent's **inertia**, which is that agent's perception of the accuracy of the poll result. This perception is the result of each agent's belief about such sources of error as coverage bias, miscounting, roundoff error, and noise in the announcement of results. This concept is rather general and seems realistic enough to be used for both human society and for designed systems of autonomous agents. This article focuses on the plurality rule, places some restrictions on agent behaviour, and considers some particular distributions of inertia. We present some numerical and analytic results on convergence to equilibria, both duvergerian and non-duvergerian. For example, a duvergerian equilibrium often occurs when all agents have the same value of inertia.

5.2 Game model

We have a set of agents whose set of allowable actions is to vote for a single alternative (not necessarily their most desirable alternative). Abstention is not allowed.

Each agent has a total order on the set of alternatives (indifference is not allowed) but as the voting rule is plurality, they vote for one alternative. Agents participate in a sequence of pre-election polls before the real election. In our model, these polls include all agents and alternatives in real election, not just a random sample. The information that these polls reveal does not have any effect on the agents' sincere preference order. In fact, we are interested in the strategic voting effect of polls rather than the so-called bandwagon or underdog effects considered in some papers [105]. In those papers, agents do not have a fixed preference order and their preference for an alternative is influenced by the popularity of that alternative.

We now discuss the assumptions in our model regarding the information and strategic behaviour of agents.

The information available to agents

The amount of information available to agents is a very important factor in their choice of strategy. The effect of poll information on the election result has been discussed in [106]. Complete information in plurality voting has been assumed in [107] and there is incomplete information in [108].

In the context of a repeated game, such as this sequence of polls under the plurality rule, in order to have complete information each agent would have to know how many agents of each **type** (sincere preference order) there are (this is usually called the *voting situation*). Even if this is unknown, we might expect to know the number of agents expressing each preference order in the previous poll. However, opinion polls for plurality will typically report only the number of agents ranking each alternative first, which we call the *scoreboard*. This lack of information on further preferences of other agents is crucial in the analysis below.

We use the concept of *inertia* to describe the reaction of agents toward the announced poll result. Agent coverage bias, miscounting or error and noise in announcing the result cause different values of uncertainty. This uncertainty brings about an inertia in agents. Each agent has an inertia value from the interval $[0, 1]$. An agent with inertia value of zero believes that the poll result is accurate. However, the poll result is meaningless to an agent with inertia value of one. In fact this agent does not consider the poll result in his decision making process. Other agents lie between these two extremes. Each agent's inertia value does not change during the sequence of polls. This seems reasonable because the set of participants in each poll does not change (it is always the entire set of agents), and the same system is used for counting and announcing the results in polls.

As far as we know this concept is new. The probability of miscounting has been discussed in [107], but is the same for all agents, whereas we have different values of inertia for different agents. The Poisson model of population uncertainty, in which there is uncertainty about the numbers of each type of agent, has been considered in [109]. In this paper agents have beliefs about these numbers that have been modelled as independent Poisson random variables. However, in our model, each agent just knows his own inertia and sincere preference order, and the scoreboard after each poll. This assumption makes sense for a system with no communication or coordination. This incomplete information influences the equilibrium result. Roughly speaking, it allows more alternatives to remain viable from the viewpoint of each agent.

The strategic behaviour of agents

The voting game described so far is still very general and allows for a wide range of outcomes. Voting games with more than two alternatives have many Nash equilibria and are not necessarily dominance solvable [87]. Eliminating dominated strategies is not sufficient to determine the result. Other refinements of equilibria such as strong and coalition-proof Nash equilibria do not always exist [99]. Some authors try to restrict the strategies of players by additional assumptions such as by assuming no voting for an alternative from another party [110].

In this paper, we assume agents have lexicographic preferences. Each agent infinitely prefers alternative x to alternative y, so he does not ignore any chance of winning of a more preferred alternative x [111]. Lexicographic preferences are not consistent with the idea of a cardinal utility function and probabilities are not relevant. Rather, they give a strong bias toward sincere voting which can still be overcome when an alternative is perceived to be a definite loser.

We also assume that each voter votes in each poll in the same way that he would if that poll were the actual election. One scenario in which this would occur is when voters do not know whether the current poll is the actual election. For example, the system designer may introduce this requirement. Thus voters will not attempt to vote strategically in the sense of misleading other voters, although they do vote strategically in the sense of playing their perceived best response. Note that the restricted information given by the scoreboard helps in this regard. For example, if bca voters could infer how many cab voters there were, they could vote for c in order that the cab voters do not abandon c, which might allow a to defeat b.

Therefore, agents vote for their most preferred alternative whom they perceive as having a non-zero chance of winning in further polls.

After each poll, each agent considers a set W of potential winners, consisting of all alternatives whom that agent perceives as having non-zero chance to win sometime in future. This set does not depend on the agents' preference order and only depends on the scoreboard and his inertia value. Agents update this set after the announced result of each poll. Agents start by voting sincerely in the first poll. Then, they update their votes according to their beliefs about potential winners during the sequence of polls. All these assumptions on behaviour are common knowledge as far as agents are concerned.

5.3 Game dynamics

5.3.1 Notation

There is a set C of alternatives (we use index c for alternatives) which has m members, and a set V of players with n members (we use index ν for agents). We consider a sequence of K polls indexed by k, where the last poll is the election. However, agents are not aware of the value of K. Each agent has a sincere strict preference order on alternatives. There are $m!$ different preference orders (or types) which are indexed by t. We have plurality as our scoring rule in which each agent votes for only one alternative. Therefore, we can assume that the set of possible strategies for player ν is $S_\nu = C$. We use the following notations through the paper:

- $s_k(c)$: the normalized score of alternative c in poll k, namely the proportion of agents who have voted for c at poll k,

- $c_k(h)$: the alternative who has h-th highest score in poll k (e.g. $c_k(1)$ is the winner of poll k, note that we do not consider ties in this paper as this case occurs relatively rarely in large electorates),
- v_t: the number of agents with type (or preference order) t,
- $W_{\varepsilon,k}$: the set of potential winners from the view point of player with inertia value ϵ according to the result of poll k,
- $V_{c,k}$: the set of agents who vote for alternative c in poll k.

Definition 5.1 (The concept of certain and doubtful). *Suppose that according to the poll result $s_k(i) < s_k(j)$. An agent with inertia ε is* certain *about this statement if*

$$(1+\varepsilon)s_k(i) < (1-\varepsilon)s_k(j). \tag{5.1}$$

Otherwise, he is doubtful.

Note that this formula implies that if inertia of an agent is 0, then he will always be certain that j is ahead of i provided that such a result is reported. Also, Equation (5.1) implies that an agent with inertia equal to 1 will always be doubtful of any claimed scores.

The supporters of each alternative may be certain that the score of their favoured alternative is less than the winner, yet they might still consider that alternative as a potential winner and vote for him in the next poll. We study the concept of potential winner in the next section.

Example 5.2. *Consider a 3 alternative election, and suppose the result of poll k is $s_k(c_k(1)) = 45\%$, $s_k(c_k(2)) = 30\%$ and $s_k(c_k(3)) = 25\%$. Any agent with* *inertia less than $\frac{1}{11}$ is certain that alternative 3 has fewer votes than alternative 2, but agents with inertia more than that are doubtful about this statement. In other words, those with $\varepsilon > \frac{1}{11}$ do not use this statement, while the others consider it in their strategic computations.*

5.3.2 Set of potential winners

In the initial state ($k = 0$), an agent with inertia ε does not have any information about the number of supporters of each alternative. Therefore, he sees all alternatives as potential winners, $W_{\varepsilon,0} = C$, and he votes sincerely in the first poll. For the next poll, the agent votes for the most desirable alternative who can win in future (not necessarily the next poll) according to his interpretation of the poll result and the voting strategies of other agents (the strategy of agents is common knowledge).

