Rochester Institute of Technology
Department of Computer Science

Complexity of Control of Borda Count Elections

By
Nathan F. Russell
nfr1228@cs.rit.edu

July 9, 2007

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in Computer Science

----
Edith Hemaspaandra, Chairperson

Christopher Homan, Reader

Piotr Faliszewski, Observer

Contents

CHAPTER 1

Introduction and history of the Borda count

In this thesis, we discuss some existing and new results relating to computational aspects of voting. In particular, we consider, apparently for the first time, the computational complexity of the application of certain types of control to the Borda count voting system. We use control in the formal sense of attempts by an election's administrator to make a specific candidate win or lose by various means. We consider control problems for weighted elections, as well as for unweighted elections with voter preferences input both individually and in succinct representation.

In many circumstances, groups of people need to decide between several alternatives ("candidates"). The group may contain a large number of stakeholders, who may disagree on which alternative(s) are more desirable than others, and very often some or all pairs of alternatives are mutually incompatible. This requires the group to in some manner combine (aggregate) the preferences of the individuals to reach a decision which is (informally speaking) as acceptable as possible to the members as a whole. Often, these preference aggregation systems involve some form of voting. By formalizing the decision method into a voting system, even those whose preferences are less influential in the final result will ideally feel that they were dealt with fairly. While voting is most frequently considered in connection with political elections and committees, preferences may be aggregated by voting in various other contexts as well.

Originally, there was often a single individual (for example, a king or queen) who made decisions for the entire group. This system made complicated methods of social choice unnecessary but did not ensure an ideal result from the perspective of other stakeholders. However, an increasing number of modern societies follow the political philosophy of democracy. This system has many variants, for example presidential and parliamentary democracy, unicameral vs. bicameral legislatures, et cetera. In every case, the governments of these societies utilize voting in what we term "political elections"—to make decisions for the entire group. Particularly, groups often use voting to elect legislators or other leaders, who may then vote among themselves about individual decisions.

In fact, a very strong case can be made that voting is necessary to popular rule. While small groups of a few people may be able to reach an informal consensus on all issues, in larger groups someone (or some small group) must be in charge, and voting may be the only fair way to select that leader in the largest groups (in which everyone cannot directly discuss potential leaders, and try to convince all others to form a consensus). In the smallest democracies every citizen could weigh in directly on every issue, as is done in a few private clubs. However, in larger groups leadership is increasingly complicated and becomes a full-time job.

Generally, it is a goal to have all voters be equal (or be unequal in a precisely defined fashion). In order to do this, groups use voting systems, in which the winner(s) are determined in what could be viewed as an algorithmic manner. We briefly digress to discuss what an election "is." While this may seem trivial, in the course of discussing elections, we can informally introduce some key concepts in the formalization we will use. First, in some fashion (with which we are unconcerned) a set of valid candidates (or, outside political elections, alternatives) is generated. Each voter is given the list of candidates, and reports his or her preferences over these candidates, formalized as an ordered list from highest to lowest (we follow the literature and assume that preferences are transitive, which may not always be the case, but simplifies our arguments). Once the preferences are collected, some voting system (alternatively, "social choice system") is used to combine them into a final result, a set of one or more winners. In the elections we consider, there is assumed to be a single winner except in the special case in which candidates are exactly tied, or in which for some reason no candidate wins.

In political contexts, various voting systems may be used. Probably the most common in leadership elections is plurality voting, in which voters are asked only for their first preference, and the alternative which receives the greatest number of votes wins. Other common systems include the single transferable vote (or instant runoff), which repeatedly eliminates the candidate with the least number of first preferences, and reassigns his or her votes to the voters' next-highest preferences, until only the desired number of winners are left (or, in the case of an election with a single winner, until one alternative has a majority of the votes). Another system is approval voting. In this system, rather than ranking candidates, voters indicate whether or not they approve of each candidate. The winner(s) are the candidate(s) with the highest number of voters approving of them.

This thesis considers a voting system known as the Borda count, which will be described in greater detail later, but essentially involves granting more points to candidates which a voter ranks higher. It also considers all of a voter's preferences, not merely the highest preference as in the case of plurality voting, an artificial gap between approval and non-approval as in approval voting, or some (often small) number of high preferences as in the case of the single transferable vote. This aspect of the Borda count is an advantage in some ways from a philosophical perspective, since all the information which a voter contributes to the election is used. In systems such as plurality voting, the vast majority of it (in a many-candidate election) is discarded, even though it is quite plausible that there is more difference between a voter's opinion of his or her second and last choice of ten candidates than there is between his or her opinion of the first and second. We could imagine a voting system in which voters are allowed to report how "far apart" they rank adjacent preferences, but such a system would tempt voters to vote dishonestly (and might allow them to do so without compromising the positive effects from their honest vote).

In political elections, it may often be the case that voters are not informed about every candidate, or feel relatively neutral about their "middle" preferences. In particular, they may strongly like or dislike only one, or a few, candidates. In these cases, protocols such as plurality or veto voting — in which a point is given to every preference but the last — may be reasonable models. However, it is interesting to note that the fact that the Borda count considers a voter's entire preference order is particularly important in those cases in which the entirety of a voter's preferences can reasonably can be seen as significant. For example, the Borda count was used to combine the rankings by judges in a wine tasting [**Cop06**]. In this setting, we might expect that the judges saw (for example) their third-choice wine as distinct from their fourth choice by just as much as they saw the first choice as better than the second.

Since modern societies rely on voting in so many contexts, some of which involve important matters such as who will be the next political leader or how huge amounts of money will be spent, it is reasonable to assume (and, indeed, indicated by history) that in some cases individuals will attempt to interfere with voting procedures, or to vote dishonestly, in order to gain a personal advantage or to have more influence on the results than they should. If individual voters or some subset of the voters cast votes which do not honestly represent their preferences, we refer to this as manipulation (it is also termed strategic voting in the special case when an individual voter, not colluding with others in complicated fashions, is considered). Alternatively, those responsible for the election (e.g., the chair) may alter the election procedures to change the results, without changing how voters choose to vote, a situation termed control. Finally, it is possible that an alteration of how voters vote is accomplished by bribing the voters. Bribery is a method which we do not discuss at great length, though there is an increasing amount of work on the topic. Particularly, Faliszewski, Hemaspaandra, and Hemaspaandra discuss bribery for many voting systems, including Condorcet and Plurality voting; they also find some general results for scoring protocols, such as the Borda count for fixed numbers of candidates [**FHH06**].

Clearly, from the perspective of society's integrity, preventing interference with elections is an important goal. This can be accomplished, of course, by using election inspectors, security procedures, et cetera, to make it more practically difficult for anyone to interfere (or, if they do so, to escape detection). In practice, such methodologies are very commonly used. However, at best, they enlarge the set of individuals who must be corrupted by anyone who would exert control beyond a certain point, or practice manipulation in a very large group. Alternatively, such strategies may limit the number of changes that can be made (disqualifying ten voters may be easier for a chair to get away with than removing a hundred). We can reasonably conclude, however, that if it were possible a better alternative would be to choose a system which, of and by itself, was difficult or impossible to interfere with.

It would be ideal if we could design election systems so that control and manipulation are impossible to accomplish (what we will later term immunity). However, it has long been known that is not possible for a reasonable election (for relatively broad, common-sense definitions of reasonable, for example one condition is that there be no voter with exclusive influence over the result) to be immune to any type of control or manipulation that might be attempted. This was shown early in the study of voting, with results such as the Gibbard-Satterthwaite Theorem [**Gib73, Sat75**] and Arrow's impossibility theorem [**Arr50**]. Respectively, these theorems state that no reasonable voting system is impossible to manipulate, and that none can have certain properties generally regarded as desirable. Likewise, in the case of control, generally speaking immunity only occurs when it is fundamental to the system under consideration; we discuss several specific examples in Chapter 3, but few if any practical voting systems are immune to every means of control.

Thus, at best we can seek to find a voting system for which control and manipulation are either impossible to accomplish in certain cases or computationally infeasible to determine how to accomplish (for example, we might seek to design an election in which deciding whether control can be exercised at all is an NP-complete problem, a property termed resistance).

A somewhat common objection to the use of complexity theory in evaluating voting systems is that in many cases it becomes more relevant for large, even unbounded, numbers of candidates. While terms such as "voting system" suggest political elections, as we previously touch upon it is also often the case that groups must choose between different opinions and recommendations in other contests. For example, a financial professional may have to decide between thousands of possible stock and bond purchases ("candidates"), each of which may be ranked differently by each of dozens or hundreds of experts and computer programs ("voters"). In these settings, it is reasonable to speak of an "election" of arbitrary numbers of candidates (and perhaps relatively few voters).

Along the same lines, in artificial intelligence applications, we may utilize voting systems to make decisions. For example, Jeong et al. propose a Borda count variant to combine the results of using three features in an image retrieval system, and in their trials achieve better results than those of systems using any one feature [**JKC99**]. Likewise, Richards and Seung propose applying Borda count and Condorcet elections to improve robustness in a certain type of neural networks [**RS04**]. Dwork, Kumar, Naor, and Sivakumar [**DKNS01**] propose the use of voting to combine search results from the Web (with a particular emphasis on combating attempted spam). Many search engines exist, each with a different method of ranking results. Some of the search engines are more or less subject, as might be expected, to influence by various spam techniques (and some allow paid inclusion, which may subvert the algorithmic rankings). The authors suggest applying social choice theory (including voting) to aggregate the results from various search engines, and mention the use of various voting systems, including the Condorcet and Borda systems.

The Borda count has been lauded for having a lower tendency towards the spoiler effect than plurality voting does. In other words, a weak candidate is less likely to lose the election for a similar strong candidate, since the weak candidate will be drawing away high preferences, rather than entire votes as in a plurality election (in other words, a voter who votes for the weak candidate first is likely to also rank the similar strong candidate relatively high, and will still be helping both).

Plurality voting can have problems not only with the spoiler effect, but with vote-splitting between similar strong candidates. For example, in the 2006 Canadian federal elections, the Conservative party gained control with only 36% of the vote [**CBC06**]. The conservatives were (and are) the only strong right-leaning party, whereas there are two, or possibly three, liberal parties, depending on classification, and votes were split between these.

As we will discuss in greater detail later, the Borda count is not only not affected by the spoiler effect, but in a sense exhibits it in reverse. More specifically, in many situations the "best" of a large number of similar candidates is far better off than he or she would be if none of the others ran (essentially, he or she will be ranked second by the voters for the other strong candidate, and the later will be ranked last by the supporters of the group of similar candidates).

In addition, since the Borda count considers a voter's entire preference list, a voter who knows that many other voters will vote for two "strong" candidates is able to benefit by ranking his or her favored one of the strong candidates first and the other last, even if he or she might honestly rank them (e.g.) third and fourth respectively. This allows him or her to have a much greater influence on the race between the two strong candidates than if he or she voted honestly (in the example given, if the election has ten candidates total, the manipulating voter will contribute nine points to the race between the strong candidates, rather than only one).

The Borda count has a long history of use in various settings. Particularly, it is used in parliamentary elections and in nomination of candidates for President in the Pacific island nations of Nauru and Kiribati, respectively [**Rei02**]. It is (perhaps inaccurately) named after Jean-Charles de Borda, who independently developed it in 1770, though similar systems had been proposed by others (including one that was used in the Roman senate [**Str80**], and another proposed by Nicolaus Casanus in the 13th century to elect the Holy Roman Emperor [**ML06**]).

We note in passing that there are other ways to apply complexity theory to elections, aside from considering ease of interference (of whatever type). For example, Conitzer and Sandholm [**CS05**] consider the amount of communication capacity needed to gather information from the voters to determine the results of an election, and point out that attempting to use minimal communication may allow voters to maintain some privacy or to compute only part of their preferences (which may be expensive to fully compute). They find that the Borda count has one of the higher requirements for communication capacity;

perhaps unsurprisingly, plurality voting (which takes into account only the first preference) has the lowest of all.

CHAPTER 2

Definitions

2.1. Elections in General

Formally, following [**BTT92**], we let an election consist of a set V of voters, each identified with a preference order, a strict, complete and transitive order on the set C of candidates.

As mentioned in the introduction, elections can be conducted using any of several voting systems. For example, many political elections in the world use plurality voting, in which voters vote only for a single candidate, their first preference. The candidate who is voted for by the greatest number of voters is the winner. In another system, approval voting, each voter either approves or disapproves of each candidate. The candidate of whom the largest number of voters approve is the winner. We do not go into much detail about specific voting systems where it is not relevant, but it is important to note that dozens of systems exist, and a fair number are in practical use for one purpose or another.

