

Manipulative Voting Dynamics

Manipulative Voting Dynamics

By

Neelam Gohar

Cambridge
Scholars
Publishing



Manipulative Voting Dynamics

By Neelam Gohar

This book first published 2017

Cambridge Scholars Publishing

Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

Copyright © 2017 by Neelam Gohar

All rights for this book reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the copyright owner.

ISBN (10): 1-4438-9879-1

ISBN (13): 978-1-4438-9879-9

Dedicated to my family specially my parents and my grandfather who have always stood by me and supported me throughout my life. They have been a constant source of love, concern, support and strength all these years. I warmly appreciate their generosity and understanding.

CONTENTS

Abstract	xi
Acknowledgments	xiii
List of Figures	xv
Chapter One.....	1
Introduction	
1.1 Background.....	1
1.1.1 Manipulative Dynamics	5
1.1.2 Tactical Voting Dynamics.....	8
1.2 Related Work	10
1.3 Problem Statement	13
1.3.1 Contribution and Comparison with Previous Work	13
1.3.2 Significance and Importance of the Problem	15
1.3.3 Specific Research Questions	16
1.4 Structure of Book.....	16
Chapter Two.....	17
Preliminaries	
2.1 Notation and Assumptions	17
2.2 Definitions	19
2.2.1 Manipulations.....	23
2.2.1.1 Types of Moves	24
2.2.1.2 Types of Manipulations	25
2.2.1.3 Weights Settings.....	27
2.2.2 Existence of Potential Functions and Pure Nash Equilibria ..	29
2.3 Summary	31
Chapter Three	33
Tactical Voting Dynamics	
3.1 Tactical Voting	33
3.1.1 Process Termination for Plurality Rule	34
3.1.2 Process Termination for other Positional Scoring Rules	42
3.1.2.1 Borda	43
3.1.2.2 Veto and k-approval Voting Rule.....	44

3.2 Weighted Votes.....	44
3.2.1 Plurality Rule.....	45
3.2.2 Borda.....	46
3.3 Conclusions.....	46
Chapter Four.....	47
Manipulation Dynamics	
4.1 Increased Support Manipulative Dynamics	
with Weighted Votes.....	47
4.1.1 A Few Examples of Manipulative Dynamics with Increased	
Support for the Winning Candidate at Each State.....	48
4.1.2 Upper Bound for General Weight Setting.....	53
4.1.3 Bound for a Small Number of Voters.....	54
4.1.3.1 Upper Bound for Bounded Real Weight Setting.....	64
4.1.4 Upper Bound when the Smallest Weight is $\varepsilon < 1$	74
4.1.5 An Upper Bound under Bounded Integer Weight Setting.....	80
4.1.6 Efficient Process.....	80
4.2 Other Voting Rules like Copeland.....	83
4.2.1 Process Termination.....	83
4.2.2 A Few Examples of Manipulative Dynamics	
with Copeland Voting Scheme.....	85
4.3 Decreased Support Manipulative Dynamics.....	90
4.3.1 How Long is the Sequence of Moves?.....	93
4.4 Conclusions.....	95
Chapter Five.....	97
Manipulative Voting Dynamics	
5.1 Mixture of Different Moves.....	97
5.1.1 Combination of Move Types that can lead to Cycles.....	99
5.2 Bounds in Terms of the Number of Distinct Weights.....	108
5.2.1 Manipulation Dynamics with Un-weighted Voters.....	108
5.3 Conclusions.....	110
Chapter Six.....	111
Cycles in Manipulation Dynamics	
6.1 Termination with Tie-breaking Rule.....	111
6.1.1 Veto Rule.....	112
6.1.2 Borda Rule.....	114
6.1.3 k-Majority Rule or k-Approval Voting Rule.....	114
6.1.4 Copeland's Rule.....	115
6.1.5 Bucklin Scheme.....	117

6.1.6 Plurality with Run-off	118
6.2 Process Termination when in Initial Settings, True and Declared Preferences of Voters are the same	119
6.2.1 Borda Rule.....	120
6.2.2 k-Approval Voting Rule.....	121
6.2.3 Copeland's Rule	122
6.2.4 Bucklin Scheme.....	123
6.2.5 Veto Rule	124
6.3 Conclusions.....	125
Chapter Seven.....	127
Summary	
7.1 Summary of Major Findings.....	127
7.2 Implications of the Findings	129
7.3 Suggestions for Further Research	129
Endnotes	131
Bibliography	133