Each agent's set of potential winners should satisfy some basic properties. The key necessary properties that we require are as follows. These are all common knowledge.

- non-emptiness: Any agent with any inertia value ε believes that there exists at least one candidate with a positive chance of winning. W should clearly be nonempty for every voter, and contain the highest scoring candidate in the current poll.
- upward closure: if an agent with inertia ε believes that $c_k(x) \in W_{\varepsilon,k}$, then he believes $c_k(x-1) \in W_{\varepsilon,k}$. This seems reasonable: if an agent believes that some alternatives have a chance to win in future in the best case, then

that agent also believes that all alternatives with higher current poll support also have a chance to win in future.

- overtaking: a possible winner must be able to overtake a higher scoring candidate who is also a possible winner. Overtaking the next higher scoring alternative is a necessary condition for winning, because the only chance an alternative has for attracting more support is that he improves his ranking position in the scoreboard. This is justified by the belief of agents about the upper closure of set of potential winners. For overtaking, alternative $c_k(x)$ needs extra support, and this support can only be obtained from the supporters of alternatives with a lower score than alternative $c_k(x)$. This is because agents who have already voted for higher scoring alternatives than $c_k(x)$ will change their votes to $c_k(x)$ if they perceive that their current choice does not have any chance to win. Upper closure of $W_{\varepsilon,k}$ would then lead to inconsistent beliefs.

  If $c_k(x)$ cannot overtake $c_k(x-1)$ in the next poll, in the most favourable case, then $x \notin W_{\varepsilon,k}$. We describe this case precisely in Proposition 5.5.

We first give an example to give the intuition behind our definitions.

Example 5.3. *Consider scoreboard $(a, b, c, d) = (40\%, 29\%, 21\%, 10\%)$ and agent ν with $\varepsilon = 0$. Voter ν reasons as follows: for each agent with inertia ε, either alternative $d \in W_{\varepsilon,k}$ or not. If yes, then also alternatives $a, b, c \in W_{\sigma,k}$ (upward closure). The agents whose most desirable potential winner is alternative d have already voted for him, and the other agents prefer to vote for alternatives a, b or c in the next poll. Thus, the score of d cannot be increased and $d \notin W_{0,k}$. However, alternative $c \in W_{0,k}$ because it is possible that all supporters of alternative d*

switch to c, yielding scoreboard $(40\%, 29\%, 31\%, 0)$, and c can overtake alternative b, and in the next round all b-supporters may switch to alternative c, and he can overtake alternative a. Because of upward closure $b, a \in W_{0,k}$.

The basic properties above show that the currently highest-scoring alternative is always considered a potential winner by each agent. The necessary conditions do not define W uniquely. Because of lexicographic preferences, voters do not abandon candidates easily, and so it makes sense that W should be as large as possible. Of course if voters voted differently in the polls and the election (for example if they know that the next round is the election and have no other constraints on strategic action), W might be smaller. For example, a candidate may be able to win by successively attracting support from others, but the number of rounds remaining may not be enough for this to occur. We are ruling out this case by our assumptions on voter behaviour. For example, uncertainty about the time of the actual election allied to lexicographic preferences implies that W should be as large as possible. Thus we argue that the necessary conditions are sufficient.

We now show how to define the set of potential winners recursively starting from the top scoring alternative.

Definition 5.4. *For $2 \leq i \leq m$, define condition $C_{ik\varepsilon}$ by*

$$(1+\varepsilon)\sum_{h \geq i} s_k(c_k(h)) > (1-\varepsilon) s_k(c_k(i-1)). \tag{$C_{ik\varepsilon}$}$$

Proposition 5.5 (The conditions for being a potential winner). *After the announced result of poll k, $c_k(x) \in W_{\varepsilon,k}$ if and only if all conditions $C_{ik\varepsilon}$ for $2 \leq i \leq x$ hold. Algorithm 1 computes the set $W_{\varepsilon,k}$.*

Algorithm 1 Function for constructing $W_{\varepsilon,k}$

Require: $k \geq 1$
  $W_{\varepsilon,k} = \{c_k(1)\}$
  for $i = 2$ to m **do**
    if Condition $C_{ik\varepsilon}$ holds **then**
      $W_{\varepsilon,k} = W_{\varepsilon,k} \cup \{c_k(i)\}$
    else
      break
    end if
  end for

Proof. Upward closure shows that the best chance of $c_k(x)$ overtaking $c_k(x-1)$ consists of attracting all supporters of agents currently voting for alternatives $c_k(h)$ with $h > x$, and retaining all current supporters. This yields condition $C_{xk\varepsilon}$, and so Algorithm 1 is clearly correct. Since overtaking of even higher alternatives must occur also, unrolling the loop in Algorithm 1 yields the result. □

Remark 5.6. *In the majority case from the viewpoint of an agent with inertia value ε, in which*

$$(1-\varepsilon)s_k(c_k(1)) > (1+\varepsilon) \sum_{c \neq c_{k(1)}} s_k(c),$$

alternative $c_k(2)$ and consequently all other alternatives except $c_k(1)$ do not have any chance to win in the future. Thus, $W_{\varepsilon,k} = \{c_k(1)\}$.

Example 5.7. *Suppose the result of poll k is $s_k(a) = 55\%$, $s_k(b) = 30\%$ and $s_k(c) = 15\%$. According to Proposition 5.5,*

$$W_{\varepsilon,k} = \begin{cases} \{a\} & 0 \leq \varepsilon \leq \frac{1}{11}; \\ \{a,b\} & \frac{1}{11} < \varepsilon \leq \frac{1}{3}; \\ \{a,b,c\} & \frac{1}{3} < \varepsilon \leq 1. \end{cases}$$

Therefore, we have 3 different sets for $W_{\varepsilon,k}$ based on the inertia value of agents. In the first inertia value interval, agents perceive the result of poll k as a majority case. Therefore, their set of potential winners is a singleton and they vote for a in poll $k+1$. In the second inertia value interval, they vote for a or b in poll $k+1$ based on their preference order. For example, an agent with preference order cab votes for a and an agent with preference order cba votes for b in poll $k+1$. In the third case where agents have high inertia, they do not care about the announced result of the poll. In fact, they believe each candidate to be viable and they just vote sincerely in poll $k+1$. An agent with inertia value of 1 always votes sincerely, regardless of the poll result.

5.4 Equilibrium results for some special cases

5.4.1 Zero inertia

In the special case where inertia is identically zero for all agents, the set of potential winners is identical for all agents. We show that in this case the sequence of polls converges to a duvergerian equilibrium, i.e., a two party competition. Note that the inertia value is fixed in all polls and also we assume there is no majority case.

