

Rochester Institute of Technology

RIT Digital Institutional Repository

Theses

6-15-2017

Election-Attack Complexity for More Natural Models

Zack Fitzsimmons
zmf6921@rit.edu

Follow this and additional works at: <https://repository.rit.edu/theses>

Recommended Citation

Fitzsimmons, Zack, "Election-Attack Complexity for More Natural Models" (2017). Thesis. Rochester Institute of Technology. Accessed from

This Dissertation is brought to you for free and open access by the RIT Libraries. For more information, please contact repository@rit.edu.

Election-Attack Complexity for More Natural Models

by

Zack Fitzsimmons

A dissertation submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in Computing and Information Sciences

Ph.D. Program in Computing and Information Sciences
B. Thomas Golisano College of Computing and
Information Sciences

Rochester Institute of Technology

Rochester, New York

June 15, 2017

Election-Attack Complexity for More Natural Models

by
Zack Fitzsimmons

Committee Approval:

We, the undersigned committee members, certify that we have advised and/or supervised the candidate on the work described in this dissertation. We further certify that we have reviewed the dissertation manuscript and approve it in partial fulfillment of the requirements of the degree of Doctor of Philosophy in Computing and Information Sciences.

Dr. Edith Hemaspaandra, Dissertation Advisor Date

Dr. Ivona Bezáková, Dissertation Committee Member Date

Dr. Lane A. Hemaspaandra, Dissertation Committee Member Date

Dr. Robert Parody, Dissertation Committee Member Date

Dr. Richard Zanibbi, Dissertation Committee Member Date

Dr. James Vallino, Dissertation Chair Date

Certified by:

Dr. Pengcheng Shi, Ph.D. Program Director Date

Election-Attack Complexity for More Natural Models

by
Zack Fitzsimmons

ABSTRACT

Elections are arguably the best way that a group of agents with preferences over a set of choices can reach a decision. This can include political domains, as well as multiagent systems in artificial-intelligence settings. It is well-known that every reasonable election system is manipulable, but determining whether such a manipulation exists may be computationally infeasible. We build on an exciting line of research that considers the complexity of election-attack problems, which include voters misrepresenting their preferences (manipulation) and attacks on the structure of the election itself (control). We must properly model such attacks and the preferences of the electorate to give us insight into the difficulty of election attacks in natural settings. This includes models for how the voters can state their preferences, their structure, and new models for the election attack itself.

We study several different natural models on the structure of the voters. In the computational study of election attacks it is generally assumed that voters strictly rank all of the candidates from most to least preferred. We consider the very natural model where voters are able to cast votes with ties, and the model where they additionally have a single-peaked structure. Specifically, we explore how voters with varying amounts of ties and structure in their preferences affect the computational complexity of different election attacks and the complexity of determining whether a given electorate is single-peaked.

For the representation of the voters, we consider how representing the voters succinctly affects the complexity of election attacks and discuss how approaches for the nonsuccinct case can be adapted.

Control and manipulation are two of the most commonly studied election-attack problems. We introduce a model of electoral control in the setting where some of the voters act strategically (i.e., are manipulators), and consider both the case where the agent controlling the election and the manipulators share a goal, and the case where they have competing goals.

The computational study of election-attack problems allows us to better understand how different election systems compare to one another, and it is important to study these problems for natural settings, as this thesis does.

Acknowledgments

First and foremost I would like to thank my advisor Edith Hemaspaandra, who has been the source of endless advice, encouragement, and patience. Over the years she has become not just an advisor, but a good friend. She has shaped me into the academic that I am today, and from her I have gained a deep appreciation for the structure of hard problems.

I would like to thank Lane A. Hemaspaandra for his guidance and helpful insight on research and on writing. I have really enjoyed collaborating on our joint papers.

Thank you to Pengcheng Shi for all of his valuable advice, and for helping to guide me through my graduate studies.

Thank you to Ivona Bezáková, Robert Parody, and Richard Zanibbi for serving on my dissertation committee and for all of the discussions that have helped me in my research and to improve this thesis.

I would like to thank the Computer Science Department at the Rochester Institute of Technology for the opportunity to teach during my graduate studies, which helped me to develop as an educator.

For all of our interesting discussions that have made my time at RIT a great experience, I would also like to thank Zack Butler, Hadi Hosseini, Peizhao Hu, Matthew Le, Erika Mesh, David Narváez, Stanisław Radziszowski, and Carlos Rivero.

For the many miles we've both traveled, by train and by car, thank you to Kristen Perrella for all of her love and support throughout my graduate studies.

And finally, I would like to thank my parents Jack and Linda, my brother John, and my sister Heather, for everything they have done for me, and for always being available for advice, encouragement, and support.

My graduate study was supported in part by NSF grant CCF-1101452 and an NSF Graduate Research Fellowship under grant no. DGE-1102937.

To my parents.

Contents

1	Introduction	1
1.1	Background	1
1.2	Computational Study of Election-Attack Problems	3
1.3	More Natural Models	4
1.4	List of Contributions	7
1.5	List of Publications	8
2	Preliminaries	9
2.1	Elections and Preferences	9
2.2	Election Systems	10
2.2.1	Scoring Rules	10
2.2.2	Extensions for Votes with Ties	11
2.2.3	Pairwise Rules	12
2.2.4	Approval Voting	14
2.3	Manipulative Attacks	14
2.3.1	Manipulation	14
2.3.2	Control	15
2.3.3	Bribery	17
2.4	Computational Complexity	17
3	Models for the Votes: Votes with Ties	21
3.1	Introduction	21
3.2	Complexity Goes Up	23
3.3	Complexity Goes Down	24
3.4	Complexity Remains the Same	26
3.4.1	Majority-Graph Result	29
3.4.2	Irrational-Voter Copeland Results	30
3.5	Axioms	32
3.6	Related Work	33
3.7	Conclusions	34
4	Models for the Electorate: Single-Peakedness with Ties	35
4.1	Introduction	35
4.1.1	Single-Peakedness with Ties	36
4.1.2	Contributions	37

4.2	Models of Single-Peakedness for Votes with Ties	38
4.2.1	Social-Choice Properties	41
4.3	Detecting Single-Peakedness	43
4.3.1	Possibly Single-Peaked Consistency	43
4.3.2	Single-Plateaued Consistency	46
4.3.3	Single-Peaked Consistency with Outside Options	50
4.3.4	Single-Peaked Consistency	53
4.4	Effect of Different Models on the Complexity of Manipulation	55
4.4.1	Complexity Goes Up	56
4.4.2	Complexity Remains the Same	59
4.5	Conclusions	65
5	Models for the Representation: Succinct Elections	66
5.1	Introduction	66
5.2	Adapting Approaches	69
5.2.1	Dynamic Programming	72
5.2.2	Fixed Numbers of Candidates	73
5.3	Kemeny Elections	74
5.4	Conclusions	82
6	Models for the Attack: Control with Manipulation	84
6.1	Introduction	84
6.2	Specification of the Model	85
6.3	Inheritance Results	87
6.4	General Upper Bounds and Matching Lower Bounds	88
6.5	Specific Systems	101
6.5.1	Plurality	106
6.5.2	Condorcet	110
6.5.3	Approval	114
6.6	Weighted Voters	117
6.7	Related Work	119
6.8	Conclusions	120
7	Conclusions	122
	Bibliography	123

Chapter 1

Introduction

1.1 Background

Elections are a flexible, widely-used framework for preference aggregation. They are used to seek fair outcomes in domains ranging from the human (often in political elections) to the electronic (where they play a central role in multiagent systems, a subfield of artificial intelligence). Examples of such applications include a meta-search engine that combines the webpage rankings of several search-engines to return the best decision on a search [DKNS01], an agent-based movie recommender system [GMHS99], and aggregating information in human computation systems [MPC13]. There is a rich literature in economics and political science that studies the use of elections, and different properties that different election systems have (see, e.g., [Bla58, RO73]).

Let's first discuss how an election is modeled. Below we have an election with the candidates a , b , c , and d , and the following four voters, where each voter strictly ranks all of the candidates from most to least preferred.

- v_1 voting ($b > c > d > a$).
- v_2 voting ($a > b > d > c$).
- v_3 voting ($d > b > c > a$).
- v_4 voting ($a > d > b > c$).

One of the most commonly used ways of scoring an election is to award one point to each voter's top preference, and the candidate(s) with the highest score win. When we score the election above in this way a wins with a score of 2. However, notice that although a is the most-preferred candidate for voters v_2 and v_4 , a is the least-preferred candidate for voters v_1 and v_3 . If we instead score this election with the Borda count, where for an election

with four candidates, each voter’s most-preferred candidate receives 3 points, their second-preferred candidate receives 2 points, their third-preferred candidate receives 1 point, and their least-preferred candidate receives 0 points, then b wins with a score of 8. And we can see that b seems to better represent the preferences of the voters, since voter v_1 gets her top choice, voters v_2 and v_3 get their second choice, and voter v_4 gets her third choice.

It is natural to wonder whether there exists an election system that will always elect winners that properly represent the preferences of the electorate. Unfortunately, an important result by Arrow, referred to as Arrow’s Impossibility Theorem [Arr50], states that given a very basic set of desirable properties, there exists no election system that satisfies all of them—there will always be tradeoffs. Since this early work by Arrow, there has been a long line of research in social-choice theory that studies different properties for election systems (see, e.g., [Smi73, You75, YL78, Tid87, FBC14]).

Another very desirable property for an election system is strategyproofness, i.e., it is not possible for a voter to cast a strategic vote that results in a more-preferred outcome for that voter.

Consider the scenario where voter v_4 has knowledge of the preferences of the other voters. Naturally, she wants a better outcome for herself, and can accomplish this by casting the strategic vote ($d > a > c > b$) so that d wins with a score of 8. So v_4 gets a personally better outcome by voting strategically (by getting her second choice instead of her third choice), but notice that this is a worse outcome overall, since now two voters (v_1 and v_2) get their third choice.

It is obvious to suggest that we should use election systems where this is not possible. Unfortunately, an important negative result from social-choice theory, the Gibbard-Satterthwaite Theorem [Gib73, Sat75] states that no reasonable election system over three or more candidates is manipulable, i.e., is not strategyproof. And when we say here that an election system is reasonable, we mean that every candidate can win and the outcome of the election is not decided by a single fixed voter (i.e., a dictator), which are very basic properties that one would want in an election system. We mention here that the Gibbard-Satterthwaite theorem only holds for so-called resolute election systems that elect a single winner, but in many cases we want an election system to be able to elect multiple winners. However, later work by Duggan and Schwartz generalizes this theorem to such irresolute election systems [DS00].

1.2 Computational Study of Election-Attack Problems

Although every reasonable election system can be manipulated, it may be computationally infeasible to determine if a successful manipulation exists. Bartholdi, Tovey, and Trick [BTT89a] (see also, Bartholdi and Orlin [BO91]) introduced this notion of examining the computational complexity of the manipulation problem, where an election system that is computationally hard to be manipulated is said to be “resistant” to manipulation, and this is an example of a computational property of an election system. This seminal work on the computational study of manipulation was very influential in the creation of the research area of computational social choice, which is an interdisciplinary field that looks at problems in social choice theory, such as elections, through a computational lens (see [BCE⁺16] for a general overview of the field). And the increasing use of elections in artificial intelligence applications encourages us to study these computational properties in the same way as the aforementioned social-choice properties.

The models of manipulation, control, and bribery are the three most-commonly studied election attacks, and we describe each of these problems below.

The influential work by Bartholdi, Tovey, and Trick considered the case of manipulation for a single manipulative voter [BTT89a]. However, in real-world settings it is likely that several voters collude to manipulate the outcome of an election. Conitzer, Sandholm, and Lang extended the single-manipulator model to this coalitional case, where a coalition of manipulators wants to set their votes to ensure that a preferred candidate wins. They also considered the destructive case, where instead of wanting to ensure that a preferred candidate wins, the manipulators want to ensure that a despised candidate loses, and the case where voters can have associated weights (where a weight ω voter can be thought of as a coalition of ω unweighted voters with the same vote) [CSL07]. A large line of research followed which explores the computational complexity of manipulation for various different election systems, and considers variants of the manipulation problem itself to better model natural scenarios that may occur.

Bartholdi, Tovey, and Trick expanded on their work on manipulation by introducing and analyzing the complexity of a family of manipulative attacks referred to as control [BTT92]. Electoral control models the actions of an election organizer, referred to as the election chair, who has control over the structure of the election (e.g., the voters) and wants to ensure that a preferred candidate wins. Different electoral control actions model real-world situations such as get-out-the-vote drives (adding voters) and gerrymandering (partitioning voters), and resistance to certain types of control actions may motivate the choice of an election system. Hemaspaandra, Hemaspaandra, and Rothe were the first to consider the case of

destructive control, where the chair wants to ensure that a despised candidate does not win [HHR07]. And like manipulation, control has also been studied for the case of weighted elections [FHH15].

Faliszewski, Hemaspaandra, and Hemaspaandra introduced the model of bribery, which is closely related to manipulation, but instead of asking if voters can cast strategic votes to ensure a preferred outcome, bribery asks whether it is possible to change the votes of a subcollection of the voters to ensure a preferred outcome [FHH09]. Bribery is often motivated in a more positive way where the cost of changing the votes of the selected voters represents the campaign costs of an election organizer that seeks to change their votes.

Overall, the study of the complexity of such election-attack problems has been one of the most important directions of research in computational social choice (see, e.g., [FHH10, FP10]).

1.3 More Natural Models

It is important that we study these election-attack problems under natural assumptions, since this may affect the computational complexity of a given attack. And the focus of this thesis is on studying just that. *We examine how different natural models for the votes, the electorate, the representation, and even the attack itself affect the computational complexity of election-attack problems.*

A prominent direction of this thesis is the study of elections that allow voters to state ties between candidates in their preferences. Allowing votes with ties is very natural and such votes are seen in real-world preference data (see, e.g., the datasets available on PREFLIB [MW13]), and election systems such as Kemeny [Kem59] and Schulze [Sch11] are defined for votes with ties.

Even though it is very natural to consider votes with ties, the computational study of elections typically assumes that voters strictly rank all of the candidates from most to least preferred, i.e., tie-free votes. In Chapter 3 we examine how allowing voters to cast votes with varying amounts of ties affects the computational complexity of the election-attack problems of manipulation, bribery, and control, extending the definitions of election systems (when necessary) to handle votes with ties. The main contribution of this chapter is that we are the first to computationally study manipulation for more general votes with ties, and the first to computationally study the standard models of bribery and control for votes with ties. The results in this chapter previously appeared in Fitzsimmons and Hemaspaandra [FH15] with the discussion of which axioms our extensions to scoring rules satisfy appearing in the later workshop version [FH16a].

Another natural model that we look at, this time for the preferences of the electorate, is a restriction on preferences called single-peakedness [Bla48]. Restricting the preferences of an electorate may sound counter-intuitive, and results such as Arrow’s theorem consider allowing unrestricted preferences to be a reasonable and desirable property. But single-peakedness is not a restriction in the sense of preventing voters from voting their true preferences, but rather modeling how preferences are structured in certain natural scenarios. Intuitively, single-peaked preferences can be thought of as modeling the preferences of an electorate with respect to a single issue, where there exists a one-dimensional ordering of the candidates (an axis) and candidates farther to the leftmost and rightmost points on the axis represent the extremes of the issue.

In Chapter 4 we consider single-peaked preferences in the setting where voters can cast votes with ties. The standard model of single-peakedness due to Black [Bla48] is defined for such votes, but other models have also been introduced that generalize single-peakedness in different ways for votes with ties. In addition to Black’s standard model, we consider the models of single-plateaued preferences [Bla58], single-peaked preferences with outside options [Can04], and possibly single-peaked preferences [Lac14], and we compare how these models relate to one another for different types of votes with ties.

Single-peaked electorates have desirable social-choice properties. For example, there are reasonable strategyproof election systems when the voters in an election are single-peaked (see, e.g., [Bar01]). So one could argue that we should use only strategyproof election systems when we have single-peaked electorates. However, as mentioned by Faliszewski et al. [FHHR11], it is not always the case that one can choose the election system in a given situation. In addition, there may be other properties that one wants in an election system even more than strategyproofness, and furthermore strategyproofness does not imply that different types of electoral control are not possible [FHHR11]. Thus we also study the complexity of election-attack problems for single-peaked votes. For tie-free votes, it has been shown that different computational problems often become easier when the votes in an election are single-peaked [FHHR11, BBHH15].

Our contributions in Chapter 4 include showing that it is easy to determine if a given collection of votes with ties satisfies each of the models of single-peakedness that we consider. In particular, we show that the consistency problem for single-peaked, single-plateaued, possibly single-peaked, and single-peaked preferences with outside options are each in P for votes with ties. Additionally, we expand on the work in Chapter 3 that considers the complexity of manipulation for votes with ties by considering the complexity of manipulation for single-peaked votes with ties. And we show that for a family of election systems called scoring rules that the complexity of manipulation does not increase when moving from the

case of tie-free single-peaked votes to single-peaked votes with ties, when the standard model of single-peakedness is used. However, when the model of possibly single-peaked preferences is used we observe an anomalous increase in complexity. Most of these results previously appeared in Fitzsimmons [Fit15] and Fitzsimmons and Hemaspaandra [FH16c].

Since our results concern computational complexity, it is important that we consider the representation of our problems. In most of the computational study of elections, the voters in an election are represented as a list of their individual votes. However, many voters may have the same vote and it is natural to represent them in a succinct way, where the voters are represented by the distinct votes cast and their corresponding counts. Though this representation can be exponentially smaller, we find in Chapter 5 that the computational complexity of different election-attack problems rarely increases from the nonsuccinct to the succinct case, which is in contrast to the case of unweighted to weighted voters. We explain this behavior by showing that several common proof techniques that show that election-attack problems are in P can be adapted for the case of succinct votes. These results previously appeared in Fitzsimmons and Hemaspaandra [FH17] and in its corresponding technical report [FH16b].

So far we have discussed different natural models that concern the voters in an election, but the model of the election-attack itself is important to consider as well.

In Chapter 6 we discuss a very natural generalization of the standard attack models. Specifically, we consider electoral control in the setting where there are also manipulators. We consider both the case when the chair and the manipulators have the same preferred outcome and cooperate to achieve their goal, and the competitive case where they have directly conflicting goals. We find that the order of who goes first, the chair or the manipulators, can have interesting effects on the complexity. We provide general upper bounds for each of the standard control actions that hold for every election system with a polynomial-time winner problem, with some of these bounds at the second and third levels of the polynomial hierarchy, which we show to be tight by constructing (admittedly artificial) election systems for which control with manipulation is complete for the corresponding case. In contrast to these high upper bounds, we find that for the important natural systems of plurality, Condorcet, and approval, that the complexity of control problems whose without-manipulators complexity is in P does not increase in the setting with manipulators. These results previously appeared in Fitzsimmons, Hemaspaandra, and Hemaspaandra [FHH13a] and its corresponding technical report [FHH13b].

1.4 List of Contributions

We briefly summarize the main contributions of this thesis below.

- We consider the natural model of allowing voters to state varying amounts of ties in their preferences, and examine how this can affect the computational complexity of election-attack problems. (Chapter 3)
- We show that for natural election systems, allowing votes with ties can both increase and decrease the complexity of bribery, and we state a general result on the effect of votes with ties on the complexity of control. (Chapter 3)
- We consider the four most natural models of single-peakedness for votes with ties and show that for each model it is in P to determine when a given collection of votes satisfies that model. (Chapter 4)
- We expand our results on the complexity of manipulation for votes with ties by considering the complexity of single-peaked votes with ties, and find that the complexity can depend on the model used. (Chapter 4)
- We consider how the succinct representation of the voters can affect the complexity of different election problems. Even though the succinct representation can be exponentially smaller than the nonsuccinct, we find that the complexity of election attacks (in the length of the input) rarely increases, and explain this behavior by showing how to adapt different techniques for showing election problems to be in P from the nonsuccinct to the succinct case. (Chapter 5)
- We find one natural case where the complexity increases when moving from the nonsuccinct to the succinct representation of the voters, namely the complexity of winner determination for Kemeny elections. (Chapter 5)
- We model the setting of control attacks on elections in which there are manipulators. We consider both the case where the chair and the manipulators have the same goal and where they have directly conflicting goals. (Chapter 6)
- We prove upper bounds for our model of control in the presence of manipulators that hold over every election system with a polynomial-time winner problem and show that most of these bounds are tight for the second and third levels of the polynomial hierarchy. (Chapter 6)

- We show for the important election systems approval, Condorcet, and plurality that the complexity of control in the presence of manipulators can be much lower than those upper bounds, even falling as low as polynomial time. (Chapter 6)

1.5 List of Publications

These are the publications that form much of the material in this thesis.

- Z. Fitzsimmons and E. Hemaspaandra. The Complexity of Succinct Elections. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (Student Abstract)*, pages 4921–4922, February 2017.
- Z. Fitzsimmons, E. Hemaspaandra, and L. Hemaspaandra. Manipulation Complexity of Same-System Runoff Elections. *Annals of Mathematics and Artificial Intelligence*, 77(3–4): 159–189, 2016.
- Z. Fitzsimmons and E. Hemaspaandra. Modeling Single-Peakedness for Votes with Ties. In *Proceedings of the 8th European Starting AI Researcher Symposium*, pages 63–74, August 2016.
 - Also appears in *Workshop Notes of the 10th Workshop on Advances in Preference Handling*, July 2016.
- Z. Fitzsimmons and E. Hemaspaandra. Complexity of Manipulative Actions When Voting with Ties. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 103–119, September 2015.
 - Also appears in *Workshop Notes of the 6th International Workshop on Computational Social Choice*, June 2016.
- Z. Fitzsimmons. Single-Peaked Consistency for Weak Orders Is Easy. In *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 127–140, June 2015.
- Z. Fitzsimmons. Realistic Assumptions for Attacks on Elections. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (Doctoral Consortium)*, pages 4235–4236, January 2015.
- Z. Fitzsimmons, E. Hemaspaandra, and L. Hemaspaandra. Control in the Presence of Manipulators: Cooperative and Competitive Cases. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 113–119, August 2013.
 - Also appears in *Workshop Notes of the 5th International Workshop on Computational Social Choice*, June 2014.

Chapter 2

Preliminaries

2.1 Elections and Preferences

An *election* is a pair (C, V) where C is a finite set of candidates and V is a finite collection of voters (a preference profile). Each voter in an election has a corresponding vote (preference order) over the set of candidates. In most of the computational study of elections, votes are assumed to be a *total order*, i.e., a strict ordering of the candidates from most to least preferred. Formally, a total order is a complete, reflexive, transitive, and antisymmetric binary relation. We use “ $>$ ” to denote strict preference between two candidates, e.g., given the candidate set $\{a, b, c\}$ a vote could be $(b > a > c)$, which means that b is strictly preferred to a , b to c , and a to c .

In this thesis we also consider other types of preference orders, and it will be clear from context when we use each type of preference order.

In Chapters 3 and 4 we consider voters with varying amounts of ties in their preferences; the most general being weak orders. (See Example 2.1.1 for an example of each of the preference orders we consider.) A *weak order* is a total order without antisymmetry. So, a voter with weak-order preferences can state transitive indifference (“ \sim ”) among the candidates, in addition to strict preference. We use “ \sim ” to denote the indifference relation. In general, a weak order can be viewed as a total order with ties. So we sometimes refer to weak orders as votes with ties, and informally refer to indifference as ties throughout this thesis.

We consider two restrictions to weak orders: top orders and bottom orders. A *top order* is a weak order with all tied candidates ranked last. Similarly, a *bottom order* is a weak order with all tied candidates ranked first. We also will sometimes discuss partial orders, where a partial order over a set of candidates is transitive, reflexive, and antisymmetric.

In Example 2.1.1 below we present an example of each of the four preference orders described above.

Example 2.1.1 *Given the candidate set $\{a, b, c, d\}$, $(a > b \sim c > d)$ is a weak order, $(a \sim b > c > d)$ is a bottom order, $(a > b > c \sim d)$ is a top order, $(a > b)$ is a partial order, and $(a > b > c > d)$ is a total order. Notice that every weak order is also a partial order, every bottom order and every top order is also a weak order and partial order, and that every total order is also a top order, bottom order, weak order, and partial order.*

Many of our results are for weighted elections, where each voter has an associated positive integral weight, and a voter with weight ω counts as a coalition of ω unweighted voters that all vote the same. Weighted elections are a very natural scenario for the real-world use of elections. For example, in an election among shareholders for a given company, the weight of a shareholder's vote may correspond to the number of shares that she holds.