ABSTRACT

In Artificial Intelligence (AI), multi-agent decision problems are of central importance. Such problems may arise when independent agents aggregate their heterogeneous preference orders among all alternatives and the result of this aggregation can be a single alternative, corresponding to the group's collective decision, or a complete aggregate ranking of all the alternatives. Voting is a general method for aggregating the preferences of multiple agents. An important technical issue that arises is the manipulation of voting schemes: a voter may be able to make the outcome what they find most favorable to himself (with respect to their own preferences) by reporting their preferences incorrectly. Unfortunately, the Gibbard-Satterthwaites theorem shows that no reasonable voting rule is completely immune to manipulation, recent literature has focused on making voting schemes computationally hard to manipulate. In contrast to most prior work, Meir et al. [40] have studied this phenomenon as a dynamic process in which voters may repeatedly alter their reported preferences until either no further manipulations are available, or else the system goes into a cycle. We develop this line of enquiry further, showing how potential functions are useful for showing convergence in a more general setting. We focus on the dynamics of weighted plurality voting under sequences made up of various types of manipulation by the voters. In cases where we have exponential bounds on the length of sequences, we identify conditions under which upper bounds can be improved. In accordance with the Nash equilibrium for plurality voting rules, we use the lexicographic tie-breaking rule that selects the winner according to a fixed priority ordering of the candidates. We also study convergence to pure Nash equilibria in plurality voting games in un-weighted settings. We are mainly concerned with polynomial bounds on the length of manipulation sequences, which depend on the types of manipulation allowed. We also consider other positional scoring rules like Borda, Veto, k-approval voting, and non-positional scoring rules like the Copeland and Bucklin voting systems.

ACKNOWLEDGMENTS

My thanks and appreciation go to Professor Paul W. Goldberg whose encouragement and guidance from start to finish have enabled me to develop an understanding of the subject. His insightful comments and constructive criticisms at different stages of my research were thought provoking and they helped me focus my ideas.

I am also thankful to my advisor Piotr Krysta for his encouragement. I must acknowledge as well many friends and colleagues for their encouragement and emotional support.

Most importantly, none of this would have been possible without the love and patience of my family to whom I would like to express my heart-felt gratitude.

Lastly, I offer my regards and blessings to all of those who supported me in any respect during the completion of this book.

Some of the results from this book are included in the proceedings of ICAI 2013.

Neelam Gohar

LIST OF FIGURES

4.1 The heaviest voter.....	56
4.2 The heaviest voter.....	57
4.3 The heaviest voter.....	58
4.4 The second heaviest voter moves	60

CHAPTER ONE

INTRODUCTION

This introductory chapter contains the following sections: Section 1.1 presents the background knowledge, some relevant recent work and a short overview of results. In Section 1.2 we summarize the related work. Section 1.3 gives a brief problem statement which also describes the contributions and significance of the problem. Section 1.4 outlines the structure of the book's chapters.

1.1 Background

One of the newer areas explored in artificial intelligence is that of multi-agent systems. This involves the analysis of interactions between multiple agents, where each has its own personal objectives. For example, each router on the internet might be an agent, and when a packet is forwarded from source to destination, each router prefers to do as little work as possible. Another example would be dividing processes between processors.

One of the actively growing subareas explored in multi-agent systems is computational social choice theory that provides a theoretical foundation for preference aggregation and collective decision-making in multi-agent domains. Computational social choice is concerned with the application of techniques developed in computer science, such as complexity analysis or algorithm design, to the study of social choice mechanisms, such as voting. It seeks to import concepts from social choice theory into AI and computing. For example, social welfare orderings developed to analyze the quality of resource allocations in human society can be applied equally well to problems in multi-agent systems or network design.

People often have to reach a joint decision even though they have conflicting preferences over the alternatives. A joint decision can be reached by an informal negotiating process or by a carefully specified

protocol. Over the course of the past decade or so, computer scientists have also become deeply involved in this study. Social choice theory investigates many kinds of multi-person decision-making problems. Multi-person decision-making problems are important, frequently encountered processes and many real-world problems involve multiple decision-makers.