Theorem 2 (**duvergerian equilibrium**). *In a plurality voting game with common* *inertia value $\varepsilon = 0$, the polling sequence yields a duvergerian equilibrium in a non-majority case after at most $m - 2$ polls.*

Proof. Let m be the number of alternatives and $\varepsilon = 0$. As agents have the same value of inertia, either all agents perceive the result as majority case or all of them perceive it as a non-majority case. As we explained before, in the majority case, agents vote for the highest scoring alternative (refer to Remark 5.6). In a non-majority case, we have $(s_k(c_k(1)) \leq \sum_{c \neq c_k(1)} s_k(c)$. According to Proposition 5.5, $c_k(2) \in W_{0,k}$, therefore, $\mid W_{0,k} \mid \geq 2$.

For all $\nu \in V_{c,k}$ for which $c \in C \setminus W_{0,k}$, ν changes his vote to his most desirable alternative in $W_{0,k}$. Thus, $s_{k+1}(c) = 0$, for each $c \in C \setminus W_{0,k}$. According to Proposition 5.5, $c_k(m) \notin W_{0,k}$. Therefore, in each poll, at least the last scored alternative is eliminated and after at most $m - 2$ polls, we have a duvergerian equilibrium. $\square$

Remark 5.8. *There is a connection with the voting method instant-runoff (IRV). When $m = 3$, if inertia is identically zero then our assumptions mean that the plurality election is actually just IRV. For general inertia and general m, we could fix some $\beta > 0$ and require that the election system automatically deletes the alternative whose support becomes less than β for the next poll. If we assume that 2 alternatives do not reach this boundary β simultaneously, we again simulate IRV. However, our procedure is more general, as several alternatives may be eliminated at one step.*

5.4.2 Constant non-zero inertia

Suppose that all agents have the same value of inertia θ, with $0 < \theta \leq 1$. Again note that the set of potential winners is identical for all agents at all times and the inertia value is fixed in all polls. This case is similar to the setup of Messner and Polborn [107] where the probability of miscounting is positive but small. Messner and Polborn introduce the concept of robust equilibrium and show that for plurality games with 3 alternatives, all such equilibria are duvergerian. However, in that paper, the value of θ is common knowledge between all agents, and this is not the case in our model. The behavioural assumptions of agents also differ. Paper [107] shows that duvergerian equilibrium happens in all robust equilibria of plurality games with 3 alternatives.

We consider a 3-alternative election with a large number of agents, with a fixed inertia value θ which is the same for all agents. W.l.o.g. we may assume that $s_1(c) < s_1(b) < s_1(a)$. We also assume there is no majority case (refer to Remark 5.6).

Proposition 5.9. *Let*

$$\theta' = \max\{\frac{s_1(a) - s_1(b) - s_1(c)}{s_1(a) + s_1(b) + s_1(c)}, \frac{s_1(b) - s_1(c)}{s_1(b) + s_1(c)}\}. \tag{5.2}$$

A c supporter with inertia $\theta \leq \theta'$ will change his vote to a or b in the second poll.

Proof. According to Proposition 5.5,

$$c \in W_{\theta,1} \Leftrightarrow \begin{cases} (1+\theta)(s_1(b) + s_1(c)) > (1-\theta)s_1(a) \\ (1+\theta)s_1(c) > (1-\theta)s_1(b) \end{cases}$$

Therefore, $c \in W_{\theta,1} \Leftrightarrow \theta > \theta'$, and $c \in C \setminus W_{\theta,1} \Leftrightarrow \theta \leq \theta'$. □

Theorem 3. *Consider a plurality voting game with $m = 3$, and fixed inertia value θ which is the same for all agents. Assuming a non-majority case, the polling sequence yields a duvergerian equilibrium after 1 poll if $\theta \leq \theta'$.*

Proof. Similar to previous case, as agents have the same value of inertia, either all agents perceive the result as majority case or all of them perceive it as a non-majority case. As we explained before, in the majority case, agents vote for the highest scoring alternative (refer to Remark 5.6). In a non-majority case, according to Proposition 5.9, as the inertia values of all agents are equal, c supporters abandon c immediately, and a duvergerian equilibrium is reached after one poll. □

Remark 5.10. *Note that same constant non-zero inertia cases do not yield duvergerian equilibrium, depending on the value of θ. If $\theta > \theta'$, then every agent continues voting sincerely and the poll results will not change in the sequence.*

Example 5.11. *Consider plurality rule with 3 alternatives where the the scoreboard of the first poll is $(40\%, 35\%, 25\%)$. If the inertia value of all agents are θ and $\theta \leq \frac{1}{6}$, we have a duvergerian equilibrium.*

5.4.3 Uniform distribution of inertia

We consider a 3-alternative election with a large number of agents, with a uniform inertia distribution on [0,1]. We describe the initial setup via a quadruple which is based on the first poll result $(s_1(a), s_1(b), s_1(c))$ and the true percentage v_6 of type cba agents (note this value is not known to any agent). W.l.o.g., we may assume that $s_1(c) < s_1(b) < s_1(a)$ and we approximate the discrete uniform distribution across agents by a continuous one for purposes of computation.

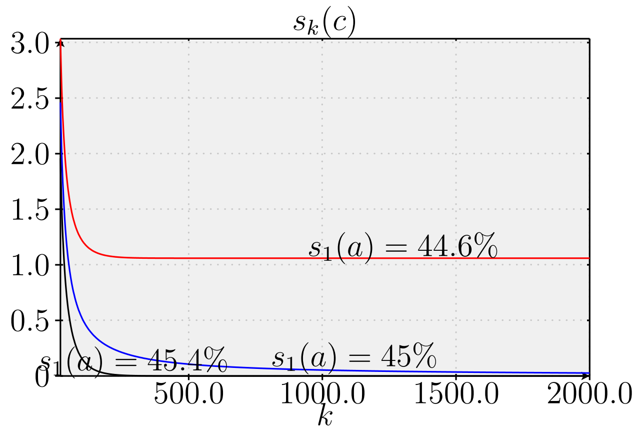

Figure 5.1: Score of the last alternative (c) as a function of k with uniform inertia distribution for three different cases where $V = (s_1(a), 35\%, 100\% - s_1(a) - 35\%, 5\%)$

All c supporters who believe that c is a loser change their votes in favour of their second alternative. The percentage of type t agents (cab and cba) who vote in favour of alternative i (a and b respectively) in poll $k+1$ is denoted by $\alpha_{t,i,k}$. Note that the assumption of a common inertia distribution implies that for all k, $\alpha_{cab,a,k} = \alpha_{cba,b,k} \equiv \alpha_k$ and $\alpha_0 = 0$.

Proposition 5.12. *For a uniform distribution of inertia for all agents during the sequence of polls and initial result* $V = (s_1(a), s_1(b), s_1(c), v_6)$, *we have*

$$\alpha_k = \frac{1}{1 + \frac{2^k\left(\frac{s_1(c)-v_6}{s_1(b)+v_6}\right)^k (s_1(b)+v_6-2s_1(c))}{(s_1(b)-s_1(c))\left(-2^k\left(\frac{s_1(c)-v_6}{s_1(b)+v_6}\right)^k + \left(1-\frac{v_6}{s_1(c)}\right)^k\right)}} \tag{5.3}$$

Proof. According to the order of alternatives in the first poll and Proposition 5.5, a c supporter concludes that c is a loser and changes his vote if $(1+\varepsilon)s_k(c) < (1-\varepsilon)s_k(b)$.