In a Borda count election, which is the focus of the present thesis, each of the candidates is given $\|C\| - n$ points for each voter who ranks him or her in nth place (so, $\|C\| - 1$ points for first, $\|C\| - 2$ for second, and so forth until the candidate a voter ranked last receives no points). In weighted elections, some voters may count more than others (for example, some voters' preferences might count for double points); more precisely, each voter is associated with a weight, and contributes a number of points to each candidate equal to the product of the weight and the number of points that would be given in an unweighted election. Weighted elections achieve real-world significance in certain non-political settings. For example, when shareholders vote for a corporate board, their votes may be weighted in proportion to the number of shares owned. Likewise, in certain committees, delegates may have votes proportionate to the number of people they represent.

In the system used in Nauru, the number of points assigned for consecutive preferences decreases hyperbolically rather than linearly; 1 point is assigned for the first preference, 1/2 for the second, 1/3 for the third, and so forth [**Rei02**]. This embodies the plausible assumption, in the case of political elections, that voters may strongly prefer their first one or two choices, but not be familiar with, or feel more or less neutral towards, others.

As a basic example of a Borda count election utilizing the Borda rule, consider three candidates, a, b, and c, and four voters, with respective preference orders $a > b > c$, $a > c > b$, $c > b > a$, and $c > a > b$. Candidate a is ranked first twice, second once, and third once, and so is given $2 + 2 + 1 + 0 = 5$ points. Candidate b, ranked second twice and third twice, is given two points, and candidate c earns five. At this point, candidates a and c are tied for winner.

In formalizing the problem, the input can be thought of as a list of voter's preference orders much like that above. This means that the input length is proportionate to the product of the numbers of voters and candidates.

The Borda count may be formulated in other ways. For example, in an alternative formulation used by among others Straffin [**Str80**], the first candidate in a preference order may be given n points and the last one. It is immediate that changing from our method to this one takes an equal number of points from each candidate, not changing the result, and so the two are equivalent: in either, the candidate with the most total points becomes the winner, and the identity of the winner will be unchanged..

In the case of a weighted election, each voter is associated with a positive integer "weight," assumed to be expressed in binary or some other reasonable representation. Before the computation of the final result, the points awarded by each voter are multiplied by that voter's weight (hence, a voter with power n has the same power as n voters each with power one). Clearly, an unweighted election is equivalent to a weighted election in which all weights are equal to one.

In some unweighted elections, and particularly where the number of voters is much greater than the number of candidates, many voters may have identical preference orders (in fact, when the disparity is great, the pigeonhole principle may require them to do so). In this case, it may make sense to think of the voters as represented in a succinct representation. Succinct representations are discussed by [**FHH06**] among others, though it is unclear where the concept originated. In such a representation, the preference order is stated once, along with the number of voters sharing that order. We assume that the number of voters is expressed in binary or some other reasonable representation. For example, we could have an input such as $3567xa > b > c; 521xb > c > a; 293xc > a > b$ to represent 3567 voters, each with preference order $a > b > c$, and likewise for the other entries. In the case where there is a small number of distinct preferences compared to the number of voters, this may render the control problem asymptotically more difficult (potentially turning vulnerability into resistance in some cases), since a short segment of the succinct input may represent an exponential number of identical voters.

Control problems for succinct representations are fundamentally different from those for weighted elections, and different methods must be used in evaluating them. In the case of a succinct representation, it is possible to add or delete some identical voters, but not others. Conversely, in the case of weighted voters, we assume that each individual voter must either cast a vote with their full weight, or not vote (not making this assumption, of course, renders much of the difference unimportant—in addition, it is difficult to imagine a practical mechanism for control that would remain undetected while allowing only part of a vote to be cast and counted).

Typically, non-succinct representations are discussed; there is a good practical reason for this. In a political election, or even in many AI applications, ballots are submitted individually. This means that even reading the ballots will take time proportionate to the product of the number of voters and the length of each ballot. In few practical cases would it be possible to determine voter preferences any faster than that, much less exponentially faster.

That said, succinct representations are potentially relevant to real-world political elections, which often have relatively few candidates, but potentially tens or hundreds of millions of voters, many of whom will share a preference order due to the pigeonhole principle alone. Often, at least portions of the preference order of voters will be deducible from such data as party memberships. In such cases, it is possible that an election chair who is able to determine if control is possible in time polynomial in a standard input length would find that it was not feasible to do so on a real computer. On the other hand, control problems requiring time which is polynomial in the length of a succinct representation might remain easy. That said, the control algorithms we present are generally relatively efficient, and in some cases they could likely be made even more efficient by allowing approximate solutions. In any case, the succinct representation case remains of theoretical interest, especially since it expresses the notion that in control of very large elections it may not be feasible to consider voters individually.

A property of the Borda count is that it encourages compromise. In, for example, a four-candidate race, a candidate who is loved by half the voters, and loathed by the other half, will lose to a candidate who is the second choice of almost everyone. In fact, as Straffin points out

[**Str80**], the (unweighted) Borda count effectively ranks candidates by their average ranking among all voters, rather than by the number of voters who rank them first. Whether favoring moderate candidates is a desirable property is a philosophical question which we do not attempt to address.

Consistent with the literature, we use the term control to mean interference with the election from outside (for example, by a chair's decision to disqualify specific candidates), and manipulation to mean interference by voters (for example, tactical voting, in which voters do not express their true preferences, and thereby achieve a more favorable result from the perspective of those preferences than if they did) [**BTT92, HHR07a, CLS07**]. A third possible method, bribery, has recently been considered [**FHH06**].

Applying complexity theory to the difficulty of controlling elections is an approach which was first introduced by Bartholdi, Tovey, and Trick [**BTT92**]. They consider attempts by an election's chair to control the plurality and Condorcet voting systems by a variety of possible means. The use of the term "destructive" for manipulation aimed at causing a candidate to lose was introduced by Conitzer, Lang, and Sandholm [**CLS03**]. Destructive control was then considered by Hemaspaandra, Hemaspaandra, and Rothe, who also analyzed approval voting as well as the influence of tie-breaking systems for certain specific means of control [**HHR07a**]. We follow the prior use of the terms constructive and destructive control; in the former, the chair attempts to ensure that a specified candidate is the unique winner, while in the later he or she attempts to ensure that a specified candidate is not the unique winner, in both cases after breaking any ties that may occur in sub-elections; more formal definitions follow.

Following prior work, we consider an election to consist of a set of voters V, each of whom has a preference order on the candidate set C. Following [**BTT92**] in particular, we assume that all preference orders are transitive, complete, and strict (less formally, each voter must unambiguously rank every candidate, without tied preferences, and the order must be transitive, which any ordinary "list" of preferences would inevitably be). In addition, we consider only the order of preferences, not "how much" voters prefer candidates to each other. There may be situations where a voter despises his or her third-ranked of three candidates, but only slightly prefers his or her first to his or her second (or, conversely, greatly prefers the first to both the second and third, but does not feel a strong distinction between those two). No common variant of the Borda count takes this sort of variability of preferences into account.

Briefly, we consider control by means of adding or deleting either voters or candidates, as well as by partition of the candidate or voter set into sub-elections. These means of control are defined formally in Section 2 of this chapter. Where sub-elections are considered, occasionally a sub-election will end in a tie. In this case, there are two possible ways to proceed, as discussed by [**HHR07a**]. In model ties eliminate, all tied candidates are eliminated from the election; in model ties promote, all are advanced to the next round.

In order to simplify proofs, we introduce the notion of a score function:

DEFINITION 2.1.1. For all voters v and all candidates c, we define $score(c,\ v)$ as the number of points which c gains from v's vote (in other words, the difference between the total number of candidates and the ordinal at which v ranks c). In the case of a weighted election, $score(c,\ v)$ is multiplied by v's weight.

In some cases it is also useful to speak of preference distance, informally how much a voter prefers one candidate to another. Thanks to the definition of the score function, this definition takes the weight of the voter into account.

DEFINITION 2.1.2. For all voters v and candidates c and d, we define $prefdist(c,\ d,\ v)$ to be $score(c,\ v)$ - $score(d,\ v)$, that is, the number of places, adjusted for weight, by which v ranks c above d.

2.2. Means of control to be considered

These formal definitions are taken, with minor clarifications and rephrasing, from the definitions used in [**BTT92**] and [**HHR07a**]. One change we make to some of these definitions is to ask whether a change of no more than k voters or candidates suffices to establish control. This formulation more naturally incorporates the notion that in some cases the chair may wish to make a relatively small change to reduce the likelihood of apprehension. Additionally, it could allow future work to consider whether control methods in which a change affects a small number of voters or candidates (for example, a change to no more than the square root of the original total) maintain the same computational properties. If we wish to consider all possible changes, k can be set to a sufficiently high value to allow this, and we can use these definitions as given.

In the case of weighted elections, we assume that k is the total weight of voters with which we are allowed to interfere (and not the number of voters).

Again, we speak of unique winners. Many if not all of our results hold whether or not winners are required to be unique, and we specify when a result holds for winners allowed to be non-unique.

2.2.1. Alteration of the Voter Set. First, we consider altering the set of voters. Perhaps most obviously, we could persuade some voters who are not presently registered to do so, and to vote.

DEFINITION 2.2.1. Control by Adding Voters

Given: A set of candidates C and a distinguished candidate $c \in C$, a set V of registered voters, an additional set W of yet unregistered voters (both V and W having preferences over C), and a positive integer $k < ||W||$.

Question (constructive): Is there a set of k or fewer voters from W (in the weighted case, a set of voters with total weight at most k) whose entry into the election would make c the unique winner?

Question (destructive): Is there a set of k or fewer voters from W (in the weighted case, a set of voters with total weight at most k) whose entry into the election would prevent c from being the unique winner?

We can also, of course, speak of deleting voters from those initially registered. By deleting, we mean preventing them from voting (or, at least, having their votes counted) by any means.

DEFINITION 2.2.2. Control by Deleting Voters

Given: A set of candidates C, a distinguished candidate $c \in C$, a set V of voters, and a positive integer $k < ||V||$.

Question (constructive): Is there a set of k or fewer voters in V (in the weighted case, a set of voters with total weight at most k) whose removal from the election would make c the unique winner?

Question (destructive): Is there a set of k or fewer voters in V (in the weighted case, a set of voters with total weight at most k) whose removal from the election would prevent c from being the unique winner?

Finally, we might, without adding or deleting voters, partition the voters into sub-elections. Some (possibly empty) set of candidates will win the sub-elections. We can then hold an election among those candidates, with all voters participating, and declare the winners of that election to win.

DEFINITION 2.2.3. Control by Partition of Voters

Given: A set of candidates C, a distinguished candidate $c \in C$, and a set V of voters.

Question (constructive): Is there a partition of V into V_1 and V_2 such that c is the unique winner in the hierarchical two-stage election in which the survivors of $(C,\ V_1)$ and $(C,\ V_2)$ run against each other with voter set V?

Question (destructive): Is there a partition of V into V_1 and V_2 such that c is not the unique winner in the hierarchical two-stage election in which the survivors of $(C,\ V_1)$ and $(C,\ V_2)$ run against each other with voter set V?

Of course, in considering partition of voters, we have to deal with handling ties among winners of sub-elections. Following [**HHR07a**], we consider two common-sense alternatives, which were informally mentioned earlier. Firstly, we could decide that if there is a tie in a sub-election, all tied candidates are eliminated and any winner comes from the other sub-election.

Definition 2.2.4. In model Ties-Eliminate (TE), all tied candidates in a sub-election are eliminated.

Of course, we may not want to eliminate candidates who were one winner of a sub-election.

Definition 2.2.5. In model Ties-Promote (TP), all tied candidates in a sub-election proceed to the next round.

2.2.2. Alteration of the Candidate Set. Now we arrive at the problems relating to altering the candidate set, rather than the voter set. One possibility is to add new candidates, perhaps in the hopes of displacing the preferred candidate upwards in the voters' reported preference lists. We begin with adding candidates who are not presently running.

Definition 2.2.6. Control by Adding Candidates

Given: A set C of qualified candidates and a distinguished candidate $c \in C$, a set B of presently non-participating candidates, a set V of voters with preferences over $C \bigcup B$, and a positive integer $k < \|B\|$ of allowed additional candidates.

Question (constructive): Is there a choice of no more than k candidates from B whose entry into the election would make c the unique winner?

Question (destructive): Is there a choice of no more than k candidates from B whose entry into the election would prevent c from being the unique winner?