Within computer science, there are a number of settings where a decision must be made based on the conflicting preferences of multiple parties. For example, determining whose job gets to run first on a machine, whose network traffic is routed along a particular link, or what advertisement is shown next to a page of search results. The paradigms of computer science give a different and useful perspective on some of the classic problems in economics and related disciplines.

For example, various results in economics prove the existence of an equilibrium but do not provide an efficient method for reaching such an equilibrium. Also greater computing power and better algorithms, have made it possible to run computationally demanding protocols that lead to much better outcomes. Preference aggregation has been extensively studied in social choice theory and voting is the most general preference aggregation scheme.

A natural and very general approach for deciding among multiple alternatives is to vote for them. Voting is one of the most popular ways of reaching common decisions. The study of elections is a showcase area where interests come from computer science specialists in theory, systems, and AI and such other fields as economics, business, operations research, and political science. Social choice theory deals with voting scenarios, in which a set of individuals must select an outcome from a set of alternatives. In the general theory of voting, agents can do more than vote for a single alternative, usually each individual ranks the possible alternatives and a voting rule selects the winning alternative based on the voters' preferences. A voting rule takes as its input a collection of votes, and as the output returns the winning alternative. For example, a simple rule known as the Plurality rule chooses the alternative that is ranked first the most often. In this case, the agents do not really need to give a full ranking, it suffices to indicate one's most preferred alternative, so each voter is in fact just voting for a single alternative.

Voting is a well-studied method of preference aggregation, in terms of its theoretical properties, as well as its computational aspects [11, 54]; various practical, implemented applications that use voting exist [18, 32, 35]. Voting is an essential element of mechanism design (how the privately known preferences of many people can be aggregated toward a social choice) for multi-agent systems, and applications built on such systems, which include ad hoc networks, virtual organizations, and a crucial aspect of decision support tools implementing online deliberative assemblies.

S. Ghosh et al. [32] present the architecture and implementation status of an agent-based movie recommender system. In particular, how the agent stores and uses user preferences to find recommendations that are likely to be useful to the user. They have adapted methods developed in the voting theory literature to find compromises between possibly disparate preference as voting is a well understood mechanism for reaching consensus. T. Haynes et al. [35] highlighted the usage of user preferences in an automated meeting scheduling system (software that automates and shares the information processing tasks of associated human users). In this modern world of processes and agents, it is not just people whose preferences must be aggregated but the preferences of computational agents must as well. In both artificial intelligence and system communities, a great array of issues has been proposed as appropriate to approach via voting systems. These issues range from spam detection to web search engines to planning in multi-agent systems and much more, e.g. [19, 20, 23, 51].

Recent work in the AI literature has studied the properties of voting schemes for performing preference aggregation [11, 23, 54]. A social choice function is a function that takes lists of people's ranked preferences and produces a single alternative (the "winner" of the election). A good social choice function represents the "will of the people". Rather than just choosing a winning alternative, most of the voting rules can also be used to find an aggregate ranking of all the alternatives. For example, we can sort the alternatives by their Borda score, thereby deciding not only on the "best" alternative but also on the second-best, and so on. There are numerous applications of this that are relevant to computer scientists, for example, one can pose the same query to multiple search engines and combine the resulting rankings of pages into an aggregate ranking.

Researchers in social choice theory have studied extensively the properties of various families of voting rules, but have typically neglected

computational issues. Sincere voting assumes that voters always choose their most preferred candidates and/or parties. It has been argued in both the formal and empirical literature, however, that voters may not always vote for their most preferred candidates. Sincere voting is voting in accordance with one's true preferences over alternatives. While strategic voting is voting over assumed outcomes, in which a voter uses their skills to determine an action that secures what is, in their view, the best possible outcome. This is the trade-off a rational voter faces in an election. They must balance their relative preference for the different candidates against the relative likelihood of influencing the outcome of the election [7]. However, in voting, one of the major technical issues is the manipulation of voting schemes. Elections are endangered not only by the organizers but also by the voters (manipulation), who might be tempted to vote strategically (that is, not according to their true preferences) to obtain their preferred outcome.