Therefore, $\alpha_k = p\{\varepsilon < \frac{s_k(b)-s_k(c)}{s_k(b)+s_k(c)}\}$. The score of alternatives a, b and c in poll k

is given by:

$$s_k(a) = s_1(a) + \alpha_{k-1}v_5 \qquad s_k(b) = s_1(b) + \alpha_{k-1}v_6 \tag{5.4}$$

$$s_k(c) = s_1(c) - \alpha_{k-1}v_6 - \alpha_{k-1}v_5 \tag{5.5}$$

Therefore,

$$\alpha_k = p\{\varepsilon < \frac{s_1(b) - s_1(c) + \alpha_{k-1}(s_1(c) + v_6)}{s_1(b) + s_1(b) - \alpha_{k-1}(s_1(c) - v_6)}\} \text{ for all } k \geq 1. \tag{5.6}$$

The stated solution formula for this recurrence is readily established by induction.

□

Proposition 5.13. *The score of the last alternative in the first poll (which we denote by c) satisfies*

$$\lim_{k\to\infty} s_k(c) = \begin{cases} 0 & \textit{if } s_1(b) + v_6 \geq 2s_1(c) \\ \left(\frac{2s_1(c) - v_6 - s_1(b)}{s_1(c) - v_6}\right) s_1(c) & \textit{if } s_1(b) + v_6 < 2s_1(c) \end{cases} \tag{5.7}$$

Proof. The score of alternative c after $k+1$ polls is

$$s_{k+1}(c) = (1 - \alpha_k)s_1(c) \tag{5.8}$$

According to Proposition 5.12, if we converge k to infinity, we have

$$\lim_{k\to\infty} \alpha_k = \begin{cases} 1 & s_1(b) + v_6 \geq 2s_1(c); \\ \frac{s_1(b) - s_1(c)}{s_1(c) - v_6} & s_1(b) + v_6 < 2s_1(c). \end{cases}$$

The result follows immediately.

□

Remark 5.14. *The convergence to zero is exponentially fast with the exponential rate decreasing as we approach the boundary between the two cases, and at the boundary it is subexponential. Figure 5.1 shows three special cases (the boundary case and 2 different cases in its neighbourhood).*

Theorem 4. *In a plurality voting game with 3 alternatives and initial result* $V = (s_1(a), s_1(b), s_1(c), v_6)$ *and uniform distribution of inertia, the polling sequence yields a duvergerian equilibrium if and only if* $s_1(b) + v_6 \geq 2s_1(c)$.

Proof. Follows immediately from Proposition 5.13. □

Fig 5.1 illustrates this inequality when $v_6 = 5\%$ and $s_1(b) = 35\%$. For $s_1(a) \geq 45\%$, we have a duvergerian equilibrium.

5.4.4 Other distributions of inertia

The above results are for very special inertia distributions; explicit analysis of this type is not possible for general distributions. In this subsection, we investigate some different distributions via numerical simulations. Intuitively, we expect that distributions skewed to the left (with more agents of low inertia) will converge to the $\varepsilon \equiv 0$ case more quickly.

We consider the continuous triangular distribution $T(p)$ whose density function's graph has vertices at $(0, 0)$, $(p, 2)$ and $(1, 0)$.

Example 5.15 (The effect of inertia distribution: Triangular vs. Uniform). *Consider the initial result* $V = (s_1(a), s_1(b), s_1(c), v_6) = (45\%, 35\%, 20\%, 5\%)$.

According to Theorem 4, we have a limiting duvergerian equilibrium for uniform inertia distribution. Numerical results in Figure 5.1 (the line for $s_1(a) = 45\%$*) also confirm this result. When we change the inertia distribution to be triangular with apex 0.5, we have the result in Figure 5.2. As we see in Figure 1, the convergence is very slow but changing the inertia distribution to* $T(0.5)$ *accelerates the process.*

Example 5.16 (The effect of voting situation). *In Figure 5.2, we have* 5% *cba agents. Figure 5.3 shows the result of the same situation with* 10% *cba agents which leads to a faster convergence. Note that the voting situation is not known to agents.*

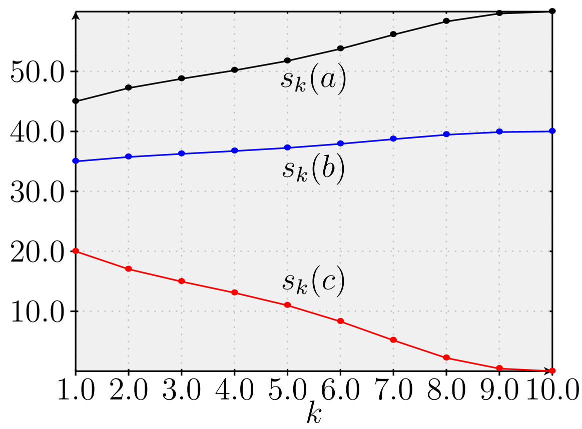

Figure 5.2: $V = (45\%, 35\%, 20\%, 5\%)$ and $T(0.5)$ inertia distribution

Example 5.17 (The effect of skewness of inertia distribution). *Consider* $V = (40\%, 35\%, 25\%, 10\%)$ *with an inertia distribution of* $T(0.5)$*. This yields a non-duvergerian equilibrium, and it appears that the score of* c *converges to 22, as shown in Figure 5.4. However, the same voting situation with an inertia distribution* $T(0.3)$ *results in a duvergerian equilibrium as shown in Figure 5.5. In this*

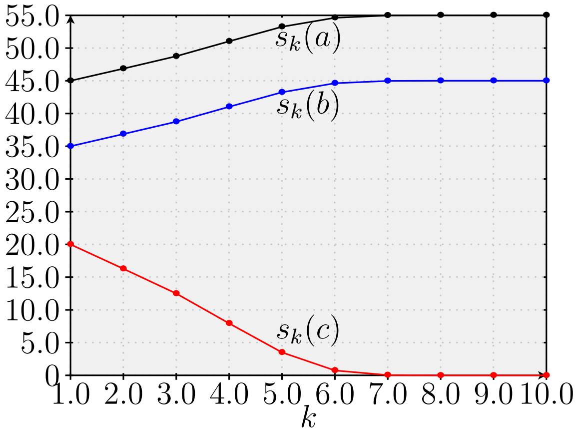

Figure 5.3: $V = (45\%, 35\%, 20\%, 10\%)$ and $T(0.5)$ inertia distribution

case, more agents validate the poll result, and we have a duvergerian equilibrium after 10 polls.

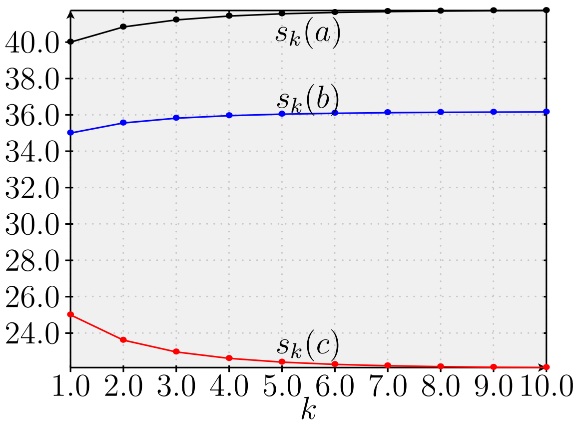

Figure 5.4: $V = (40\%, 35\%, 25\%, 10\%)$ and $T(0.5)$ inertia distribution

5.5 Conclusion and future directions

In this paper we tried to study a repeated game with unknown number of rounds and incomplete information. The strategy of each player depends on his belief