2.2 Election Systems

An *election system* (an election rule), \mathcal{E} , is a mapping from an election to a set W , referred to as the winner(s), where W can be any subset of the candidate set. This is referred to as the nonunique winner model. In the unique winner model the winner of an election must be a single candidate. When it is not otherwise specified we use the nonunique winner model.

The winner problem for an election system \mathcal{E} is defined by the following decision problem.

Name: \mathcal{E} -winner

Given: An election (C, V) and a candidate $p \in C$.

Question: Is p a winner of the election (C, V) using election system \mathcal{E} ?

One reasonable computational property that we can require from an election system is that it is computationally easy to determine the winner, i.e., the winner problem is in P. However, there exist election systems with desirable properties in social choice theory, e.g., the Kemeny rule [Kem59, KS60], which have an NP-hard winner problem [BTT89b] and thus is not in P (unless $P = NP$). Note that with the exception of the results in Section 5.3, we consider election systems with polynomial-time winner problems.

We now discuss three important families of election systems, the first being scoring rules.

2.2.1 Scoring Rules

A *scoring rule* denotes a set of scoring vectors of the form $\langle s_1, \dots, s_m \rangle$, where for each i , $s_i \in \mathbb{Q}$ and $s_i \geq s_{i+1}$, and when given an election with m candidates, uses the vector of length m to determine each candidate's score. When the preferences are all total orders,

a candidate at position i in the preference order of a voter receives a score of s_i from that voter. The candidate(s) with the highest total score win. Four important scoring rules are plurality, veto, Borda, and t -Approval.

Plurality: with scoring vector $\langle 1, 0, \dots, 0 \rangle$.

Veto: with scoring vector $\langle 1, \dots, 1, 0 \rangle$.

Borda: with scoring vector $\langle m-1, m-2, \dots, 1, 0 \rangle$.

t -Approval: with scoring vector $\langle \underbrace{1, \dots, 1}_t, 0, \dots, 0 \rangle$.

2.2.2 Extensions for Votes with Ties

To use a scoring rule to determine the outcome of an election containing votes with ties we must extend the definition of scoring rules given in the previous section. The scoring-rule extensions for weak orders defined below generalize the extensions introduced by Baumeister et al. [BFLR12] and by Narodytska and Walsh [NW14] which in turn generalizes extensions used for the Borda count (see [Eme13] for a discussion of such extensions).

Write a preference order with ties as $G_1 > G_2 > \dots > G_r$ where each G_i is a set of tied candidates. For each set G_i , let $k_i = \sum_{j=1}^{i-1} \|G_j\|$ be the number of candidates strictly preferred to every candidate in the set. See the caption of Table 2.1 for an example.

We now introduce the following scoring-rule extensions, which as stated above, generalize previously used scoring-rule extensions [BFLR12, NW14]. In Table 2.1 we present an example of each of these extensions for Borda.

Min: Each candidate in G_i receives a score of $s_{k_i + \|G_i\|}$.

Max: Each candidate in G_i receives a score of s_{k_i+1} .

Round-down: Each candidate in G_i receives a score of s_{m-r+i} .

Average: Each candidate in G_i receives a score of

$$\frac{\sum_{j=k_i+1}^{k_i+\|G_i\|} s_j}{\|G_i\|}.$$

The optimistic and pessimistic models from the work by Baumeister et al. [BFLR12] are the same as our max and min extensions respectively, for top orders. All of the scoring-rule extensions for top orders found in the work by Narodytska and Walsh [NW14] can be

Borda	$score(a)$	$score(b)$	$score(c)$	$score(d)$
Min	3	1	1	0
Max	3	2	2	0
Round-down	2	1	1	0
Average	3	1.5	1.5	0

Table 2.1: The score of each candidate for preference order ($a > b \sim c > d$) using Borda with each of our scoring-rule extensions. We write this order as $\{a\} > \{b, c\} > \{d\}$, i.e., $G_1 = \{a\}$, $G_2 = \{b, c\}$, and $G_3 = \{d\}$. Note that $k_1 = 0$, $k_2 = 1$, and $k_3 = 3$.

realized by our definitions above, with our round-down and average extensions yielding the same scores for top orders as their round-down and average extensions. With the additional modification that $s_m = 0$ our min scoring-rule extension yields the same scores for top orders as round-up in the work by Narodytska and Walsh [NW14].

2.2.3 Pairwise Rules


In addition to scoring rules, election systems can be defined by the pairwise majority elections between the candidates, i.e., for a pair of candidates $a, b \in C$ a beats b by majority if more than half of the voters state $a > b$.

One important example is the Copeland rule [Cop51], where given an election, each candidate receives one point for each pairwise majority election she wins and receives 0.5 points for each tie. The Copeland rule was later parameterized by Faliszewski et al. [FHHR09] as Copeland $^\alpha$ (where α is a rational number between 0 and 1), and instead of each candidate receiving 0.5 points for each tie, they receive α points. We mention here that it was somewhat recently discovered that an election system that is the same as Copeland 1 was proposed in the thirteenth-century by Ramon Llull, a Catalan mystic and philosopher (see [HP01]). So, as is now common, we refer to Copeland 1 as Llull.

As with the case of scoring rules, we also consider the case of Copeland $^\alpha$ elections for votes with ties. We extend the definition for Copeland $^\alpha$ in the obvious way (i.e., $a > b$ by majority if more voters state $a > b$ than $b > a$), as was done for the case of top orders by Baumeister et al. [BFLR12] and Narodytska and Walsh [NW14].


When discussing elections defined by pairwise majority elections we sometimes refer to the *induced majority graph* of an election. An election (C, V) can be represented by an induced majority graph where each candidate in the election corresponds to a vertex in the graph, and for every pair of candidates $a, b \in C$ the graph contains the edge $a \rightarrow b$ if $a > b$ by majority. See Example 2.2.1 for an example of an election and its corresponding induced majority graph.

Example 2.2.1 Given the election (C, V) where $C = \{a, b, c, d\}$ and the collection of votes, V , consists of one voter voting $(a > b > c > d)$, one voter voting $(c > b > a > d)$, and one voter voting $(a > c > b > d)$, we have the following induced majority graph.



When discussing pairwise election systems it is important to mention the notion of the Condorcet winner of an election, i.e., the candidate that beats every other candidate pairwise [Con85]. Notice that a is the Condorcet winner in Example 2.2.1, since for each candidate $b \in C - \{a\}$ there is the edge $a \rightarrow b$ in the induced majority graph. The Condorcet winner of an election seems like a good choice for an outcome that represents the electorate; unfortunately a Condorcet winner is not guaranteed to exist for a given election. And it is not difficult to construct an election that realizes this situation. In Example 2.2.2 we present such an example. There is also the notion of a weak Condorcet winner, which ties-or-beats every other candidate pairwise. So, in contrast to the case of the Condorcet winner, an election can have multiple weak Condorcet winners, though it is possible to have neither.

Example 2.2.2 Given the election (C, V) where $C = \{a, b, c, d\}$ and the collection of votes, V , consists of one voter voting $(b > a > c > d)$, one voter voting $(c > b > a > d)$, and one voter voting $(a > c > b > d)$, we have the following induced majority graph, which contains a cycle, and thus does not have a Condorcet winner or even a weak Condorcet winner.



An election system that elects exactly the Condorcet winner when one exists is called Condorcet consistent (or weak Condorcet consistent for the case of electing exactly the weak Condorcet winner(s) when they exist). Copeland ^{α} is an example of a Condorcet consistent election system, and many election systems based on pairwise comparisons are as well. The notion of the Condorcet winner of an election was used by Bartholdi, Tovey, and Trick in their work that introduces electoral control [BTT92] to define an election system, called Condorcet, that elects the Condorcet winner when one exists, and otherwise no one wins.

2.2.4 Approval Voting

Approval voting is an election system that is somewhat differently structured than the two families stated above. In approval voting, instead of each voter voting a total-order preference, each voter votes a 0-1 vector of length $\|C\|$ and indicates approval (1) or disapproval (0) for each candidate, and the candidates with the most approvals win [BF83].

2.3 Manipulative Attacks

A major research direction discussed in this thesis is the computational study of how hard different manipulative attacks are for a given election system. In this section we define the three main families of manipulative attacks: manipulation, bribery, and control. We mention here that in each of the manipulative attacks discussed we assume that the manipulative agent(s) (either the manipulators, the chair, or the briber) has complete information of the preferences of the voters.¹ We briefly describe each general family of attacks and present a formal definition of each below. Note that we define each of these problems for the nonunique winner model (our standard model). For the unique winner model it will generally be enough to change instances of “a winner” to “a unique winner” in the definitions.

In each of our manipulative attacks we formally define the constructive unweighted cases, where “constructive” means that the preferred outcome of the strategic agent(s) is to ensure that a preferred candidate wins, and “unweighted” means that we are considering elections where voters do not have corresponding weights. In the corresponding destructive cases the preferred outcome of the strategic agent(s) is to ensure that a despised candidate does not win, and in the corresponding weighted cases each voter has a corresponding weight.

2.3.1 Manipulation

The manipulation problem, first introduced by Bartholdi, Tovey, and Trick [BTT89a], asks when given an election, a manipulator, and a preferred candidate, if the manipulator can set her vote to ensure that her preferred candidate wins. It is reasonable to assume that in an election, especially when there are many voters, that there are multiple manipulators who act as a coalition. Conitzer, Sandholm, and Lang [CSL07] introduced the coalitional manipulation problem, which we formally define below.

¹This may not seem to fit our general focus of natural models, but since many of our results show NP-hardness, considering the case where the manipulative agent(s) has complete information strengthens our results; we can show hardness for a setting with partial information by a reduction from the complete-information setting.

Name: \mathcal{E} -Constructive Unweighted Coalitional Manipulation (CUCM)

Given: An election (C, V) , a collection of manipulative voters W , and a preferred candidate $p \in C$

Question: Is there a way to set the votes of the manipulators such that p is a winner of the election $(C, V \cup W)$ using election system \mathcal{E} ?

We will also consider the cases of destructive unweighted coalitional manipulation (DUCM) and the cases of manipulation for weighted elections (CWCM and DWCM), which were both introduced by Conitzer, Sandholm, and Lang [CSL07].

2.3.2 Control

Electoral control is the problem of determining if it is possible for an election organizer with control over the structure of an election, whom we refer to as the election chair, to ensure a preferred outcome. This preferred outcome can either be ensuring that a preferred candidate wins (the constructive case [BTT92]) or that a despised candidate loses (the destructive case [HHR07]).

The standard models of control can be split into two groups. The first being the nonpartition control types where the election chair can control the election by adding or deleting the voters or candidates. The second being the partition control types where either the candidates or voters are partitioned and subelections are held restricted to these partitions before a final runoff is held among the candidates that survive. In these cases we consider different models for what it means to survive, the ties-eliminate (TE) model and the ties-promote (TP) model. In the ties-eliminate (TE) model, only a unique winner of a subelection proceeds to the runoff (if multiple candidates tie as winners then no candidates proceed to the runoff), and in the ties-promote (TP) model all of the winning candidates in a subelection proceed to the runoff.²

We formally define the constructive versions of each of the standard control actions below.

Name: \mathcal{E} -Constructive Control by Adding Candidates (CCAC)³

Given: An election (C, V) , a set of unregistered candidates D , a preferred candidate $p \in C$, and a limit $k \in \mathbb{N}$.

²Recent work by Hemaspaandra, Hemaspaandra, and Menton shows that in the nonunique winner model two pairs of the standard control models collapse. Specifically, the models of destructive control by partitioning candidates and destructive control by runoff partitioning candidates, in each of the tie-breaking models [HHM13].

³Note that the seminal control paper by Bartholdi, Tovey, and Trick [BTT92] defines control by adding candidates without a limit k . This is sometimes referred to as control by unlimited adding candidates. We use the now more commonly used model where there is a limit to the number of candidates that the election chair can add.

Question: Does there exist a subset of the unregistered candidates $D' \subseteq D$ such that $\|D'\| \leq k$ and p is a winner of the election $(C \cup D', V)$ using election system \mathcal{E} ?

Name: \mathcal{E} -Constructive Control by Deleting Candidates (CCDC)

Given: An election (C, V) , a preferred candidate $p \in C$, and a limit $k \in \mathbb{N}$.

Question: Does there exist a subset of the candidates $C' \subseteq C$ such that $\|C'\| \leq k$ and p is a winner of the election $(C - C', V)$ using election system \mathcal{E} ?

Name: \mathcal{E} -Constructive Control by Adding Voters (CCAV)

Given: An election (C, V) , a collection on unregistered voters U , a preferred candidate $p \in C$, and a limit $k \in \mathbb{N}$.

Question: Does there exist a subcollection of the unregistered voters $U' \subseteq U$ such that $\|U'\| \leq k$ and p is a winner of the election $(C, V \cup U')$ using election system \mathcal{E} ?

Name: \mathcal{E} -Constructive Control by Deleting Voters (CCDV)

Given: An election (C, V) , a preferred candidate $p \in C$, and a limit $k \in \mathbb{N}$.

Question: Does there exist a subcollection of the voters $V' \subseteq V$ such that $\|V'\| \leq k$ and p is a winner of the election $(C, V - V')$ using election system \mathcal{E} ?

Name: \mathcal{E} -Constructive Control by Partitioning Candidates (CCPC)

Given: An election (C, V) and a preferred candidate $p \in C$.

Question: Does there exist a partition of the candidates (C_1, C_2) such that p is a winner of the runoff election between the winners of (C_1, V) under the given tie-handling model and the candidates in C_2 , all using election system \mathcal{E} ?

Name: \mathcal{E} -Constructive Control by Runoff Partitioning Candidates (CCRPC)

Given: An election (C, V) and a preferred candidate $p \in C$.

Question: Does there exist a partition of the candidates (C_1, C_2) such that p is a winner of the runoff election that consists of the winners of (C_1, V) and (C_2, V) under the given tie-handling model, all using election system \mathcal{E} ?

Name: \mathcal{E} -Constructive Control by Partitioning Voters (CCPV)

Given: An election (C, V) and a preferred candidate $p \in C$.

Question: Does there exist a partition of the voters (V_1, V_2) such that p is a winner of the runoff election that consists of the winners of (C, V_1) and (C, V_2) under the given tie-handling model, all using election system \mathcal{E} ?

We will also consider the corresponding destructive and weighted cases of these control problems.

For the destructive cases, simply change each instance of “preferred candidate” to “despised candidate,” and “ p is a winner” to “ p is not a winner.”

For the weighted cases of partitioning and candidate control it is clear to see how weighted elections are taken into account. It is not as clear how to adapt the definitions of voter control, since we could interpret the limit (in either adding or deleting voters) as the number of voters to add/delete *or* the total vote weight to add/delete. We use the former, where the limit remains the number of voters, and this is precisely the definition used by Faliszewski, Hemaspaandra, and Hemaspaandra [FHH15].

2.3.3 Bribery

Bribery is the problem of determining if it is possible to change the votes of a subcollection of the voters, within a given limit, to ensure that a preferred candidate wins [FHH09].

Name: \mathcal{E} -Bribery

Given: A candidate set C , a collection of voters V , a preferred candidate $p \in C$, and a limit $k \in \mathbb{N}$.

Question: Is there a way to change the votes of at most k of the voters in V so that p is a winner under election system \mathcal{E} ?

In the corresponding problem of weighted bribery where each voter has an associated positive integer weight, the limit still denotes the number of voters to bribe. Though none of our results use prices, we mention here that bribery is also considered for the case where voters have an associated price, where the limit is then the total budget used [FHH09].

2.4 Computational Complexity

The focus of this thesis is on computational complexity results. Most of our results will concern the classes P and NP. We assume that the reader is familiar with these two classes, and the notions of what it means to be “hard” and what it means to be “complete” for a given class. We also assume that the reader is familiar with polynomial-time many-one

reductions (\leq_m^p), and since all of the reductions in this thesis are polynomial-time many-one reductions, we often simply refer to them as reductions. (See [GJ79] for a general introduction to P, NP, and polynomial-time many-one reductions.)

As is standard in computational complexity theory, the election attacks defined above are decision problems, and our results generally concern whether a given election system is easy to attack, i.e., it is in P to determine if the attack is possible, or if it is difficult to attack, i.e., it is NP-hard to determine if the attack is possible. Most of our results that show problems to be NP-hard also show them to be NP-complete, where membership in NP is generally trivial to show for these problems.

Our proofs of NP-hardness are all due to reductions from NP-complete problems such as Partition and Exact Cover by 3-Sets, which we define below. We have some results that use variants of these problems, but we will introduce them locally to where they are used.

Name: Partition [Kar72]

Given: Given a nonempty set of positive integers $\{k_1, \dots, k_t\}$ such that $\sum_{i=1}^t k_i = 2K$.

Question: Does there exist a subset A of $\{k_1, \dots, k_t\}$ such that $\sum A = K$?

Name: Exact Cover by 3-Sets [Kar72]

Given: Given a set $B = \{b_1, \dots, b_{3k}\}$, and a collection $\mathcal{S} = \{S_1, \dots, S_n\}$ of three-element subsets of B .

Question: Does there exist a subcollection \mathcal{S}' of \mathcal{S} such that every element of B occurs in exactly one member of \mathcal{S}' ?

Our polynomial-time results range from simple greedy approaches, to reductions to more complex algorithms such as network flow and matrix permutation problems, and such problems will be defined locally for the results that use them.

In Chapters 5 and 6 we discuss results that include completeness results for higher classes, that are widely believed to be strictly larger than NP. Most notably, the classes Θ_2^p , Δ_2^p , NP^{NP} , coNP^{NP} , and $\text{coNP}^{\text{NP}^{\text{NP}}}$ in the polynomial hierarchy. (We introduce these classes here, but Chapters 5 and 6 each contain further discussion of the classes that appear in their results.)

The polynomial hierarchy is defined by the classes Σ_k^p , Π_k^p , and Δ_k^p , where $\Delta_0^p = \Pi_0^p = \Sigma_0^p = \text{P}$ and, for all $k \geq 0$, $\Sigma_{k+1}^p = \text{NP}^{\Sigma_k^p}$, $\Pi_{k+1}^p = \text{coNP}^{\Sigma_k^p}$, and $\Delta_{k+1}^p = \text{P}^{\Delta_k^p}$, where the superscript is used to denote access to an oracle for a set of one's choice from that class. [MS72, Sto76]. In Chapter 6, we use the commonly used notation $\text{NP}^{\text{NP}} = \Sigma_2^p$, $\text{coNP}^{\text{NP}} = \Pi_2^p$, and $\text{coNP}^{\text{NP}^{\text{NP}}} = \Pi_3^p$ to refer to these classes.

The class $\Theta_2^p = P^{\text{NP}^{\lceil \log \rceil}}$ was first studied by Papadimitriou and Zachos [PZ83], and denotes the class of problems solvable by a P-machine that can ask $O(\log n)$ queries to an NP oracle. Hemachandra showed that this class is equivalent to $P_{\parallel}^{\text{NP}}$, the class of problems solvable by a P-machine that can ask one round of parallel queries to an NP oracle [Hem89].

Note the following relationships between these classes:

$$P \subseteq NP \cap \text{coNP} \subseteq NP \cup \text{coNP} \subseteq \Theta_2^p \subseteq \Delta_2^p \subseteq \Sigma_2^p \cap \Pi_2^p \subseteq \Sigma_2^p \cup \Pi_2^p \subseteq \Pi_3^p.$$

There are far fewer completeness results for higher levels of the polynomial hierarchy than for NP (see [SU02a, SU02b]), and it is particularly interesting to find natural problems at such high levels of complexity.

The computational complexity classes mentioned above concern the worst-case time complexity of a given problem. A given NP-hard problem may have many more easy instances than hard ones, and in practice a heuristic algorithm may be able to perform quite well. However, there are known theoretical limits to the performance of heuristics for NP-hard problems. A recent survey by Hemaspaandra and Williams [HW12] discusses such limitations and shows that, due to the work by Buhrman and Hitchcock [BH08] and Cai et al. [CCHO05], no polynomial-time heuristic algorithm can err on only subexponentially many instances—unless the polynomial hierarchy collapses.

Usually not all instances of a problem are equally likely, so the distribution of the instances must be taken into account to move beyond worst-case analysis to the average case. The notion of average-case complexity due to Levin [Lev86] takes the distribution into account, but this is difficult to work with and few problems are complete for the class.

Some recent work on election attacks has studied the performance of heuristics to examine how hard these problems are in practice. Most of the empirical study of the hardness of election attacks follows a similar design as the influential work by Walsh [Wal11] by examining the performance of algorithms for an election-attack problem for elections with votes sampled from either theoretical distributions such as the impartial culture model [GK68] or from real-world data, e.g., the datasets available on PREFLIB [MW13]. The analysis done in these studies has generally consisted of descriptive statistics, e.g., the observed runtime of an algorithm for the instances in the experiment, and some patterns and trends can be discussed using such data, which can be valuable to motivate theoretical study. For example, Walsh observed a smooth phase transition for manipulation of STV and veto [Wal11], and Mossel, Procaccia, and Racz later proved this for independent and identically distributed votes [MPR13].

The rigorous study of the hardness of problems beyond worst-case complexity is a valuable

and very difficult direction of research. But even showing that a problem is NP-hard gives strong insight into the structure of the problem, and shows that polynomial-time algorithms for the problem do not exist—unless $P = NP$.

Chapter 3

Models for the Votes: Votes with Ties

3.1 Introduction

In this chapter we consider the effect of different generalizations of total-order preferences on the complexity of the election attacks of manipulation, bribery, and control. The majority of this chapter will focus on preferences that allow voters to state ties among the candidates in addition to strict preference, but we will also consider the case where the preferences of the voters do not need to be transitive, which is referred to as the case of irrational votes.

The computational study of the problems of manipulation, control, and bribery has largely been restricted to elections where the voters have tie-free votes. Recent work by Narodytska and Walsh [NW14] (see also [ML15]) studies the computational complexity of the manipulation problem for top orders, i.e., votes where ties are allowed, but only among a voter's least-preferred candidates and are otherwise tie free. Some of our manipulation results can be seen as expanding on the manipulation results from Narodytska and Walsh [NW14] by solving open questions and considering more general types of votes with ties.

It is important that we understand the complexity of election problems for elections that allow votes with ties, since in practical settings voters often have ties between some of the candidates. For example, the online preference repository PREFLIB contains several election datasets containing votes with ties [MW13], and it is natural to allow agents to state votes with ties when they have utilities over a set of candidates.

Election systems in use are sometimes defined for votes with ties. For example, both the Kemeny rule [Kem59] and the Schulze rule [Sch11] are defined for votes that contain ties. Also, there exist variants of the Borda count that are defined for top-order votes (see [Eme13]). As described in Section 2.2.2, when necessary, we extend the definitions of election systems to properly handle votes with ties.

Note that we are not considering the case of incomplete votes, but rather votes with ties.

It is tempting to want to break ties in a voter’s preferences when they have votes with ties, but it is very natural for a voter to have equal preference among some of the candidates and so they should not be forced to state a tie-free vote. We mention in passing the line of work on incomplete preferences that seeks to determine the “best extension” of a voter’s preferences and use this to determine the winner of a given election (see, e.g., [XC11]).

For the computational study of manipulation, we are the first to consider orders that allow a voter to state ties at each position of her preference order, i.e., weak orders. In contrast to the work by Narodytska and Walsh [NW14], we give an example of a natural case where manipulation becomes hard when given votes with ties, while it is in P for tie-free votes. And we are the first to study the complexity of the standard models of control and bribery for votes with ties. However, we mention here that Baumeister et al. [BFLR12] consider a different version of bribery called extension bribery, for top orders (there called top-truncated votes) [BFLR12].

This chapter is organized as follows. Our results are distributed among three sections, each of which deals with a different behavior of votes with ties. In Section 3.2 we consider cases where the complexity of an election attack can increase when moving from votes without ties to votes with ties. In Section 3.3 we present examples where the complexity decreases for votes with ties (and state a general observation on control). And in Section 3.4 we present cases where the complexity remains the same, a general result on two-voter majority graphs, and discuss axioms that our scoring-rule extensions defined in Section 2.2.2 satisfy. We discuss related work in Section 3.6 and our general conclusions and some directions for future work in Section 3.7.