Manipulation in voting is considered to be any scenario in which a voter reveals false preferences in order to improve (with respect to their own preferences) the outcome of the election. A manipulative vote leads to successful manipulation if it changes the election outcome to one preferred by that particular voter. Since voters are considered rational agents, who want to maximize their own utility, their best strategy may be to manipulate an election if this will gain them a higher utility. This has various negative consequences: not only do voters spend valuable computational resources determining which lie to employ, but worse, the outcome may not be one that reflects the social good. The Gibbard-Satterthwaite result [33, 58] states that any non-dictatorial voting scheme is vulnerable to manipulation, that is, there will always be a preference profile in which at least one of the individuals has an incentive not to elicit their true preferences. Gibbard-Satterthwaite, Gardenfors, and other such theorems open doors to strategic voting, which makes voting a richer phenomenon. In order to achieve some standard of non-manipulability in voting schemes, in all the previous work the complexity of the manipulation has been considered where one could try to avoid manipulation by using protocols where determining a beneficial manipulation is hard; for a survey, see [25].

Complexity offers a powerful tool to frustrate manipulators who seek to manipulate or control election outcomes. The motivation for studying complexity issues comes from the Gibbard-Satterthwaite theorem which shows that every reasonable election system can be manipulated [33, 58].

So better design of election systems cannot prevent manipulation. Computational complexity can serve as a barrier to dishonest behavior by the voters, and Bartholdi et al. [4] proposed classifying voting rules according to how difficult it is to manipulate them. They argued that well-known voting rules such as Plurality, Borda, Copeland, and Maximin are easy to manipulate. Since then, the computational complexity of manipulation under various voting rules received considerable attention in the literature, both from the theoretical and from the experimental perspective (see, [61, 63]) and the recent surveys [13, 24, 60]. The complexity of the manipulation problem for a single voter is quite well understood and this problem is efficiently solvable for most common voting rules with the notable exception of the single transferable vote (STV) [4, 5]. The more recent work has focused on coalitional manipulation, i.e. manipulation by multiple, possibly weighted voters.

We have not dealt with computational complexity issues here, we are considering bounds on the length of sequences of manipulations that voters can perform. Despite the basic manipulability of reasonable voting systems, it would still be desirable to find ways to reach a stable result, which no agent will be able to change. One possibility is the convergence of myopic improvement dynamics, where strategic voters change their votes step by step in order to get a better outcome. A voting profile is in equilibrium when no voter can secure the election of their preferred candidate. This iterative voting is used, in the real world, in various situations, such as elimination decisions in various “reality shows.” The study of dynamics in strategic voting is very interesting and highly relevant to the multi-agent systems, as it helps to tackle multi-agent decision-making problems, where autonomous agents (that may be distant, self-interested and/or unknown to each other) have to choose a joint plan of action or allocate resources or goods. We work with different types of moves that lead to successful manipulation.

1.1.1 Manipulative Dynamics

Meir et al. [40] have studied this phenomenon as a dynamic process in which voters may repeatedly alter their reported preferences until either no further manipulations are available, or else the system goes into a cycle. Here we develop this line of enquiry further, showing how potential functions are useful for showing convergence in a more general setting. We focus on Plurality voting with weighted voters, and obtain bounds on the lengths of sequences of manipulations, that depend on which types of

manipulation are allowed. We analyze the sequences of changes of votes that may result from various voters performing manipulations and we bound the length of sequences of votes with the help of potential functions. Potential functions are valuable for proving the existence of pure Nash equilibria and the convergence of best response dynamics. Even Dar et al. [21] introduced the idea of using a potential function to measure closeness to a balanced allocation, and used it to show convergence for sequences of randomly selected “best response” moves in a load-balancing setting in which tasks may have variable weights, and resources may have variable capacities. We study convergence to pure Nash equilibria in Plurality voting games. In such a game, the voters strategically choose a candidate to vote for, and the winner is determined by the Plurality rule. A voting profile is in equilibrium, when no voter can change their vote to ensure that a more preferable candidate gets elected. In our model, we assume the elementary stepwise system (ESS), i.e. at every state only one voter is allowed to move. Thus, a voter switches their support to another candidate in response to the moves of other voters so that a sequence of moves occurs. This sequence may stop at a steady state where no voter wishes to switch, or may continue indefinitely. This steady state is called the Nash equilibrium. The concept of Nash equilibria has become an important mathematical tool in analyzing the behavior of selfish users in non-cooperative systems [50], i.e. games where players act in an independent and selfish way. Such iterative games reach an equilibrium point from either an arbitrary or a truthful initial state. We focus more on weighted voting setting, where voters may have different weights in elections. The topic of convergence to stable outcomes in strategic voting settings is interesting to artificial intelligence. We are mainly concerned with polynomial bounds on the length of manipulation sequences.