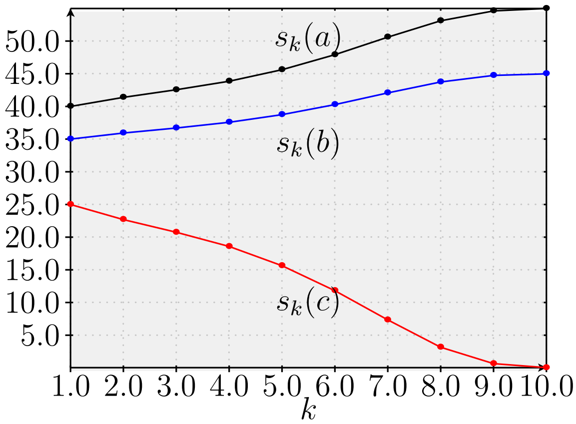

Figure 5.5: $V = (40\%, 35\%, 25\%, 10\%)$ and $T(0.3)$ inertia distribution

about the belief of other players. The sequence of opinion polls helps agents to coordinate on an equilibrium in an environment with some uncertainties about the accuracy of these polls. The amount of information available to agents has a critical role in influencing the strategic choices of agents. In this paper, we try to simplify the model with some assumptions about the strategy of players as a starting point for studying this game. Even in this simplified model, there are too many special cases that can happen depending on the inertia distribution or preference distribution of agents. We try to explain the model by some examples that give insight into different scenarios.

As a future direction, it is interesting to study how the strategy of agents will change if they have more information or in a more complicated model, each agent has different amounts of information. For example, some agents may have extra information than others regarding the inertia distribution of other agents or their preference order or the number of rounds ahead. Therefore, they may have different belief about the strategy of each agent.

Another interesting direction would be to to allow inertia to change from one poll to the next. For example, if random sampling is used instead of polling all voters, the sample size might vary between polls. More generally we want to explore the effect of inertia in other models with different behavioural assumptions for example, when voters use some simple heuristic strategies. We expect to observe substantial differences in equilibrium outcomes when non-zero inertia is introduced into the model.

Chapter 6

Conclusions

In this chapter, we first present a brief summary of the research which was discussed throughout the thesis, and also we mention some relevant papers which have been published recently (after writing of our papers) or have not been discussed during the relevant chapters. In the second section, we have a brief discussion of the other work in the general area which we did not study during this thesis, and can be future directions of this research. The future directions of each chapter are discussed separately at the end of that chapter.

6.1 Summary

In the first part of this thesis we discussed that although there are preference profiles that do not admit any strategic manipulation, it has become clear to researchers that rules that are never manipulable must be very hard to find. This negative result has inspired several strands of research. One strand proceeds by weakening the assumptions of single-valuedness, leading to many results, most of which have the same negative character. Another direction ("domain conditions") is to sacrifice universality, and for a given set of preference profiles, to attempt to find strategy-proof social choice function on this set. The other strand is to quantitatively measure the manipulability of each rule, with the aim of discovering the rules with minimum manipulability.

In this strand we studied the manipulability of scoring rules and Copeland's method by introducing a new measure in Section 1.4.

Xia recently has generalized the asymptotic behaviour study of strategic manipulation under a general distribution of preferences and the fixed number of candidates. In this model, all types of strategic behaviour are unified as vote operations [12].

In this strand the computational hardness of manipulation has also been studied with the worst-case and typical-case analysis, approximability, and heuristic approaches for various definitions of manipulation such as control, bribery and possible winner.

The probability of safe manipulation was studied in Chapter 2. [112] has discussed the complexity of safe manipulation under scoring rules.

The complexity of optimal manipulation, i.e., finding a strategic vote that brings about the manipulator's goal yet deviates as little as possible from his sincere preference order is studied in [113]. They have obtained polynomial-time algorithms for all scoring rules.

As we discussed earlier limiting the domains of voters' preference orders has been known as another way for decreasing the possibility of manipulation and control. However recent results show that in some cases the single peaked preferences are more vulnerable to manipulation and control [114].

In the second part of the thesis we concentrated on non-cooperative study of strategic manipulation voting games. The equilibrium result and its convergence was discussed at this part. As we saw in Chapters 4 and 5, the behavioural assumptions of agents and available amount of information affect the outcomes of games considerably. Best reply dynamics for scoring rules was studied in Chapter 4. In Chapter 5, we concentrated on plurality voting rule and studied a new model where voters have partial information via pre-election polls and also have some value of uncertainty regarding the result of these polls.

Recently, Elkind and Erdèlyi have studied manipulation under voting rule uncertainty where manipulators have uncertainty regarding the voting rule, and should choose their strategies independent of the voting rule [115].

In this thesis, we concentrated on simultaneous voting games. Sequential voting game is one possible future direction for this research such as games with multiple binary issues that are sequentially voted on by the voters [116] or Stackelberg voting games [117].

6.2 Future directions

During this thesis, we just concentrated on the problem of strategic manipulation. In this section, we intend to summarise some of the other topics in this field briefly.

By reviewing the papers in computational social choice area, we can find a large number of papers considering

- The mechanism design of social choice functions with some desirable behaviours,

- Communication and privacy complexity which analyzes the communication requirements in voting systems for making a decision,
- The computational aspects of fair division such as cake-cutting and allocations of indivisible goods,
- The computational aspects of coalitional voting games,
- Social choice theory in combinatorial domains such as the reasoning of combinatorial preferences, the compact representation of preferences for multi-issue topics and preference aggregations.

Belief and judgement aggregation are other topics that have been considered in some papers such as [118, 119].

Another topic in this area is matching problem where we should find a pair for each element of two groups by considering some preferences regarding the elements of each group. For example finding a match in a marriage decision or finding a correct match in kidney donations.

Vote elicitation in multiagent systems is another topic which has not been discussed in this thesis. The information elicited from an agent depends on what other agents have revealed about their preferences. Depending on the elicitation costs across voters and number of candidates, complexity and strategy-proofness of vote elicitation differ. Walsh has considered the complexity issues in preference elicitation and manipulation in [120].

Bibliography

[1] R. Reyhani, G. Pritchard, and M. C. Wilson, "Manipulation of voting rules," Tech. Rep. 357, University of Auckland CDMTCS Research Report Series, 2009.

[2] G. Pritchard, R. Reyhani, and M. C. Wilson, "Power measures derived from the sequential query process," *Mathematical Social Sciences*, vol. 65, no. 3, pp. 174–180, 2013.

[3] M. C. Wilson and R. Reyhani, "The probability of safe manipulation," in *Proceedings of Third International Workshop on Computational Social Choice (COMSOC 2010), Düsseldorf, Germany*, pp. 423–431, Düsseldorf university press, Sep 2010.

[4] R. Reyhani and M. C. Wilson, "Best reply dynamics for scoring rules," in *Proceedings of 20th European Conference on Artificial Intelligence, (ECAI 2012)*, vol. 242, pp. 672–677, Montpellier, France: IOS Press, 2012.

[5] R. Reyhani, M. C. Wilson, and J. Khazaei, "Coordination via polling in plurality voting games under inertia," in *Proceedings of Fourth International Workshop on Computational Social Choice, (COMSOC 2012), Kraków, Poland*, Sep 2012.

[6] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet, "A short introduction to computational social choice," in *SOFSEM 2007: Theory and Practice of Computer Science* (J. van Leeuwen, G. Italiano, W. van der Hoek, C. Meinel, H. Sack, and F. Plil, eds.), vol. 4362 of *Lecture Notes in Computer Science*, pp. 51–69, Springer Berlin, Heidelberg, 2007.