Alternatively, we may instead delete candidates whose presence may be preventing the desired result. Following [**HHR07a**], we do not allow deleting the distinguished candidate in destructive control, since this would make the control problem trivial.

Definition 2.2.7. Control by Deleting Candidates

Given: A set C of candidates, a distinguished candidate $c \in C$, a set V of voters, and a positive integer $k < ||C||$ of allowed deletions.

Question (constructive): Are there k or fewer candidates in C whose disqualification would make c the unique winner?

Question (destructive): Are there k or fewer candidates in C, excluding c, whose disqualification would prevent c from being the unique winner?

We could also consider having two elections, in which one set of candidates first compete for the right to meet a second set of candidates.

Definition 2.2.8. Control by Partition of Candidates

Given: A set C of candidates, a distinguished candidate $c \in C$, and a set V of voters.

Question (constructive): Is there a partition of C into C_1 and C_2 such that c is the unique winner in the sequential two-stage election in which the winners in the sub-election (C_1, V) who survive the tie-handling rule move forward to face the candidates in C_2 (with voter set V)?

Question (destructive): Is there a partition of C into C_1 and C_2 such that c is not the unique winner in the sequential two-stage election in which the winners in the sub-election (C_1, V) who survive the tie-handling rule move forward to face the candidates in C_2 (with voter set V)?

Rather than an election in which the first set of candidates must run for the privilege of competing on equal footing with the second set, it may make more sense to have two sub-elections the winners of which then compete against one another.

Definition 2.2.9. Control by Runoff Partition of Candidates

Given: A set C of candidates, a distinguished candidate $c \in C$, and a set V of voters.

Question (constructive): Is there a partition of C into C_1 and C_2 such that c is the unique winner in the election in which those candidates surviving (with respect to the tie-handling rule) sub-elections on voter set V and respective candidate sets C_1 and C_2 have a runoff with voter set V?

Question (destructive): Is there a partition of C into C_1 and C_2 such that c is not the unique winner in the election in which those candidates surviving (with respect to the tie-handling rule) sub-elections $(C_1,\ V)$ and $(C_2,\ V)$ have a runoff with voter set V?

CHAPTER 3

Prior work

3.1. Manipulation

Manipulation, in contrast to control, is the practice of voters (either individually or in cooperation) submitting votes which do not honestly reflect their preferences in order to achieve an election outcome which better fits those preferences than if they had voted honestly. Ideally, we would want all elections to be impossible to manipulate as well as to control. At least in the case of manipulation, as alluded to earlier, this is not possible. The Gibbard-Satterthwaite Theorem states, informally, that in an election of three or more candidates, any voting system has one of several properties [**Gib73, Sat75**]. The system must either be dictatorial, in other words have some voter who is solely in control of the outcome, have some candidate who can never win (thus, effectively have a smaller number of candidates) or be susceptible to manipulation (that is, have a scenario in which a voter with total knowledge of other voters' preferences will have an incentive to vote dishonestly). This is paraphrased from the formulation used by [**CLS07**]; several alternate formulations exist.

The Gibbard-Satterthwaite Theorem can be generalized as well. For example, Aswal, Chatterji and Sen show that any non-manipulable social choice function defined on certain "linked domains" must be dictatorial [**ACS03**].

This inevitable potential for strategic voting, in addition to allowing voters to have more influence than they "should" in view of their respective weights, can pose a practical problem in that, if many voters vote dishonestly, there may be outcomes that few if any voters actually would desire. For example, in the realistic case of incomplete information, we can imagine a plurality election in which a candidate believed to be strong is actually the first preference of relatively few voters, but is elected thanks to strategic voters who vote for him or her rather than their first choice, who they believe cannot win. In the case of the Borda count, even with complete information, there is a strong incentive to "bury" a candidate whom a voter dislikes and expects to win.

In general, the Borda count is widely known to have poor properties against manipulation, and particularly against tactical voting; not only is manipulation possible, but in many cases determining a manipulation strategy is an easy task [**Ser02**]. As a simple example, consider a four-candidate election in which a dishonest voter prefers $a > b > c > d$. If the voter in question knows that b and c are liked by a large proportion of the other voters, and as such are the only candidates with a real possibility of winning, he or she will want to have the maximum possible impact on the race between them. In this scenario, he or she can vote $b > a > d > c$, and thereby give b an edge of three points in the b vs. c race, instead of only one point if he or she had voted honestly. Interestingly, however, Barbie, Puppe, and Tasnádi find that the Borda count for preference sets that satisfy certain criteria is strategy-proof. [**BPT06**]

It is fair to note that similar concerns apply to plurality voting, in which not voting for a favored weak candidate may be advantageous—but in a plurality election, clearly, a voter can help only one candidate, whereas a manipulative voter in a Borda count election may be able to over-rank several preferred candidates as well as burying several strong candidates below weak candidates who he or she may dislike even more.

Of course, even though manipulation is in a certain sense inevitable, we can still attempt to make it as difficult as possible (for example, by making it an NP-complete problem to determine whether manipulation by a specific group of voters is possible; as in the case of control, this property is termed "resistance").

Conitzer, Sandholm, and Lang consider the manipulation of few-candidate elections [**CLS07**]. In addition to the Borda count, they consider the plurality, Condorcet, and single transferable vote systems, as well as others which we do not discuss. Conitzer, Sandholm, and Lang additionally consider the number of candidates needed before an election gains resistance in those cases in which it is possible at all.

Conitzer, Sandholm, and Lang make an assumption (somewhat irrelevant when considering control of elections with honest voters, but significant when considering manipulation) that voters must submit their entire preference sets up front, and may not change them once submitted, even when an election has multiple rounds. They consider manipulation both in the case in which the manipulating voters have complete information about the preferences of non-manipulators, and in that in which they do not. They also consider other variations. For example, voters may be weighted or unweighted, and may be trying to manipulate constructively or destructively, either as individuals or as a coalition. One interesting result is that when manipulation is hard for coalitions with complete information, it is also hard for individuals with incomplete information.

In some cases manipulation is easy. For example, constructive manipulation of a plurality election requires us to compute whether the preferred candidate wins in the event that every manipulator votes for that candidate. If this is the case, then the manipulators can clearly do no better if they vote in any other way. A positional scoring protocol is one in which two conditions are true. Firstly, candidates receive, for each integer n no greater than the number of candidates, a fixed number of points for being ranked nth by any voter. Secondly, the sequence of point values is non-increasing (being ranked early is no worse than being ranked late). The Borda count, when restricted to a fixed number of candidates, is a positional scoring protocol. Conitzer, Sandholm, and Lang find that, with the exception of plurality voting, all such protocols for weighted elections are resistant to constructive manipulation by coalitions even with only three candidates.

Further, [**CLS07**] find that for all voting systems in which candidates are scored based on the votes, and in which the score function is monotone (that is, for a voter to change his or her vote to rank a candidate higher does not hurt that candidate), a group which includes the Borda count, destructive manipulation of weighted elections by coalitions of voters is in P. This is true because, if there is a manipulation to make the distinguished candidate lose, there must be one in which every manipulator ranks the same candidate first and ranks the distinguished candidate last; such a manipulation helps the first-ranked candidate as much as possible against the distinguished candidate. It is easy to determine the result of manipulation if every manipulator uses votes of that form, and try doing so for each possible first-ranked candidate.

Campbell and Kelly consider a class of "social welfare functions" which generate social choice rules (i.e., voting systems) which are completely immune to manipulation when certain conditions hold regarding the initial election [**CK06**]. The authors begin with what they term the "folk theorem:" in Condorcet elections, when a Condorcet winner exists, for any given non-winning candidate, each manipulator either prefers the winner to that candidate, or not. If a manipulator prefers the current winner, he or she will vote honestly and rank the current winner earlier, thereby not changing the result to favor the candidate under consideration. If he or she prefers the candidate under consideration, he or she will rank that candidate first even when voting honestly. As such, in either case, manipulation by that voter will not make some non-winning candidate the unique Condorcet winner, again assuming that such a winner existed initially (note that the Gibbard-Satterthwaite Theorem does not set any conditions on the outcome of the pre-manipulation election). Campbell and Kelly extend this result to certain other systems.

The Borda count is vulnerable to destructive manipulation, that is, the manipulation problem is polynomial, as shown by [**CLS07**]. However, as previously noted, vulnerability applies to other common voting systems (and particularly to the commonly used plurality system, which in fact is even more trivial to manipulate). In fact, again, the Gibbard-Satterthwaite Theorem implies that any reasonable voting system can be manipulated (though it may be resistant in some or all cases). We need to consider other factors in deciding which system to use in which situations.

3.2. Control

The intent of this thesis is to determine the computational complexity of control of Borda count voting systems (weighted and otherwise) through a variety of potential means of control, including some of those considered for other voting systems in [**BTT92**] and [**HHR07a**] (namely, control by addition and removal of candidates and voters, partition of candidates, partition of voters, and runoff partition of candidates, though we do not find results for all these cases). Following the reasoning of those papers, we aim to determine whether the Borda Count is immune to exercising constructive and destructive control by various means of control (it is fundamentally impossible to establish control), is resistant to control (determining whether it is possible to establish control is NP-complete), or is vulnerable (determining whether it is possible to establish control is in P).

Consistent with the Gibbard-Satterthwaite Theorem, and despite the specific advantage of avoiding the spoiler effect, the Borda count is vulnerable to certain other types of practical manipulation and control. For example, Serais argues that a losing candidate may be able to win by encouraging a similar candidate slightly weaker than herself to run (the opposite of the outcome in plurality voting) [**Ser02**]. This practice is referred to as cloning of candidates.

As an example of the risk from cloning, consider an election with two candidates, a and b, and voters who are in favor of the two in equal numbers (or equivalently, two voters with preference order $a > b$ and $b > a$). Clearly, the candidates are tied. Now let us introduce a new candidate c, who is a clone of a, in other words ranked similarly by the voters, so that all voters either prefer b to both a and c, or rank b after both. Further assume that c is less known than a, and so is ranked behind him or her by all voters.

In the modified election, clearly, we will have half the voters with preference order $a > c > b$ and the other half with order $b > a > c$. Here a clearly wins, and this result holds even if a small number of voters have preferences inconsistent with the described properties. On the other hand, if those properties hold for all voters, any two-candidate election in which even one voter initially favors a can be controlled by adding a sufficient number of clones of a to give a a large number of points from that voter. Serais considers the situations in which the Borda count may be controlled via cloning in much more detail [**Ser02**].

It appears that no work has been done specifically on computational aspects of control of the Borda count (although there has been work on manipulation and bribery). In particular, our concern will be with the asymptotic difficulty of control of elections, even for an election administrator with very great power (and knowledge of the preferences of all voters). This contrasts with practical control, in which the concern may be whether an election with a small number of candidates can be controlled without discovery (though our notion of making at most k changes covers this as well). We also consider elections in which voters are weighted, or in which unweighted voters are expressed in succinct representation. In the remainder of this section, all elections are unweighted unless specified otherwise.

Bartholdi, Tovey, and Trick begin by making assumptions that they state are conservative for real elections [**BTT92**]. For example, they assume that an election's chair is able to know each voter's preferences (for all candidates, presently eligible or not), and that the voters will vote "sincerely"—that is, will not use manipulation strategies such as tactical voting, individually or otherwise (we make the same assumptions). Additionally, they simplify their analysis by not considering tie-handling rules. Section 2 of their paper is devoted to summarizing the field of computational complexity, and particularly NP theory, however one point of interest is that they introduce certain terminology relating to the computational difficulty of establishing control, specifically the notions of immunity, susceptibility (non-immunity), resistance, and vulnerability. As noted by Hemaspaandra, Hemaspaandra, and Rothe [**HHR07a**], for those systems which we discuss in depth and all those discussed by [**BTT92**], the control problems for which resistance holds are in fact NP-complete. This holds because it is easy to determine the winner of these elections, and a nondeterministic control algorithm can try each possible alteration to the election and check whether the desired outcome is achieved.

Bartholdi, Tovey, and Trick go on to specify the means of control they will consider, which are among the constructive cases we define in Chapter 2, Sections 2 and 3, namely addition and deletion of voters, addition and deletion of candidates, and partition and runoff partition of candidates. In every case, they limit themselves to what others term constructive control—control in which the chair attempts to make a specific "distinguished" non-winning candidate the unique winner. In all cases, preferences are assumed to be constant - the set of voters, presently registered or otherwise, is finite, and all voters have a preference for all candidates, presently running or not, and will vote honestly. This differs from the situation in which manipulation is considered, which requires its own set of assumptions (for example, the assumption that all voter preferences are specified up front and cannot be changed).