For our model, we consider an election with m alternatives, and with n voters, each of whom has a total ordering of the alternatives. A system is comprised of a finite number of states and transitions occur from state to state when voters change their mind and support an alternative candidate. Every state is mapped into a real value by the potential function and transitions cause the potential to increase or decrease. States can be defined as the profiles of “declared preferences” of voters. A transition is a manipulation move by a single voter. We focus on the Plurality voting rule because Plurality has been shown to be particularly susceptible to manipulation, both in practice and theory [29, 57]. We consider other voting rules as well. We assume that voters have knowledge regarding the candidates currently supported by the other voters in the case of Plurality

voting. For other positional scoring rules, voters have knowledge regarding the total scores of all candidate at a state. Complete information is not needed in such a set-up. Voters manipulate according to their true preferences. Voters change their vote (make manipulation) after observing the current state and outcome. If voters have their true preferences then a manipulator changes their preference list in favor of a less preferred candidate and makes him a new winner if they do not like the current winner and it results in a better winner (for that voter) than the current winner. In the case where voters declared preferences that are different to their true preferences and the outcome is not favorable, then they change preference list in favor of their most preferred candidate who can win. If a voter cannot affect the outcome at some state, they simply keep their current preference list. This process of manipulation proceeds in turns, where a single voter changes preference list at each step/turn. Voters take turns modifying their votes; these manipulations are according to the way in which they affect the outcome of the election. The process ends when no voter has objections and the outcome is set by the last state.

In manipulation dynamics, voters change their mind to make a “manipulative vote” that changes the outcome of the election. We are considering bounds on the length of sequences of manipulations that can take place. We also consider voting rules with lexicographic tie-breaking rule that depends strictly on linear preference orders to choose a winner in case of ties. In most of our results we use a weighted voting system. A weighted voting system is the one in which the preferences of some voters carry more weight than the preferences of other voters. Some of our results have dependence on the voters’ weights. We have results for different weight settings. We used weighted votes as the introduction of weights generalizes the usability of voting schemes.

If voters are allowed to vote simultaneously, then this iterative process may never converge to an equilibrium [40]. For this reason, in our model, only one voter is allowed to move at every state. The system is modeled as a sequence of steps and in each step one voter switches from one candidate to another. We establish bounds on the length of sequences of manipulations that voters can perform. We consider these with respect to the different types of moves that lead to successful manipulation. We do not concern ourselves with the impact of manipulation on social welfare; we treat manipulation as an “occupational hazard” and ask the question: In a system where manipulation may occur, when can we guarantee that the voters will end up satisfied with their (possibly manipulative) votes, in the

context of the votes offered by the others? Put another way, we posit that in various real-world situations, it may be better to reach a poor decision than no decision at all. We can regard the voting system as a game in which each voter has, as pure strategies, the set of all votes they may make. (In Plurality voting, a vote is just the choice of a single “preferred” candidate.) Each voter has a type, consisting of a ranking of the candidates that represents their real preferences. We ask whether pure Nash equilibria exist for any set of voter types, and more importantly whether such an equilibrium can be reached by the players via a sequence of myopic changes of vote. This can be regarded as a very simplistic model of a negotiation process among the voters, and we would like to ensure that it does not end in deadlock.

Our main issue is the proof of termination, and the bounds on the length of sequences of manipulations that can take place. We are interested in bounds on the number of possible steps that are purely in terms of number of candidates m and number of voters n (and independent of the total size of the weight which can be quite large). An important property of the voting rules discussed in [4] is that they may produce multiple winners. In real-life settings, when an election ends in a tie, it is not uncommon to choose the winner using a tie-breaking rule that is non-lexicographic in nature. When an election under a particular voting rule ends in a tie, we use a lexicographic tie-breaking rule that uses a fixed linear order on candidates to break ties.