[7] A. Gibbard, "Manipulation of voting schemes: a general result," *Econometrica*, vol. 41, pp. 587–601, 1973.

[8] M. A. Satterthwaite, "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of Economic Theory*, vol. 10, no. 2, pp. 187–217, 1975.

[9] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe, "Copeland voting fully resists constructive control," in *Algorithmic Aspects in Information and Management* (R. Fleischer and J. Xu, eds.), vol. 5034 of *Lecture Notes in Computer Science*, pp. 165–176, Springer Berlin / Heidelberg, 2008.

[10] I. McLean, H. Lorrey, and J. Colomer, "Voting in the medieval papacy and religious orders," in *Modeling Decisions for Artificial Intelligence* (V. Torra, Y. Narukawa, and Y. Yoshida, eds.), vol. 4617 of *Lecture Notes in Computer Science*, pp. 30–44, Springer Berlin / Heidelberg, 2007.

[11] K. Arrow, A. Sen, and K. Suzumura, *Handbook of Social Choice and Welfare*. Elsevier Science, 2002.

[12] L. Xia, “How many vote operations are needed to manipulate a voting system?,” *Arxiv preprint arXiv:1204.1231*, 2012.

[13] F. Brandt, V. Conitzer, and U. Endriss, “Computational social choice,” *To appear in G. Weiss (Ed.), Multiagents systems, MIT Press*, 2012.

[14] W. Gehrlein and D. Lepelley, “Probability calculations in voting theory: an overview of recent results,” *Reasoned Choices: Essays in Honor of Academy Professor Hannu Nurmi*, 2004.

[15] E. Stensholt, “Voces populi and the art of listening,” *Social Choice and Welfare*, vol. 35, pp. 291–317, 2010.

[16] K. Dowding and M. van Hees, “In praise of manipulation,” *British Journal of Political Science*, vol. 38, no. 01, pp. 1–15, 2008.

[17] A. Lehtinen, “The welfare consequences of strategic behaviour under approval and plurality voting,” *European Journal of Political Economy*, vol. 24, no. 3, pp. 688–704, 2008.

[18] A. D. Taylor, *Social choice and the mathematics of manipulation*. Outlooks, Cambridge: Cambridge University Press, 2005.

[19] T. Walsh, “Where are the really hard manipulation problems? the phase transition in manipulating the veto rule,” in *In Proceedings of the 21th IJCAI*, pp. 324–329, 2009.

[20] G. Pritchard and M. C. Wilson, “Exact results on manipulability of positional voting rules,” *Social Choice and Welfare*, vol. 29, no. 3, pp. 487–513, 2007.

[21] G. Pritchard and A. Slinko, “On the average minimum size of a manipulating coalition,” *Social Choice and Welfare*, vol. 27, no. 2, pp. 263–277, 2006.

[22] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra, “The complexity of bribery in elections,” *Proceedings of the 20th AAAI Conference on Artificial Intelligence, (AAAI 2006)*, vol. 21, pp. 641–646, 2006.

[23] E. Elkind, P. Faliszewski, and A. Slinko, “Swap bribery,” in *Algorithmic Game Theory* (M. Mavronicolas and V. Papadopoulou, eds.), vol. 5814 of *Lecture Notes in Computer Science*, pp. 299–310, Springer Berlin / Heidelberg, 2009.

[24] R. Christian, M. Fellows, F. Rosamond, and A. Slinko, “On complexity of lobbying in multiple referenda,” *Review of Economic Design*, vol. 11, no. 3, pp. 217–224, 2007.

[25] G. Erdélyi, L. Hemaspaandra, J. Rothe, and H. Spakowski, “On approximating optimal weighted lobbying and frequency of correctness versus average-case polynomial time,” in *Fundamentals of Computation Theory* (E. Csuhaj-Varjú and Z. Ésik, eds.), vol. 4639 of *Lecture Notes in Computer Science*, pp. 300–311, Springer Berlin / Heidelberg, 2007.

[26] I. Schlotter, P. Faliszewski, and E. Elkind, “Campaign management under approval-driven voting rules,” in *Proceedings of the 25th AAAI Conference on Artificial Intelligence, (AAAI 2011)*, pp. 726–731, 2011.

[27] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra, "Multimode control attacks on elections," *Journal of Artificial Intelligence Research*, vol. 40, pp. 305–351, 2011.

[28] E. Elkind, P. Faliszewski, and A. Slinko, "Cloning in elections," in *Proceedings of 24th AAAI Conference on Artificial Intelligence, (AAAI 2010)*, vol. 10, pp. 768–773, 2010.

[29] N. Betzler, S. Hemmann, and R. Niedermeier, "A multivariate complexity analysis of determining possible winners given incomplete votes," in *Proceedings of the 21th International Joint Conference on Artificial Intelligence (IJCAI 2009)*, vol. 2, pp. 53–58, 2009.

[30] L. Xia and V. Conitzer, "Determining possible and necessary winners under common voting rules given partial orders," *Journal of Artificial Intelligence Research*, vol. 41, pp. 25–67, May 2011.

[31] M. Zuckerman, A. D. Procaccia, and J. S. Rosenschein, "Algorithms for the coalitional manipulation problem," *Artificial Intelligence*, vol. 173, no. 2, pp. 392–412, 2009.

[32] P. Pattanaik, "Counter-threats and strategic manipulation under voting schemes," *The Review of Economic Studies*, vol. 43, no. 1, pp. 11–18, 1976.

[33] G. Pritchard and M. C. Wilson, "Asymptotics of the minimum manipulating coalition size for positional voting rules under impartial culture behaviour," *Mathematical Social Sciences*, vol. 58, no. 1, pp. 35–57, 2009.

[34] M. C. Wilson and G. Pritchard, "Probability calculations under the IAC hypothesis," *Mathematical Social Sciences*, vol. 54, no. 3, pp. 244–256, 2007.

[35] D. Lepelley, A. Louichi, and H. Smaoui, "On Ehrhart polynomials and probability calculations in voting theory," *Social Choice and Welfare*, vol. 30, no. 3, pp. 363–383, 2008.

[36] P. Favardin, D. Lepelley, and J. Serais, "Borda rule, Copeland method and strategic manipulation," *Review of Economic Design*, vol. 7, pp. 213–228, 2002.

[37] L. S. Shapley and M. Shubik, "A method for evaluating the distribution of power in a committee system," in *The Shapley value*, pp. 41–48, Cambridge: Cambridge Univ. Press, 1988.

[38] D. S. Felsenthal and M. Machover, *The measurement of voting power*. Cheltenham: Edward Elgar Publishing Limited, 1998. Theory and practice, problems and paradoxes.

[39] I. Lindner, "The power of a collectivity to act in weighted voting games with many small voters," *Social Choice and Welfare*, vol. 30, no. 4, pp. 581–601, 2008.

[40] F. Carreras, "A decisiveness index for simple games," *European Journal of Operational Research*, vol. 163, no. 2, pp. 370–387, 2005.

[41] D. Lepelley and B. Mbih, "The proportion of coalitionally unstable situations under the plurality rule," *Economics Letters*, vol. 24, no. 4, pp. 311–315, 1987.

[42] D. Lepelley and B. Mbih, “Strategic manipulation in committees using the plurality rule: alternative concepts and frequency calculations,” *Group Decision and Negotiation*, vol. 6, no. 2, pp. 119–138, 1997.