Bartholdi, Tovey, and Trick consider two voting systems. Plurality voting, which is used in most elections in the United States and several other countries, gives each voter one vote, which he or she is assumed to cast for the most preferred candidate. The winner is simply the candidate who receives more votes than any other. The second system, Condorcet voting, requires the winner to be preferred to every other candidate by a majority of voters. This of course leads to a potential pitfall in the system—in many cases there will not be a winner at all, the so-called Condorcet paradox. As a simple example, if there are 3 candidates a, b, and c, and 3 voters with preference orders $a > b > c$, $b > c > a$, and $c > a > b$, then it is easy to verify that a is preferred to c only once, b is preferred to a only once, and c is preferred to b only once, so no candidate can win this Condorcet election. We do not go into great detail regarding results specific to Condorcet voting. The reason is that this system is essentially unlike both the Borda count and plurality voting. As such, results related to the Condorcet system often use methods not relevant to other systems. The Condorcet system depends on pairwise comparisons by individual voters between all candidates, not on exactly how the voters regard each candidate. Conversely, the Borda count, like plurality voting, depends on a ranking of all candidates, compiled from the voters' preferences (both are scoring protocols when restricted to a fixed number of candidates; the Condorcet system is not). These systems also always have a (possibly not unique) winner. This contrasts with the Condorcet system, which has no winners in many possible elections (a simple example being the election earlier in this paragraph).

Bartholdi, Tovey, and Trick show that plurality voting is computationally resistant to constructive control by adding candidates. This is done by reduction from the NP-complete hitting set problem, which informally requires determining whether, given a finite set and a set of its subsets, it is possible to choose a subset of less than a specified size such that it shares at least one element with every given subset. Briefly, the reduction is done by creating a correspondence between voters and set members. This is done in such a way that the preferred candidate wins ground against the original winner for each subset a member of which is chosen, but loses ground against another opponent every time any set item is chosen. By setting the voter preferences properly, it is possible to construct an election in which control is possible if and only if the appropriate hitting set exists.

Likewise, [**BTT92**] find that plurality voting is computationally resistant to control by deleting candidates. This is shown by reduction from exact cover by 3-sets, an NP-complete problem which involves determining whether, given a set B whose size is divisible by 3 and a collection S of 3-element subsets of B, a set S' of subsets of S can be chosen such that each element of B appears in exactly one element of S'. The authors design an election such that, if manipulation by deleting candidates is possible, an exact cover exists. This involves an intricate balancing of votes won and lost by various candidates. The preferred candidate gains exactly one vote for each member added to the covering set, gaining enough votes to win if and only if each subset is covered. If an element of B is chosen more than once, a non-preferred candidate will gain enough votes to win, and the same will happen if there is no cover.

[**BTT92**] find plurality voting to be resistant to control by partition of candidates as well as runoff partition of candidates. The runoff partition case is done by reducing to the problem of control by deleting candidates. While intricate, the transformation is somewhat intuitive, since the favored candidate can only win if the earlier round in which he or she does not take part places a set of candidates against whom he or she can win in the runoff. The case of partition of candidates is proved by adding filler candidates to the previously constructed runoff partition case, and showing that it is still necessary to solve the original deletion problem in order to exercise control over the election by partition of candidates.

Plurality voting is computationally vulnerable to control by adding and by deleting voters. Essentially, this is true because it considers only each voter's favorite candidate. This means that it is easy to see which voters should be added to help the preferred candidate, or deleted to hurt those candidates who are presently beating him or her.

Plurality voting is computationally vulnerable to control by partition of voters. If any candidate will lose to the preferred candidate c, there is a candidate c' who is the strongest of those who do so. Now, if there is no candidate who defeats c and c' combined by a margin of than two votes, it is possible to assign the voters for c and c' to separate partitions, and then to partition the remaining voters so that no candidate beats the current winners in their respective partitions. This allows c to face and defeat c' in the final election.

One significant result mentioned in [**BTT92**] is the Weak Axiom of Revealed Preference (WARP), which is true of voting systems in which the winner of an election is always the winner of every subset of candidates containing him or her. In voting systems which satisfy WARP, constructive control by adding candidates is impossible—any original candidate who will win an augmented election must have been winning already. This is true, for example, of Condorcet elections, since the winner must be preferred by a majority to every candidate in the set, and so to every candidate in his or her subset. It is not, however, true of plurality voting. It is quite possible that, in a plurality election, the winner does not win among a subset of candidates because voters who would be spread between candidates outside the subset instead vote for a less-preferred candidate, giving that candidate enough votes to defeat the winner of the augmented election.

Hemaspaandra, Hemaspaandra, and Rothe introduce a notion of a system being computationally vulnerable to a specific method of control, which means that not only does the polynomial algorithm decide whether control is possible, but it tells the chair exactly how to go about establishing control whenever it is [**HHR07a**]. There are few if any practical examples in which a voting system in which recognizing winners is in P is vulnerable, but not computationally vulnerable. In any event, [**HHR07a**] prove computational vulnerability in every case in which they prove vulnerability.

In passing, we note that voting systems do exist in which determining the winner is an NP-hard problem. These include the Dodgson, Kemeny, and Young systems, all essentially variants on Condorcet voting. The first two systems were shown to have NP-hard winner-determination problems by [**BTT89**], and the second by [**RSV03**]. Clearly, such systems, while often interesting, are impractical for real elections.

Hemaspaandra, Hemaspaandra, and Rothe additionally introduce, as previously mentioned, the notion of destructive control. In some real-world situations, the chair may not have a strong preference for who wins, as long as a specific despised candidate loses. This may free the chair to not consider interactions between multiple candidates, making the control problem easier. In fact, when determining winners is in P, it is impossible for a system to be resistant to destructive control unless it is also resistant to constructive control. This is true because determining whether destructive control is possible is the same as determining whether constructive control is possible for some non-distinguished candidate (who will thus become a winner, making the distinguished candidate not the unique winner). As a special case, destructive control by deleting candidates does not allow the chair to delete the distinguished candidate (thus preventing a trivial, uninteresting method of control).

Hemaspaandra, Hemaspaandra, and Rothe consider three voting systems: plurality, approval, and Condorcet voting. We will not further discuss Condorcet voting, simply because plurality voting more closely resembles the Borda count, and it is reasonable to assume that that results regarding plurality voting will be more often directly useful in analyzing the Borda case. In addition, we do not go into much detail regarding the constructive results for the case of approval voting, which is first considered by [**HHR07a**]; the arguments used to show these results are often intricate and not easily applicable to the Borda count, and in many cases immunity trivially holds, which is likewise not the case for the Borda count.

Hemaspaandra, Hemaspaandra, and Rothe begin by showing susceptibility in many cases. Often, this is proved by finding a small election in which a change allowed by the means of control in question changes the winner. We will not discuss these proofs in general; when we mention resistance or vulnerability, it implies susceptibility. Immunity, generally speaking, is found in cases in which the voting system in question has fundamental properties preventing it from being possible to control. For example, Condorcet voting is immune to destructive control by deleting candidates for the simple reason that it satisfies unique-WARP (the version of WARP considering unique winners), and so if the despised candidate wins the original election, he or she will win an election among any subset containing him or her. As previously mentioned, deleting the despised candidate is not allowed as it would permit trivial control by deleting candidates, and so any subset of remaining candidates must contain the distinguished candidate; due to unique-WARP, the distinguished candidate will still win among this subset.

[**HHR07a**] find approval voting to be immune to destructive control by deletion of candidates, partition of candidates, and runoff partition of candidates. This is true because it fulfills the version of WARP (as defined earlier) for unique winners, and so the chosen candidate must be winner in any subset containing him or her—and thus must win the final election, as well as the sub-elections involved in the partition methods. Interestingly, no immunity results are found for plurality voting at all (which supports our intuition that none exist for the similar Borda count).

Approval voting is computationally vulnerable to destructive control by partition of voters regardless of tie-handling rules. [**HHR07a**] consider two tie-breaking models, Ties-Eliminate and Ties-Promote, as defined previously. As in many such cases, if the distinguished candidate c already is not the unique winner, control is trivially possible by leaving one partition empty. If there are exactly one or two candidates, it is impossible to change the winner by partition of voters, so control is possible or impossible depending on whether the distinguished candidate presently is the unique winner. In the case of model Ties-Eliminate, they go on to consider whether some two non-distinguished candidates can be made to each tie or beat c in one partition. This is done by allocating voters who approve of each intended winner to that winner's partition, using arithmetical reasoning to ensure that this strategy can succeed for the pair of candidates in question (and continuing to search pairs if not). If no such pair exists, control is impossible. In model Ties-Promote, similar reasoning is used, except that they ensure that the the intended winners strictly defeat c in each partition. This is important since at this point in the algorithm, c is known to be the original winner, and since relative standings in approval voting do not depend on the candidate set chosen, c would still be the final winner if not prevented from winning in either partition (this is also the case in model Ties-Eliminate, but there of course a tied result suffices to ensure that c does not reach the final round).

Approval voting is computationally vulnerable to destructive control by deletion of voters. This is shown by reasoning very similar to that for partition of voters; in the case of approval voting, for each potential winner i, it is possible to count the number of voters who approve of c but not of i. If the number of such voters is sufficient, enough can be deleted to allow i to unseat c. In the related case of addition of voters, [**HHR07a**] likewise propose adding enough eligible voters who approve a chosen candidate, but not c. There are other details involved in their argument for both cases, which actually counts voters for each candidate considered as a potential "unseater"; they do not actually loop across all candidates, though such a strategy would succeed.

Plurality voting is computationally vulnerable to constructive as well as destructive control by partition of voters in model Ties-Eliminate. In the constructive case, [**HHR07a**] argue if the distinguished candidate can be made to win, this can be done in several ways. In the case in which a partition contains no voters, clearly the winner remains the winner of the original election, since all voters are in the second partition. In the second case, where the distinguished candidate is nominated by both partitions, he or she is clearly also the overall winner. In the third case, the distinguished candidate wins one sub-election while the other produces no winner due to a tie involving the distinguished candidate and one or more others. Since the distinguished candidate wins one sub-election and ties the other, he or she must already win the original election. If the distinguished candidate is not already the winner, it is only necessary to determine whether one of two other cases can be made to hold. Specifically, either the distinguished candidate can be made to win one of the two sub-elections and then win the runoff, or can be made to win a sub-election, while the other sub-election produces no winner due to a tie between non-distinguished candidates.

In the destructive case under model Ties-Eliminate, plurality voting is again vulnerable. As in all cases of destructive control, when the distinguished candidate c is already not the unique winner, there is a trivial solution. When he or she is the unique winner, and the two next-highest candidates have vote totals which sum to at least that of c, partitions can be arranged so that one of the two ties or beats c in each partition. When this is not the case, c clearly must win in at least one partition, but destructive control may still be possible. Essentially, the method used in this case is to find some candidate who ties or defeats c in a two-candidate plurality election, and who can be made to win the second sub-election by assigning his or her opponents' supporters to the first partition. To be clear, such a candidate may exist even when the distinguished candidate originally wins, because the candidate may gain support in the final runoff by voters whose first choice (possibly the distinguished candidate, possibly not) is not in that runoff. In this case, control is possible if and only if such a candidate exists.

[**HHR07a**] use a single reduction in resistance results for partition of voters in model Ties-Promote as well as other means of control for plurality voting. In considering resistance results, they first establish a rather broad result—if the winner of an election can be determined in polynomial time (as is true in the Borda count), and a specific control problem is NP-hard, the system cannot be immune to control by that means unless P=NP. This is true because if immunity holds, the decision problem is in P (it is easy to state that control is impossible for any given case, if it is never possible). When the problem is already known to be NP-hard, this of course implies that P=NP.

In the case of plurality voting, [**HHR07a**] use one reduction from the Hitting Set problem for several cases in each of constructive and destructive control. In the case of destructive control, the reduction involves ensuring that the distinguished candidate loses ground to another candidate for each set member added, but gains ground against a third candidate when a set member is added. Given this construction, relatively short arguments show that plurality elections are resistant to destructive control by adding candidates, deleting candidates, partition of candidates in both models Ties-Eliminate and Ties-Promote, and runoff partition of candidates, again in both models.

For partition of voters, a separate reduction to a new problem, Restricted Hitting Set, is used to show that plurality voting is resistant to destructive as well as constructive control by this method in model Ties-Promote. Other cases of constructive control were previously shown by [**BTT92**].