1.1.2 Tactical Voting Dynamics

Sincere voting is voting in accordance with one’s true preferences rather than the alternatives, while strategic voting is where a voter assumes likely outcomes and uses skill to determine an action that secures what they believe to be the best possible outcome. Strategic voting under the Plurality rule refers to a voter deserting a more preferred candidate with a poor chance of winning in favor of a less preferred candidate with a better chance of winning [26]. The logic of tactical/strategic voting, of course, is that of Duverger’s law, which states that the supporters of a small party would not “waste” their votes by voting for their most preferred party (candidate) because they do not have a chance of winning under a Plurality system with single member districts. Instead, they vote for the major party that is most acceptable to them, and that has a chance of winning.

Let us suppose a voter believes that their most preferred candidate has little chance of competing for the lead in the election. Voting for such a candidate may be a “waste.” The voter may decide to switch their vote to the one among the leading candidates she most prefers in order to make her vote “pivotal” in determining a more preferable outcome. This is the trade-off a rational voter faces in an election. Strategic voting is an important component of Duverger’s law: If voters are rational, they end up voting for one of the two leading candidates [6].

Another voting dynamic we consider is that of tactical voting in which a voter changes their mind to make a tactical vote according to the mind-changing rule (defined later). The purpose of making a tactical vote is to increase the score of a preferred candidate which may or may not lead to changing an election outcome. The set of all voters’ declared preferences is summarized in the concept of a state. A transition occurs from current state to a new state when a voter changes their mind and chooses a different candidate to support (under Plurality). In a state of a system, each voter determines whether they can improve (with regard to their own true preferences) the outcome by altering their own vote while assuming that all other votes remain the same. A mind-changing rule is as follows: A voter considers all alternative candidates that they ranked higher than the current winner of the state. They can then change their support to the alternative who currently has the most votes, breaking away from their own preferred candidate. With this mind-changing, a transition occurs and the system moves from the current state into a new one. At each iteration, the state of the system associates each voter with the candidate currently supported by that voter under Plurality rule. Tactical voting is different to manipulation dynamics because it simply raises the votes of the leading candidate that the voter prefers. In this kind of voting, a voter, instead of wasting vote by voting for his most preferred candidate who does not have a chance of winning, thinks it better to vote for a candidate who is more acceptable and has a chance of winning and, thus, raise the score of that alternative. We analyze the sequences of votes that may result from various voters voting tactically in both weighted and un-weighted settings.

Our results. We focus on Plurality voting with weighted voters, in which each voter reports a single preferred candidate. A voter’s weight is fixed throughout. The score of a candidate is the total weight of voters who support that candidate, hence the winning candidate is the one with highest score, and we assume a standard lexicographic tie-breaking rule in which candidates have a given total order on them that determines the winner if

two or more of them have maximal score. Meir et al. [40] also consider this tie-breaking rule, and compares it with a randomized one.

We investigate the rate of convergence, i.e. the number of steps of manipulation that may be needed to reach a pure Nash equilibrium. We focus on types of manipulation where there are no cycles in the state/transition graph, where convergence is guaranteed, and we analyze bounds on the number of steps required. The rate of convergence will be expressed as a function of the number of voters, n ; the number of candidates, m ; the ratio w_{\max} of maximum to minimum weights; and the number of distinct weights, K . Guaranteed convergence may also depend on types of manipulation available; a classification is given in Chapter 2.

We identify combinations of types of moves that are able to lead to cycles of manipulation moves. We consider combinations of move types where convergence is guaranteed, and exhibits various potential functions to obtain upper bounds on the number of manipulation steps possible. Alternative types of moves seem to require alternative potential functions, and we give upper bounds as expressions in terms of the parameters, n , m , and K .

1.2 Related Work

Meir et al. [40] studied the convergence of pure strategy Nash equilibria in Plurality games. They showed that myopic best response dynamics may cycle, even when they start from a truthful voting profile, for both deterministic and randomized tie-breaking schemes. Our work extends that of Meir et al. [40] in that we consider weighted voters as well. The notion of voting dynamics as well as the convergence of voting games, particularly Plurality, exist in previous research. For deterministic tie-breaking schemes, we demonstrate that if one excludes certain deviations then an improvement path is guaranteed. There is also a number of very recent papers, apart from Meir et al. [40], that analyze strategic behavior in voting using tools of non-cooperative game theory [14, 62]. Desmedt and Elkind [14] consider the setting where all voters are strategic, where an election can be viewed as a game, and the election outcomes correspond to the Nash equilibria of this game. They analyze two variants of Plurality voting, namely, simultaneous voting, where all voters submit their votes at the same time, and sequential voting, where the voters express their preferences one by one. Sequential voting always has an equilibrium in pure strategies. They take the approach suggested by Farquharson [26] and view manipulation as an unavoidable attribute of an electoral system with

rational voters.