We now present our results, many of which are about scoring rules. In Section 2.2.2 we presented the extensions min, max, round-down, and average to properly handle votes with ties, which extend scoring rules to score “groups” of tied candidates instead of each candidate individually. To recall the definitions of each of our scoring-rule extensions, we repeat the example from Table 2.1. Consider the candidate set $\{a, b, c, d\}$ and the vote $(a > b \sim c > d)$. We show the score assigned to each candidate using Borda ($\langle 3, 2, 1, 0 \rangle$) using each of our scoring-rule extensions.

Borda using min: $score(a) = 3$, $score(b) = score(c) = 1$, and $score(d) = 0$.

Borda using max: $score(a) = 3$, $score(b) = score(c) = 2$, and $score(d) = 0$.

Borda using round-down: $score(a) = 2$, $score(b) = score(c) = 1$, and $score(d) = 0$.

Borda using average: $score(a) = 3$, $score(b) = score(c) = 1.5$, and $score(d) = 0$.

3.2 Complexity Goes Up

The related work on the complexity of manipulation for top orders [NW14] did not find a natural case where manipulation complexity increases when moving from total orders to top orders. We found such cases for Constructive Weighted Coalitional Manipulation (CWCM) by considering votes that are single-peaked (in the possibly single-peaked model), so the proof of the following theorem is deferred to Chapter 4, namely the proofs of Theorems 4.4.2, 4.4.3, and 4.4.4, where we discuss models of single-peakedness.

Theorem 3.2.1 *The complexity of 3-candidate Borda CWCM for possibly single-peaked preferences goes from P to NP-complete for top orders using max, round-down, and average.*

We now present cases where we observe an increase in the computational complexity of control and of bribery when moving from tie-free votes to votes with ties.

Consider the complexity of Constructive Control by Adding Voters (CCAV), which asks if the election chair can ensure that a preferred candidate wins by adding voters to the election. (See Section 2.3.2 for a formal definition.) This problem is known to be in P for plurality for total orders [BTT92].

Theorem 3.2.2 [BTT92] *Plurality CCAV for total orders is in P.*

Plurality using max for bottom orders is essentially the same as approval voting (where each voter indicates either approval or disapproval of each candidate and the candidate(s) with the most approvals win). For example, given the set of candidates $\{a, b, c, d\}$, an approval vector that approves of a and c , and a preference order $(a \sim c > b > d)$ yield the same scores for approval and plurality using max respectively. So the theorem below immediately follows from the proof of Theorem 4.43 from Hemaspaandra, Hemaspaandra, and Rothe [HHR07], which shows that CCAV for approval voting is NP-complete, and so we see an increase in the complexity with respect to the case of total orders.

Theorem 3.2.3 *Plurality CCAV for bottom orders (and thus also for weak orders) using max is NP-complete.*

We now show that the case of plurality for bottom orders and weak orders using average is NP-complete.

Theorem 3.2.4 *Plurality CCAV for bottom orders (and thus also for weak orders) using average is NP-complete.*

Proof. Let $B = \{b_1, \dots, b_{3k}\}$ and a collection $\mathcal{S} = \{S_1, \dots, S_n\}$ of 3-element subsets of B where each $S_j = \{b_{j_1}, b_{j_2}, b_{j_3}\}$ be an instance of Exact Cover by 3-Sets, which asks if there exists a subcollection \mathcal{S}' of \mathcal{S} such that each $b \in B$ occurs in exactly one member of \mathcal{S}' . Without loss of generality let k be divisible by 4 and let $\ell = 3k/4$. We construct the following instance of control by adding voters.

Let the candidates $C = \{p\} \cup B$. Let the addition limit be k . Let the collection of registered voters consist of the following $(3k^2 + 9k)/4 + 1$ voters. (When “ \dots ” appears at the end of a vote here, the remaining candidates from C are ranked lexicographically. For example, given the candidate set $\{a, b, c, d\}$, the vote $(b > \dots)$ denotes the vote $(b > a > c > d)$.)

- For each i , $1 \leq i \leq \ell$, $k + 3$ voters voting $(b_i \sim b_{i+\ell} \sim b_{i+2\ell} \sim b_{i+3\ell} > \dots)$.
- One voter voting $(p > \dots)$.

Let the collection of unregistered voters consist of the following n voters.

- For each $S_j \in \mathcal{S}$, one voter voting $(p \sim b_{j_1} \sim b_{j_2} \sim b_{j_3} > \dots)$.

Notice that from the registered voters, the score of each b_i candidate is $(k - 1)/4$ greater than the score of p . Thus the chair must add voters from the collection of unregistered voters so that no b_i candidate receives more than $1/4$ more points, while p must gain $k/4$ points. Therefore the chair must add the voters that correspond to an exact cover. \square

The complexity of bribery for plurality also goes from P for total orders to NP-complete for votes with ties.

Theorem 3.2.5 [FHH09] *Unweighted bribery for plurality for total orders is in P.*

The proof that bribery for plurality for bottom orders and weak orders using max is NP-complete immediately follows from the proof of Theorem 4.2 from Faliszewski, Hemaspaandra, and Hemaspaandra [FHH09], which showed bribery for approval to be NP-complete.

Theorem 3.2.6 *Unweighted bribery for plurality for bottom orders and weak orders using max is NP-complete.*

3.3 Complexity Goes Down

Narodytska and Walsh [NW14] show that the complexity of coalitional manipulation can go down when moving from total orders to top orders. In particular, they show that

the complexity of coalitional manipulation (weighted or unweighted) for Borda goes from NP-complete to P for top orders using min (which they refer to as round-up). This is because when using min an optimal manipulator vote is to put p first and have all other candidates tied for last.

In contrast, notice that the complexity of a (standard) control action cannot decrease when more lenient votes are allowed. This is because the votes that create hard instances of control are still able to be cast when more general votes are possible. The election chair is not able to directly change votes, except in a somewhat restricted way in candidate control cases, but it is clear to see how this does not affect the statement below.¹

Observation 3.3.1 *If a (standard) control problem is hard for a type of vote with ties, it remains hard for votes that allow more ties.*

What about bribery? Bribery can be viewed as a two-phase action consisting of control by deleting voters followed by manipulation. Hardness for a bribery problem is typically caused by hardness of the corresponding deleting voters problem or the corresponding manipulation problem. If the deleting voters problem is hard, this problem remains hard for votes that allow ties, and it is likely that the bribery problem remains hard as well. Our best chance of finding a bribery problem that is hard for total orders and easy for votes with ties is a problem whose manipulation problem is hard, but whose deleting voters problem is easy. Such problems exist, e.g., all weighted m -candidate t -approval systems except plurality and triviality.²

Theorem 3.3.2 [FHH09] *Weighted bribery for m -candidate t -approval for all $t \geq 2$ and $m > t$ is NP-complete.*

For m -candidate t -approval elections (except plurality and triviality) the corresponding weighted manipulation problem was shown to be NP-complete by Hemaspaandra and Hemaspaandra [HH07] and the corresponding deleting voters problem was shown to be in P by Faliszewski, Hemaspaandra, and Hemaspaandra [FHH15].

Theorem 3.3.3 *Weighted bribery for m -candidate t -approval for weak orders and for top orders using min is in P.*

Proof. To perform an optimal bribery, we cannot simply perform an optimal deleting voter action followed by an optimal manipulation action. For example, if the score of b is already

¹A similar argument is used to explain the relationship between the easiness of control with respect to general and single-peaked votes (see footnote 14 of Brandt et al. [BBHH15]).

²By triviality we mean a scoring rule with a scoring vector that gives each candidate the same score.

at most the score of p , it does not make sense to delete a voter with vote $(b > p \sim a)$. But in the case of bribery, we would change this voter to $(p > a \sim b)$, which could be advantageous.

However, the weighted version of constructive control by deleting voters (CCDV) algorithm from [FHH15] still basically works. Since m is constant, there are only a constant number of different votes possible. And we can assume without loss of generality that we bribe only the heaviest voters of each vote-type and that each bribed voter is bribed to put p first and have all other candidates tied for last. In order to find out if there exists a successful bribery of k voters, we look at all the ways we can distribute this k among the different types of votes. We then manipulate the heaviest voters of each type to put p first and have all other candidates tied for last, and see if that makes p a winner. \square

3.4 Complexity Remains the Same

Narodytska and Walsh [NW14] show that 4-candidate Copeland^{0.5} CWCM remains NP-complete for top orders. They conjecture that this is also the case for 3 candidates and point out that the reduction that shows this for total orders from Faliszewski, Hemaspaandra, and Schnoor [FHS08] won't work. We will prove their conjecture using the following variation of Partition, which we define as Partition' and show to be NP-complete below.

Name: Partition'

Given: A nonempty set of positive even integers $\{k_1, \dots, k_t\}$ and a positive even integer \widehat{K} .

Question: Does there exist a partition (A, B, C) of $\{k_1, \dots, k_t\}$ such that $\sum A = \sum B + \widehat{K}$?

Theorem 3.4.1 *Partition' is NP-complete.*

Proof. The construction here is similar to the first part of the reduction to a different version of Partition from Faliszewski, Hemaspaandra, and Hemaspaandra [FHH09].

Given $\{k_1, \dots, k_t\}$ such that $\sum_{i=1}^t k_i = 2K$, corresponding to an instance of Partition, we construct the following instance $\{k'_1, \dots, k'_t, \ell'_1, \dots, \ell'_t\}, \widehat{K}$ of Partition'. Let $k'_i = 4^i + 4^{t+1}k_i$, $\ell'_i = 4^i$, and $\widehat{K} = 4^{t+1}K + \sum_{i=1}^t 4^i$. (Note that in Faliszewski, Hemaspaandra, Hemaspaandra [FHH09] "3"s were used, but we use "4"s here so that when we add a subset of $\{k'_1, \dots, k'_t, \ell'_1, \dots, \ell'_t, \widehat{K}\}$, we never have carries in the last $t + 1$ digits base 4, and we set the last digit to 0 to ensure that all numbers are even.)

If there exists a partition (A, B, C) of $\{k'_1, \dots, k'_t, \ell'_1, \dots, \ell'_t\}$ such that $\sum A = \sum B + \widehat{K}$, then $\forall i, 1 \leq i \leq t$, $\lfloor (\sum A)/4^i \rfloor \bmod 4 = \lfloor (\sum B + \widehat{K})/4^i \rfloor \bmod 4$. Note that $\lfloor (\sum A)/4^i \rfloor \bmod 4 = \|A \cap \{k'_i, \ell'_i\}\|$, $\lfloor (\sum B)/4^i \rfloor \bmod 4 = \|B \cap \{k'_i, \ell'_i\}\|$, and $\lfloor \widehat{K}/4^i \rfloor \bmod 4 = 1$. So, $\|A \cap$

$\{k'_i, \ell'_i\} = \|B \cap \{k'_i, \ell'_i\}\| + 1$. It follows that exactly one of k'_i or ℓ'_i is in A and neither is in B . Since this is the case for every i , it follows that $B = \emptyset$. Now look at all k_i such that k'_i is in A . That set will add up to K , and so our original Partition instance is a positive instance.

For the converse, it is immediate that a subset D of $\{k_1, \dots, k_t\}$ that adds up to K can be converted into a solution for our Partition' instance, namely, by putting k'_i in A for every k_i in D , putting ℓ'_i in A for every k_i not in D , letting $B = \emptyset$, and putting all other elements of $\{k'_1, \dots, k'_t, \ell'_1, \dots, \ell'_t\}$ in C . \square

We can now use Partition' to prove the following theorem, which in turn proves the conjecture by Narodytska and Walsh [NW14].³

Theorem 3.4.2 *3-candidate Copeland ^{α} CWCM remains NP-complete for top orders, bottom orders, and weak orders, for all rational $\alpha \in [0, 1)$ in the nonunique winner case (our standard model).*

Proof. The proof for bottom-order votes follows from the proof of the case for total-order votes due to Faliszewski, Hemaspaandra, and Schnoor [FHS08]. We prove the top-order case below, and it is clear to see that this proof also holds for weak orders.

Let $\{k_1, \dots, k_t\}, \widehat{K}$ be an instance of Partition', which asks whether there exists a partition (A, B, C) of $\{k_1, \dots, k_t\}$ such that $\sum A = \sum B + \widehat{K}$.

Let k_1, \dots, k_t sum to $2K$ and without loss of generality assume that $\widehat{K} \leq 2K$. We now construct an instance of CWCM. Let the candidate set $C = \{a, b, p\}$ and let the preferred candidate be p . Let there be two nonmanipulators with the following weights and votes.

- One weight $K + \widehat{K}/2$ nonmanipulator voting $(a > b > p)$.
- One weight $K - \widehat{K}/2$ nonmanipulator voting $(b > a > p)$.

From the votes of the nonmanipulators, $score(a) = 2$, $score(b) = 1$, and $score(p) = 0$. In the induced majority graph, there is the edge $a \rightarrow b$ with weight \widehat{K} , the edge $a \rightarrow p$ with weight $2K$, and the edge $b \rightarrow p$ with weight $2K$. Let there be t manipulators with, weights k_1, \dots, k_t .

Suppose that there exists a partition of $\{k_1, \dots, k_t\}$ into (A, B, C) such that $\sum A = \sum B + \widehat{K}$. Then for each $k_i \in A$, have the manipulator with weight k_i vote $(p > b > a)$, for each $k_i \in B$, have the manipulator with weight k_i vote $(p > a > b)$, and for each $k_i \in C$ have the manipulator with weight k_i vote $(p > a \sim b)$. From the votes of the nonmanipulators and manipulators, $score(a) = score(b) = score(p) = 2\alpha$.

³Menon and Larson independently proved the top-order case of the following theorem [ML15].

For the other direction, suppose that p can be made a winner. When all of the manipulators put p first then $score(p) = 2\alpha$ (the highest score that p can achieve). Since $\alpha < 1$, the manipulators must have voted such that a and b tie. This means that a subcollection of the manipulators with weight K voted $(p > b > a)$, a subcollection with weight $K - \widehat{K}$ voted $(p > a > b)$, and a subcollection with weight \widehat{K} voted $(p > a \sim b)$. No other votes would cause b and a to tie. Notice that the weights of the manipulators in the three different subcollections form a partition (A, B, C) of $\{k_1, \dots, k_t\}$ such that $\sum A = \sum B + \widehat{K}$. \square

3-candidate Copeland ^{α} CWCM is unusual in that the complexity can be different if we look at the unique winner case instead of the nonunique winner case (our standard model). We prove that the only 3-candidate Copeland CWCM case that is hard for the unique winner model remains hard using a very similar approach.

Theorem 3.4.3 *3-candidate Copeland⁰ CWCM remains NP-complete for top orders, bottom orders, and weak orders, in the unique winner case.*

Proof. The proof for bottom-order votes follows from the proof of the case for total-order votes due to Faliszewski, Hemaspaandra, and Schnoor [FHS08]. We prove the top-order case below, and it is clear to see that this proof also holds for weak orders.

Let $\{k_1, \dots, k_t\}, \widehat{K}$ be an instance of Partition', which asks whether there exists a partition (A, B, C) of $\{k_1, \dots, k_t\}$ such that $\sum A = \sum B + \widehat{K}$.

Let k_1, \dots, k_t sum to $2K$ and without loss of generality assume that $\widehat{K} \leq 2K$. We now construct an instance of CWCM. Let the candidate set $C = \{a, b, p\}$. Let the preferred candidate be $p \in C$. Let there be two nonmanipulators with the following weights and votes.

- One weight $K + \widehat{K}/2$ nonmanipulator voting $(a > p > b)$.
- One weight $K - \widehat{K}/2$ nonmanipulator voting $(b > a > p)$.

From the votes of the nonmanipulators $score(a) = 2$, $score(b) = 0$, and $score(p) = 1$. The induced majority graph contains the edge $a \rightarrow b$ with weight \widehat{K} , the edge $a \rightarrow p$ with weight $2K$, and the edge $p \rightarrow b$ with weight \widehat{K} . Let there be t manipulators, with weights k_1, \dots, k_t .

Suppose that there exists a partition of $\{k_1, \dots, k_t\}$ into (A, B, C) such that $\sum A = \sum B + \widehat{K}$. Then for each $k_i \in A$ have the manipulator with weight k_i vote $(p > b > a)$, for each $k_i \in B$ have the manipulator with weight k_i vote $(p > a > b)$, and for each $k_i \in C$ have the manipulator with weight k_i vote $(p > a \sim b)$. From the votes of the nonmanipulators and the manipulators $score(p) = 1$ and $score(a) = score(b) = 0$.

For the other direction, suppose that p can be made a unique winner. When all of the manipulators put p first then $\text{score}(p) = 1$. So the manipulators must have voted so that a and b tie, since otherwise either a or b would tie with p and p would not be a unique winner. Therefore a subcollection of the manipulators with weight K voted $(p > b > a)$, a subcollection with weight $K - \widehat{K}$ voted $(p > a > b)$, and a subcollection with weight \widehat{K} voted $(p > a \sim b)$. No other votes would cause a and b to tie. \square

The remaining 3-candidate Copeland ^{α} CWCM cases remain in P when moving from total orders to votes with ties. The theorem below follows using the same arguments as in the proof of the case without ties from Faliszewski, Hemaspaandra, and Schnoor [FHS08].

Theorem 3.4.4 *3-candidate Copeland ^{α} CWCM remains in P for top orders, bottom orders, and weak orders, for $\alpha = 1$ for the nonunique winner case and for all rational $\alpha \in (0, 1]$ in the unique winner case.*

3.4.1 Majority-Graph Result

We now state a general theorem on two-voter majority graphs for votes with ties. See Brandt et al. [BHKS13] for related work on majority graphs constructed from a fixed number of voters with total orders. Recall that a majority graph can be constructed from an election (C, V) by representing each candidate as a vertex in the graph and for every pair of candidates $a, b \in C$ the graph contains the edge $a \rightarrow b$ if $a > b$ by majority.

Theorem 3.4.5 *A majority graph can be induced by two weak orders if and only if it can be induced by two total orders.*

Proof. Given two weak orders v_1 and v_2 that describe preferences over a candidate set C , we construct two total orders, v'_1 and v'_2 iteratively as follows.

For each pair of candidates $a, b \in C$ and $i \in \{1, 2\}$, if $a > b$ in v_i then set $a > b$ in v'_i .

For each pair of candidates $a, b \in C$, if $a > b$ in v_1 (v_2) and $a \sim b$ in v_2 (v_1) then the majority graph induced by v_1 and v_2 contains the edge $a \rightarrow b$. To ensure that the majority graph induced by v'_1 and v'_2 contains the edge $a \rightarrow b$ we must set $a > b$ in v'_2 (v'_1).

After performing the above steps there may still be a set of candidates $C' \subseteq C$ such that v_1 and v_2 are indifferent between each pair of candidates in C' . For each pair of candidates $a, b \in C'$, $a \sim b$ in v_1 and v_2 , which implies the majority graph does not contain an edge between a and b . To ensure that majority graph induced by v'_1 and v'_2 does not contain an edge between a and b , without loss of generality set v'_1 to strictly prefer the lexicographically smaller to the lexicographically larger candidate and the reverse in v'_2 .

The process described above constructs two orders v'_1 and v'_2 and ensures that the majority graph induced by v_1 and v_2 is the same as the majority graph induced by v'_1 and v'_2 . Since for each pair of candidates $a, b \in C$ and $i \in \{1, 2\}$ we consider each possible case where $a \sim b$ is in v_i and set either $a > b$ or $b > a$ in the corresponding order v'_i , it is clear that v'_1 and v'_2 are total orders. \square

Observe that as a consequence of Theorem 3.4.5 we get a transfer of NP-hardness from total orders to weak orders for two manipulators when the result depends only on the induced majority graph. The proofs for Copeland $^\alpha$ unweighted manipulation for two manipulators for all rational α for total orders depend only on the induced majority graph [FHS08, FHS10], so we can state the following corollary to Theorem 3.4.5.

Corollary 3.4.6 *Copeland $^\alpha$ unweighted manipulation for two manipulators for all rational $\alpha \neq 0.5$ for weak orders is NP-complete.*

3.4.2 Irrational-Voter Copeland Results

Another way to give more flexibility to voters is to let the voters state preferences that are not necessarily transitive, which are referred to as “irrational.” This simply means that for every unordered pair $a, b \in C$ of distinct candidates, the voter has $a > b$ or $b > a$. For example, a voter’s preferences could be $(a > b, b > c, c > a)$. As mentioned by Faliszewski et al. [FHHR09], a voter is likely to have preferences that are not transitive when making a decision based on multiple criteria.

Additionally, the preferences of voters can include ties as well as irrationality, and we will also consider such votes.

It is known that unweighted Copeland $^\alpha$ manipulation is NP-complete for total orders for all rational α except 0.5 [FHS08, FHS10]. For irrational voters, this problem is in P for $\alpha = 0, 0.5$, and 1, and NP-complete for all other α [FHS10]. *Weighted* manipulation for Copeland $^\alpha$ has not been studied for irrational voters. We will do so here.

Theorem 3.4.7 *3-candidate Copeland $^\alpha$ CWCM remains in P for irrational voters with or without ties, for $\alpha = 1$ for the nonunique winner case and for all rational $\alpha \in (0, 1]$ in the unique winner case.*

Theorem 3.4.8 *3-candidate Copeland $^\alpha$ CWCM remains NP-complete for irrational voters with or without ties, for $\alpha = 0$ in the unique winner case and for all rational $\alpha \in [0, 1)$ in the nonunique winner case.*

The proofs of the above two theorems follow from the arguments in the proofs of the corresponding rational cases, i.e., the proofs of Theorem 4.1 and 4.2 from Faliszewski, Hemaspaandra, and Schnoor [FHS08] for the case of voters without ties and the proofs of Theorems 3.4.2, 3.4.3, and 3.4.4 above for the case of voters with ties.

When $\alpha = 1$, which (as mentioned in Section 2.2.3) is also known as Llull, interesting things happen. It is known that 4-candidate Llull CWCM is in P for the unique and nonunique winner cases [FHS12]. For larger fixed numbers of candidates, this is open. Though it is known that unweighted manipulation for Llull (with an unbounded number of candidates) is NP-complete [FHS10]. In contrast, we will show now that for irrational voters, all these problems are in P, and we show this result by a reduction to maximum network flow, which is well-known to be in P (see [AMO93]).

Theorem 3.4.9 *Llull CWCM is in P for irrational voters with or without ties, in the nonunique winner case and in the unique winner case.*

Proof. Given a set of candidates C , a collection of voters V , k manipulators, and a preferred candidate $p \in C$, the preferences of the manipulators will always contain $p > a$ for all candidates $a \neq p$. This determines the score of p . In addition, let the initial preferences of the manipulators be $a > b$ for each pair of candidates $a, b \in C - \{p\}$ such that a defeats b in V or such that a ties b in V and a is lexicographically smaller than b . Note that, if $k > 0$, there are no pairwise ties in the election with the manipulators set in this way and that the manipulators all have strict preferences between every pair of candidates (i.e., no ties in their preferences). For every $a \neq p$, let $\text{score}_0(a)$ be the score of a with the manipulators set as above.

Construct the following flow network. The nodes are: a source s , a sink t , and all candidates other than p . For every $a \in C - \{p\}$, add an edge with capacity $\text{score}_0(a)$ from s to a and add an edge with capacity $\text{score}(p)$ from a to t . For every $a, b \in C - \{p\}$, add an edge from candidate a to candidate b with capacity 1 if, when all manipulators set $b > a$, the score of a decreases by 1 (and the score of b increases by 1).

If there is manipulation such that p is a winner, then for every candidate $a \in C - \{p\}$, $\text{score}(a) \leq \text{score}(p)$ so there is a network flow that saturates all edges that go out from s .

If there is a network flow that saturates all edges that go out from s then for every $a, b \in C - \{p\}$ such that there is a unit of flow from a to b , change $a > b$ to $b > a$ in all manipulators.

This construction can be adapted to the unique winner case by letting the capacity of the edge from a to t be $\text{score}(p) - 1$ instead of $\text{score}(p)$. \square

3.5 Axioms

Scoring rules on total orders satisfy a number of important axioms (explained below) such as neutrality, anonymity, consistency, continuity, and monotonicity.