CHAPTER 4

Analysis of the Borda count

We consider the possible means of control one at a time, considering constructive control, destructive control, or both (that is, both causing a specified "distinguished" candidate to win, and causing that candidate to lose, following the terminology of [**HHR07a**]). We endeavor to find general principles involved in control of the Borda count, and particularly a close relationship between this task and known NP-complete problems to reduce the degree to which novel reasoning needed in specific proofs. In particular, we find that certain weighted problems closely correspond with the partition problem, as has been previously seen in other voting systems. It should be noted that in the case of constructive control by addition of weighted voters, our reasoning (while arrived at independently) is very similar to that of Conitzer, Lang, and Sandholm, who analyze the manipulation of few-candidate elections (including Borda elections) [**CLS07**].

4.1. Addition and deletion of voters

4.1.1. Destructive control by addition and deletion of voters. Destructive control by adding and deleting voters is not difficult. We must make sure that some candidate defeats the distinguished candidate, and can do this by first adding or deleting the voters who will contribute the most to the goal. If deletion or addition of any voters allow us to replace the winner with a given new candidate, those will.

Theorem 4.1.1. *The unweighted Borda count is computationally vulnerable to destructive control by adding and by deleting voters.*

Proof. We show this result by presenting a polynomial-time algorithm as follows:

To establish destructive control by deleting no more than k voters, for each candidate c' other than the distinguished candidate c:

- For each voter v in V, determine $\mathit{prefdist}(c, c', v)$ as defined in Definition 2.1.2.
  - Sort the voters by this distance.

  - Determine whether deleting the k voters with the largest positive preference distance (or all, if there are fewer than k with positive preference distance) suffices to allow c' to at least tie c.
- If that is true for some c', output a set of voters to be deleted for one such c'. If it is true for none, output control impossible.

In the case of addition of voters, a similar algorithm suffices. Again, for each candidate c' other than c:

- For each voter v in W, determine $\mathit{prefdist}(c', c, v)$ as defined in Definition 2.1.2.
  - Sort the voters by this distance.
  - Determine whether adding the k voters with the largest positive preference distance (or all, if there are fewer than k with positive preference distance) suffices to allow c' to at least tie c.
- If that is true for some c', output a set of voters to be added for one such c'. If it is true for none, output control impossible.

We may note that each opposing candidate c' can be considered separately, since in order for c to no longer be the unique winner some such candidate must defeat him or her, and this is both a necessary and a sufficient condition. Clearly, the algorithms run in polynomial time. □

4.1.2. Constructive control by addition and deletion of unweighted voters. In the case of constructive control by addition of unweighted voters, we observe that the control problem is hard because adding a voter or voters may benefit some candidates while harming others.

THEOREM 4.1.2. *The unweighted Borda count is resistant to constructive control by adding voters.*

PROOF. We show this by reduction from exact cover by 3-sets (X3C). The X3C problem involves determining whether, given a set $S = \{s_1, s_2, ..., s_{3n}\}$ for some integer n and a collection $T = \{T_1, T_2, ..., T_m\}$ of 3-elements subsets of S, T contains an exact cover for S, that is, some subset of T whose elements collectively contain each member of S exactly once (or, equivalently, so that every element of S appears in exactly one element of T). This problem is NP-complete as per [**GJ79**].

To begin with, if m is at least $3n^2$, we must pad S by adding triples of new elements, each of which appear in exactly one new subset in T, until m is less than $3n^2$. The reason for this transformation will become obvious later, but it is clear that it does not exponentially enlarge the input, and will not change the existence of an exact cover.

When we refer to a set T_i within a preference order, we mean that that set's member candidates appear in the preference order at that point in arbitrary order.

The candidate set C contains buffer candidates $b_1 ... b_{4\|S\|^3}$ (that is, a total of $4\|S\|^3$ buffer candidates), a candidate for each set member s_i (which we also term s_i) and the distinguished candidate c. We will set up the transformation such that the election is controllable when, and only when, there is an exact cover. Note that $\|C\| = 4\|S\|^3 + \|S\| + 1$.

The number of allowed additions, k, should be such that only the appropriate number of 3-sets can be added, so we set $k = n = \frac{\|S\|}{3}$.

The set of presently qualified voters V includes several groups of voters. The first group consists of $4\|S\|^3$ voters, half having preference order

$$s_1 > s_2 > ... > s_{\|S\|} > c > b_1 > b_2 > ... > b_{4\|S\|^3}$$

and the other half with preference order

$$c > s_{\|S\|} > s_{\|S\|-1} > ... > s_2 > s_1 > b_{4\|S\|^3} > b_{4\|S\|^3-1} > ... > b_2 > b_1.$$

We note that the different ordering ensures that none of the s_i and b_i candidates has an advantage over any other member of the same group, since the arithmetic means of the rankings within each group are constant. More precisely, for each s_j and each b_l, each voter v from the first group has $\mathit{prefdist}(v,\ s_j,\ b_l) = \|S\| - j + 1 + l$ and each voter from the second group has $j + 4\|S\|^3 - l$. A pair of voters, one from each group, combine to give each s_j an advantage over each b_l of $4\|S\|^3 + \|S\| + 1$, so all these voters combined give a total advantage of $8\|S\|^6 + 2\|S\|^4 + 2\|S\|^3$. This also leaves c tied with all the s_i candidates, since c acts like an s_j with $j = 0$ for the purpose of these calculations.

V additionally contains $k - 1$ voters whose purpose is to give c a disadvantage relative to the s_1 candidates. Of these voters, half have preference order

$$s_1 > s_2 > ... > s_{\|S\|-1} > s_{\|S\|} > b_1 > b_2 > ... > b_{4\|S\|^3-1} > b_{4\|S\|^3} > c$$

and the other half have preference order

$$s_{\|S\|} > s_{\|S\|-1} > ... > s_2 > s_1 > b_{4\|S\|^3} > b_{4\|S\|^3-1} > ... > b_2 > b_1 > c.$$

We note that c is now behind each of the s_i candidates by a total of

$$(k-1)(4\|S\|^3) = (\frac{\|S\|}{3} - 1)(4\|S\|^3) = \frac{4\|S\|^4}{3} - 4\|S\|^3.$$

We set the set W of unqualified voters to contain one voter v_i for each set T_i in T, and each such voter has preference order

$$c > b_1 > b_2 > ... > b_{\|S\|^3} > T_i > b_{2\|S\|^3+1} > ... > b_{4\|S\|^3} > S - T_i.$$

Adding one of these voters to the election will aid c by approximately $4\|S\|^3$ points over the candidates corresponding to $S - T_i$, but will also give the candidates in T_i an advantage of approximately $3\|S\|^3$ points over those in $S - T_i$ (the rough equivalence holds because buffer candidates greatly outnumber all other candidates; we will show this formally later).

We note again that, before any voters from V' are added to the election, c is ahead of all the buffer candidates by $8\|S\|^6+2\|S\|^4+2\|S\|^3$ but is behind all of the s_i candidates by $\frac{4\|S\|^4}{3} - 4\|S\|^3$. For any voter v_j in W, the s_i candidates which are not in the subset T_j corresponding to that voter can gain no more than $\|S\|$ (actually, no more than $\|S\| - 3$) points relative to one another as a result of v_j's vote. From all these voters combined, no s_i candidate can gain more than $m\|S\|$ points beyond the $3\|S\|^3$ it gains when chosen. This means that no such candidate can reach the total "offset" of $\|S\|^3$ needed to defeat c when a correct cover is chosen, or to fail to defeat c when one is not chosen, as long as m is less than $\|S\|^2$. That inequality is guaranteed by the padding performed at the beginning of this proof. This justifies our statement that we can speak approximately about the number of points given to the "unchosen" s_i by voters from W and the correctness of the reduction will not be affected, so long as all margins of victory are at least $\|S\|^3$.

Additionally, we note that no buffer candidate can ever win: their disadvantage of $8\|S\|^6 + 2\|S\|^4 + 2\|S\|^3$ points is far too large for any points given by voters in W to overcome.

We are now ready to prove this theorem. To begin with, we show that that the existence of an exact cover of size $\frac{\|S\|}{3} = k$ implies that control is possible. To exert control, we add the voters corresponding to the members of the cover, thereby adding $k \cdot \|C\|$, or (within the approximation bounds described above) $\frac{4\|S\|^4}{3}$, points to c and (to the same approximation) $\frac{3\|C\|}{4}$, or $3\|S\|^3$, to each s_i (depending on the order of candidates within the subset). But c had an initial disadvantage of only $\frac{4\|S\|^4}{3} - 4\|S\|^3$ points, and is now ahead of each other candidate.

On the other hand, if no exact cover exists, control is impossible. If we attempt to choose a number x fewer than k voters to add, then c gains $4x\|S\|^3$ points, while each s_i whose subset is chosen gains $3\|S\|^3$. Since the s_i whose subsets are chosen had an initial advantage over c

of $\frac{4\|S\|^4}{3} - 4\|S\|^3$ points, and x is at most $k-1 = \frac{\|S\|}{3} - 1$, c can gain at most $\frac{4\|S\|^4}{3} - 4\|S\|^3$ points, and cannot win; he or she is behind each of those s_i by $3\|S\|^3$ points.

If we choose k voters, but choose a set of additions corresponding to adding more than one subset containing some element s_j, in other words hitting s_j more than once, s_j will gain at least $6\|S\|^3$ points from these additions and again will defeat c. $\square$

Finally, we note that this reasoning holds whether we consider unique winners, or a winner, as the standard of victory for the election. The large number of buffer candidates ensures that any addition of a voter makes c either win or not win; no ties are possible.

The above result aside, we can solve the control problem for small numbers of candidates. More specifically:

Theorem 4.1.3. *The unweighted Borda count is computationally vulnerable to constructive control by adding or deleting voters for elections of three candidates.*

Proof. The control algorithm proceeds as follows:

- Sort the voters into three categories: Category 1 is those who rank the distinguished candidate first, category 2 being those who rank the distinguished candidate second, and category 3 being those who rank the distinguished candidate last.
- Without exceeding k, add as many non-qualified voters of category 1 in the adding voters case, or delete as many qualified voters of category 3 in the deleting voters case, as possible.
- If there are additions or removals left, and exactly one candidate leads or ties the distinguished candidate, add or remove voters from category 2 to the extent necessary to make the distinguished candidate strictly defeat that candidate. (We note that if the distinguished candidate is last at the beginning of this phase, nothing can be done as any change to category 2 voters will help one non-distinguished candidate and hurt the other equally).
- If the distinguished candidate is now the unique winner, control is possible; output the list of voters to be deleted or added. Otherwise, it is impossible; output accordingly.

$\square$

While this explicit algorithm works for the case of three candidates, it is important to note that for *any* bounded number of candidates, and unweighted voters, there will be a constant number of unique voters, since the number of possible preference orders is at most the factorial of the number of candidates. As such, we conjecture that it is possible to design a polynomial-time algorithm for any fixed number of candidates. This would imply that the Borda count is computationally vulnerable to control by adding or deleting unweighted voters *for constant numbers of candidates*. However, as we show above, control is hard for the general case for adding voters (and very likely is also hard for deleting voters). Intuitively, this is true since for large numbers of candidates, the number of possible interactions becomes very high.

We considered at some length the possibility that the unweighted Borda count is resistant to constructive control by deleting, as well as by adding, voters. We attempted to modify the reduction used to show the unweighted Borda count resistant to constructive control by adding voters so that it would work in the case of deleting voters as well; however, this work ran into difficulty.

The basic idea we attempted was to prevent spurious control possibilities involving deleting voters other than those corresponding to set members s_i; in an effort to do this, we made several alterations to the reduction above. We attempted to set V to include both V' and V from the existing reduction, and added two new candidates, c_1 and c_2. Then we let r be the number of voters to be undeletable, that is, the members of the initial V, and p be $3\|C\|$. We gave each existing voter a preference order with c_1 and c_2 near one another, as well as near c. Next we created a voter with preference order $c_1 >$ (*pr buffer candidates*) $> c >$ (*pr buffer candidates*) $> c_2 >$ (*all* s_i) $>$ (*remaining buffer candidates*). All voters to be made undeletable then had c replaced with $c_2 >$ (*p buffer candidates*) $> c >$ (*p buffer candidates*) $> c_1$ in their preference orders. All voters initially in V' were given a preference order with $c > c_1 > c_2$ in place of c . The intent of this mechanism was that any spurious control attempt that removed an undeletable voter(s) would have c beat by c_1, and any that broke that mechanism by removing the new voter would have c defeated by c_2.