Xia and Conitzer [62] consider a voting process in which voters vote one after another as an extensive-form game. They study equilibria of sequential voting for a number of voting rules (including Plurality). However, they use a deterministic tie-breaking rule. Feddersen et al. [27] study a Plurality voting game in which voters are strategically rational and search for different equilibria choices. However, in order to reach an equilibrium, they limit the possible preference choices to single-peaked preferences. Also the model assumes that both voters and candidates possess complete information and that voters use only pure strategies. There have been several studies applying the Nash equilibrium to dynamics process, particularly in the allocation of public goods. Much of the work is summarized in [38]. They characterize all Nash equilibria, using different approaches under the restriction that preferences are single-peaked preferences as in Feddersen et al. [27]. Hinich et al. [37] change single-peaked preferences to a specific probabilistic model of voters over a Euclidean space of candidates, where individuals vote randomly according to probability functions based on their preferences, and where the candidates maximize expected votes. Another relevant work is by Messner and Polborn [41]. They focus on the existence and uniqueness of strong equilibria in Plurality games. Strong equilibrium is a weaker concept, still stronger than the Nash equilibrium. No coalitional manipulation can get an incentive by making a coordinated diversion in a case of strong equilibrium. The same approach is used in Dhillon and Lockwood [15] by considering dominant strategies in Plurality voting. To reach an equilibrium, de Trenquayle [59] uses a specific voting rule and Euclidean preferences, proving that under the Euclidean preferences the majority rule converges and that there is a unique equilibrium. All the above papers assume that voters have some knowledge of the other voters' preferences.

Another model was suggested by Myerson and Weber [44] who found Nash equilibrium for positional scoring rules like approval, Borda, and Plurality. They assumed voters have some knowledge about the preferences of other voters but not every election converges. Iterative voting with Plurality was examined by Chopra et al. [9], limiting voters' information about others voters' preferences, e.g. when voters are myopic, and also assuming that voters have sufficient information about all voters. The focus of this study is on the role played by the state of knowledge of the agents. A related dynamical model was considered by Airiau and Endriss [2] who examined the conditions needed to achieve an equilibrium

in iterative games using the Plurality voting rule. In the model a group of agents make a sequence of collective decisions on whether to remain in the current state of the system or switch to an alternative state. At each step, a voter is selected at random and may propose a single alternative to the one currently winning the election; a pairwise vote takes place between the current winner and the new alternative. As in [40], cycles may arise; indeed, the ability of the chosen voter to select a single alternative for a pairwise election makes it possible to exhibit cycles which cannot be escaped. An iterative procedure for reaching a solution was also used by [17], but they use money like value among voters/agents. They consider how agents can come to a consensus without needing to reveal full information about their preferences, and without needing to generate alternatives prior to the voting process.

The study of manipulability of various voting rules, i.e. understanding the algorithmic complexity of individual or coalitional manipulation, is an active research area. Much of this work views manipulation as a type of adversarial behavior that needs to be prevented, either by imposing restrictions on voter preferences, or by identifying a voting rule for which manipulation is computationally hard, preferably in the average case rather than in the worst case. Making manipulation difficult to compute is a way followed recently by several authors [4, 5, 10, 11, 12] who address the computational complexity of manipulation for Plurality and other voting rules.

Bartholdi et al. [4] showed how computational complexity protects the integrity of social choice. While many standard voting schemes can be manipulated with only polynomial computational effort, they exhibit a voting rule that efficiently computes winners but is computationally resistant to manipulation. Bartholdi and Orlin [5] showed that Single Transferable Vote (STV) is apparently unique among voting schemes in actual use today in that it is computationally resistant to manipulation. Under STV each voter ranks all the candidates in order of preference. STV tallies votes by reallocating support from weaker candidates to stronger candidates and excess support from elected candidates to remaining contenders. Conitzer et al. [10] asked the question: How many candidates are needed to make elections hard to manipulate? They answered this question for the voting protocols: Plurality, Borda, STV, Copeland, Maximin, regular Cup, and randomized Cup.