It is interesting to see whether these axioms still hold if we allow weak orders using our scoring-rule extensions. Informally, an anonymous election system treats all of the voters the same, and a neutral election system treats all of the candidates the same. It is immediate that all of our scoring-rule extensions satisfy anonymity and neutrality.

An election system is consistent if for all elections (C, V) and (C, V') if there are common winners, then these are exactly the winners of $(C, V \cup V')$. An election system is continuous if for all elections (C, V) and (C, V') if $x \in C$ is the unique winner of (C, V) then there exists an $n \in \mathbb{N}$ such that x is the unique winner of $(C, kV \cup V')$ for all $k \geq n$. Consistency and continuity each follow from the same arguments used for the case for total orders (see [You75] and [Zwi16]). Basically, consistency follows because the score of a candidate in the union of two elections is the sum of the scores in the individual elections. Similarly, continuity holds since we can multiply scores.

That leaves monotonicity, which for total orders and resolute election systems (a single candidate wins) requires that for all elections (C, V) , if $x \in C$ is the winner then if x is “lifted” in a vote $v_i \in V$ (i.e., the preference of x is increased in the vote v_i and the relative positions of candidates in $C - \{x\}$ remain unchanged) to yield the vote v'_i then x is the winner of $(C, \{v'_i\} \cup V - \{v_i\})$.

To extend this definition for irresolute election systems Pegel [Pel81] states that when a winning candidate x is lifted in a vote that x must remain a winner and no new winners are added. (See also footnote 58 of [Zwi16] and [SZ12] for a discussion of monotonicity for irresolute election systems.) We consider the obvious analogue for weak orders, and the following is an example of the lifting operation. For an election $(\{x, y, z\}, V)$ where x is a winner, consider the vote $v_i \in V$ and $v_i = (y > x \sim z)$, so x can be lifted to yield $v'_i = (y > x > z)$, $v'_i = (y \sim x > z)$, or $v'_i = (x > y > z)$.

It is easy to see that the max, min, and average scoring-rule extensions all satisfy monotonicity by a similar argument. For each of these extensions, when a winning candidate x is lifted in a vote the score of x can only increase or remain the same, while the scores of every other candidate can only decrease or remain the same. The round-down scoring-rule extension does not satisfy monotonicity. Consider the following counterexample.

Example 3.5.1 Consider the candidate set $\{a, b, c, d, e\}$, the votes $v_1 = (a > c \sim b > d > e)$ and $v_2 = (c > a > d > b > e)$, and the 5-candidate scoring rule $\langle 16, 8, 4, 2, 0 \rangle$ using round-down.

Notice that $\text{score}(a) = 16$, $\text{score}(b) = 6$, $\text{score}(c) = 20$, $\text{score}(d) = 6$ and $\text{score}(e) = 0$. If we lift the winning candidate c in vote v_1 to yield $v'_1 = (a > c > b > d > e)$ we increase the number of groups in the vote, so the score of a increases. So given v'_1 and v_2 $\text{score}(a) = \text{score}(c) = 24$. Since a became a winner when c was lifted, round-down does not satisfy monotonicity.

Notice that if we used the scoring vector $\langle 4, 3, 2, 1, 0 \rangle$ (5-candidate Borda) our counterexample would not work, since by lifting a winning candidate we can increase the number of groups in v_1 by at most one, and for Borda this means that no other candidate's score will increase by more than one (while a winning candidate's score increases by at least one), so no candidate will become a winner when x is lifted if she was not already a winner.

We mention here that Gärdenfors [Gär73] considers extensions to Borda for votes with ties that are equivalent to Borda using average, min, and round-down and finds that monotonicity holds (though Gärdenfors looks at election systems that produce a social ordering instead of a set of winners and uses a somewhat different definition of monotonicity) [Gär73].

3.6 Related Work

The recent work by Narodytska and Walsh [NW14] studied the complexity of manipulation for top orders and is very influential to our computational study of more general votes with ties. Baumeister et al. [BFLR12] and Narodytska and Walsh [NW14] studied several extensions for election systems for top orders, which we further extend for weak orders. Menon and Larson [ML15] expanded on the work by Narodytska and Walsh and characterized the complexity of weighted manipulation for all three-candidate scoring rules for top orders.

We also mention related work that considered other more general types of votes. Konczak and Lang [KL05] introduced the possible and necessary winner problems, where given an election where voters have incomplete votes, i.e., votes that may contain unspecified preferences among (some of) the candidates, a possible winner is a candidate that wins in at least one extension of the votes to total orders and a necessary winner is a candidate that wins in every extension [KL05]. Note that the votes considered by Konczak and Lang are tie-free. However, they stated how their work could be extended to also consider ties. Pini et al. [PRVW11] also studied the possible and necessary winner problems, and they considered votes with ties, incompleteness, and incomparability (where incomparability means that a voter refuses to state a preference between two candidates).

Baumeister et al. [BFLR12] studied the possible winner problem for different types of truncated votes (i.e., votes with incompleteness), where all of the unranked candidates appear at the top, bottom, or middle of each vote [BFLR12]. Baumeister et al. also introduced the

model of extension bribery, where given voters with top-truncated preferences, voters are paid to extend their vote to ensure that a preferred candidate wins [BFLR12]. They additionally extended the definitions of the election systems used to handle top-truncated votes, which we generalize for votes with ties.

3.7 Conclusions

Allowing voters to state ties in their preferences is an important step to model more natural scenarios for the computational study of elections. One would expect that the complexity of many election attack problems remains the same when moving from tie-free votes to votes with ties, but we were able to find interesting results on what can happen.

We were the first to find a case where the complexity of manipulation increases when moving from tie-free votes to votes with ties. We found cases where the complexity of bribery can increase and cases where it can decrease, and we stated a general observation about the complexity of control.

For 3-candidate Copeland^α elections, we found that in all cases for votes with ties and even irrational votes with and without ties, the complexity of weighted manipulation remains the same. However, though the complexity remains the same, in many cases a different argument is required for votes with ties.

Overall, we presented a general overview of what can happen with respect to the worst-case complexity of election-attack problems when voters are allowed to state votes with ties. It would be interesting to see how the complexity of other election problems is affected by votes with ties, which we consider to be a natural model for preferences in practical settings.

Chapter 4

Models for the Electorate: Single-Peakedness with Ties

4.1 Introduction

For a given election it is natural to consider that the preferences of the electorate have some structure. In this chapter we consider votes that are single-peaked, and building upon the work in the previous chapter, we consider different ways to model single-peakedness for votes with ties.

The model of single-peaked preferences was introduced by Black [Bla48] and models the preferences of a collection of voters with respect to a one-dimensional axis, i.e., a total ordering of the candidates, where each voter in a single-peaked electorate has a single most-preferred candidate (peak) on the axis and the farther that a candidate is from the voter's peak the less preferred they are by the voter [Bla48]. Figure 4.1 includes two preference orders that are single-peaked and one that is not single-peaked with respect to the axis $A = a < d < b < e < c$.

Single-peakedness is arguably the most important and commonly examined restriction on preferences in economics and social choice, and the notion of single-peakedness nicely models how an electorate views candidates with respect to a single polarizing issue (e.g., liberal vs. conservative), where candidates farther to the leftmost and rightmost points on the axis represent the extremes of the issue.

This standard model of single-peaked preferences has many desirable social-choice properties. When the voters in an election are single-peaked, the majority relation is transitive [Bla48] (i.e., for each pair of candidates a, b if $a > b$ and $b > c$ then $a > c$, by majority) and there exist voting rules that are strategyproof [Mou80] (i.e., a voter cannot misrepre-




Figure 4.1: The solid line represents the single-peaked preference order ($b > e > d > a > c$), the dashed line represents the single-peaked preference order ($c > e > b > d > a$), and the dotted line represents the preference order ($a > c > d > e > b$), which is *not* single-peaked.

sent her preferences to achieve a personally better outcome). Additionally, computational problems often become easier when preferences are single-peaked, when compared to the general case. For tie-free votes, the complexity of election-attack problems often decreases when the voters in an election are single-peaked [FHHR11, BBHH15] (or even *nearly* single-peaked [FHH14a]), and the winner problems for Kemeny, Dodgson, and Young elections are in P [BBHH15] when they are Θ_2^P -complete in general [HHR97, HSV05, RSV03].

4.1.1 Single-Peakedness with Ties

In computational social choice the study of single-peaked preferences is typically restricted to tie-free preferences, but, as discussed in Chapter 3, in practical settings voters often have some degree of ties in their preferences and election systems exist that properly consider these preferences with ties. We introduce four important models for single-peakedness with ties here, and formal definitions will follow in Section 4.2.

The standard model of single-peaked preferences introduced by Black [Bla48] is naturally defined for votes with ties. For a voter to have single-peaked preferences in this model they must have strictly increasing and then strictly decreasing preferences or only strictly increasing (decreasing) preferences along a given axis. So when votes with ties are allowed voters can only have tied preference between candidates that appear on opposite sides of their peak.

Black later extended this model to single-plateaued preferences, which model the preferences of a collection of voters in a similar way, but allow voters to have multiple most preferred candidates (which form a plateau instead of a peak) in their preferences [Bla58,

Chapter 5].¹

Recently, two additional models of single-peakedness have been introduced that extend the standard model: the model of possibly single-preferences (introduced by Lackner in the context of partial votes [Lac14]),² where a voter is possibly single-peaked with respect to a given axis if they have weakly increasing and then weakly decreasing preferences or only weakly increasing (decreasing) preferences along the axis, and the model of single-peaked preferences with outside options [Can04], where a voter is single-peaked with outside options with respect to a given axis if they have single-peaked preferences over a closed interval of the axis, candidates outside this interval are strictly less preferred, and they are tied among all of the candidates outside this interval.

4.1.2 Contributions

The aforementioned social choice and computational results on elections with single-peaked preferences lead us to ask the question: How can we determine what axis (if any) makes a given electorate single-peaked? This is known as the single-peaked consistency problem. For tie-free votes, Bartholdi and Trick [BT86] showed that this problem was in P by a reduction to the problem of determining if a 0-1 matrix has the consecutive ones property. We extend this result to show that the consistency problem for single-peaked, single-plateaued, possibly single-peaked, and single-peaked preferences with outside options for votes with ties are each in P. For the model of possibly single-peaked preferences, this solves the main open problem in the work by Lackner [Lac14] that introduces this model.

We show that each of these problems are in P by a reduction to the problem of determining if a 0-1 matrix has the consecutive ones property, and so we get a nice correspondence between the axis and the permutation of the columns of the matrix in the reduction. A reduction to this consecutive ones problem was used by Bartholdi and Trick [BT86] to determine the single-peaked consistency for tie-free votes, and Bredereck, Chen, and Woeginger [BCW13] present a reduction to this same problem to determine the single-crossing consistency for tie-free votes.³ We also mention that after single-peaked consistency for tie-free votes was shown to be in P, both Doignon and Falmagne [DF94] and Escoffier, Lang, and Öztürk [ELÖ08] independently found faster direct algorithms.

In Chapter 3 we considered the effect of allowing votes with ties on the complexity of manipulation, control, and bribery. In this chapter we extend this work for the case of

¹We mention that for weak orders the definition of single-peaked preferences from Fishburn [Fis73, Chapter 9] is the same as Black's definition of single-plateaued preferences [Bla58, Chapter 5].

²The possibly single-peaked model is also known as the existentially single-peaked model.

³The model of single-crossing preferences is another domain restriction [Mir71]. Its corresponding consistency problem for tie-free votes was first shown to be in P by Elkind, Faliszewski, and Slinko [EFS12].

single-peaked preferences for votes with ties, and examine the effect of different models of single-peakedness on the complexity of manipulation. We present the first natural case where there is an increase in complexity of manipulation from tie-free votes to votes with ties, by considering votes that are possibly single-peaked. Menon and Larson [ML16] later systematically examined the complexity of manipulation for single-peaked preferences with outside options for top orders (which is equivalent to the model of possibly single-peaked preferences for top orders) and found that the complexity often increases with respect to the tie-free case. In contrast, we show for scoring rules and other natural systems that for the standard model of single-peakedness, and for single-plateauedness, allowing votes with ties does not increase the complexity of manipulation.

This chapter is organized as follows. In Section 4.2 we introduce each of the four models of single-peakedness for votes with ties, and compare the social-choice properties of these different models and how they each relate to one another. Section 4.3 is split into four subsections, and in each subsection we redefine one of the four models of single-peakedness using forbidden substructures and describe a transformation from its consistency problem to the problem of determining if a 0-1 matrix has the consecutive ones property. In Section 4.4 we discuss the computational impact of single-peaked votes with ties on the complexity of manipulation and how this is further impacted by the choice of the model of single-peakedness. We conclude in Section 4.5 by summarizing our results.

4.2 Models of Single-Peakedness for Votes with Ties

We consider four important models of single-peaked preferences for votes with ties. For the following definitions, for a given axis A (a total ordering of the candidates) and a given preference order v , we say that v is strictly increasing (decreasing) along a segment of A if each candidate is preferred to the candidate on its left (right) with respect to A . In Figure 4.2 we present an example of each of the four models of single-peaked preferences, and in Figure 4.3 we show how the four models relate to each other. We begin by stating the definition of the standard model of single-peakedness from Black [Bla48].

Definition 4.2.1 *Given a preference profile V of weak orders over a set of candidates C , V is single-peaked with respect to a total ordering of the candidates A (an axis) if for each voter $v \in V$, A can be split into three segments X , Y , and Z (X and Z can each be empty) such that Y contains only the most-preferred candidate of v , v is strictly increasing along X and v is strictly decreasing along Z .*

We can restate the above definition informally by stating that each preference order is

strictly increasing to a peak and then strictly decreasing, only strictly increasing, or only strictly decreasing, all along the given axis.

Observe that for a preference profile of votes with ties to be single-peaked with respect to an axis, each voter can have a tie between at most two candidates at each position in her preference order since the candidates must each appear on separate sides of her peak. Otherwise the preference order would not be *strictly* increasing/decreasing along the given axis.

We define the corresponding problem of single-peaked consistency for weak orders below.

Name: Single-Peaked Consistency for Weak Orders

Given: A preference profile V of weak orders and a set of candidates C .

Question: Does there exist an axis A such that V is single-peaked with respect to A ?

The model of single-plateaued preferences extends single-peakedness by allowing voters to have multiple most-preferred candidates (a single plateau at the top) instead of a single peak [Bla58]. This is defined by extending Definition 4.2.1 so that Y can contain multiple candidates.

Definition 4.2.2 *Given a preference profile V of weak orders over a set of candidates C , V is single-plateaued with respect to a total ordering of the candidates A (an axis) if for each voter $v \in V$, A can be split into three segments X , Y , and Z (X and Z can each be empty) such that Y contains only the most-preferred candidates of v , v is strictly increasing along X and v is strictly decreasing along Z .*

We define the corresponding problem of single-plateaued consistency for weak orders below.

Name: Single-Plateaued Consistency for Weak Orders

Given: A preference profile V of weak orders and a set of candidates C .

Question: Does there exist an axis A such that V is single-plateaued with respect to A ?

Another extension to the standard model of single-peakedness is the model of possibly single-peaked preferences introduced by Lackner [Lac14]. A preference profile is possibly single-peaked with respect to a given axis if there exists an extension of each preference order to a total order such that the new preference profile of total orders is single-peaked. This can be stated without referring to extensions, for votes with ties, in the following way.

Definition 4.2.3 *Given a preference profile V of weak orders over a set of candidates C , V is possibly single-peaked with respect to a total ordering of the candidates A (an axis) if for each voter $v \in V$, A can be split into three segments X , Y , and Z (X and Z can each be empty) such that Y contains the most-preferred candidates of v , v is weakly increasing along X and v is weakly decreasing along Z .*

We define the corresponding problem of possibly single-peaked consistency for weak orders below.

Name: Possibly Single-Peaked Consistency for Weak Orders

Given: A preference profile V of weak orders and a set of candidates C .

Question: Does there exist an axis A such that V is possibly single-peaked with respect to A ?

Notice that the above definition extends the definition of single-plateaued preferences to allow for multiple plateaus on either side of the peak. So for votes with ties, possibly single-peaked preferences model when voters have weakly increasing and then weakly decreasing or only weakly increasing/decreasing preferences along an axis.

Another generalization of single-peaked preferences for votes with ties was introduced by Cantala [Can04]: the model of single-peaked preferences with outside options. When preferences satisfy this restriction with respect to a given axis, each voter has a segment of the axis where they have single-peaked preferences and candidates appearing outside of this segment on the axis are strictly less preferred and the voter is tied among them. Similar to how single-plateaued preferences extend the standard single-peaked model to allow voters to state multiple most-preferred candidates, the notion of single-peaked preferences with outside options extends the standard model to allow voters to state multiple least-preferred candidates.

Definition 4.2.4 *Given a preference profile V of weak orders over a set of candidates C , V is single-peaked with outside options with respect to a total ordering of the candidates A (an axis) if for each voter $v \in V$, A can be split into five segments O_1 , X , Y , Z , and O_2 (O_1 , X , Z , and O_2 can each be empty) such that Y contains only the most-preferred candidate of v , v is strictly increasing along X and v is strictly decreasing along Z , for all candidates $a \in X \cup Y \cup Z$ and $b \in O_1 \cup O_2$, v states $a > b$, and for all candidates $x, y \in O_1 \cup O_2$, v states $x \sim y$.*

We define the corresponding consistency problem for single-peaked preferences with outside options for weak orders below.

Name: Single-Peaked Consistency with Outside Options for Weak Orders

Given: A preference profile V of weak orders and a set of candidates C .

Question: Does there exist an axis A such that V is single-peaked with outside options with respect to A ?




Figure 4.2: Given the axis $A = a < b < c < d < e < f < g$, the solid line is the single-peaked order ($e > d > c \sim f > b > a > g$), the dashed line is the single-plateaued order ($b \sim c > d > e > f > g > a$), the dashed-dotted line is the single-peaked with outside options order ($d > c \sim e > a \sim b \sim f \sim g$), and the dotted line is the possibly single-peaked order ($f \sim g > e > d \sim c > b > a$). See Figure 4.3 for the relationships between the four models.

4.2.1 Social-Choice Properties

We now briefly state some general observations on single-peaked preferences for votes with ties, including how the models relate to each other, as well as their social-choice properties. See the related work by Barberà [Bar07] for a more in-depth comparison of the different social-choice properties of the models of single-peaked, single-plateaued, and single-peaked preferences with outside options for votes with ties.

It is easy to see that for total-order preferences each of the four models of single-peakedness with ties that we consider are equivalent. In Figure 4.3 we show how the models are related for weak orders, top orders, and bottom orders.

One of the most well-known and desirable properties of an election with single-peaked preferences is that the majority relation is transitive [Bla48]. A majority relation is transitive if for all distinct candidates $a, b, c \in C$ if $a > b$ and $b > c$ by majority then $a > c$ by majority. When the majority relation is transitive then we know that weak Condorcet winners exist.

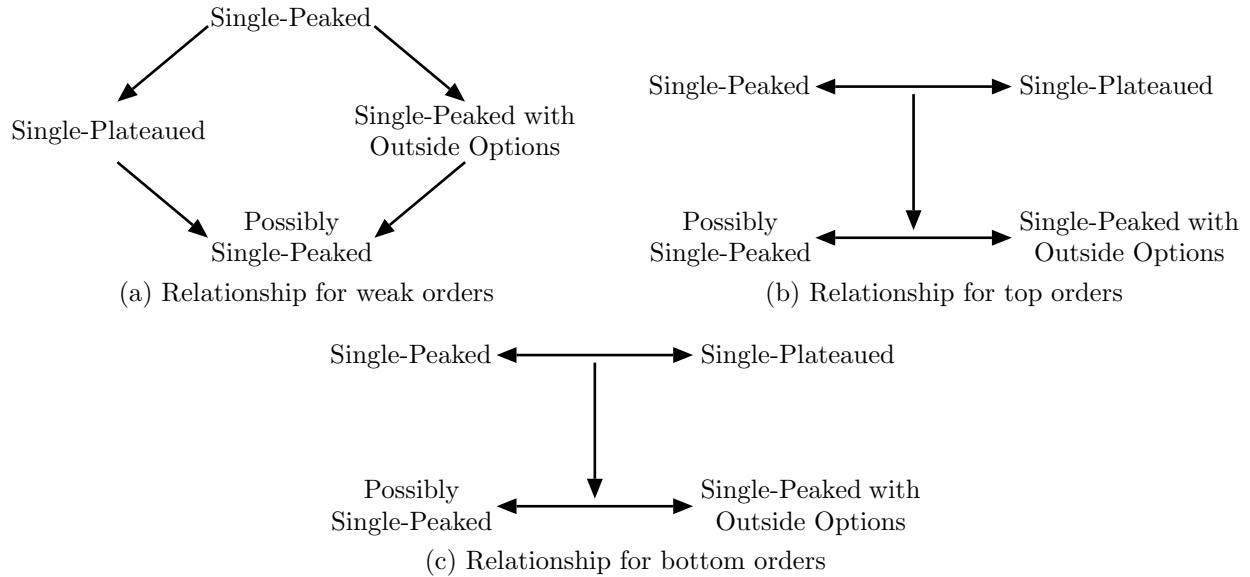


Figure 4.3: Relationships between the four models of single-peaked preferences for different types of votes with ties, where $A \rightarrow B$ indicates that a preference profile satisfying model A also satisfies model B .

(Recall that a weak Condorcet winner is a candidate that ties-or-beats every other candidate pairwise.) The transitivity of the majority relation also holds when the voters have single-plateaued preferences [Bla58].

It is not necessarily the case that the majority relation is transitive when preferences in an election are single-peaked with outside options or are possibly single-peaked. Cantala [Can04] provides an example of a profile that is single-peaked with outside options that does not have a weak Condorcet winner, and so does not have a transitive majority relation. To illustrate this: Consider the preference profile V comprised of the following five voters from Table 9.1 in Fishburn [Fis73].

- One voter voting $(b > a > c)$.
- Two voters voting $(c > b > a)$.
- Two voters voting $(a > b \sim c)$.

When we evaluate this preference profile under the majority rule where $x > y$ by majority if more voters state $x > y$ than $y > x$, then V has the majority cycle $a > c > b > a$ [Fis73]. Clearly V is single-peaked with outside options (and thus also possibly single-peaked) with respect to the axis $A = a < b < c$.

The possibly single-peakedness model considers the existence of an extension of the preferences of all of the voters to single-peaked total orders. We briefly consider the case where

every extension to the votes are single-peaked total orders, and make two observations.

Observation 4.2.5 *If a preference profile of weak orders is single-peaked then all extensions of the preferences to total orders are also single-peaked.*

Observation 4.2.6 *If a preference profile of weak orders is single-plateaued and each preference order has at most two most-preferred candidates, then all extensions of the preferences to total orders are single-peaked.*

4.3 Detecting Single-Peakedness

Given a preference profile it is natural to ask how to determine if an axis exists such that the profile satisfies one of the above restrictions. This is referred to as the consistency problem for a restriction. Bartholdi and Trick [BT86] showed that single-peaked consistency for total orders can be determined in polynomial time (i.e., is in P) by reducing to the following problem of determining if a 0-1 matrix has the consecutive ones property, and we extend this result to show that the consistency problem for single-peaked, single-plateaued, possibly single-peaked, and single-peaked preferences with outside options, for weak orders are each in P.

Name: Consecutive-ones matrix.

Given: A 0-1 matrix M .

Question: Does there exist a permutation of the columns of M such that in each row all of the 1's are consecutive?

The above problem was shown to be in P by Fulkerson and Gross [FG65]. Booth and Lueker [BL76] improved on this result by finding a linear-time algorithm through the development and use of the novel PQ-tree data structure, which contains all (possibly exponentially many) permutations of the columns of a matrix such that all of the 1's are consecutive in each row.

4.3.1 Possibly Single-Peaked Consistency

The most general of the four models that we consider is the model of possibly single-peaked preferences. The construction and corresponding proof will be the basis for showing that the consistency problems for single-peaked, single-plateaued, and single-peaked preferences with outside options for weak orders are each also in P.

Given an axis A and a preference order v , if v is possibly single-peaked with respect to A then v cannot have strictly decreasing and then strictly increasing preferences with respect to A . Following the terminology used by Lackner [Lac14], we refer to this as a v -valley.