However, this strategy failed. The set of voters to make undeletable included the large number of voters whose role was to make sure the buffer candidates did not uniquely win. However, there are as many such voters as there are buffer candidates (and so are not enough buffer candidates to form the construction in the previous paragraph). Additionally, the new voters gave c an advantage over the s_i, and if we ranked the s_i near c, the balance between c, c_1, and c_2 could be distorted. Adding more buffer candidates to overcome that effect would not help, since we would still need the voters to make sure that buffer candidates did not uniquely win, and these voters would still need to be undeletable.

Obvious methods of repairing this flaw also ran into problems. Making only some "buffer candidate killer" voters undeletable could allow only those having c higher than the s_i to remain after deletion, which might introduce spurious control possibilities (and determining whether such a situation existed might make the reduction outside P, since it itself would be a special case of control by adding voters).

Making fewer buffer candidate killers would make it harder to guarantee that some of them didn't sufficiently hurt all of the b_i to prevent their winning, making it difficult to allow the parts of the reduction depending on approximate totals to succeed.

If we used no buffer candidate killers at all, or an insufficient number, there could be cases where control should be possible but wasn't (and hence the reduction failed) because a buffer candidate won. If we "built in" killing by ranking the buffer candidates in a balanced fashion within the s_i, with the average low, there would be a risk that the control strategy that "should" work will fail because some "lucky" buffer candidate beat c.

It is possible that future work could fix these problems, but it is clear that for this particular voting system, showing resistance to deletion of voters will often be a more complex case than that for addition of voters. Showing resistance for the case of partition of voters would likely be still more intricate, since it involves both adding and deleting voters.

4.1.3. Constructive control by addition and deletion of weighted voters. In this case, it is possible that we will only barely be able to establish control, since causing the distinguished candidate to gain ground against the current winner may help some third candidate even more. The exact balancing involved makes the control problem hard if weights can be large. The case of adding voters of course follows from that unweighted case, but this proof is more elegant and more directly applicable to weighted elections. In addition, we find a resistance result for elections of only three candidates, unlike the method used for unweighted elections, which appears to require an unbounded number of candidates, and even if it were improved would still require many buffer candidates.

THEOREM 4.1.4. *The weighted Borda count, even with 3 candidates, is resistant to constructive control by adding or deleting voters.*

PROOF. It should be noted that, while we arrived at this proof independently, the reduction from the partition problem, and some technical details of the preference lists used, resemble the proof used by [**CLS07**] in showing that manipulation of positional scoring protocols for three candidates is hard (their Theorem 6, which specifically mentions the Borda count case in its first corollary).

We show this by reduction from the partition problem. In this problem, given a set S of positive integers, we must decide whether there exists some subset of S that sums to half the overall sum t. We note that, as pointed out by Garey and Johnson, this problem is NP-complete even when restricted to positive integers; in fact, this is how they specify the problem [**GJ79**].

In the case of adding voters, let the candidate set be $\{c, c_1, c_2\}$, with c the distinguished candidate. For each element $s_{\rm i}$ of S, create a voter v_i with weight $6s_{\rm i}$ and preference order $c_1 > c > c_2$. Further create two special voters, v' with weight $3t$ and preference order $c_2 > c > c_1$, and $\hat{v}$ with weight one and preference order $c > c_1 > c_2$. Make the two special voters the initial V, add the v_i to W and set k to $3t$ to allow adding voters with a total weight equal to half the sum. If it is possible to exert constructive control, the partition instance has a solution; simply form one side of the partition from the set members corresponding to the voters added to establish control.

We construct the reduction such that, in order for c to win, he or she must gain a sufficient number of points against c_2 without losing too many points against c_1. We design the preference orders so that it is impossible for adding a v_i to accomplish the former goal without a partial compromise of the latter by the same amount. It is possible that a control strategy might attempt to add voters corresponding to a total score which is too large or small to form a valid partition. The role of special voter v' is to set up the desired sum by unbalancing the vote totals so that a set of voters corresponding to a valid partition must be added. When a partition exists, the best possible control strategy will result in a tie. $\hat{v}$ will break this tie in c's favor. We note that, since $\hat{v}$ can help c's standings against any other candidate by at most two points (whereas an imbalance involving any set of other voters will hurt by a multiple of three), $\hat{v}$ will not contribute enough points to save c when the partition is unsuccessful. We design the reduction so that, in order to establish control, the chosen set of voters to be added must exactly balance the contribution of the vote cast by v' to the races between c and each non-distinguished candidate. To do this, those voters must total a weight equal to that of v', which we will set to $3t$. Since the weights are six times the original set elements, the voters to be added will correspond to a subset with a sum of $\frac{t}{2}$, which is of course a valid partition.

If we attempt to add a subset summing to a total x which is less than $\frac{t}{2}$, then c will gain $6x$ points from the v_i, $3t$ points from v', and two points from $\hat{v}$, giving a total score of $6x + 3t + 2$. Since x is an integer, and can be at most $\frac{t}{2} - 1$, c can gain no more than $6t - 4$ points. Meanwhile, c_1 will gain $6t$ points from v' and 1 from $\hat{v}$ for a total of $6t + 1$ and will defeat c by at least 5 points.

We cannot choose to add voters corresponding to a set summing more than $\frac{t}{2}$, since $k = 3t$ and the weights are six times the values of the corresponding set elements. It follows that control cannot established unless we add voters corresponding to elements with a total of exactly $\frac{t}{2}$ points.

In the reverse direction, if there is a valid partition, c can be made to win by adding the voters corresponding to the members of one side of the partition. This gives c a total of $6t + 2$ points against $6t + 1$ for c_1 and $6t$ for c_2.

In the case of constructive control by deleting voters, essentially the same construction can be used. We initially add all voters from both V and W in the addition of voters case to V, and set k to $\|S\| = \|V\| - 2$ voters.

Essentially the same argument as for the addition of voters case applies. However we must ensure that neither special voter can be deleted to establish control in a case in which it is impossible to establish control by deleting only from the v_i. If $\hat{v}$ is deleted, clearly this can only harm the standing of c, whom $\hat{v}$ ranks first. On the other hand, if v' is deleted, we cannot then delete all of the v_i since there are $2k$ total weight among the v_i and deleting v' alone used all the possible deletions. Any of the v_i who are allowed to vote in this case will overwhelm $\hat{v}$ and give victory to v_1, since $\hat{v}$ has power one, but none of the v_i has power below six. It follows that neither special voter may be deleted, so the reduction from the prior result holds. □

We note in passing that this reduction works even if we consider non-unique winners. This is true because a tie as the best possible result can never happen: $\hat{v}$ is the only voter with a power not a multiple of three. As such, the point total of each candidate is guaranteed to be congruent to a unique value modulo three in the addition of voters case. The deletion of voters case requires a bit more thought, but clearly the same property holds unless $\hat{v}$ is deleted. If deleting $\hat{v}$ causes a tie then a partition existed. In this case, constructive control without the tie will always be possible by deleting the same set but not deleting $\hat{v}$.

4.1.4. Addition and deletion of unweighted voters specified in succinct representation. If voters are provided in succinct representation, the result for the non-succinct case holds:

THEOREM 4.1.5. *The unweighted Borda count is computationally vulnerable to destructive control by adding and by deleting voters if voters are specified in succinct representation.*

PROOF. A variation of the proof of Theorem 4.1.1 works here. In the case of deletion of voters, if the distinguished candidate c already is not the unique winner, control is possible; output an empty set of voters to be deleted. Otherwise, for each non-distinguished candidate c', while c is still the unique winner and we have deletions remaining under k, use a simple greedy algorithm to determine the result if we delete as many voters as possible from the "block" (group of voters) which has the greatest positive preference distance between c and c'. If, before c is no longer the unique winner, no remaining block ranks c ahead of c', or if we exhaust k, control is impossible using this c'. If control is possible using some c', output the set of voters to be deleted for one such c'. Otherwise, output "control impossible": if c can be made to not uniquely win, this can be done by giving c' the optimal number of points for each deletion used, which can be done by first taking the greatest preference distances, so this algorithm would have succeeded.

In the case of addition of voters, the algorithm proceeds along much the same lines: if the distinguished candidate c already is not the unique winner, control is possible; output an empty set of voters to be added. Otherwise, for each non-distinguished candidate c', while c is still the unique winner and we have additions remaining under k, use a simple greedy algorithm to determine the result if we add as many voters as possible from the block which has the greatest positive preference distance between c' and c. If, before c is no longer the unique winner, no unadded group ranks c' ahead of c, or if we exhaust k, control is impossible using this c'. If control is possible using some c', output the set of voters to be added for one such c'. Otherwise, output "control impossible": if c can be made to not uniquely win, this can be done by giving c' the optimal number of points for each deletion used, which can be done by first taking the greatest preference distances, so this algorithm would have succeeded.

These algorithms proceed in polynomial time since they will always either add (delete) every voter in a group, or as many voters as possible within the group, and that number is easy to calculate by computing the room remaining under k. Rather trivially, a variant which produced results with the absolute minimum number of voters rather than finding an arbitrary solution with at most k alterations would also be polynomial-time; such a solution could simply add the minimum number of voters needed from the last group, and can determine this number in polynomial time using basic arithmetic. □

THEOREM 4.1.6. *The Borda count is resistant to constructive control by adding voters if voters are specified in succinct representation.*

PROOF. This result follows immediately from Theorem 4.1.2. □

We note that, contrary to one possible intuition, there is no obvious connection to the weighted equivalent discussed in Section 4.3. This is because while a partition instance can be represented by a weighted election, doing the direct equivalent with a succinct representation would be something closer to partition with divisible objects (which is related to the continuous knapsack problem, an easy problem).

4.2. Partition of voters

THEOREM 4.2.1. *The unweighted Borda count is vulnerable to destructive control by partition of voters in case ties eliminate.*

PROOF. If the distinguished candidate c originally is not the unique winner, control is trivially possible by choosing one sub-election to be empty, and the other to contain every voter. If there are no more than two candidates, and c initially is the unique winner, control is impossible since c must win at least one partition (having the highest number of points among the one or two candidates), and then must win the runoff. This result, true for partition of voters in many scoring protocols, is due to [**HHR07a**].

In all other cases, we must assign voters to sub-elections such that c either is not the unique winner of either sub-election, or uniquely wins exactly one, and is then tied or defeated in the runoff by the other semifinalist, who must then at least tie him or her in a two-candidate majority election (to which the two-candidate Borda count is equivalent). In the former case, we must ensure that some candidates c_1 and c_2, distinct from one another (if they are not, no solution exists since c wins overall) as well as from c, respectively at least tie c in the sub-elections V_1 and V_2 (whether or not they themselves win those sub-elections). For each such pair of candidates, we initially proceed as follows:

- Assign all voters who prefer $c_1 > c > c_2$ (possibly with additional candidates inserted) to V_1, and all who prefer $c_2 > c > c_1$

to V_2. Clearly, these voters can only hurt c in those groups, and only help c in the opposite groups.

- For each remaining voter v, compute the ratio of preference distances $\mathit{prefdist}(v, c, c_1)/\mathit{prefdist}(v, c, c_2)$. We term the voters who prefer both c_1 and c_2 to c "good voters," and those who prefer neither to c "bad voters." The others were assigned in the last step. Note that the bad voters with the highest ratio are those that give c the most points over c_1 for each point they give c against c_2. Conversely, the good voters with a high ratio help c_1 over c as much as possible for each point they give c_2 over c.
- Tentatively assign all good voters to V_1 and all bad voters to V_2. If c_1 does not defeat c in V_1, control is impossible using this pair of candidates. If c_1 defeats c in V_1, and c_2 defeats c in V_2, control is possible; output the current partition.
- Otherwise, while c_1 is still defeating or tying c in V_1, and c is defeating or tying c_2 in V_2:
  - Among all the good voters in V_1 and bad voters in V_2, move the voter with the lowest ratio into the other sub-election. That voter will, intuitively, do the least to hurt, or the most to help, the desired winner in the sub-election into which he or she is moved.

When that loop terminates, if c is not the unique winner in either sub-election, control is possible, and we can output the current partition. If not, control may still be possible: the numbers of points contributed by each voter may be relatively large integers (up the number of candidates). The above algorithm assigned good voters who give each intended winner c_1 or c_2 as many points as possible over c in their respective partitions, relative to the number of points those voters could give the other intended winner if assigned differently (and likewise, bad voters were assigned to take as few as possible). However, the above algorithm may fail to establish control in a situation where the desired winners in both sub-elections are nearly tied with c, and where moving some set of voters with a suboptimal ratio may allow the desired result in both sub-elections but moving the voter with the optimal ratio does not (informally, because it "overshoots" the desired result in one sub-election, and leaves too many behind in the other). If c is uniquely winning in both groups, or uniquely winning in one and tying in the other, control is not possible with this pair of intended winners (except possibly in the case where one intended winner defeats c in a two-candidate majority election, dealt with later). At this point, if control can be established at all, it can be done by moving a set of voters totaling at most $\|C\|$, since this is the greatest number of points any single voter can contribute; any larger differential will have been handled by moving individual optimal voters earlier if it can be fixed at all.