The main manipulation question studied in Conitzer et al. [11] is that of coalitional manipulation by weighted voters. They characterize the exact

number of candidates for which manipulation becomes hard for the Plurality, Borda, STV, Copeland, Maximin, Veto, Plurality with run-off etc. They show that for simpler manipulation problems, manipulation cannot be hard with few candidates. Some earlier work showed that a high complexity of manipulation relies on both the number of candidates and the number of voters being unbounded. Conitzer and Sandholm [12] derived hardness results for the more common setting where the number of candidates was small but the number of voters could be large. They show that with complete information about the others' votes, individual manipulation is easy, and coalitional manipulation is easy with un-weighted voters.

1.3 Problem Statement

We study the convergence to pure strategy Nash equilibria in Plurality voting games. We also study other positional scoring rules and some non-positional scoring rules as well. We consider election with m alternatives and with n voters, each of whom has a total ordering of the alternatives. In such a game, the voters strategically choose a candidate to vote for, and the winner is determined by the Plurality/other voting rules. Voters take turns modifying their votes; these manipulations are classified according to the way in which they affect the outcome of the election. We focus on achieving a stable outcome, taking strategic behavior into account. A voting profile is in equilibrium when no voter can change their vote so that their preferred candidate gets elected. We investigate bounds on the number of iterations that can be made for different voting rules. We focus on the weighted voting settings, where voters may have different weights in elections. We consider equi-weighted votes also. An important property of the voting rules is that they may produce multiple winners, i.e. they are, in fact, voting correspondences. When an election ends in a tie, we choose the winner using a tie-breaking rule that is lexicographic in nature.

1.3.1 Contribution and Comparison with Previous Work

Most of the previous work about manipulation dealt with computational complexity issues, where one could try to avoid manipulation by using protocols where determining a beneficial manipulation is hard [11, 25]. The well-known Gibbard-Satterthwaite theorem [33, 58] states that a reasonable voting rule is completely immune to strategic manipulation. This makes the analysis of elections a complicated and challenging task. One approach to understanding voting is the analysis of solution concepts

such as Nash equilibria (NE). Several studies exist in prior research that apply game-theoretic solution concepts to the voting games. But the most recent and relevant work is that of Meir et al. [40] who suggested the framework of voting as a dynamic process in which voters repeatedly change their reported preferences one at a time (if voters are allowed to change their preferences simultaneously, the process will never converge). This iterative process continues until either no further manipulations are available or the system goes into a cycle. In the paper they study different versions of iterative voting, varying tie-breaking rules, weights and policies of voters, and the initial profile. Their results show that, in order to guarantee convergence, it is necessary and sufficient that voters restrict their actions to natural best responses. They also showed that with weighted voters, or when better replies are used, convergence is not guaranteed. Hence, myopic better response dynamics may cycle, even when starting from a truthful voting profile, for both deterministic and randomized tie-breaking schemes. This topic of convergence to stable outcomes in a strategic voting setting is interesting to artificial intelligence. It tackles the fundamental problem of decision-making where agents are considered to be autonomous entities and they have to choose a joint plan of action or allocation of resources. A related dynamical model was considered by Airiau and Endriss [2], in which, at each step, a voter is selected at random and may propose a single alternative to the one currently winning the election; a pairwise vote takes place between the current winner and the new alternative. As in [40] cycles may arise; the ability of the chosen voter to select a single alternative for a pairwise election makes it possible to exhibit cycles which cannot be escaped. We expand this framework further, concerning the dynamics of weighted Plurality voting under sequences made up by various types of manipulations by the voters. We also consider other voting rules apart from Plurality. We use the idea of using a potential function for studying the rate of convergence to equilibria in a more general setting. For lexicographical tie-breaking scheme under different weight settings, we demonstrate that convergence to equilibria can be guaranteed considering different types of moves that lead to successful manipulation. Polynomial bounds are obtained and proofs are based on constructing a potential function with a guaranteed value of increase/decrease at each step. Our results suggest different choices of potential functions can handle different versions of the problem. We also show that a cycle exists if we allow all types of moves, that is the reason we obtain bounds for different subsets of moves where the voting dynamics converge. Our results and the results obtained in [40] provide a relatively complete knowledge of what