Definition 4.3.1 *A preference order v over a candidate set C contains a v -valley with respect to an axis A if there exist candidates $a, b, c \in C$ such that $a < b < c$ in A and v states $a > b$ and $c > b$.*

Using the v -valley substructure we can state the following lemma, which will simplify the argument used in the proof of Theorem 4.3.5.

Lemma 4.3.2 [Lac14] *Let V be a preference profile of weak orders. V is possibly single-peaked with respect to an axis A if and only if no preference order $v \in V$ contains a v -valley with respect to A .*

To construct a matrix from a preference profile of weak orders, we extend the transformation from Bartholdi and Trick [BT86] so that it works for weak orders. We describe the construction below.

Construction 4.3.3 *Let V be a preference profile of weak orders over candidate set C . For each $v \in V$ construct a $(\|C\| - 1) \times \|C\|$ matrix X_v . Each column of X_v corresponds to a candidate in C . For each candidate $c \in C$ let k be the number of candidates that are strictly preferred to c by v and let the corresponding column in matrix X_v contain k 0's starting at row one, with the remaining entries filled with 1's. All $\|V\|$ of the matrices are row-wise concatenated to yield the $(\|V\| \cdot (\|C\| - 1)) \times \|C\|$ matrix X .*

The main difference in our construction is that we have to handle the case of ties, so we determine the number of 1's in each column using the number of candidates ranked strictly less than a given candidate, so this results in one fewer row in each of the individual preference matrices. In the construction used by Bartholdi and Trick [BT86], given a preference order v over a set of candidates C , for all $a, b \in C$, v states $a > b$ if and only if the number of 1's in the column corresponding to a is greater than the number of 1's in the column corresponding to b in X_v . Notice that this still holds for our construction.

Below we show how Construction 4.3.3 is applied to a preference profile of weak orders that is possibly single-peaked.

Example 4.3.4 Consider the preference profile V that consists of the preference orders v and w . Let the preference order v be $(a \sim c > b > e \sim d > f)$ and the preference order w be

($a > b > c > e \sim d > f$). Notice that V does not contain a guiding order,⁴ which is required by the algorithm for weak orders found in Lackner [Lac14].

$$X_v = \begin{array}{c} \begin{array}{cccccc} a & b & c & d & e & f \end{array} \\ \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \end{array}$$

$$X_w = \begin{array}{c} \begin{array}{cccccc} a & b & c & d & e & f \end{array} \\ \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \end{array}$$

We then row-wise concatenate X_v and X_w to construct X .

$$X = \begin{array}{c} \begin{array}{cccccc} a & b & c & d & e & f \end{array} \\ \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \end{array}$$

$$X' = \begin{array}{c} \begin{array}{cccccc} b & a & c & d & e & f \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \end{array}$$

Next, we permute the columns of X so that in each row all of the 1's are consecutive to yield X' . Observe that V is possibly single-peaked with respect to $b < a < c < d < e < f$, the ordering of the columns of X' as its axis.

We now show that possibly single-peaked consistency for weak orders and the problem of determining if the constructed 0-1 matrix has the consecutive ones property are equivalent using Lemma 4.3.2 and Construction 4.3.3.

⁴Lackner defines a *guiding order* as an implicitly specified total order constructed in the following way.

If there exists a voter $v \in V$ such that v has a unique last-ranked candidate, then that candidate is appended to the top of the guiding order. This is then repeated on the preference profile restricted to the candidates not yet added to the guiding order until either the guiding order is a total order or there is no $v \in V$ with a unique last-ranked candidate, the case where no guiding order exists [Lac14]. Note that if a given preference profile is possibly single-peaked then it remains possibly single-peaked if a guiding order is added as an additional preference order [Lac14].

Theorem 4.3.5 *A preference profile V of weak orders is possibly single-peaked consistent if and only if the matrix X , constructed using Construction 4.3.3, has the consecutive ones property.*

Proof. Let V be a preference profile of weak orders. Essentially the same argument as used by Bartholdi and Trick [BT86] holds.

If V is possibly single-peaked with respect to an axis A then by Lemma 4.3.2 we know that no preference order $v \in V$ contains a v-valley with respect to A . When the columns of the matrix X are permuted to correspond to the axis A no row will contain the sequence $\dots 1 \dots 0 \dots 1 \dots$ since this corresponds to a preference order that strictly decreases and then strictly increases along the axis A (a v-valley). Therefore X has the consecutive ones property.

For the other direction suppose that V is not possibly single-peaked. Then by Lemma 4.3.2 we know that for every possible axis there exists a preference order $v \in V$ such that v contains a v-valley with respect to that axis. So every permutation of the columns of X will correspond to an axis where some preference order has a v-valley. As stated in the proof of the other direction, a v-valley corresponds to a row containing the sequence $\dots 1 \dots 0 \dots 1 \dots$ so clearly X does not have the consecutive ones property.

The only difference from the argument used by Bartholdi and Trick [BT86] for total orders is that in our case the preference orders can remain constant at the peak and at points on either side of the peak. The same argument still applies since by Lemma 4.3.2 the absence of v-valleys with respect to an axis is equivalent to a profile of weak orders being possibly single-peaked with respect to that axis. \square

Corollary 4.3.6 *Possibly single-peaked consistency for weak orders is in P.*

4.3.2 Single-Plateaued Consistency

Single-plateaued preferences are a much more restrictive model than possibly single-peaked preferences since essentially they are single-peaked except that each preference order can have multiple most-preferred candidates [Bla58, Chapter 5].

Since a preference order must be strictly increasing and then strictly decreasing with respect to an axis (excluding its most-preferred candidates) we can again use the v-valley substructure. However we will need another substructure to prevent two candidates that are tied in a voter's preference order from appearing on the same side of that voter's peak (plateau).

Definition 4.3.7 A preference order v over a candidate set C contains a nonpeak plateau with respect to A if there exist candidates $a, b, c \in C$ such that $a < b < c$ in A and v states either $a > b \sim c$ or $c > b \sim a$.

We use the v-valley and nonpeak plateau substructures to state the following lemma.

Lemma 4.3.8 Let V be a preference profile of weak orders. V is single-plateaued with respect to an axis A if and only if no preference order $v \in V$ contains a v-valley with respect to A and no preference order $v \in V$ contains a nonpeak plateau with respect to A .

Proof. Let C be a candidate set, V be a preference profile of weak orders, and A be an axis.

If V is single-plateaued with respect to A then for every preference order $v \in V$, A can be split into segments X , Y , and Z such that v is strictly increasing along X , remaining constant along Y , and strictly decreasing along Z . Since v is only ever strictly decreasing along Z and Z is the rightmost segment of A , v cannot contain a v-valley with respect to A . For a nonpeak plateau to exist with respect to A there must exist candidates $a, b, c \in C$ such that $a < b < c$ in A and v states either $a > b \sim c$ or $c > b \sim a$.

We first consider the case where v states $a > b \sim c$. Since v strictly prefers a to b and a to c , and both b and c are to the right of a on the axis, we know that both b and c must be in segment Z . However, v is strictly decreasing along Z , so v cannot have a nonpeak plateau of this form.

We now consider the case where v states $c > b \sim a$. Since c is strictly preferred to a and b by v and both a and b are to the left of c on the axis we know that both a and b must be in segment X . However, v is strictly increasing along X , so v cannot have a nonpeak plateau of this form.

For the other direction we consider the case when no preference order $v \in V$ contains a v-valley with respect to A and no preference order $v \in V$ contains a nonpeak plateau with respect to A .

Since no preference order $v \in V$ contains a v-valley with respect to A , we know from Lemma 4.3.2 that V is possibly single-peaked with respect to A . Since we also know that no preference order $v \in V$ contains a nonpeak plateau with respect to A it is easy to see that V is single-plateaued with respect to A . □

Since the nonpeak plateau substructure is needed in addition to the v-valley substructure, we need to extend Construction 4.3.3 so that if a preference order contains a nonpeak plateau with respect to an axis A , then when the columns of its corresponding preference matrix are permuted according to A the matrix will contain a row with the sequence $\dots 1 \dots 0 \dots 1 \dots$.

Construction 4.3.9 Let V be a preference profile of weak orders over candidate set C . For each $v \in V$ construct a $(\|C\| - 1) \times \|C\|$ matrix X_v . Each column of X_v corresponds to a candidate in C . For each candidate $c \in C$ let k be the number of candidates that are strictly preferred to c by v and let the corresponding column in matrix X_v contain k 0's starting at row one, with the remaining entries filled with 1's (as in Construction 4.3.3). The following extensions to Construction 4.3.3 ensure that if v has a nonpeak plateau with respect to an axis A then when the columns of X_v are permuted according to A it will not have consecutive ones in rows.

For each pair of candidates $a, b \in C$ such that v states $a \sim b$, they are not among v 's most-preferred candidates, and there is no candidate $c \in C - \{a, b\}$ such that v is tied among a, b , and c , then append three additional rows to the matrix X_v where the column corresponding to a is $\begin{bmatrix} 0 & 1 & 1 \end{bmatrix}'$, the column corresponding to b is $\begin{bmatrix} 1 & 1 & 0 \end{bmatrix}'$, each column corresponding to a candidate strictly preferred to a and b is $\begin{bmatrix} 1 & 1 & 1 \end{bmatrix}'$, and each column corresponding to a remaining candidate is $\begin{bmatrix} 0 & 0 & 0 \end{bmatrix}'$.

If there exist four candidates $a, b, c, d \in C$ such that v states $d > a \sim b \sim c$, then let X be a matrix with no solution and halt.

After constructing a matrix X_v for each $v \in V$, all $\|V\|$ of the matrices are row-wise concatenated to yield a matrix X .

We now show how Construction 4.3.9 is applied to a preference profile of weak orders that is single-plateaued.

Example 4.3.10 We consider the same preference profile as in Example 4.3.4 and we bold the additional rows in this example. Let the preference order v be $(a \sim c > b > e \sim d > f)$ and the preference order w be $(a > b > c > e \sim d > f)$.

$$X_v = \begin{array}{c} \begin{array}{cccccc} a & b & c & d & e & f \\ \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{0} \end{bmatrix} \end{array} \end{array}$$

$$X_w = \begin{array}{c} \begin{array}{ccccc} a & b & c & d & e & f \\ \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{0} \end{bmatrix} \end{array} \end{array}$$

We then row-wise concatenate X_v and X_w to construct X .

$$X = \begin{array}{c} \begin{array}{cccccc} a & b & c & d & e & f \end{array} \\ \left[\begin{array}{cccccc} 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{0} \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{0} \end{array} \right] \end{array} \qquad
 X' = \begin{array}{c} \begin{array}{cccccc} e & b & a & c & d & f \end{array} \\ \left[\begin{array}{cccccc} 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} & \mathbf{0} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{1} & \mathbf{0} \end{array} \right] \end{array}$$

Next, we permute the columns of X such that in each row all of the ones are consecutive to yield X' . Observe that V is single-plateaued with respect to this new ordering $e < b < a < c < d < f$ as its axis. Also notice that an axis containing d and e adjacent to each other (as seen in Example 4.3.4) would not correspond to an ordering of the columns of X with consecutive ones in rows due to the additional rows from the extensions made to Construction 4.3.3 in Construction 4.3.9.

Construction 4.3.3 ensures that no preference order contains a v-valley and the extensions made in Construction 4.3.9 ensure that no preference order contains a nonpeak plateau. So the proof of the following theorem uses a similar argument to the proof of Theorem 4.3.5. Now the presence of v-valleys *or* nonpeak plateaus, not just v-valleys, is equivalent to a row containing the sequence $\cdots 1 \cdots 0 \cdots 1 \cdots$.

Theorem 4.3.11 *A preference profile V of weak orders is single-plateaued consistent if and only if the matrix X , constructed using Construction 4.3.9, has the consecutive ones property.*

Proof. Let V be a preference profile of weak orders. We extend the argument used by Bartholdi and Trick [BT86] and the proof of Theorem 4.3.5 except in this case we use Lemma 4.3.8 instead of Lemma 4.3.2.

If V is single-plateaued with respect to an axis A then by Lemma 4.3.8 we know that no $v \in V$ contains a v-valley with respect to A and no $v \in V$ contains a nonpeak plateau with respect to A . When the columns of the matrix X are permuted to correspond to the axis A no row will contain the sequence $\dots 1 \dots 0 \dots 1 \dots$ since this would correspond to a preference order that strictly decreases and then strictly increases along the axis A (a v-valley) or it would correspond to a preference order that has two tied candidates appearing on the same side of its peak (a nonpeak plateau). Therefore X has the consecutive ones property.

Conversely, suppose that V is not single-plateaued consistent. We know from Lemma 4.3.8 that for every possible axis there exists a preference order $v \in V$ such that v contains a v-valley or v contains a nonpeak plateau with respect to that axis.

If there exists a preference order $v \in V$ and candidates $a, b, c, d \in C$ such that v states $d > a \sim b \sim c$ then we know that X has no solution.

Else, we know that every permutation of the columns of X directly corresponds to an axis where a preference order has a v-valley or a nonpeak plateau. And as stated in the other direction, the presence of a v-valley or a nonpeak plateau corresponds to a row containing the sequence $\dots 1 \dots 0 \dots 1 \dots$.

Therefore X does not have the consecutive ones property. □

Corollary 4.3.12 *Single-plateaued consistency for weak orders is in P.*

4.3.3 Single-Peaked Consistency with Outside Options

It is easy to see that a given preference profile of top orders is single-peaked with outside options if and only if it is possibly single-peaked. So it is immediate from the result by Lackner [Lac14] that shows that possibly single-peaked consistency for top orders is in P, that the consistency problem for single-peaked preferences with outside options for top orders is in P.

Recall that a preference order is single-peaked with outside options if and only if it is single-peaked for the standard model of single-peakedness in a closed interval of the axis, tied among all candidates outside of that interval, and all candidates outside of that interval are strictly less preferred. We can think of this model as similar to single-plateaued preferences except instead of a plateau of most-preferred candidates we allow multiple least-preferred candidates. So, instead of using a nonpeak plateau substructure, we use a nonleast plateau substructure.

Definition 4.3.13 *A preference order v over a candidate set C contains a nonleast plateau with respect to an axis A if there exist candidates $a, b \in C$ such that a and b are adjacent in A and either there exists a candidate $c \in C$ such that v states $a \sim b > c$ or a and b are among v 's most-preferred candidates.*

We can now use this substructure and the v-valley substructure to state the following lemma.

Lemma 4.3.14 *Let V be a preference profile of weak orders. V is single-peaked with outside options with respect to an axis A if and only if no preference order $v \in V$ contains a v-valley with respect to A and no preference order $v \in V$ contains a nonleast plateau with respect to A .*

Proof. Let C be a candidate set, V be a preference profile of weak orders, and A be an axis.

If V is single-peaked with outside options with respect to A then for every preference order $v \in V$, A can be split into five segments O_1, X, Y, Z , and O_2 , such that v is strictly increasing along X , strictly decreasing along Z , remains constant along O_1 and O_2 , Y contains only the most-preferred candidate of v , and for all candidates $a \in X \cup Y \cup Z$ and $b \in O_1 \cup O_2$, v states $a > b$, and for all candidates $x, y \in O_1 \cup O_2$, v states $x \sim y$. Since v only strictly decreases along Z and no v only remains constant in O_2 (the rightmost segment of A), v cannot contain a v-valley with respect to A . For a nonleast plateau to exist with respect to A there must exist candidates $a, b \in C$ such that a and b are adjacent in A , and a and b are among v 's most-preferred candidates or there exists a candidate $c \in C$ such that v states $a \sim b > c$. We know that the former is not possible, since v must have a single most-preferred candidate to be single-peaked with outside options. For the latter, since a and b are both strictly preferred to c we know that $a, b \in X \cup Y \cup Z$, but since a and b are adjacent on A and v states $a \sim b$ it must be the case that $a, b \in O_1 \cup O_2$. Therefore it is clear that v cannot contain a nonleast plateau.

For the other direction we consider the case where no preference order $v \in V$ contains a v-valley with respect to A and no preference order contains a nonleast plateau with respect to A .

Since no preference order $v \in V$ contains a v-valley with respect to A we know from Lemma 4.3.2 that V is possibly single-peaked with respect to A . Since we also know that no preference order $v \in V$ contains a nonleast plateau with respect to A it is easy to see that V is single-peaked with outside options with respect to A □

We adapt Construction 4.3.9 so that if a preference order contains a nonleast plateau with respect to an axis A , then when the columns of its preference matrix are permuted according to A the matrix will contain the sequence $\dots 1 \dots 0 \dots 1 \dots$. Since Construction 4.3.9 already ensures this for v-valleys and nonpeak plateaus, our extended construction below only needs to change what happens for plateaus that contain most-preferred candidates or least-preferred candidates in a given preference order.

Construction 4.3.15 *Let V be a preference profile of weak orders over candidate set C . For each $v \in V$ construct a $(\|C\| - 1) \times \|C\|$ matrix X_v . Each column of X_v corresponds to a candidate in C . For each candidate $c \in C$ let k be the number of candidates that are strictly preferred to c by v and let the corresponding column in matrix X_v contain k 0's starting at row one, with the remaining entries filled with 1's (as in Construction 4.3.3). The following extensions to Construction 4.3.3 ensure that if v has a nonleast plateau with respect to an axis A then when the columns of X_v are permuted according to A it will not have consecutive ones in rows.*

For each pair of candidates $a, b \in C$ such that v states $a \sim b$, they are not among the most-preferred candidates in v , they are not among the least-preferred candidates in v , and there is no candidate $c \in C - \{a, b\}$ such that v is tied among a, b , and c , then append three additional rows to the matrix X_v where the column corresponding to a is $\begin{bmatrix} 0 & 1 & 1 \end{bmatrix}'$, the column corresponding to b is $\begin{bmatrix} 1 & 1 & 0 \end{bmatrix}'$, each column corresponding to a candidate strictly preferred to a and b is $\begin{bmatrix} 1 & 1 & 1 \end{bmatrix}'$, and each column corresponding to a remaining candidate is $\begin{bmatrix} 0 & 0 & 0 \end{bmatrix}'$.

If there exist two candidates $a, b \in C$ such that v states $a \sim b$ and a and b are among v 's most-preferred candidates, then let X be a matrix that has no solution and halt.

If there exist three candidates $a, b, c \in C$ such that v states $a \sim b \sim c$ and they are not among v 's least-preferred candidates, then let X be a matrix that has no solution and halt.

After constructing a matrix X_v for each $v \in V$, all $\|V\|$ of the matrices are row-wise concatenated to yield a matrix X , except in the case where X has already been determined.

Construction 4.3.3 ensures that no preference order contains a v-valley and the extensions made in Construction 4.3.15 ensure that no preference order contains a nonleast plateau similarly to how Construction 4.3.9 extends Construction 4.3.3 to ensure that no preference order contains a nonpeak plateau. So the proof of the following theorem uses essentially the same argument as the proof of Theorem 4.3.11.

Theorem 4.3.16 *A preference profile V of weak orders is single-peaked with outside options consistent if and only if the matrix X , constructed using Construction 4.3.15 has the consecutive ones property.*

Corollary 4.3.17 *Single-peaked consistency with outside options for weak orders is in P.*

4.3.4 Single-Peaked Consistency

We now present our results for the strongest domain restriction for weak orders that we examine, single-peaked preferences. Recall that a preference order is single-peaked with respect to an axis A if and only if it is strictly increasing to a single most-preferred candidate (peak) and then strictly decreasing with respect to A . So we again use the v-valley substructure, but like the previous two cases of single-plateaued preferences and single-peaked preferences with outside options, we need an additional substructure. Even if no preference order has a v-valley with respect to A it may not be single-peaked because it is indifferent between two candidates on the same side of its peak or has more than one most-preferred candidate.

We can handle the first condition just mentioned with the nonpeak plateau substructure used in Section 4.3.2, but the second condition requires us to view *any* plateau as a forbidden substructure.

Definition 4.3.18 *A preference order v over a candidate set C contains a plateau with respect to an axis A if there exist candidates $a, b \in C$ such that a and b are adjacent in A and v states $a \sim b$.*

We can now use the plateau substructure and the v-valley substructure to state the following lemma.

Lemma 4.3.19 *Let V be a preference profile of weak orders. V is single-peaked with respect to an axis A if and only if no preference order $v \in V$ contains a v-valley with respect to A and no preference order $v \in V$ contains a plateau with respect to A .*

Proof. Let C be a candidate set, V be a preference profile of weak orders, and A be an axis.

If V is single-peaked with respect to A then for each preference order $v \in V$ the axis A can be split into segments X , Y , and Z such that v is strictly increasing along X , Y contains only the most-preferred candidate of v , and v is strictly decreasing along Z . Since v is only ever strictly decreasing along Z and Z is the rightmost segment of A , v cannot contain a v-valley with respect to A . For a plateau to exist with respect to A there must

exist candidates $a, b \in C$ such that a and b are adjacent on A , but clearly this is not possible since v is only every strictly increasing/decreasing along A .

For the other direction we consider the case when no preference order $v \in V$ contains a v-valley with respect to A and no preference order $v \in V$ contains a plateau with respect to A .

Since no preference order $v \in V$ contains a v-valley with respect to A we know from Lemma 4.3.2 that V is possibly single-peaked with respect to A . Since we also know that no preference order $v \in V$ contains a plateau with respect to A it is easy to see that V is single-peaked with respect to A . \square

We extend Construction 4.3.9 so that if a preference order contains a plateau with respect to an axis A , then when the columns of its preference matrix are permuted according to A the matrix will contain the sequence $\dots 1 \dots 0 \dots 1 \dots$. Since Construction 4.3.9 already ensures this for the case of nonpeak plateaus, our extended construction below only needs to add a condition for plateaus that contain the most-preferred candidates in a given preference order.

Construction 4.3.20 *Follow Construction 4.3.9 except add the following condition while constructing a preference matrix X_v for each preference order $v \in V$.*

If there exist two candidates $a, b \in C$ such that v states $a \sim b$ and they are among the most-preferred candidates in v , then output a matrix that has no solution.

Clearly the extension to Construction 4.3.9 above ensures that if there are multiple most-preferred candidates in a preference order then the preference matrix constructed from that order does not have the consecutive ones property.

When a preference order has a unique most-preferred candidate and is single-plateaued, it is clearly also single-peaked. Construction 4.3.20 ensures that no preference order contains more than one most-preferred candidate the same way that Construction 4.3.9 ensures that no preference order contains three or more candidates that are all tied and that are not among its most-preferred candidates, since this always results in a nonpeak plateau. So the proof of the following theorem follows from the proof of Theorem 4.3.11, but using Lemma 4.3.19 instead of Lemma 4.3.8.

Theorem 4.3.21 *A preference profile V of weak orders is single-peaked consistent if and only if the matrix X , constructed using Construction 4.3.20 has the consecutive ones property.*

Corollary 4.3.22 *Single-peaked consistency for weak orders is in P.*

4.4 Effect of Different Models on the Complexity of Manipulation

In this section we consider the computational impact of allowing ties in single-peaked votes on the complexity of Constructive Weighted Coalitional Manipulation (CWCM), compare this with the case of single-peaked total orders, and show that the complexity can depend on the model of single-peakedness for votes with ties used.

As is standard in the study of manipulation for single-peaked electorates, we follow the model introduced by Walsh [Wal07] where the societal axis is given as part of the input to the problem and the manipulators must state votes that are single-peaked with respect to this axis (in our case, for the corresponding model of single-peakedness). This is a natural setting, since it can be assumed that it is known how the candidates fall between the two extremes of the issue under consideration, and that the votes of the manipulators should agree with how the candidates are viewed by the electorate.

For total-order preferences, weighted manipulation for any fixed number of candidates is known to be NP-complete for every scoring rule that is not isomorphic to plurality or triviality [HH07]. For single-peaked total orders, Faliszewski et al. completely characterized the complexity of weighted manipulation for 3-candidate scoring rules [FHHR11] and this result was generalized by Brandt et al. for any fixed number of candidates [BBHH15]. These results both showed that the complexity of weighted manipulation for scoring rules often decreases when voters have tie-free single-peaked votes. We show that for single-peaked votes with ties the complexity of manipulation can either increase or remain the same depending on the model used.

This section on manipulation is split into two subsections that each examine a behavior of the complexity of manipulation when moving from tie-free votes to votes with ties for a given model of single-peakedness. In Section 4.4.1 we show that the complexity of manipulation can increase from the tie-free case to the case of votes with ties by considering possibly single-peaked top orders (and so these results also hold for single-peaked preferences with outside options). In contrast, in Section 4.4.2 we show that the complexity of manipulation remains the same for votes with ties for scoring rules and other natural systems for the standard model of single-peaked preferences and the model of single-plateaued preferences.