[43] D. Lepelley and B. Mbih, “The vulnerability of four social choice functions to coalitional manipulation of preferences,” *Social Choice and Welfare*, vol. 11, no. 3, pp. 253–265, 1994.

[44] J. R. Chamberlin, “An investigation into the relative manipulability of four voting systems,” *Behavioral Science*, vol. 30, no. 4, pp. 195–203, 1985.

[45] F. Aleskerov, D. Karabekyan, R. Sanver, and V. Yakuba, “Evaluating the degree of manipulability of certain aggregation procedures under multiple choices,” *Journal of the New Economic Association*, pp. 37–61, 2009.

[46] A. Slinko, “How the size of a coalition affects its chances to influence an election,” *Social Choice and Welfare*, vol. 26, no. 1, pp. 143–153, 2006.

[47] A. Slinko, “How large should a coalition be to manipulate an election?,” *Mathematical Social Sciences*, vol. 47, no. 3, pp. 289–293, 2004.

[48] L. Xia and V. Conitzer, “Generalized scoring rules and the frequency of coalitional manipulability,” in *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, pp. 109–118, ACM, 2008.

[49] D. G. Saari, “Susceptibility to manipulation,” *Public Choice*, vol. 64, pp. 21–41, 1990.

[50] B. Peleg, “A note on manipulability of large voting schemes,” *Theory and Decision*, vol. 11, no. 4, pp. 401–412, 1979.

[51] E. Friedgut, G. Kalai, and N. Nisan, "Elections can be manipulated often," *Annual IEEE Symposium on Foundations of Computer Science*, pp. 243–249, 2008.

[52] S. Dobzinski and A. Procaccia, "Frequent manipulability of elections: The case of two voters," in *Internet and Network Economics* (C. Papadimitriou and S. Zhang, eds.), vol. 5385 of *Lecture Notes in Computer Science*, pp. 653–664, Springer Berlin / Heidelberg, 2008.

[53] M. Isaksson, G. Kindler, and E. Mossel, "The geometry of manipulation: A quantitative proof of the Gibbard-Satterthwaite theorem," in *Foundations of Computer Science (FOCS 2010)*, pp. 319–328, IEEE, 2010.

[54] E. Elkind and H. Lipmaa, "Small coalitions cannot manipulate voting," in *Financial Cryptography and Data Security* (A. Patrick and M. Yung, eds.), vol. 3570 of *Lecture Notes in Computer Science*, pp. 578–578, Springer Berlin / Heidelberg, 2005.

[55] E. Elkind and H. Lipmaa, "Hybrid voting protocols and hardness of manipulation," in *Algorithms and Computation* (X. Deng and D.-Z. Du, eds.), vol. 3827 of *Lecture Notes in Computer Science*, pp. 206–215, Springer Berlin / Heidelberg, 2005.

[56] P. Faliszewski, E. Hemaspaandra, and H. Schnoor, "Copeland voting: ties matter," in *Proceedings of AAMAS 2008*, (Richland, SC), pp. 983–990, International Foundation for Autonomous Agents and Multiagent Systems, 2008.

[57] V. Conitzer, T. Sandholm, and J. Lang, "When are elections with few candidates hard to manipulate?," *Journal of the ACM (JACM)*, vol. 54, no. 3, 2007.

[58] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra, "Using complexity to protect elections," *Communications of the ACM*, vol. 53, no. 11, pp. 74–82, 2010.

[59] L. Xia and V. Conitzer, "A sufficient condition for voting rules to be frequently manipulable," in *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, pp. 99–108, ACM, 2008.

[60] L. Xia, M. Zuckerman, A. Procaccia, V. Conitzer, and J. Rosenschein, "Complexity of unweighted coalitional manipulation under some common voting rules," in *Proceedings of IJCAI 2009*, vol. 9, pp. 348–353, 2009.

[61] J. Bartholdi, C. Tovey, and M. Trick, "The computational difficulty of manipulating an election," *Social Choice and Welfare*, vol. 6, no. 3, pp. 227–241, 1989.

[62] P. Faliszewski and A. Procaccia, "AI's war on manipulation: Are we winning?," *AI Magazine*, vol. 31, no. 4, 2010.

[63] V. Conitzer and T. Sandholm, "Nonexistence of voting rules that are usually hard to manipulate," in *Proceedings of the National Conference on Artificial Intelligence, (AAAI 2006)*, vol. 21, pp. 627–634, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006.

[64] C. Menton and P. Singh, “Manipulation can be hard in tractable voting systems even for constant-sized coalitions,” *Arxiv preprint arXiv:1108.4439*, 2011.

[65] A. Slinko, “On asymptotic strategy-proofness of classical social choice rules,” *Theory and Decision*, vol. 52, no. 4, pp. 389–398, 2002.

[66] A. Slinko, “On asymptotic strategy-proofness of the plurality and the run-off rules,” *Social Choice and Welfare*, vol. 19, no. 2, pp. 313–324, 2002.

[67] E. Mossel, A. Procaccia, and M. Racz, “A smooth transition from powerlessness to absolute power,” *Arxiv preprint arXiv:1205.2074*, 2012.

[68] J. Rothe and L. Schend, “Typical-case challenges to complexity shields that are supposed to protect elections against manipulation and control: A survey,” in *International Symposium on Artificial Intelligence and Mathematics*, 2012.

[69] G. Erdélyi, M. Nowak, and J. Rothe, “Sincere-strategy preference-based approval voting broadly resists control,” *Mathematical Foundations of Computer Science 2008*, pp. 311–322, 2008.

[70] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe, “A richer understanding of the complexity of election systems,” *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*, pp. 375–406, 2009.

[71] H. Liu, H. Feng, D. Zhu, and J. Luan, “Parameterized computational complexity of control problems in voting systems,” *Theoretical Computer Science*, vol. 410, no. 27-29, pp. 2746–2753, 2009.

[72] L. Xia, "Computing the margin of victory for various voting rules," in *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, (New York, NY, USA), pp. 982–999, ACM, 2012.

[73] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe, "Llull and Copeland voting computationally resist bribery and constructive control," *Journal of Artificial Intelligence Research*, vol. 35, no. 1, pp. 275–341, 2009.

[74] A. Slinko and S. White, "Is it ever safe to vote strategically?," Tech. Rep. 563, University of Auckland Department of Mathematics Report Series, 2008.

[75] P. Favardin and D. Lepelley, "Some further results on the manipulability of social choice rules," *Social Choice and Welfare*, vol. 26, no. 3, pp. 485–509, 2006.

[76] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2008.

[77] S. Brains and P. Fishburn, "Approval voting," *Birkh/iuser, Boston*, 1983.

[78] R. Farquharson, *Theory of voting*. New Haven, Conn.: Yale University Press, 1969.

[79] V. Conitzer and T. Sandholm, "New complexity results about Nash equilibria," *Games and Economic Behavior*, vol. 63, no. 2, pp. 621–641, 2008.

[80] H. Aziz and M. Paterson, "Computing voting power in easy weighted voting games," *Arxiv preprint arXiv:0811.2497*, Jan 2008.

[81] M. Zuckerman, P. Faliszewski, V. Conitzer, and J. Rosenschein, “An NTU cooperative game theoretic view of manipulating elections,” *Internet and Network Economics*, pp. 363–374, 2011.