In order to deal with this potential problem, we must work out exactly how many points are needed, or can be spared, in each sub-election, and find a "fixing set" of voters to move to make both sub-elections work. At first it appears that this prevents this algorithm from being polynomial, since finding a set of voters to move will require solving a variant of the knapsack problem, which is widely known to be NP-complete. However, the most points any voter can give or take from the result in any two-candidate race is the number of candidates $\|C\|$, and so moving a voter of optimal ratio cannot create an imbalance greater than this number. If the input is expressed as a listing of individual voters' rankings of candidates, the input size is proportionate to the products of the voter and candidate counts. This means that the input size is not less than $\|C\|$, which when expressed in unary is shorter than our input, since each individual preference order is of length greater than $\|C\|$. The unary knapsack problem is in P, a fact we will take advantage of.

Now, we can construct a unary knapsack instance. Assume without loss of generality that c is presently uniquely winning in V_1. Let m be the "extra" victory margin, that is, the gap in points between $c_2$and c in V_2. For each voter v_i we create an item with weight equal to $\mathit{prefdist}(v_i,\ c,\ c_1)$ and value $\mathit{prefdist}(v_i,\ c,\ c_2)$; note that one or both values may be negative. Since the goal is not to move too many voters and cause c to win in V_2, we set the capacity to be m. If the knapsack instance has a solution with total value equal to at least the number of points by which c is defeating c_1 in V_1, control is possible by moving the voters whose corresponding items were added to the knapsack. This hurts c_2 in V_2 by those voters' total weight, but since this is at most m, c_2 still defeats or ties c in V_2. As such, c uniquely wins neither partition, and so cannot win the election; control is possible; we output the partitions accordingly. If the knapsack instance has no such solution, control is not possible for the reasons outlined above.

If a fixing set exists or is not needed for some non-distinguished candidates c_1 and c_2, control is possible; output the final partitions.

We note in passing that by setting a larger capacity (but still smaller than the input size, since no margin of victory can be larger than combined lengths of all preference orders) we could use this reduction without the initial steps.

Finally, we note that it is possible to establish control by allowing c to win one sub-election, but not uniquely win in the runoff election against the winner of the second sub-election. The run-off election is of course a a two-candidate Borda election, which is equivalent to a majority-vote election. For each candidate c' who defeats c in such an election (i.e., is preferred to c by a majority of voters, subject to our tie-breaking criterion), we can run the algorithm above with intended sub-winners of c and c', and a distinguished candidate (intended loser) of each other candidate. If and only if the algorithm fails again in every case, control is impossible. This algorithm could of course be designed to be more efficient by eliminating some redundant computations, but in its current form suffices to establish that the control problem is in P. □

In passing, we conjecture that since the unweighted Borda count is resistant to *constructive* control by adding unweighted voters as per Theorem 4.1.2, it is also resistant to constructive control by partition of voters, which in a sense involves an extension of that method.

We conjecture that the unweighted Borda count is resistant to destructive control in model Ties-Promote. This appears likely to be true because a potentially unlimited number of candidates may be tied and proceed to the run-off, making control of the run-off in this special case similar to control for addition and deletion of candidates. As discussed in the next section, it is unclear how a polynomial-time algorithm for control by addition and deletion of candidates in *any* case could be designed. Additionally, Hemaspaandra, Hemaspaandra, and Rothe [**HHR07a**] find that plurality voting—in some ways a simpler version of the Borda count—is resistant to destructive and constructive control by partition of voters in model Ties-Promote, suggesting that the same is likely to be true for the Borda count system.

4.3. Addition and deletion of candidates

In the weighted as well as unweighted Borda count, we do not analyze the case of adding or deleting candidates. However, it seems reasonable to conjecture that destructive control by adding or deleting candidates is easy. In order to make the distinguished candidate c lose, he or she must be beaten by some other candidate c', and adding or deleting any third candidate will change the relative ranking of these two in an easily computable manner. More specifically, it is immediate that deleting a candidate ranked by a specific voter below c, but above c', will hurt c in the two-candidate race by that voter's weight. It is true that in some cases, other candidates will benefit and achieve a total still higher than the candidate whose relative performance is under consideration, but when considering destructive control this is unimportant.

In the case of adding or deleting candidates for constructive control, we expect the Borda count to be resistant to these methods. This appears to be true because any single change in the set of candidates will potentially affect how well the preferred or despised candidate performs against all of the others (for example, inserting a candidate whom a voter places between two original candidates will give one additional point to the higher of these—as previously alluded to, this is one reason why cloning of candidates can be beneficial, if the clone will narrowly trail the original favored candidate in the eyes of many voters). The interactions are complex, and the problem intuitively seems likely to be hard. If it is easy, we would not expect it to be solvable by simple algorithms like the one presented above. If any progress is made with these cases, we anticipate that it will require methods unlike those previously applied to any voting system (though some insights from such systems as k-plurality voting would likely apply). We discuss these methods in our conclusions, but do note one important result:

THEOREM 4.3.1. *The Borda count, even for unweighted elections, does not satisfy the weak axiom of revealed preferences (WARP).*

PROOF. As defined above, WARP is the property under which the candidate winning an election is the winner of any subset containing him or her. Consider an election with 3 (unweighted) voters and 4 candidates. The voters rank the candidates in the orders $b > a > c > d$, $b > a > d > c$, and $a > c > d > b$, respectively. Candidate a wins. However, if we consider the subset of candidates a, b then b wins; thus, the winner a is not the winner of all subsets that contain him or her, and the Borda count does not satisfy WARP. □

More generally, it is possible to create an election in which a potential candidate (here, candidates c and d) outside of the subset under consideration will displace some candidates (here, a) into high rankings more often than others (b), and hence give points to the displaced candidates, who may not have won already if they were ranked low by voters who ranked the potential candidate highly. Clearly, since unweighted elections are a subset of weighted elections (that subset in which all weights are equal to one), this also holds for weighted elections.

4.4. Summary of new results

We find the following:

- The unweighted Borda count is computationally vulnerable to destructive control by adding and by deleting voters, even with voters presented in succinct representation.
- The unweighted Borda count is computationally vulnerable to destructive control by partition of voters in case Ties-Eliminate. (We may reasonably conjecture that it is resistant in case Ties-Promote, based on the results of [**HHR07a**] for plurality voting, which is in a sense a "simpler" system).
- The weighted Borda count, even with only three candidates, is resistant to constructive control by adding and by deleting voters.
- The unweighted Borda count is resistant to control by adding voters at least in the case of unbounded numbers of candidates whether or not voters are specified in succinct representation.

CHAPTER 5

Conclusions

We find that in the case of adding unweighted voters the Borda count is computationally vulnerable to destructive control, but resistant to constructive control. In the case of adding and deleting weighted voters, it is resistant to constructive control with only three candidates. These are the cases in which we might reasonably expect it to be easiest to establish control (intuitively, partition of voters is a combination of adding and deleting voters, while coordinating between two sub-elections, and the methods of control involving alteration to the candidate set are very complicated to consider for the Borda count in general). As such, we might reasonably conjecture that further work will show the weighted Borda count to be resistant to constructive control by every means (except in the case of two or fewer candidates, where it degrades to a majority-vote election). Even in the unweighted case, we show resistance for unbounded numbers of candidates, but this result appears unlikely to hold with a fixed number of candidates, particularly since the Borda count is vulnerable to constructive control by adding and also by deleting voters with three candidates.

As a general rule, as we might expect, control is more difficult in weighted Borda count elections. Clearly, it cannot be less difficult, since unweighted elections are simply a special case of weighted elections in which all voters have equal weight. Succinct representations appear to be intermediate in difficulty in a sense. That is, in the case in which a number of weighted candidates is changed to the same number of groups of voters, it appears that the control problem becomes easier in that we can use greedy algorithms rather than solving the partition instance directly. As one concrete example of this, we find that in the case of constructive control by adding or deleting weighted voters with three candidates, the Borda count is resistant; on the other hand, in the case of succinct representations, it is vulnerable.

Despite the Borda count's high resistance to constructive control, it is unlikely to be a practically useful election system due to the common and obvious strategies for manipulation. Particularly, as mentioned earlier, it is vulnerable to a manipulation strategy—viable even for individual voters—of reporting the gap between strong candidates in one's preference order as wider than it in fact is. As we mentioned, these strategies can potentially allow a single voter to have power much out of proportion with his or her actual preference distance between the strong candidates. Worse, it is possible that many voters in political elections will use this type of tactical voting. An obvious strategy for tactical voting is to place weak candidates in the middle of the reported preference order, with the liked and disliked strong candidates near the beginning and end, respectively. Some manipulators might be unfamiliar with the weak candidates, and rank them in an "obvious" order, for example alphabetically or in the order in which candidates appear on the ballot. This might give an advantage to the first weak candidate in this order, who would gain second-place points from voters tactically voting to aid *any* strong candidate, perhaps leading to a candidate with very little actual support or recognition winning. On the other hand, a "typical" (informally speaking) manipulation strategy for plurality voting, for example, will benefit a voter's favorite among the strong candidates and nobody else. There is, then, much more chance in a Borda count election that a relative unknown could win due to interacting attempts at manipulation (whether or not the election reaches the "ideal" result, from the perspective of the voters' preferences, is another issue, and of course all reasonable voting systems are manipulable).

The manipulation strategies for many commonly used voting systems, such as plurality voting, approval voting, and instant runoff voting generally do not allow amplification of effectiveness for individual dishonest voters. In these systems, the total influence of any voter on a pairwise contest between any pair of candidates is limited to that voter's power. The systems do not consider preference distance. In limiting the worst-case effectiveness of manipulation, this weakness becomes a strength in a sense. Along the same lines, these other systems do not suffer from the same extreme vulnerability to cloning of candidates found in the Borda count. Of course, plurality voting at least does have problems with a spoiler effect. In fact, all the most common systems are vulnerable to some forms of control or manipulation (and, again, no reasonable system can ever be completely immune to manipulation).

While it has significant problems with manipulation, it is interesting to note that the Borda count is resistant to constructive control by adding and by deleting weighted voters even with small numbers of candidates. This could potentially make it a practical system for weighted elections in which this is a concern, such as perhaps corporate shareholder meetings in which it may be possible to discourage certain shareholders from attending the meeting, such as by delaying mail to some, and conversely possible to remind others to attend.

As another advantage, the Borda count tends to choose a candidate with a broad base of support. In the context of political elections in particular, all voters may be pleased with a candidate who is somewhat moderate. A plurality election in the same may become essentially a race between two extreme candidates, either of whom about half the population might be strongly displeased with. More concretely, if there are five candidates, and roughly half the voter population has a ranking of $a > b > c > d > e$, while most of the remainder has $e > d > c > b > a$, a plurality election is certain to choose one of the "extreme" candidates. Even in this case with "extremist" voters, however, the Borda count will be about equally likely to choose any candidate, depending on exactly how the votes are cast (if the election had two voters, one of each of the above forms, it would actually be a five-way tie). Which system this tendency is an "advantage" for is of course highly subjective, and depends on the goals for which the voting system is applied: if we want to minimize the number of strongly displeased voters, the Borda count may be better, while if we want to make sure that as many voters as possible get their very first choice, the plurality system will be better (and, clearly, is the best possible system here).

It is also interesting to note that the fact that a control or manipulation problem is NP-complete does not necessarily mean that the election is safe from practical manipulation and control. First of all, in many political elections there are only two or three candidates with a reasonable chance of winning in practice (even in nations with several major parties, all individual candidates may not be equally qualified, for example). Additionally, Procaccia and Rosenschein establish that certain NP-hard manipulation problems can be solved in polynomial time on average across certain probability-weighted "junta distributions" [**PR06**]. More exactly, they find that (apparently only unweighted) scoring protocols—such as the Borda count for fixed numbers of candidates—are "sensitive" to manipulation by coalitions (i.e., cooperating groups) of voters averaged across these distributions when the number of candidates is constant. This result is of potential concern for political elections, in which the number of candidates will frequently be very small. Of course, other systems, including plurality voting, are scoring protocols for fixed numbers of voters just as the Borda count is.