Most of our results in this section will be about scoring rules. We again repeat the example from Table 2.1 to recall the definitions of each of our scoring-rule extensions. Consider the candidate set $\{a, b, c, d\}$ and the vote $(a > b \sim c > d)$. We show the score assigned to each candidate using Borda (i.e., $\langle 3, 2, 1, 0 \rangle$) using each of our scoring-rule extensions.

Borda using min: $score(a) = 3$, $score(b) = score(c) = 1$, and $score(d) = 0$.

Borda using max: $score(a) = 3$, $score(b) = score(c) = 2$, and $score(d) = 0$.

Borda using round-down: $score(a) = 2$, $score(b) = score(c) = 1$, and $score(d) = 0$.

Borda using average: $score(a) = 3$, $score(b) = score(c) = 1.5$, and $score(d) = 0$.

4.4.1 Complexity Goes Up

Since for total orders, the standard model of single-peakedness and the model of possibly single-peaked preferences are equivalent, the following theorem from Faliszewski et al. [FHHR11] also holds for possibly single-peaked votes.

Theorem 4.4.1 [FHHR11] *3-candidate Borda CWCM for single-peaked total orders is in P.*

We now consider the complexity of 3-candidate Borda CWCM for top orders that are possibly single-peaked, and we observe a surprising increase in complexity. It is important to note that the related work by Narodytska and Walsh [NW14] did not find a natural case where the complexity of manipulation increases when moving from total orders to a vote with ties.

Theorem 4.4.2 *3-candidate Borda CWCM for possibly single-peaked top orders using max is NP-complete.*

Proof. Let $\{k_1, \dots, k_t\}$ be an instance of Partition, which asks if there exists a subset A of $\{k_1, \dots, k_t\}$ such that $\sum A = K$.

We construct the following instance of manipulation. Let the candidate set $C = \{a, b, p\}$ and set $a < p < b$ as the single-peaked axis. We have two nonmanipulators with the following weights and votes.

- One weight $3K$ nonmanipulator voting $(a > p \sim b)$.
- One weight $3K$ nonmanipulator voting $(b > p \sim a)$.

From the nonmanipulators, $score(p) = 6K$, while $score(a)$ and $score(b)$ are both $9K$.

Let there be t manipulators, with weights k_1, \dots, k_t , and let p be the preferred candidate. Without loss of generality, all of the manipulators put p first. Then p receives a score of $10K$ overall. However, a and b can score at most K each from the votes of the manipulators, for p to be a winner. So the manipulators must split their votes so that a subcollection of manipulators with weight K votes $(p > a > b)$ and a subcollection with weight K votes $(p > b > a)$. Notice that these are the only votes possible to ensure that p wins and that

the manipulators cannot simply all vote ($p > a \sim b$) since both a and b receive a point from that vote (since we are using max) and we have no points to spare. \square

The above argument for max does not immediately apply to the other scoring-rule extensions. In particular, for min the optimal vote for the manipulators is always to rank p first and to rank the remaining candidates tied and less preferred than p (as in the proof of Proposition 3 of Narodytska and Walsh, which considered the general case of top orders [NW14]). So that case is in P, with an optimal manipulator vote of ($p > a \sim b$).

However, it is not hard to modify the proof to show that the reduction used in the proof of Theorem 4.4.2 also works for the round-down case.

Theorem 4.4.3 *3-candidate Borda CWCM for possibly single-peaked top orders using round-down is NP-complete.*

Proof. Let $\{k_1, \dots, k_t\}$ be an instance of Partition, which asks if there exists a subset A of $\{k_1, \dots, k_t\}$ such that $\sum A = K$.

We construct the following instance of manipulation. Let the candidate set $C = \{a, b, p\}$ and set $a < p < b$ as the single-peaked axis. We have two nonmanipulators with the following weights and votes.

- One weight $3K$ nonmanipulator voting ($a > p \sim b$).
- One weight $3K$ nonmanipulator voting ($b > p \sim a$).

From the nonmanipulators, $score(p)$ is 0, while $score(a)$ and $score(b)$ are both $3K$.

Let there be t manipulators, with weights k_1, \dots, k_t , and let p be the preferred candidate. Without loss of generality, all of the manipulators put p first. Then p receives a score of $4K$ overall and a and b must score at most K each from the votes of the manipulators, for p to win. So the manipulators must split their votes so that a subcollection of manipulators with weight K votes ($p > a > b$) and a subcollection with weight K votes ($p > b > a$) so that their votes do not increase the scores of a and b too much. \square

The average scoring-rule extension case is more complicated since it is less close to Partition than the previous cases. We are still be able to show NP-completeness, but we have to reduce from the special, restricted version of Partition that we defined previously in Section 3.4 as Partition'.⁵

⁵A similar situation occurred in the proof of Proposition 5 in Narodytska and Walsh [NW14], where a (very different) specialized version of Subset Sum was constructed to prove that 3-candidate Borda CWCM (in the non-single-peaked case) for top orders using average remained NP-complete.

Theorem 4.4.4 *3-candidate Borda CWCM for possibly single-peaked top orders using average is NP-complete.*

Proof. Let $\{k_1, \dots, k_t\}, \widehat{K}$ be an instance of Partition'. We are asking whether there exists a partition (A, B, C) of $\{k_1, \dots, k_t\}$ such that $\sum A = \sum B + \widehat{K}$. Recall that all numbers involved are even. Let k_1, \dots, k_t sum to $2K$. Without loss of generality, assume that $\widehat{K} \leq 2K$.

We construct the following instance of manipulation. Let the set of candidates be $\{a, b, p\}$ and set $a < p < b$ as the single-peaked axis. We have two nonmanipulators with the following weights and votes.

- One weight $6K + \widehat{K}$ nonmanipulator voting $(a > p \sim b)$.
- One weight $6K - \widehat{K}$ nonmanipulator voting $(b > p \sim a)$.

From the nonmanipulators, $score(p)$ is $6K$, $score(a) + score(b) = 30K$ and $score(a) - score(b) = 3\widehat{K}$.

Let there be t manipulators, with weights $3k_1, \dots, 3k_t$, and let p be the preferred candidate.

First suppose there exists a partition (A, B, C) of $\{k_1, \dots, k_t\}$ such that $\sum A = \sum B + \widehat{K}$. For every $k_i \in A$, let the weight $3k_i$ manipulator vote $(p > b > a)$. For every $k_i \in B$, let the weight $3k_i$ manipulator vote $(p > a > b)$. For every $k_i \in C$, let the weight $3k_i$ manipulator vote $(p > a \sim b)$. Notice that after this manipulation that $score(p) = 18K$, $score(a) = score(b)$, and $score(a) + score(b) = 30K + 6K$. It follows that $score(p) = score(a) = score(b) = 18K$.

For the converse, suppose that p can be made a winner. Without loss of generality, assume that p is ranked uniquely first by all manipulators. Then $score(p) = score(a) = score(b) = 18K$. Let A' be the set of weights of manipulators that vote $(p > b > a)$, let B' be the set of weights of manipulators that vote $(p > a > b)$, and let C' be the set of weights of manipulators that vote $(p > a \sim b)$. No other votes are possible. Let $A = \{k_i \mid 3k_i \in A'\}$, $B = \{k_i \mid 3k_i \in B'\}$, and $C = \{k_i \mid 3k_i \in C'\}$. Therefore (A, B, C) corresponds to a partition of $\{k_1, \dots, k_t\}$. Note that $\sum A = \sum B + \widehat{K}$. \square

The results stated above for 3-candidate Borda for possibly single-peaked top orders also hold for the case of single-peaked preferences with outside options, since the two models are equivalent for top orders. Menon and Larson [ML16] later systematically studied the complexity of weighted manipulation for single-peaked preferences with outside options for top orders and found that for 3-candidate scoring rules and other natural systems that the

complexity generally increases [ML16]. In the next section we show that no such increase occurs for votes with ties for scoring rules and other natural systems for the standard model of single-peaked preferences, and for single-plateaued preferences.

4.4.2 Complexity Remains the Same

In this section we show that for the standard model of single-peakedness and for single-plateauedness that the complexity of weighted manipulation for m -candidate scoring rules using max, min, round-down, and average does not increase when moving from total orders to top orders, bottom orders, or weak orders.

We first present the following lemma, which will be used in the proofs of our results. Note that a similar lemma is shown by Brandt et al. [BBHH15] for manipulation of single-peaked total orders, but it requires some care to see that our similar result holds for votes with ties for each of the four single-peaked models.

Lemma 4.4.5 *If p can be made a winner by a manipulation of top-order, bottom-order, or weak-order votes that are single-peaked, single-plateaued, possibly single-peaked, or single-peaked with outside options for a scoring rule using max, min, round-down, or average, then p can be made a winner by a manipulation (of the same type) in which all manipulators rank p uniquely first.*

Proof. Suppose that a manipulator has a vote of the form $(X_{\text{wo}} > P > Y_1 > \dots > Y_\ell)$ where X_{wo} is a weak order over X (where X can be empty), $\ell \geq 0$, each of P, Y_1, \dots, Y_ℓ are nonempty groups of tied candidates, and $p \in P$. We have the following two cases.

Case 1: If all of the candidates in P appear on the same side of the candidates in X with respect to the axis (or X is empty), then we can change the vote of the manipulator to $(p > (P - \{p\})_{\text{to}} > X_{\text{to}} > Y_1 > \dots > Y_\ell)$ where $(P - \{p\})_{\text{to}}$ and X_{to} are total orders such that the vote satisfies the given notion of single-peakedness with ties with respect to the axis.

Case 2: Otherwise, we let P' be the candidates to the left of X and P'' be the candidates to the right of X with respect to the axis. Without loss of generality let $p \in P'$. We can then change the vote of the manipulator to $(p > (P' - \{p\})_{\text{to}} > X_{\text{to}} > P''_{\text{to}} > Y_1 > \dots > Y_\ell)$.

Notice that in all of the above cases and for all candidates $c \in C - \{p\}$ that $\text{score}(p) - \text{score}(c)$ does not decrease, regardless of which scoring-rule extension is used. So if p is a winner with the initial vote, p is still a winner in the vote that ranks p uniquely first. \square

The following results carefully adapt the arguments used by Brandt et al. [BBHH15] for manipulation of single-peaked total orders to the case of votes with ties for single-peaked and for single-plateaued votes.

Theorem 4.4.6 *Let $\alpha = \langle \alpha_1, \alpha_2, \dots, \alpha_m \rangle$ be a scoring vector. If α -CWCM is in P for single-peaked total orders, then α -CWCM is in P for single-peaked and single-plateaued top orders, bottom orders, and weak orders for all our scoring rule extensions.*

Proof. If the given axis has p in the leftmost or rightmost position, then, by Lemma 4.4.5, we can assume that all manipulators rank p uniquely first. There is exactly one such single-peaked or single-plateaued vote with ties, namely to rank p first and strictly rank the remaining candidates according to their position on the axis.

So, let the axis be $a_{m_1} < \dots < a_1 < p < b_1 < \dots < b_{m_2}$.

We first consider the P cases for single-peaked total orders described in Lemma 6.6 of Brandt et al. [BBHH15] where for all $i, j > 1$ such that $i + j \leq m + 1$ it holds that $(\alpha_1 - \alpha_i)(\alpha_1 - \alpha_j) \leq (\alpha_i - \alpha_{i+1})(\alpha_j - \alpha_{j+1})$.

The crucial observation in the proof of this lemma is that for every set of single-peaked total order votes S , we are in one of the following two cases:

1. For all i , $1 \leq i \leq m_1$, $score_S(a_i) \leq score_S(p)$.
2. For all i , $1 \leq i \leq m_2$, $score_S(b_i) \leq score_S(p)$.

We will show that this is also the case for single-peaked and single-plateaued orders with ties. This then implies that the optimal vote for all manipulators is $(p > b_1 > \dots > b_{m_2} > a_1 > \dots > a_{m_1})$ (in case 1) or $(p > a_1 > \dots > a_{m_1} > b_1 > \dots > b_{m_2})$ (in case 2).

For the sake of contradiction let S be a collection of nonmanipulators and i_1 and i_2 be integers such that $1 < i_1 \leq m_1 + 1$ and $1 < i_2 \leq m_2 + 1$ such that $score_S(a_{i_1}) > score_S(p)$ and $score_S(b_{i_2}) > score_S(p)$.

We now create an \widehat{S} such that $score_{\widehat{S}}(a_{i_1}) > score_{\widehat{S}}(p)$ and $score_{\widehat{S}}(b_{i_2}) > score_{\widehat{S}}(p)$ and the argument in the proof of Lemma 6.6 in Brandt et al. [BBHH15] still holds.

- First, remove from S all voters that have p tied for first. Notice that if we remove all votes with p tied for first from S then $score(p)$ does not increase and neither $score(a_{i_1})$ nor $score(b_{i_2})$ decreases. So we can assume that p is never tied for first in a vote (although this does not mean that other candidates cannot be tied for first in a vote when we are in the single-plateaued case).

- We know that p is not tied for first, so p is tied with at most one other candidate in each vote. So in every vote break the tie against p . Notice that this can lower only the score of p regardless of whether we are using the min, max, round-down, or average scoring-rule extension.

Let ℓ_i be the total weight of the voters in \widehat{S} that rank some candidate in $\{a_1, \dots, a_{m_1}\}$ first and p in the position of α_i (in every extension other than round-down this means that p is ranked i th; for round-down we mean that p is ranked in the α_i group (i.e., p is followed by $m - i$ groups)). Using a similar argument to the proof of Lemma 6.6 from Brandt et al. [BBHH15] with \widehat{S} as described above we can reach a contradiction.

We now present the remaining P cases, which can be seen as the with-ties analogue of the P cases of Lemma 6.7 from Brandt et al. [BBHH15].

We examine each of the cases individually. For convenience, we normalize the scoring vector so that $\alpha_m = 0$.

If $\alpha_2 = 0$ then clearly the optimal action for the manipulators is to vote a total order with p uniquely first and the remaining candidates ranked arbitrarily such that the vote is single-peaked.

If $\alpha_1 = \alpha_{\lfloor \frac{m-1}{2} \rfloor + 2}$, then, as pointed out in [BBHH15], Lemma 6.6 applies, which has been handled above.

The last P-time case is that $\alpha_1 \leq 2\alpha_2$, $\alpha_1 > \alpha_2 > 0$, and $\alpha_2 = \alpha_{m-1}$.

Notice that when the votes are single-peaked or when they are single-plateaued then a_{m_1} and b_{m_2} are the only candidates that can occur last in a vote, either uniquely or tied (except for the case where all candidates are tied, but we can easily ignore these votes since they do not change the relative scores between candidates).

Since we know from Lemma 4.4.5 that all of the manipulators can put p strictly ranked first, we can have manipulator votes of the following forms.

- $(p > \dots > a_{m_1})$.
- $(p > \dots > b_{m_2})$.
- $(p > \dots > a_{m_1} \sim b_{m_2})$.

Notice that for all votes with p first that $score(p) \geq \alpha_2$.

For all votes, $score(a_{m_1}) = 0$ or $score(b_{m_2}) = 0$, or a_{m_1} and b_{m_2} are tied for last, in which case $score(a_{m_1}) + score(b_{m_2}) \leq 2\alpha_2$ (the worst case occurs with max where they each receive a score of α_{m-1}).

So, $score_S(a_{m_1}) + score_S(b_{m_2}) \leq 2score_S(p)$. If $score_S(a_{m_1}) \leq score_S(p)$, $(p > a_1 > \dots > a_{m_1} > b_1 > \dots > b_{m_2})$ is the optimal manipulator vote. Otherwise, $score_S(b_{m_2}) \leq score_S(p)$ and $(p > b_1 > \dots > b_{m_2} > a_1 > \dots > a_{m_1})$ is the optimal manipulator vote.

We have proven that the above cases hold for single-peaked and for single-plateaued weak orders. It is clear that the case for single-peaked bottom orders follows from the case of single-peaked total orders (since they are equivalent) and how to adapt the above arguments to hold for single-peaked top orders, and for single-plateaued top orders and bottom orders. \square

It is known that 3-candidate Copeland $^\alpha$ CWCM for all rational $\alpha \in [0, 1)$ is NP-complete for total orders in the nonunique-winner model [FHS08], and when $\alpha = 1$ (also known as Llull) CWCM is in P for $m \leq 4$ [FHS08, FHS12] and the cases for $m \geq 5$ remain open. In Chapter 3 we showed that the NP-completeness of the 3-candidate case holds for top orders, bottom orders, and weak orders, and Menon and Larson [ML15] independently showed the top-order case. We also showed in Chapter 3 that 3-candidate Llull CWCM is in P for top orders, bottom orders, and weak orders.

Recall that weak Condorcet winners always exist when preferences are single-peaked with ties, and when they are single-plateaued. Thus the result that Llull CWCM for single-peaked total orders is in P from Brandt et al. [BBHH15] also holds for the case of single-peaked and single-plateaued top orders, bottom orders, and weak orders.

Copeland $^\alpha$ for all rational $\alpha \in [0, 1)$ was shown to be in P by Yang [Yan15] for single-peaked total orders.

Theorem 4.4.7 [Yan15] *Copeland $^\alpha$ CWCM for all rational $\alpha \in [0, 1)$ is in P for single-peaked total orders.*

In contrast, for top orders that are single-peaked with outside options, it is NP-complete even for three candidates [ML16].

Theorem 4.4.8 [ML16] *3-candidate Copeland $^\alpha$ CWCM for all rational $\alpha \in [0, 1)$ is NP-complete for top orders that are single-peaked with outside options.*

For single-peaked and single-plateaued weak orders, bottom orders, and top orders we again inherit the behavior of single-peaked total orders.

Theorem 4.4.9 *Copeland $^\alpha$ CWCM for all rational $\alpha \in [0, 1)$ is in P for single-peaked and single-plateaued top orders, bottom orders, and weak orders.*

Proof. Let A be our axis. Consider a set of nonmanipulators with single-plateaued weak orders. Replace each nonmanipulator v of weight w with two nonmanipulators v_1 and v_2

of weight w . The first nonmanipulator breaks the ties in the vote in increasing order of L and the second nonmanipulator breaks the ties in the vote in decreasing order of A , i.e., if v states $a \sim b$ and $a < b$ in A , then set v_1 to state $a > b$ and v_2 to state $b > a$. Note that v_1 and v_2 are single-peaked total orders and that the weighted majority graph induced by the nonmanipulators after replacement can be obtained from the weighted majority graph induced by the original nonmanipulators by multiplying each weight in the graph by 2. When we also multiply the manipulator weights by 2, we have an equivalent Copeland ^{α} CWCM problem, where all nonmanipulators are single-peaked total orders and all manipulators have even weight.

Suppose p can be made a winner by having the manipulators cast single-plateaued votes with ties. Now replace each manipulator of weight $2w$ by two weight- w manipulators. The first manipulator breaks the ties in the vote in increasing order of L and the second manipulator breaks the ties in the vote in decreasing order of L . Now the replaced manipulator votes are single-peaked total orders and p is still a winner. We need the following fact from the proof of Theorem 4.4.7: *if p can be made a winner in the single-peaked total order case, then p can be made a winner by having all manipulators cast the same p -time computable vote.* It follows that p can be made a winner by having all replaced manipulators cast the same single-peaked total order vote. But then p can be made a winner by having all original manipulators cast the same single-peaked total order vote. Since this vote is P-time computable, it follows that Copeland ^{α} CWCM for all rational $\alpha \in [0, 1)$ is in P for single-plateaued weak orders, and this also holds for single-peaked weak orders since every single-peaked profile of weak orders is also single-plateaued.

It is clear to see that similar arguments hold for single-peaked top orders, and for single-plateaued top orders and bottom orders. The case for single-peaked bottom orders follows from the case for single-peaked total orders. \square

The most surprising result in the related work by Menon and Larson [ML16] on the complexity of manipulation for single-peaked preferences with outside options for top orders concerns the elimination extension to the veto scoring rule.

For elimination veto for total orders, a candidate with the lowest score (using the veto scoring rule) is eliminated, and the rule is repeated on the votes restricted to the remaining candidates until there is one candidate left (see [CT07]). For comparison with related work, we use the unique winner model for these cases, and for votes with ties we use the min extension to the veto scoring rule.

Menon and Larson [ML16] find that the complexity of 3-candidate CWCM for elimination veto for single-peaked preferences with outside options for top orders using min is NP-complete, whereas for single-peaked total orders [ML16] and even for total orders in the

general case [CT07] it is in P.

Theorem 4.4.10 [CT07] *m-candidate elimination veto CWCM for total orders is in P in the unique-winner model.*

Theorem 4.4.11 [ML16] *m-candidate elimination veto CWCM for single-peaked total orders is in P in the unique-winner model.*

Theorem 4.4.12 [ML16] *3-candidate elimination veto CWCM for top orders that are single-peaked with outside options using min is NP-complete in the unique-winner model.*

Menon and Larson [ML16] state this case as a counterexample to the conjecture by Faliszewski et al. [FHHR11], which states that the complexity for a natural election system will not increase when moving from the general case to the single-peaked case. (Though Menon and Larson do qualify that the conjecture from Faliszewski et al. concerned total orders.) However, for the standard model of single-peaked preferences and for single-plateaued preferences, elimination veto CWCM for top orders, bottom orders, and weak orders using min is in P for any fixed number of candidates, thus the counterexample crucially relies on using a nonstandard definition.

Theorem 4.4.13 *m-candidate elimination veto CWCM for single-peaked and for single-plateaued top orders, bottom orders, and weak orders using min is in P in the unique-winner model.*

Proof. The proof of this theorem follows from an argument similar to the proof of the theorem for single-peaked total orders from Menon and Larson [ML16], which follows from the fact that for total orders the candidate eliminated after each round is on the leftmost or rightmost location of the axis, and the reverse of an elimination order is single-peaked with respect to the axis, where an elimination order is the total order created by appending the candidate eliminated after each round.

It is easy to see that both of these statements also hold for single-peaked votes with ties, since the only candidates that can be vetoed in each round are still the leftmost and rightmost candidates on the axis.

It is also clear that Lemma 12 from Coleman and Teague [CT07] still holds, which states that if there exists a collection of votes that can induce an elimination order then it can be induced by all of the manipulators voting the reverse of the elimination order, and so all of the manipulators can vote the same vote. So since we know that the reverse of an elimination order is single-peaked with respect to the axis, and that there are only polynomially many

possible elimination orders with p first, the manipulators simply try each elimination order with p first.

Since the elimination order is always a total order, the above argument clearly holds for single-peaked top orders, bottom orders, and weak orders.

For single-plateaued preferences it is possible for voters with multiple most-preferred candidates to veto more than two candidates (after all of the candidates ranked below their most-preferred candidates have been eliminated). But this case only becomes an issue when there are more than two candidates that all of the voters have among their most-preferred candidates, and since we can assume without loss of generality that the manipulators all vote a single-peaked total order with p first, this holds for the single-plateaued case as well. \square

4.5 Conclusions

The standard model of single-peakedness is naturally defined for votes with ties. We presented three additional variants of single-peaked preferences for votes with ties, and showed that each corresponding consistency problem is in P, by reductions to the problem of determining if a 0-1 matrix has the consecutive ones property.

We also considered the complexity of manipulation for single-peaked votes with ties, and though we found an anomalous increase in complexity from votes without ties to votes with ties for the possibly single-peaked model, we find that for scoring rules and other important natural systems, the complexity of weighted manipulation does not increase when moving from total orders to votes with ties in the standard single-peaked and in the single-plateaued cases. Single-peaked and single-plateaued preferences for votes with ties also retain the important social-choice property of the existence of weak Condorcet winners. This is not to say that possibly single-peaked and single-peaked preferences with outside options are without merit, since they both model easily understood structure in preferences.

Single-peaked preferences are studied because they are a simply stated and important domain restriction that gives insight into how the voters view the candidates, and elections with single-peaked voters have nice properties. However, a recent experimental study suggests that in real-world settings voters are often not single-peaked [MFG12], but in this study the single-peaked results only used Black's definition for total orders. It would be interesting to see if real-world datasets of weak orders contain voters that are single-peaked, single-plateaued, single-peaked with outside options, or possibly single-peaked.