[82] M. Zuckerman, P. Faliszewski, Y. Bachrach, and E. Elkind, “Manipulating the quota in weighted voting games,” *Artificial Intelligence*, vol. 180, pp. 1–19, 2012.

[83] V. Conitzer and T. Sandholm, “Complexity of determining nonemptiness of the core,” *In Proceedings of the 18th International Joint Conference on Artificial Inteligence(IJCAI-03)(Acapulco, Mexico)*, pp. 613–618, 2003.

[84] R. Myerson and R. Weber, “A theory of voting equilibria,” *The American Political Science Review*, vol. 87, pp. 102–114, Mar. 1993.

[85] M. Fey, “Stability and coordination in Duverger’s law: A formal model of preelection polls and strategic voting,” *American Political Science Review*, vol. 91, no. 1, pp. 135–147, 1997.

[86] D. P. Myatt, “On the theory of strategic voting,” *Review of Economic Studies*, vol. 74, no. 1, pp. 255–281, 2007.

[87] A. Dhillon and B. Lockwood, “When are plurality rule voting games dominance-solvable?,” *Games and Economic Behavior*, vol. 46, no. 1, pp. 55–75, 2004.

[88] Y. Desmedt and E. Elkind, “Equilibria of plurality voting with abstentions,” in *Proceedings of the 11th ACM conference on Electronic commerce*, EC ’10, (New York, NY, USA), pp. 347–356, ACM, 2010.

[89] F. D. Sinopoli, “Sophisticated voting and equilibrium refinements under plurality rule,” *Social Choice and Welfare*, vol. 17, no. 4, pp. 655–672, 2000.

[90] C. Papadimitriou and T. Roughgarden, “Computing correlated equilibria in multi-player games,” *Journal of the ACM (JACM)*, vol. 55, no. 3, p. 14, 2008.

[91] J. Y. Halpern and R. Pass, “Iterated regret minimization: A new solution concept,” *Games and Economic Behavior*, vol. 74, no. 1, pp. 184 – 207, 2012.

[92] G. Christodoulou and E. Koutsoupias, “On the price of anarchy and stability of correlated equilibria of linear congestion games,” *Lecture Notes in Computer Science*, vol. 3669, pp. 59–70, Jan 2005.

[93] T. Roughgarden, “Intrinsic robustness of the price of anarchy,” in *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC ’09, (New York, NY, USA), pp. 513–522, ACM, 2009.

[94] J. Mitlöhner, D. Eckert, and C. Klamler, “Simulating the effects of misperception on the manipulability of voting rules,” in *Proceedings of the 1st International Workshop on Computational Social Choice (COMSOC 2006)*, pp. 345–355, 2006.

[95] E. Clough, “Strategic voting under conditions of uncertainty: A reevaluation of Duverger’s law,” *British Journal of Political Science*, vol. 37, no. 02, pp. 313–332, 2007.

[96] P. Bodo, “The effects of noise and manipulation on the accuracy of collective decision rules,” *Algorithmic Decision Theory*, pp. 62–73, 2009.

[97] V. Conitzer, T. Walsh, and L. Xia, “Dominating manipulations in voting with partial information,” in *Proceedings of 25th AAAI Conference on Artificial Intelligence*, pp. 638–643, 2011.

[98] R. Myerson, “Comparison of scoring rules in Poisson voting games,” *Journal of Economic Theory*, vol. 103, no. 1, pp. 219–251, 2002.

[99] M. Messner and M. K. Polborn, “Strong and coalition-proof political equilibria under plurality and runoff rule,” *International Journal of Game Theory*, vol. 35, pp. 287–314, 2007.

[100] D. Fudenberg and D. K. Levine, *The theory of learning in games*. MIT Press, May 1998.

[101] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings, “Convergence to equilibria in plurality voting,” *Proc. of 24th Conference on Artificial Intelligence (AAAI-10)*, July 2010.

[102] P. C. Fishburn, *Decision and value theory*. Wiley, Jan. 1964.

[103] O. Lev and J. Rosenschein, “Convergence of iterative voting,” in *Proceedings of AAMAS 2012*, 2012.

[104] W. H. Riker, “The two-party system and Duverger’s law: An essay on the history of political science,” *The American Political Science Review*, vol. 76, no. 4, pp. 753–766, 1982.

[105] I. McAllister and D. T. Studlar, “Bandwagon, underdog, or projection? opinion polls and electoral choice in Britain, 1979-1987,” *The Journal of Politics*, vol. 53, no. 3, pp. 720–741, 1991.

[106] A. Reijngoud and U. Endriss, “Voter response to iterated poll information,” in *Proceedings of AAMAS 2012*, 2012.

[107] M. Messner and M. Polborn, “Miscounts, Duverger’s law and Duverger’s hypothesis.” Available at: http://works.bepress.com/polborn/25, 2011.

[108] T. R. Palfrey, “A mathematical proof of Duverger’s law,” in *Models of Strategic Choice in Politics* (P. C. Ordeshook, ed.), pp. 69–91, University of Michigan Press, 1989.

[109] R. B. Myerson, “Comparison of scoring rules in Poisson voting games,” *Journal of Economic Theory*, vol. 103, pp. 219–251, Mar. 2002.

[110] D. P. Myatt and S. D. Fisher, “Everything is uncertain and uncertainty is everything: Strategic voting in simple plurality elections,” Economics Series Working Papers 115, University of Oxford, Department of Economics, 2002.

[111] P. Fishburn, “Axioms for lexicographic preferences,” *The Review of Economic Studies*, vol. 42, no. 3, pp. 415–419, 1975.

[112] E. Ianovski, L. Yu, E. Elkind, and M. C. Wilson, “The complexity of safe manipulation under scoring rules,” in *Proceedings of IJCAI 2011*, pp. 246–251, 2011.

[113] S. Obraztsova and E. Elkind, “Optimal manipulation of voting rules,” *accpeted in AAMAS 2012*, 2012.

[114] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe, "The shield that never was: Societies with single-peaked preferences are more open to manipulation and control," *Information and Computation*, vol. 209, no. 2, pp. 89–107, 2011.

[115] E. Elkind and G. Erdélyi, "Manipulation under voting rule uncertainty," in *Proceedings of AAMAS 2012*, AAMAS '12, (Richland, SC), pp. 627–634, International Foundation for Autonomous Agents and Multiagent Systems, 2012.

[116] L. Xia, V. Conitzer, and J. Lang, "Strategic sequential voting in multi-issue domains and multiple-election paradoxes," in *Proceedings of the 12th ACM conference on Electronic commerce*, EC '11, (New York, NY, USA), pp. 179–188, ACM, 2011.

[117] L. Xia and V. Conitzer, "Stackelberg voting games: Computational aspects and paradoxes," in *Proceedings of AAAI 2010*, vol. 10, pp. 921–926, 2010.

[118] U. Endriss, U. Grandi, and D. Porello, "Complexity of judgment aggregation: safety of the agenda," in *Proceedings of AAMAS 2010*, AAMAS '10, (Richland, SC), pp. 359–366, International Foundation for Autonomous Agents and Multiagent Systems, 2010.

[119] S. Konieczny and R. Pino Prez, "Logic based merging," *Journal of Philosophical Logic*, vol. 40, pp. 239–270, 2011. 10.1007/s10992-011-9175-5.

[120] T. Walsh, "Complexity issues in preference elicitation and manipulation," *ISAIM 2008*, Jan 2008.