As we state in the introduction, the Borda count has philosophical advantages in addition to its computational properties. In particular, it is one of relatively few voting systems which make use of all the information provided by the voters in a standard preference order of the sort we consider. Most other systems consider only the first preference as in the case of plurality voting, a few early preferences as in instant runoff voting, or a simplified "summary" of preferences as in approval voting. Once we decide to model preferences as a simple sorted list, using the entire list seems only reasonable. If we use more information (e.g. asking voters how much their opinion of successive candidates in the list differs), we open up fairly obvious strategies for tactical voting that do not even require voters to lie about their preference order. For example, in that case, voters could widen those gaps between strong candidates, and leave those above the first, and below the last, strong candidate narrow.

The Borda count's "reverse spoiler effect" encouraging cloning of candidates can reasonably be seen as a theoretical weakness. However, from a philosophical standpoint it may be seen as an advantage in the context of political elections. It may be ideal for voters to be offered a choice between a relatively large number of candidates, who might better represent the range of opinions among the voters. One way to meet this goal is to ensure that the stronger candidates do not have an incentive to discourage other candidates of similar views as plurality voting tends to (we can consider, for example, the tendency in U.S. presidential elections for stronger candidates to attempt to keep similar weak candidates out of the debates). Of course, in a Borda count election we run into the exact opposite problem: it may be to the benefit of strong candidates to discourage those weak candidates that oppose them. However, it is possible that the most reasonable way to discourage weak opponents would be for candidates to engage in debate with *all* opponents, even weak ones who might not be invited otherwise, which may be seen as desirable from the perspective of society.

It should be noted that this definition of "strong" and "weak" candidates is informal, but is reasonable to consider in our general context of perfect knowledge, in which voters and candidates can recognize those candidates who are presently leading. In addition, it has some application to real political elections, in which there are often one or several candidates (e.g., minor-party candidates or extremists) who cannot be reasonably expected to win, a fact which is known to all concerned. Along the same lines, in artificial intelligence applications, there may be certain alternatives which efficient heuristics suggest are only marginally worthy of consideration. It might be possible for "voters" to apply these heuristics, and use the results in a control strategy (accepting some risk that the heuristics fail).

Despite theoretical deficiencies in the Borda count, Reilly points out that Nauru and Kiribati are among relatively few nations outside the West to remain free since independence, and do so in spite of significant practical problems such as poverty [**Rei02**]. Kiribati uses the Borda count in committee elections within parliament, which decide who will stand for president (a national plurality election is then held among the three or four highest-scoring Borda count winners, who are chosen by and among the members of parliament). As we might expect, the committee elections have a small number of voters, who have some knowledge of one another's preferences and opportunity to collude; in one particular election, this allowed manipulation. More specifically, the outgoing government and a group of independents forced out the leader of the opposition, in favor of their own respective leaders and two candidates who had no real intent of running. The nationwide four-candidate election then effectively had to choose between the leaders of the two cooperating groups of manipulators [**Rei02**]. This outcome, in view of the Borda count's overall success in Kiribati, shows that the Borda count is very much a mixed blessing. But, of course, we might say the same of any other reasonable voting system, and neither our results nor any existing ones call this into doubt.

5.1. Further work

In view of our work, many open questions remain. Firstly, it is completely unclear how the Borda count performs against control via alterations to the candidate set, (including against addition, deletion, partition, and runoff partition of candidates). In the case of arbitrary numbers of candidates, we conjecture that the Borda count is resistant to control by most or all of these means. This is due to complicated interactions of gaining or losing points as candidates leave a (sub)election. That said, it is immediate that candidate-related problems are easy with any constant number of candidates since we can simply try all possibilities. However, if the Borda count is modified to only consider each voter's highest few (constant number of) preferences, it is unclear what would result.

One interesting question is whether the Borda count can be made more resistant to control via relatively minor modifications to prevent easily computing trade-offs in greedy algorithms. For example, we could change the system so that some changes of one position in preference orders are worth more points than others. For example, we could give 999 points to the first choice and 501 to the second choice of each voter in a three-candidate election without greatly distorting the results of most real-world elections. A reasonable conjecture is that this sort of change will not affect the difficulty of control, even if the new set of weights is specified in binary, unless the modifications are relatively large (in other words, often make the winner different than in a standard Borda count election). Hybridization is one technique which might make control harder than an unmodified Borda count election. Construction of modified ("hybrid") voting systems resistant to various methods is considered by [**HHR07b**], who use this technique to find an election system resistant to every standard means of control (however, the resulting voting system is much more complex than the Borda count).

We may consider how the modified Borda count used in Nauru (which gives one point to the first-ranked candidate for each voter, 1/2 to the second, and generally $1/m$ to the mth) performs against manipulation and bribery, as well as control. We conjecture that control of this variant will be no easier than the more common variant, and possibly harder because control of this sort of election could not use simple calculations based on preference distance, on which many of our results for unweighted elections depend. For example, in this case a voter with $a > b > c$ does not help a against c as much as two voters with $a > c > b$ do, and as we add more candidates matters become increasingly complicated. Of course, all point totals will be rational numbers, and comparison of rational numbers in fraction form is an easy problem (even when the numerator and denominator are unbounded, as with unbounded numbers of candidates), so we do not expect a great change in complexity, except perhaps in the succinct representation case.

It is possible that approximating control in the "Nauru count" (that is, designing a polynomial-time algorithm which successfully exerts control in all cases in which it is possible to make the desired winner lead all rivals by, for example, 1/10 of the points, but fails for some cases in which the margin of victory is narrower) would be easier than for the standard Borda count, since it might be possible to disregard or summarize certain lower preferences. For example, in a 100 candidate election, a candidate ranked 90th by a voter has an advantage of only 1/900th of a point over the candidate ranked 100th, and it would take dozens to hundreds of such voters to make up for one voter ranking the candidates in the opposite order in the top 5.

We may also consider whether constructive control of weighted Borda count elections by various methods of altering the voting set is polynomial in any special cases, e.g., when some subset of the candidates have a very small amount of support. It seems reasonable that this might allow "space" to help the distinguished candidate while also helping such weak candidates, but not enough to risk them winning, but how well this generalizes is unclear. This scenario is applicable to many real-world cases; for example, political elections often have a few strong candidates, and many proposed solutions (e.g., chess moves) are clearly bad and will have little support. If this particular special case is vulnerable to control, we might then consider whether a few weak candidates suffice, or if it is necessary that the great majority be weak. Then we could consider the average case, for example, the case in which weak candidates are randomly placed within the bottom half of voter preferences.

We might consider how the Borda count behaves with small constant numbers of voters, for example two or three, but large numbers of candidates. Intuitively, it appears that in this scenario the system would become vulnerable to all forms of control involving alteration of the voter set, since it would be easy to compute the results of the election under each possible change. However, there might still be resistance to means of control involving candidates.

The Borda count clearly has a significant weakness against manipulation as well as against certain means of control. Dummet notes that a "penalty for dissimilarity" between candidates in a given voter's presence order, effectively applied against the latter, and that this is the reason why adding a clone of a candidate may help that candidate [**Dum98**], and proceeds to modify the system to fix this. We might consider whether something similar will work against control.

Acknowledgements: This research was supported in part by grants NSF-CCR-0311021 and NSF-IIS-0713061. Thanks to Professor Hemaspaandra for feedback and continued support during my at times slow progress on this thesis, and thanks to Professor Homan for valuable feedback towards the final version. Additionally, thanks to my entire committee for valued advice, and for being willing to schedule my defense during the summer quarter. Thanks to my parents for encouragement, and for being a sounding board for my ideas.

Bibliography

[Amy00] D. Amy. Behind the Ballot Box: a Citizen's Guide to Voting Systems, Praeger Publishers, 2000

[Arr50] K. Arrow. *A Difficulty in the Concept of Social Welfare*, Journal of Political Economy 58(4), pp. 328–346, August 1950

[ACS03] N. Aswal, S. Chatterji, and A. Sen. *Dictatorial Domains,* Economic Theory, vol. 22, issue 1, pp. 45-62, 2003

[BPT06] M. Barbie, C. Puppe, and A. Tasnádi. *Non-manipulable Domains for the Borda Count,* Economic Theory, vol. 27, issue 2, pp. 411-430, 2006

[BTT89] [J. Bartholdi, C. Tovey, and M. Trick. *Voting Schemes for Which it can be Difficult to Tell Who Won the election*, Social Choice and Welfare 6, pp. 157–165, 1989

[BTT92] J. Bartholdi, C. Tovey, and M. Trick. *How Hard is it to Control an Election?*, Mathematical and Computer Modeling, 16(8/9), pp. 27-40, 1992

[CK06] Campbell, D.E.; Kelly, J.S. *Social Welfare Functions Generating Social Choice Rules that are Invulnerable to Manipulation* Mathematical Social Sciences, vol. 51, issue 1, pp. 81–89, 1 January 2006

[CBC06] CBC news. Canada Votes 2006, available `http://www.cbc.ca/canadavotes/`, accessed June 28, 2007

[CLS03] V. Conitzer, J. Lang, and T. Sandholm. *How Many Candidates are Needed to Make Elections Hard to Manipulate?,* Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge, ACM Press, pp. 201-214. July 2003

[CLS07] V. Conitzer, J. Lang, and T. Sandholm. *When are Elections With Few Candidates Hard to Manipulate?,* J. ACM, to appear

[CS05] Conitzer, V. and Sandholm, T. 2005. *Communication Complexity of Common Voting Rules*, Proceedings of the 6th ACM Conference on Electronic Commerce, 2005

[Cop06] Copia. *Judgement of Paris 30th Anniversary Tasting Results,* available `http://www.copia.org/content/node/483`, modified May 24, 2006, accessed May 29, 2007

[Dum98] M. Dummett. *The Borda Count and Agenda Manipulation,* Social Choice and Welfare, Springer, vol. 15(2), pp. 289-296, 1998.

[Dun01] P. Dunne. *An Annotated List of Selected NP-complete Problems.* The University of Liverpool, Dept of Computer Science, COMP202. Available `http://www.csc.liv.ac.uk/~ped/teachadmin/COMP202/annotated_np.html`, modified Jan. 26, 2001, accessed June 2, 2007

[DKNS01] Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. *Rank Aggregation Methods for the Web.* Proceedings of the 10th International Conference on World Wide Web, held May 1–5, 2001

[FHH06] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. *The Complexity of Bribery in Elections*, AAAI 2006, pp. 641–646, July 2006

[GJ79] M. Garey. and D. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman and co., 1979

[Gib73] A. Gibbard. *Manipulation of Voting Schemes: a General Result*, Econometrica, vol. 41, no. 4, pp. 587–601, 1973

[HHR07a] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. *Anyone but Him: The Complexity of Precluding an Alternative*, Artificial Intelligence, vol. 171, issue 5–6, 2007

[HHR07b] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. *Hybrid Elections Broaden Complexity-Theoretic Resistance to Control,* IJCAI 2007

[JKC99] S. Jeong, K. Kim, B. Chun, J. Lee, and Y. Bae. *An Effective Method for Combining Multiple Features of Image Retrieval,* Proceedings of the IEEE Region 10 Conference, 1999

[ML06] I. McLean and H. Lorrey. *Voting in the Medieval Papacy and Religious Orders,* Nuffield College Working Papers in Polticis, available `http://www.nuffield.ox.ac.uk/Politics/papers/2006/mclean02.pdf`, modified 2006, accessed April 10, 2007

[PR06] A. Procaccia and J. Rosenschein. *Junta Distributions and the Average-case Complexity of Manipulating Elections*, Journal of Artificial Inteilligence Research, Proceedings of the Fifth international Joint Conference on Autonomous Agents and Multiagent Systems 2006

[Rei02] B. Reilly. *Social Choice in the South Seas: Electoral Innovation and the Borda Count in the Pacific Island Countries*, International Political Review (2002), Vol 24, No. 4, 355-372. Available `http://rangevoting.org/ReillySCSS.pdf,` accessed Oct. 22, 2006

[RS04] W. Richards and H. Seung. *Neural Voting Machines,* MIT Technical Reports, available `http://dspace.mit.edu/handle/1721.1/30513`, accessed May 30, 2007

[RSV03] J. Rothe, H. Spakowski, and J. Vogel. *Exact complexity of the winner problem for Young elections,*Theory of Computing Systems, vol. 36, no. 5, pp.375–386, 2003.

[Sat75] M. Satterthwaite. *Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions* Journal of Economic Theory 10 (April 1975), pg. 187–217

[Ser02] J. Serais. *Cloning Manipulation of the Borda Rule*, Meeting SCW 2002

[Str80] P. Straffin. Topics in the Theory of Voting, Birkhäuser Boston, Inc., 1980