Chapter 5

Models for the Representation: Succinct Elections

5.1 Introduction

In the computational study of elections, it is generally assumed that the preferences of the voters are represented as a list of votes. Though this may be a reasonable representation for paper ballots in political elections, in artificial intelligence applications a more succinct representation where the preferences of the electorate are represented as a list of distinct votes and their counts may be more natural. For example, this representation is used by the online preference repository PREFLIB for election data [MW13].

In this chapter we consider how this succinct representation of the voters can affect the complexity of different election problems, and contrast this with the case of weighted voters.

Though the succinct representation may be exponentially smaller than the nonsuccinct representation, we find that in surprisingly few cases the complexity of election problems increases. Related work that considers succinct elections did not find a case where the complexity increases [Rus07, FHHR09, FHH09, HHR09, FHHR11]. We explain this phenomenon by showing that many common proof techniques that show that election problems are in P can be adapted to the succinct case. However, this does not mean that it is not possible for the complexity to increase. For example, consider the following election system X . Every voter votes by an approval vector (i.e., states an approval or disapproval for each candidate), and a candidate wins if they are approved by exactly half of the voters. Though simple, this election system is not particularly natural. It has been chosen to show that an increase in complexity is possible. Consider the problem of constructive control by adding candidates (CCAC) for election system X . Given an election where the voters are represented succinctly,

a set of unregistered candidates, and a preferred candidate p it is NP-complete to ensure that p wins the election using system X by adding unregistered candidates to the election. This is because this problem is basically subset sum, which is in P for unary numbers by dynamic programming and is NP-complete for binary numbers [GJ79]. It is straightforward to see that X -CCAC is in P and X -Succinct-CCAC is NP-complete.¹ (We mention in passing that X -Weighted-CCAC is also NP-complete.)

We did find a natural case where the complexity increases. The complexity of determining the winner in a Kemeny election, which is well-known to be Θ_2^P -complete for the case of nonsuccinct votes [HSV05], is Δ_2^P -complete for succinct votes. This solves an open problem from the previous work by Hemaspaandra, Spakowski, and Vogel [HSV05].

Throughout this chapter we consider succinct elections. Recall that our standard definition for elections (which we refer to as *nonsuccinct elections* in this chapter to discern between the two models) consists of a finite set of candidates C and a list of voters V , where each voter $v \in V$ has a preference order over the candidates. In a *succinct election* V is not a list of voters, but instead a list of distinct votes v (preference orders) and their positive integer count $\kappa(v)$. To see how these two definitions compare, consider the nonsuccinct election consisting of $C = \{a, b, c\}$ and the following four voters: $(a > b > c)$, $(a > b > c)$, $(a > b > c)$, and $(b > c > a)$. This list of four voters can be represented succinctly as: $3 \times (a > b > c)$ and $1 \times (b > c > a)$.

We will compare complexity results between succinct and weighted elections. In a weighted election V is still a list of voters, but each $v \in V$ has an associated positive integer weight $\omega(v)$, and can be thought of as a coalition of $\omega(v)$ voters all voting the same. It is important to note that weights are indivisible, while the counts in a succinct election are not. Unlike the previous two chapters, here we assume that the preference order of each voter is a total order, i.e., she strictly ranks each candidate from most to least preferred.

We consider how moving from the nonsuccinct case to the succinct case affects the complexity of different election problems; and so we briefly restate the nonsuccinct definitions

¹To show that X -CCAC is in P, let s_1, \dots, s_m be the approval scores of candidates c_1, \dots, c_m , let the c_1 be the preferred candidate and let c_1, \dots, c_ℓ be the registered candidates. We want to know if there exists a subset I of $\{\ell + 1, \dots, m\}$ of size at most k such that $2s_1 = \sum_{i=1}^{\ell} s_i + \sum_{i \in I} s_i$. This can easily be computed in polynomial time, using dynamic programming. Simply let, for all $\ell + 1 \leq i \leq m$, $0 \leq k' \leq k$, $0 \leq t \leq \sum_{1 \leq i \leq m} s_i$, $S[i, t, k'] = 1$ if and only if there exists a size- k' subset I of $\{\ell + 1, \dots, m\}$ such that $\sum_{i=1}^{\ell} s_i + \sum_{i \in I} s_i = t$. This can be done in polynomial time, since there are only polynomially many i , k' , and t to consider.

When the input is given succinctly, there are exponentially many values for t possible. In this case, we can show that it is NP-complete. We reduce from Partition: Given $\{k_1, \dots, k_m\}$ such that $\sum k_i = 2K$, we need to determine if a subset of these integers sums to K . We have $m + 1$ candidates, preferred candidate p and candidates c_1, \dots, c_m . We have K voters that approve of only p , and k_i voters that approve of only c_i . p is the only registered voter and the addition limit is m . It is immediate that there is a subset of k_1, \dots, k_m that sums to K if and only if p can be made a winner by adding candidates.

of these problems and describe how they change for the succinct case.

Recall the problem \mathcal{E} -Winner, which given an election (C, V) and a candidate $p \in C$, asks if p is a winner of (C, V) using election system \mathcal{E} . In the succinct case, \mathcal{E} -Succinct-Winner, V is represented succinctly, and in the weighted case, \mathcal{E} -Weighted-Winner, each voter has a corresponding positive integer weight.

In addition to the winner problem, we examine the complexity of manipulation and constructive control by adding voters. Formal definitions for each of these problems can be found in Sections 2.3.1 and 2.3.2 respectively, but we informally restate them here, and describe how the succinct and weighted cases differ.

The coalitional unweighted manipulation problem (CUCM) asks if given an election, a list of manipulative voters, and a preferred candidate $p \in C$, if there is a way to set the votes of the manipulators such that p wins [BTT89a, CSL07].

The weighted case (CWCM) is essentially the same definition as above, except each of the voters (both the nonmanipulators and the manipulators) have an associated positive integer weight [CSL07]. In the corresponding succinct case, the nonmanipulator votes are represented succinctly, and the manipulators are represented as k encoded in binary.

Constructive control by adding voters (CCAV) asks if given an election, a list of unregistered voters, a limit k , and a preferred candidate p , if the election chair can ensure that p wins by adding at most k of the unregistered voters [BTT92].

In the standard model of weighted voter control the parameter k denotes the number of weighted voters the chair can add/delete [FHH15]. In the succinct case, the only change from the \mathcal{E} -CCAV definition above is that the voters (registered and unregistered) are represented succinctly.

As with much of the other computational complexity results in this thesis, many of our results here concern the classes P and NP. However, we also have several results concerning the classes Θ_2^p and Δ_2^p . Recall that $\Delta_2^p = \text{P}^{\text{NP}}$, i.e., the class of problems solvable in P with access to an NP oracle, and that $\Theta_2^p = \text{P}^{\text{NP}[\log]}$ is the class of problems solvable in P with a log number of queries to an NP oracle (which is equivalent to the class of problems solvable by a P-machine that can ask one round of parallel queries to an NP oracle [Hem89]).

Most of our polynomial-time results for the succinct case of problems modify the proof of the nonsuccinct case, but some interesting results will use the well-known result due to Lenstra Jr., which shows that even though solving an integer linear program is NP-complete in the general case [Kar72, BT76], it is in P when the number of variables is fixed [Len83]. Informally, an integer linear program is a system of linear inequalities with integer variables and coefficients.

5.2 Adapting Approaches

The general theme of this chapter is that even though we would expect the complexity of election problems to increase (with respect to the length of the input) when the voters are represented succinctly, we find that in surprisingly few cases such an increase occurs. In this section we will discuss several common approaches used to show that election problems are in P (for the nonsuccinct case), and describe how they can be adapted for the case where voters are represented succinctly, sometimes straightforwardly and sometimes in a more complicated way.

In particular, we show how greedy approaches, limited brute-forcing, network flow, and edge matching/cover techniques can be adapted. To showcase these adaptations, we will show that the dichotomy result for constructive control by adding voters (CCAV) for pure scoring rules [HHS14a] holds for the succinct case.

A *pure scoring rule* defines a family of scoring vectors where the m -candidate scoring vector can be computed in polynomial time in m , and the $m + 1$ -candidate scoring vector can be obtained from the m -candidate scoring vector by adding a single coefficient [BD09]. Recall that a scoring rule describes a family of scoring vectors $\langle \alpha_1, \alpha_2, \dots, \alpha_m \rangle$, $\alpha_i \geq \alpha_{i+1}$ for each m -candidate input, where each candidate receives a score of α_i from each vote where she is ranked i th, and the candidate(s) with the highest total score win.

Assuming $P \neq NP$, the following pure scoring rules are asymptotically (i.e., for a large enough number of candidates) the only cases where CCAV is in P, both for the nonsuccinct and the succinct case.

- $\langle \alpha, \beta, 0, \dots, 0 \rangle$, where $\alpha > \beta$.
- t -Approval, $\langle \underbrace{1, \dots, 1}_t, 0, \dots, 0 \rangle$, where $t \leq 3$.
- t -Veto, $\langle 1, \dots, 1, \underbrace{0, \dots, 0}_t \rangle$, where $t \leq 2$.
- $\langle 2, 1, \dots, 1, 0 \rangle$.

In contrast, all weighted cases are NP-complete, except triviality (i.e., a scoring rule where every candidate scores the same), 1-approval (plurality), 2-approval, and 1-veto [FHH15, Lin12].

The following case demonstrates the adaptation of greedy algorithms as well as how to handle limited brute-forcing.

Theorem 5.2.1 *For $\alpha > \beta \geq 0$, $\langle \alpha, \beta, 0, \dots, 0 \rangle$ -CCAV is in P for succinct elections.*

Proof. We follow the proof of the claim from [HHS14a], which is found in its corresponding technical report [HHS14b], and show how to adapt this to the succinct case.

The proof shows that there is a constant ℓ such that it is never better to add ℓ votes that put p second than it is to add ℓ different votes that put p first (assuming that votes that rank the first two candidates identically also rank all other candidates identically). Let V_1 be the set of unregistered voters that put p first and let V_2 be the set of unregistered voters that put p second.

So, if p can be made a winner, p can be made a winner by adding at most ℓ voters from V_2 or in a way that does not leave ℓ different votes in V_1 unused.

In the first case, we can brute-force over all sets of V_2 voters of size $\leq \ell$. Add these voters. We are then left with the problem of checking whether we can add at most k' V_1 voters to an election to make p a winner, where k' is k minus the number of V_2 voters added. In Hemaspaandra, Hemaspaandra, and Schnoor [HHS14a], this is done by brute-forcing over every j in $\{0, \dots, k'\}$ and checking whether we can add j voters from V_1 to make p a winner. If we know j , we know what the score of p after control will be, and so the last part can then be done greedily, by for each candidate adding as many voters that put a second as possible one voter at the time (until we have added j voters total).

We have two problems in the succinct case. The first is that we cannot in polynomial time brute-force over all possible values of j , since k' is not polynomially bounded. The second problem is that we don't have the time to add voters one at the time. The second problem can of course easily be fixed by, for each candidate a , adding as many voters with a in second place as possible all at once: If s_a is the initial score of a and fs_p is the final score of p (note that $fs_p = s_p + j\alpha$), we can add at most $\lfloor (fs_p - s_a)/\beta \rfloor$ voters with a second.

To handle the first problem, note that we can formulate the problem as an integer linear program with one variable (j), namely, there is a j , $0 \leq j \leq k'$, such that

$$\sum_{a \in C - \{p\}} \min(\lfloor (s_p + j\alpha - s_a)/\beta \rfloor, v_a) \geq j,$$

where v_a is the number of voters in V_1 that rank a second.² This integer linear program can be solved in polynomial time by [Len83].

It remains to show that the second case, where p can be made a winner in a way that does not leave ℓ different votes in V_1 unused, can be determined in polynomial time in the succinct case.

²Notice that the ILP above checks that if we add the maximum number of voters such that the score of each of the non- p candidates is at most $s_p + j\alpha$, then we have added at least j voters. This implies that we can add j voters in such a way that p is still a winner.

For the nonsuccinct case, in Hemaspaandra, Hemaspaandra, and Schnoor [HHS14a] this case is handled by brute-forcing over all sets S of at most $\ell - 1$ candidates who are ranked second by unused V_1 voters, and then brute-forcing over all possible sets of unused V_1 voters consistent with S . We can describe a set of voters consistent with S as a function $u : S \rightarrow \{1, \dots, \|V_1\|\}$ such that $u(c)$ is the number of unused V_1 voters that rank c second. We can brute-force over all such functions u in polynomial time. This way, we loop over all possible sets of unused V_1 voters, and so we also loop over all possible sets of added V_1 voters. After adding the V_1 voters, we need to see if p can be made a winner by adding V_2 voters. This can be done in a similar way as adding the V_1 voters in the first case above.

In the succinct case, we can still loop over all sets S of at most $\ell - 1$ candidates. We can not brute-force all functions u , but this will turn again into an integer linear program, this time with $\ell - 1$ variables (for the u values) and one more variable to handle the V_2 voters similarly to how we handled the V_1 voters in the first case above. This integer linear program can be solved in polynomial time by [Len83]. \square

Theorem 5.2.2 *CCAV is in P for t -approval when $t \leq 3$ and for t -veto when $t \leq 2$ for succinct elections.*

Proof. The results for 1-approval, 2-approval, and 1-veto follow via simple greedy algorithms that can easily be adapted to work in the succinct case.

3-approval-CCAV was shown to be in P by Lin [Lin12] using a reduction to Simple b -Edge Matching for Multigraphs (see [Sch03]). The essence of the reduction is that every unregistered voter of the form $(\{p, x, y\} > \dots)$ (by which we mean a voter who gives a point to p , x , and y) corresponds to an edge (x, y) in the constructed multigraph. For every candidate $c \neq p$, c is a vertex in the graph and $b(c)$ is the final score of p minus the initial score of c (i.e., the number of points that we can add to c while keeping c 's score less than or equal to p 's score). In the succinct case, we would have too many edges in the graph. However, Capacitated b -Edge Matching (for graphs where edges have integer capacities) is also in P (see [Sch03]). So, we simply set the capacity of edge (x, y) equal to the number of unregistered voters of the form $(\{p, x, y\} > \dots)$, and we let $b(c)$ be the final score of p minus the initial score of c as previously.

Similarly, 2-veto-CCAV was shown to be in P by reduction to Simple b -Edge Cover for Multigraphs. In that reduction, every unregistered voter of the form $(\dots > \{x, y\})$ corresponds to an edge (x, y) in the constructed multigraph. Again, we can modify this construction to a reduction to Capacitated b -Edge Cover, which is also in P, by letting the capacity of edge (x, y) be the number of unregistered voters of the form $(\dots > \{x, y\})$. \square

Another common technique to show that election problems are in P are reductions to network flow (e.g., we use a reduction to the problem of maximum network flow in the proof of Theorem 3.4.9).

Theorem 5.2.3 $\langle 2, 1, \dots, 1, 0 \rangle$ -CCAV is in P for succinct elections.

Proof. The essence of the nonsuccinct proof is to (after some easy preprocessing) build a min-cost flow network. The capacities of the edges are computed from the scores of the candidates and the multiplicities of the voters and are polynomially bounded in the size of the input. In the succinct case, we can use the exact same network. The capacities are now binary integers, but min-cost network flow is still in P in that case. \square

5.2.1 Dynamic Programming

There is one other common technique that is used to prove that election problems are in P. This is dynamic programming. As alluded to in the introduction, dynamic programming approaches do not generalize to the succinct case. This does not mean that the succinct case is not in P. Let's for example look at the result by Hemaspaandra and Schnoor [HS16] that shows that for all pure scoring rules with a constant number of different coefficients, manipulation is in P. This is shown by dynamic programming. And this algorithm is not in P for the succinct case, not even for restricted versions of this problem.

However, as just stated, for specific instances we may not need dynamic programming. For example, it is immediate that plurality and veto can be easily solved in polynomial time, in the nonsuccinct and the succinct case. For a more involved example, for scoring rule $\langle 0, \dots, 0, -1, -2 \rangle$, one can greedily put a candidate $c \neq p$ with highest score in a -2 spot, and after filling all -2 spots, in a -1 spot. Note that this simple greedy algorithm does not directly generalize to the succinct case, since we do not have enough time to fill slots one-at-a-time.

Still, in this particular case it is not hard to see that the succinct case is in P: Greedily put a highest scoring candidate $c \neq p$ in $\lfloor (score(c) - score(p))/2 \rfloor$ -2 spots (or in as many of those as available). After filling all -2 spots, fill the -1 spots.

Things quickly get more complicated though. And there is a clear limit to how far this will generalize, since Borda manipulation is NP-complete [DKNW11, BNW11].

Still, we conjecture that all these manipulation cases are in P for succinct elections.

Again, this is in contrast to the weighted case, which is in P only for triviality and plurality. This holds even for fixed numbers of candidates, as explained in the next section.

5.2.2 Fixed Numbers of Candidates

It is reasonable to assume that the number of candidates in an election may be fixed. For the case of weighted voters many problems are hard, even when the number of candidates is fixed.

Theorem 5.2.4 1. *m -candidate α -CWCM is NP-complete for every scoring rule α that is not plurality or triviality [HH07].*

2. *m -candidate α -CCAV and m -candidate α -CCDV are each NP-complete for every scoring rule $\alpha = \langle \alpha_1, \dots, \alpha_m \rangle$, when $\|\alpha_1, \dots, \alpha_m\| \geq 3$, for weighted elections [FHH15].*

Faliszewski, Hemaspaandra, and Hemaspaandra [FHH09] showed that for succinct votes that manipulation is in P for any fixed number of candidates by describing an integer linear program with a fixed number of variables.

Theorem 5.2.5 [FHH09] *m -candidate α -CUCM is in P for every scoring rule α , for succinct elections.*

A similar approach can be used to show that for succinct votes, constructive control by adding/deleting voters is in P for every scoring rule for a fixed number of candidates.

Theorem 5.2.6 *m -candidate α -CCAV and m -candidate α -CCDV are each in P for every scoring rule $\alpha = \langle \alpha_1, \dots, \alpha_m \rangle$, for succinct elections.*

Proof. Given a set of candidates C , a list of registered votes V , a list of unregistered votes U , an addition limit $k \in \mathbb{N}$, and a preferred candidate $p \in C$ we can determine if the chair can ensure that p wins in polynomial time. The argument closely follows the approach used by Faliszewski, Hemaspaandra, and Hemaspaandra [FHH09] to show that m -candidate α -bribery and m -candidate α -CUCM, for every scoring rule α are each in P for succinct elections.

We describe an integer linear program with a fixed number of variables, which due to the result from Lenstra Jr. [Len83], can be solved in polynomial time.

Since the number of candidates m is fixed, we know that only $m!$ different preference orders are possible. We list these preference orders in order in the following way: $o_1, \dots, o_{m!}$. For $1 \leq i \leq m!$, let k_i be the number of voters in V with preference order o_i , and let k'_i be the number of voters in U with preference order o_i . Let $x_1, \dots, x_{m!}$ be variables. The constants used as input to our linear program are the coefficients of the scoring vector $\alpha_1, \dots, \alpha_m$, the counts for the registered voters $k_1, \dots, k_{m!}$, the counts for the unregistered voters $k'_1, \dots, k'_{m!}$, and the addition limit k . We have the following constraints.

We first need to ensure that all of our variables have a positive value. So, $\forall i, 1 \leq i \leq m,$

$$x_i \geq 0.$$

We also need to ensure that no more than the number of unregistered votes of each type are added. So, $\forall i, 1 \leq i \leq m!,$

$$x_i \leq k'_i.$$

The following constraint ensures that no more than k total votes are added.

$$\sum_{i=1}^{m!} x_i \leq k.$$

Our final constraint ensures that no other candidate has a score greater than p . (Note that $\text{pos}(c, j)$ below denotes the position of candidate c in preference order o_j .) So, $\forall i, 1 \leq i \leq m,$

$$\sum_{j=1}^{m!} \alpha_{\text{pos}(p, j)}(k_j + x_j) \geq \sum_{j=1}^{m!} \alpha_{\text{pos}(c_i, j)}(k_j + x_j).$$

It is easy to see that the above constraints are satisfied if and only if it is possible to add at most k unregistered voters from U such that p wins.

The above integer linear program can be easily modified for the case of deleting voters. □

5.3 Kemeny Elections

Kemeny elections were introduced in [Kem59]. It is the unique election system that is neutral, consistent, and Condorcet [YL78]. A candidate p is a Kemeny winner if p is ranked first in a Kemeny consensus. A Kemeny consensus is a linear order $>$ over the candidates that minimizes the Kendall's tau distance to V , i.e., that minimizes

$$\sum_{a, b \in C, a > b} \|\{v \in V \mid b >_v a\}\|,$$

where $>_v$ is the preference order of voter v .

Observation 5.3.1 *Kemeny-Weighted-Winner is equivalent to Kemeny-Succinct-Winner.*

Notice that the above observation holds since we do not modify the votes to score the election. This is not the case for Dodgson and Young elections, which we will discuss later in the section.

The Kemeny score of a candidate c is the minimum value of the Kendall's tau distance of all linear orders over the candidates that rank c first. And so p is a Kemeny winner if and only if p has a lowest Kemeny score.

Name: Kemeny-Score

Given: An election (C, V) , a candidate $c \in C$, and an integer k .

Question: Is the Kemeny score of c at most k ?

Kemeny-Score was shown to be NP-complete and Kemeny-Winner was shown NP-hard by a reduction from feedback arc set in [BTT89a]. Kemeny-Winner was shown Θ_2^p -complete in [HSV05]. Membership in Θ_2^p is easy: Simply compute the Kemeny scores of all candidates with one round of parallel queries to Kemeny-Score and check that p 's score is the lowest. Kemeny-Winner was shown to be Θ_2^p -complete by a chain of the following three reductions:

1. Min-Card-Vertex-Cover-Compare \leq_m^p Vertex-Cover-Member,
2. Vertex-Cover-Member \leq_m^p Feedback-Arc-Set-Member, and
3. Feedback-Arc-Set-Member \leq_m^p Kemeny-Winner.

These reductions ultimately show that the Θ_2^p -complete problem Min-Card-Vertex-Cover-Compare reduces to Kemeny-Winner.

Theorem 5.3.2 [HSV05] *Kemeny-Winner is Θ_2^p -complete.*

Footnote 3 in [HSV05] points out that Kemeny-Succinct-Winner is in Δ_2^p and explicitly leaves the exact complexity of this problem as an open question. It is easy to see that Kemeny-Succinct-Winner is in Δ_2^p : Simply compute the Kemeny scores of all candidates using binary search on queries to Kemeny-Weighted-Score (which is in NP) and check that p 's score is the lowest). We will now show that Kemeny-Succinct-Winner (and Kemeny-Weighted-Winner) are in fact Δ_2^p -complete.

We first define the following weighted version of Min-Card-Vertex-Cover-Compare.

Name: Min-Weight-Vertex-Cover-Compare

Given: Vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, where $\omega(G)$ denotes the weight of the graph.

Question: Is the weight of G 's minimum-weight vertex cover less than or equal to the weight of H 's minimum-weight vertex cover?

Lemma 5.3.3 *Min-Weight-Vertex-Cover-Compare* \leq_m^p *Kemeny-Succinct-Winner*.

Proof. This follows by careful inspection and modification of the proof from [HSV05]. In particular, since that proof consists of a chain of three reductions between Θ_2^p -complete problems, we need to define suitable Δ_2^p -complete weighted versions of the two intermediate problems and show that the weighted versions of the reductions still hold. In essence, we “lift” the constructions and proofs from Θ_2^p to Δ_2^p . This works surprisingly well.

So, we define a weighted version of Vertex-Cover-Member.

Name: Weighted-Vertex-Cover-Member

Given: Vertex-weighted graph G and vertex v in G .

Question: Is v a member of a minimum-weight vertex cover of G ?

And a weighted version of Feedback-Arc-Set-Member.

Name: Weighted-Feedback-Arc-Set-Member

Given: Irreflexive and antisymmetric edge-weighted digraph G and vertex v in G .

Question: Is there a minimum-weight feedback arc set of G that contains all arcs entering v ?

And we show that the following three reductions hold.

1. (Lemma 5.3.4) *Min-Weight-Vertex-Cover-Compare* \leq_m^p *Weighted-Vertex-Cover-Member*.
2. (Lemma 5.3.5) *Weighted-Vertex-Cover-Member* \leq_m^p *Weighted-Feedback-Arc-Set-Member*.
3. (Lemma 5.3.6) *Weighted-Feedback-Arc-Set-Member* \leq_m^p *Kemeny-Succinct-Winner*.

□

Lemma 5.3.4 *Min-Weight-Vertex-Cover-Compare* \leq_m^p *Weighted-Vertex-Cover-Member*.

Proof. We extend the construction from Lemma 4.12 from [HSV05] to vertex-weighted graphs. Given two vertex-weighted graphs G and H with $\omega(G) = \omega(H)$, construct graph F as follows. $F = (G + (\{v\}, \emptyset)) \times (H + (\{w\}, \emptyset))$, keeping the weights of the vertices in G and H , and letting $\omega(v) = \omega(w) = 1$. Map (G, H) to (F, w) . Write mvc for the minimum vertex cover weight of a vertex-weighted graph. We extend the argument from Lemma 4.12 from [HSV05] to show that this mapping reduces Min-Weight-Vertex-Cover-Compare to Weighted-Vertex-Cover-Member.

Note that V is a vertex cover of F if and only if

1. V contains all vertices in $G + (\{v\}, \emptyset)$ and a vertex cover of H . The lightest such vertex covers have weight $\omega(G) + 1 + mvc(H)$ and do not contain w .
2. V contains all vertices in $H + (\{w\}, \emptyset)$ and a vertex cover of G . The lightest such vertex covers have weight $mvc(G) + \omega(H) + 1$ and contain w .

Since $\omega(G) = \omega(H)$ it is immediate that $mvc(G) \leq mvc(H)$ if and only if there is a minimum-weight vertex cover of F that contains w . \square

Lemma 5.3.5 *Weighted-Vertex-Cover-Member \leq_m^p Weighted-Feedback-Arc-Set-Member.*

Proof. We extend the construction from the proof of Lemma 4.8 from [HSV05], which itself is an extension of the standard reduction from Vertex-Cover to Feedback-Arc-Set from [Kar72], to graphs with weights. Given a vertex-weighted graph G , define edge-weighted digraph $H = (W, A)$ as follows.

1. $W = \{v, v' \mid v \in V(G)\}$.
2. A consists of the following edges.
 - (a) For $v \in V(G)$, edge (v, v') with weight $\omega(v)$.
 - (b) For $\{v, w\} \in E(G)$, edges (v', w) and (w', v) with weight $\omega(G)$.

The reduction maps (G, \hat{v}) to (H, \hat{v}') .

If V' is a vertex cover of G that contains \hat{v} , then, by the proof of Lemma 4.9 from [HSV05], $\{(v, v') \mid v \in V'\}$ is a feedback arc set for H that contains all arcs entering \hat{v}' . Note that the weight of this feedback arc set is $\omega(V')$.

Now suppose that A' is a feedback arc set of H that contains (\hat{v}, \hat{v}') . Again by the proof of Lemma 4.9 from [HSV05], $V' = \{v \in V \mid \exists w \in W : (v, w) \in A' \text{ or } (v', w) \in A'\}$ forms a vertex cover of G that contains \hat{v} . Note that $\omega(V') \leq \omega(A')$. \square

Lemma 5.3.6 *Weighted-Feedback-Arc-Set-Member \leq_m^p Kemeny-Succinct-Winner.*

Proof. Given an irreflexive and antisymmetric edge-weighted digraph G and vertex v in G , we first multiply all edge weights by 2 (this does not change the membership of (G, v) in Weighted-Feedback-Arc-Set-Member). We then use McGarvey’s construction [McG53] to construct in polynomial time a succinct election that has G as its weighted majority graph. The remainder of the proof is identical to the proof of Lemma 4.2 of [HSV05]. \square

So, Lemmas 5.3.4, 5.3.5, and 5.3.6 show that Min-Weight-Vertex-Cover-Compare reduces to Kemeny-Succinct-Winner. It remains to show the following lemma.

Lemma 5.3.7 *Min-Weight-Vertex-Cover-Compare is Δ_2^p -complete.*

Proof. Membership in Δ_2^p is easy to see: Use binary search to compute the weight of G ’s and H ’s minimum-weight vertex covers. Hardness is much more complicated, since there are not many Δ_2^p -complete problems to choose from (significantly fewer than even Θ_2^p -complete problems). The closest available problem is MAXSATASG₌ [Wag87], which consists of pairs of 3cnf formulas with the same maximal satisfying assignment. Though Wagner does not specify what happens when the formulas are not satisfiable, it is easy to see from the proof of Δ_2^p -hardness of MAXSATASG₌ that we need look only at satisfiable formulas.

By negating the variables in the formulas, we obtain the following Δ_2^p -hard promise problem.³

Name: MINSATASG₌

Given: Two satisfiable 3cnf formulas $\phi(x_1, \dots, x_n)$ and $\psi(x_1, \dots, x_n)$.

Question: Do ϕ and ψ have the same lexicographically minimal satisfying assignment (viewed as an n -bit string)?

It is convenient to introduce the following intermediate Δ_2^p -complete problem.

Name: Min-Weight-Vertex-Cover-Equality

Given: Vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, where $\omega(G)$ denotes the weight of the graph.

Question: Is the weight of G ’s minimum-weight vertex cover equal to the weight of H ’s minimum-weight vertex cover?

³A similar problem was used to show the only previously known Δ_2^p -completeness result in computational social choice, namely of Online-Veto-Weighted-Coalition-Manipulation [HHR14].

We will show that $\text{MINSATASG}_=$ reduces to $\text{Min-Weight-Vertex-Cover-Equality}$ and that $\text{Min-Weight-Vertex-Cover-Equality}$ reduces to $\text{Min-Weight-Vertex-Cover-Compare}$, which proves Lemma 5.3.7.

Let $\phi(x_1, \dots, x_n)$ and $\psi(x_1, \dots, x_n)$ be two satisfiable 3cnf formulas. Without loss of generality, assume that ϕ and ψ have the same number m of clauses (simply pad).

Let $f(\phi)$ be the graph computed by the standard reduction from 3SAT to Vertex-Cover from [Kar72]. Then $f(\phi)$ consists of $2n + 3m$ vertices: $\{x_i, \bar{x}_i \mid 1 \leq i \leq n\} \cup \{a_i, b_i, c_i \mid 1 \leq i \leq m\}$, and the following edges: $\{x_i, \bar{x}_i\}$ for $1 \leq i \leq n$ and if the j th clause in ϕ is $\ell_1 \vee \ell_2 \vee \ell_3$, where $\ell_r \in \{x_i, \bar{x}_i \mid 1 \leq i \leq n\}$, then we have edges $\{a_j, \ell_1\}$, $\{b_j, \ell_2\}$, $\{c_j, \ell_3\}$.

The properties of f that we need here, which follow immediately from the proof of [Kar72], are as follows.

1. $f(\phi)$ does not have a vertex cover of size less than $n + 2m$.
2. If W is a vertex cover of size $n + 2m$, then $W \cap \{x_i, \bar{x}_i \mid 1 \leq i \leq n\}$ corresponds to a satisfying assignment of ϕ , in the sense that $\|W \cap \{x_i, \bar{x}_i\}\| = 1$ for $1 \leq i \leq n$, and $\alpha_1 \cdots \alpha_n$ defined as $\alpha_i = 1$ if and only if $x_i \in W$ is a satisfying assignment for ϕ .
3. If α is a satisfying assignment for ϕ , then there is a vertex cover of size $n + 2m$ such that $W \cap \{x_i, \bar{x}_i \mid 1 \leq i \leq n\}$ corresponds to this assignment, i.e., the set $\{x_i \mid \alpha_i = 1\} \cup \{\bar{x}_i \mid \alpha_i = 0\}$ can be extended to a vertex cover of size $n + 2m$ by adding $2m$ vertices from $\{a_i, b_i, c_i \mid 1 \leq i \leq m\}$.

Now set the weights of the vertices as follows: $\omega(a_i) = \omega(b_i) = \omega(c_i) = 2^n$, $\omega(x_i) = 2^n + 2^{i-1}$, and $\omega(\bar{x}_i) = 2^n$. Note that for W a set of vertices, $\|W\| = \lfloor \omega(W)/2^n \rfloor$. In particular, a minimum-weight vertex cover will also have minimum size.

We will now show that n -bit string α is the smallest satisfying assignment of ϕ if and only if $f(\phi)$ has a minimum-weight vertex cover of weight $(n + 2m)2^n + \alpha$. (We interpret n -bit string α as a binary number between 0 and $2^n - 1$.)

If α is a satisfying assignment for ϕ , then by property 3 above the set $\{x_i \mid \alpha_i = 1\} \cup \{\bar{x}_i \mid \alpha_i = 0\}$ can be extended to a vertex cover of size $n + 2m$ by adding $2m$ vertices from $\{a_i, b_i, c_i \mid 1 \leq i \leq m\}$. The weight of this vertex cover is $(n + 2m)2^n + \alpha$. Since the minimum size of a vertex cover is $n + 2m$, the weight of a minimum-weight vertex cover is $(n + 2m)2^n + \beta$, for some β such that $0 \leq \beta < 2^n$. Such a β corresponds to a satisfying assignment by property 2 above.

Since ψ is also a satisfiable 3cnf formula over x_1, \dots, x_n with m clauses, it also holds that n -bit string α is the smallest satisfying assignment of ψ if and only if $f(\psi)$ has a minimum weight vertex cover of weight $(n + 2m)2^n + \alpha$.

This then implies that ϕ and ψ have the same minimal satisfying assignment if and only if $f(\phi)$ and $f(\psi)$'s minimum-weight vertex covers have the same weight. Note that $\omega(f(\phi)) = \omega(f(\psi))$. This completes the reduction from $\text{MINSATASG}_=$ to $\text{Min-Weight-Vertex-Cover-Equality}$.

It remains to show that $\text{Min-Weight-Vertex-Cover-Equality} \leq_m^p \text{Min-Weight-Vertex-Cover-Compare}$.

This is straightforward. Given two vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, map this to the following two vertex-weighted graphs ($+$ denotes the disjoint union and \times denotes the join of two graphs).

1. $\widehat{G} = X + G + H$, where X is a 2-vertex, 1-edge graph, where both vertices have weight $2\omega(G)$.
2. $\widehat{H} = (G \times H) + (G \times H)$.

Note that $\omega(\widehat{G}) = \omega(\widehat{H})$. Let mvc be the minimum vertex cover weight of a vertex-weighted graph. Then $mvc(\widehat{G}) = 2\omega(G) + mvc(G) + mvc(H)$. $mvc(\widehat{H}) = 2 \min(\omega(G) + mvc(H), mvc(G) + \omega(H)) = 2\omega(G) + 2 \min(mvc(G), mvc(H))$. It follows that $mvc(G) = mvc(H)$ if and only if $mvc(\widehat{G}) \leq mvc(\widehat{H})$. \square

This completes the proof of the main theorem of this section.

Theorem 5.3.8 *Kemeny-Weighted-Winner and Kemeny-Succinct-Winner are Δ_2^p -complete.*

Dwork et al. show that Kemeny-Winner is already NP-hard if we have four voters [DKNS01]. In fact, one can easily combine the techniques from [HSV05] and [DKNS01] to obtain the following theorem.

Theorem 5.3.9 *Kemeny-Winner for four voters is Θ_2^p -complete.*

Proof. We combine the constructions from [HSV05] and [DKNS01], to reduce Feedback-Arc-Set-Member to Kemeny-Winner for four voters.

Given an irreflexive and antisymmetric digraph $G = (V, A)$ and vertex $v \in V$, we first compute graph G' as done in [DKNS01]. $G' = (V', A')$ such that $V' = V \cup A$ and $A' = \{(v, (v, w)), ((v, w), w) \mid (v, w) \in A\}$. Note that G has a feedback arc set of size k if and only if G' has a feedback arc set of size k . In addition, and important for our proof, (G, v) is in Feedback-Arc-Set-Member if and only if (G', v) is in Feedback-Arc-Set-Member. Also note that G' is irreflexive and antisymmetric. Now apply the reduction from Feedback-Arc-Set-Member to Kemeny-Winner from [HSV05] on (G', v) . This reduction outputs an

election $g(G')$ that corresponds to G' and (G', v) is in Feedback-Arc-Set-Member if and only if $(g(G'), v)$ is in Kemeny-Winner. The construction from [DKNS01] shows that we need only four voters in $g(G')$. \square

Do we get the same complexity jump if we look at the succinct and weighted case for four votes? Note that for the succinct case we allow many voters with the same vote, and so the Kemeny scores can be large. However, even though the scores can be large, it is easy to see that for each instance there are only a polynomial number of possible scores for each candidate, namely, if the multiplicities of the four votes are $k_1, k_2, k_3,$ and $k_4,$ the only possible scores are $\ell_1 k_1 + \ell_2 k_2 + \ell_3 k_3 + \ell_4 k_4,$ where $0 \leq \ell_i \leq \|C\|(\|C\| - 1)/2.$ For each candidate, simply query the Kemeny-Weighted-Score oracle for all these possible scores in parallel, and then determine whether p has the highest score.

Theorem 5.3.10 *Kemeny-Weighted-Winner and Kemeny-Succinct-Winner for four votes are Θ_2^p -complete.*

We end this section by looking at two other voting systems whose winner problem is Θ_2^p -complete [HHR97, RSV03] namely, Dodgson voting [Dod76] and Young voting [You77].

The Young score of a candidate c is defined as the minimum number of voters that need to be deleted in order to make c a weak Condorcet winner. The Dodgson score of a candidate c is the minimum number of switches between adjacent candidates in voters such that c becomes a Condorcet winner. The candidates with the lowest scores are the winners. Unlike Kemeny, these winner problems depend on a modification of the votes, and so the succinct case and the weighted case do not coincide. In addition, one has to carefully think about what the weighted winner problem means: Do we charge by the voter (in analogy to the definition of weighted control) or do we charge by voter weight.

Clearly, by brute-force, we have the following.

Theorem 5.3.11 *For a fixed number of voters, Young-Weighted-Winner and Dodgson-Weighted-Winner are in P, both in the model where we charge by the voter and in the model where we charge by the vote-weight.*

The succinct case for Young winner is also in P, but this seems to need much more powerful machinery. So, though this is not a complexity jump from tractable to intractable, it seems to be a jump in the complexity of the algorithm.

Theorem 5.3.12 *For a fixed number of votes, Young-Succinct-Winner is in P.*

Proof. Let v_1, \dots, v_ℓ be the different votes. Then the Young score of c is at most k if and only if there exist k_1, \dots, k_ℓ such that $k_i \leq \kappa(v_i)$ and c is a weak Condorcet winner in the election system consisting of $\kappa(v_i) - k_i$ votes v_i . All this can be formulated as an integer linear program with ℓ variables, which can be solved in polynomial time by [Len83]. \square

For the general weighted cases where we charge by the voter, it is easy to see that the scores are polynomially bounded, and so we get.

Theorem 5.3.13 *Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by the voter are Θ_2^p -complete.*

It is easy to see that Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by vote-weight as well as Young-Succinct-Winner and Dodgson-Succinct-Winner are in Δ_2^p , by using binary search to find the scores of all candidates.

We conjecture that, as in the case for Kemeny, we can “lift” the Θ_2^p -hardness proof for Dodgson and Kemeny winner to a Δ_2^p -hardness proof for the weighted case.

Conjecture 5.3.14 *Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by vote-weight are Δ_2^p -complete.*

We do not think that this “lifting” will work for the succinct case and leave the exact complexity of Young-Succinct-Winner and Dodgson-Succinct-Winner, which we know to be Θ_2^p -hard and in Δ_2^p , as an open question.

5.4 Conclusions

Overall, we found that when the voters are represented succinctly the complexity of election problems rarely increases. This is quite surprising, since the succinct representation can be exponentially smaller and so we would expect to see an increase in complexity (with respect to the length of the input). We explained this behavior by showing that common techniques for showing that election problems are in P can be adapted for the case of succinct voters.

There are several interesting directions for future work. In the cases we found where the complexity increases when moving to a succinct representation, the weighted and succinct problems were equivalent. Will this always be the case? Also, what is the exact complexity of Young-Succinct-Winner and Dodgson-Succinct-Winner?

Another direction would be to consider elections where only a fixed number of votes are allowed. This is a natural case to consider, since one can imagine scenarios where even though there are many candidates, voters all aligning with the same political party may all

vote the same. This is reminiscent to how in Chapter 4 we looked at elections where the preferences of the voters have a single-peaked structure (and so only a restricted number of the total possible votes in an election could occur), but instead we restrict the number of votes in an election without assuming structure among them.

Chapter 6

Models for the Attack: Control with Manipulation

6.1 Introduction

In this chapter we consider a new model for strategically attacking an election that combines the two well-known models of control and manipulation. Recall that control models the actions of an agent, referred to as the chair, with control over the structure of the election (e.g., the collection of voters), who seeks to ensure their preferred outcome by modifying the election (e.g., by deleting voters) [BTT92]. Manipulation models the actions of a coalition of voters that misrepresent their preferences to seek to ensure their preferred outcome [BTT89a]. These actions have been well-studied individually, and results on control and manipulation form a large part of this thesis.

Here, we make the very natural assumption that these attacks do not happen alone in a given election. There may be an election chair that seeks to ensure a preferred candidate wins by deleting voters from the election, and a coalition of manipulators that want to ensure that the chair's preferred candidate does not win. We refer to this setting of control vs. manipulation as the competitive case, and we also consider the case where the chair and the manipulators share the same goal and refer to this as the cooperative case. For the competitive case we show that the order of who goes first, the chair or the manipulators, and allowing the manipulators to change their vote for partition control cases, can have interesting effects on the worst-case complexity.

Many of our results will be about the classes P and NP, but for the competitive cases we have completeness results at higher complexity classes. Namely, the classes NP^{NP} , coNP^{NP} , $\text{coNP}^{\text{NP}^{\text{NP}}}$ in the polynomial hierarchy [MS72, Sto76], and the class DP, i.e, the class of

languages that are the difference of two languages in NP [PY84].

This chapter includes the following contributions for this new model.

We state inheritance results in Section 6.3 for our new setting, which explain which complexity results from manipulation and control cases carry over to our new setting. In Section 6.4 we state general upper bounds for the complexity of control with manipulation that hold for all election systems with p -time winner problems, and even though many of these bounds are quite high (at the second and third levels of the polynomial hierarchy), we show most of them to be tight by constructing (admittedly artificial) election systems for which the corresponding case of control with manipulation is complete for that class. And in line with our general theme of what happens in natural settings, we show in Sections 6.5 and 6.6 what happens for real-world election systems. We find that for the important election systems of approval, Condorcet, and plurality that our with-manipulation setting does not raise the complexity from P with respect to the complexity of the corresponding without-manipulation control case. And for the case of weighted voters we found a natural case where the complexity differs depending on who goes first (as long as $NP \neq coNP$). We discuss some related work in Section 6.7 and in Section 6.8 we state some general conclusions and directions for future work.

6.2 Specification of the Model

Our model of allowing control in the presence of manipulators varies the standard control definitions to allow some of the voters to be manipulators, and so their preferences are initially unspecified. We mention that for adding voters (AV), it is legal to have manipulators among the registered and/or the unregistered votes. For the cooperative cases, the question is whether the chair can choose preferences for the manipulators such that, along with using his or her legal control-decision ability for that control type, p can be made (precluded from being) a winner. We denote these types by adding in an “M+,” e.g., plurality-M+CCAV for the collaborative case of constructive control by adding voters. For the competitive cases, we can look at the case where the manipulative coalition sets its votes and then the chair chooses a control action, and we call that MF for “manipulators first.” Or we can have the chair control first and then the manipulators set their votes, which we call CF for “chair first.” Since the manipulators seek to thwart the chair, the case Borda-CCAV-MF, for example, asks whether under Borda, no matter how the manipulative voters, moving first, set their votes, there will exist some choice of at most k unregistered voters that the chair can add so that p is a winner. For partition cases, we add the string “-revoting” to indicate that after the first-round elections occur, the manipulators can change their votes in the runoff.

Notice that for a given control action the CF case is a subset of the MF case, since if there exists a control action such that for all manipulations the chair is successful, then the chair is successful with the same control action when the manipulators go first.

Recall the definitions of manipulation and control from Sections 2.3.1 and 2.3.2. Below we formally state the control plus manipulation action of constructive control by deleting voters (CCDV) for the collaborative (M+), chair-first (CF), and manipulator-first (MF) cases.

Name: \mathcal{E} -M+CCDV/ \mathcal{E} -CCDV-CF/ \mathcal{E} -CCDV-MF

Given: An election $(C, V \cup W)$ (where V and W denote the nonmanipulative and manipulative voters respectively), a preferred candidate $p \in C$, and a delete limit $k \in \mathbb{N}$.

Question (M+): Does there exist a subcollection $V' \subseteq (V \cup W)$ such that $\|V'\| \leq k$, and a way to set the votes of the manipulators, such that p is a winner of $(C, (V \cup W) - V')$ under election system \mathcal{E} ?

Question (CF): Does there exist a subcollection $V' \subseteq (V \cup W)$ such that $\|V'\| \leq k$, so that regardless of how the manipulators set their votes, p is a winner of $(C, (V \cup W) - V')$ under election system \mathcal{E} ?

Question (MF): Regardless of how the manipulators set their votes, does there exist a subcollection $V' \subseteq (V \cup W)$ such that $\|V'\| \leq k$, and p is a winner of $(C, (V \cup W) - V')$ under election system \mathcal{E} ?

To allow many things to be spoken of compactly, we use “stacked” notation to indicate every possible string one gets by reading across and taking one choice from each bracket one encounters on one’s path across the expression. So, for example, $CC \begin{bmatrix} A \\ D \end{bmatrix} V - \begin{bmatrix} CF \\ MF \end{bmatrix}$ refers to four control types, not just two, and $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix} \\ \begin{bmatrix} PC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix} \end{bmatrix}$ refers to $2 \times (2 \times 2 + 3 \times 2) = 20$ control types.

Notice that for our competitive setting, we seem to be asymmetrically focusing on things from the perspective of the chair. That is, regardless of whether the chair moves first or whether the manipulators move first, our problems are always posed in terms of the chair’s constructive or destructive goal regarding the candidate p . It would be natural to ask—and indeed, a conference referee asked us to address the issue of—whether one can interestingly study the competitive problem from the perspective of the manipulators rather than that of the chair. That is, in the MF case for example, one would ask whether the manipulators can act so as to achieve or block victory for p , regardless of the actions of the chair that follow. And one could similarly look at the CF case from the manipulators’ perspective.

After all, in many real-world settings, what one cares about may well be the perspective of the manipulators. Thus being able to address this issue would itself be an additional motivation. Fortunately, in the competitive case—and this holds in both the nonunique-winner model and the unique-winner model, and holds for all types of constructive and destructive attacks discussed here—the chair achieving his or her goal in the model where we view things from the perspective of the chair is precisely the same as the manipulators failing to meet their goal in the model where we view things from the perspective of the manipulators. This follows from the definitions. Thus we are implicitly handling the case of the manipulators’ perspective: For all our competitive cases, studying a constructive (respectively, destructive) attack problem from the perspective of the manipulators is exactly the same as studying the complement of the *destructive* (respectively, *constructive*) version of the same problem in the model we use here, that is, from the perspective of the chair. For example, the sets $\mathcal{E}\text{-DCAV-CF-ManipulatorFocus}$ and $\overline{\mathcal{E}\text{-CCAV-CF-ChairFocus}}$ are the same on all syntactically legal inputs (and they will of course differ on all syntactically illegal inputs). (We will not use “focus” suffixes in except in the previous sentence, since all our problems will implicitly be “-ChairFocus.”) We caution that the above discussion should not be interpreted as saying that the constructive and destructive problems are each other’s opposites. That is not true, although there is a partial connection between these cases, see the discussion in footnote 5 of [HHR07].

6.3 Inheritance Results

Each control type many-one reduces to each of its cooperative and to each of its competitive control-plus-manipulation variants, because for those variants the zero-manipulator cases degenerate to the pure control case. For example, $\mathcal{E}\text{-CCDV} \leq_m^p \mathcal{E}\text{-M+CCDV}$ and $\mathcal{E}\text{-CCDV} \leq_m^p \mathcal{E}\text{-CCDV-MF}$. In particular, NP-hardness results for control inherit upward to each related cooperative and competitive case.

For manipulation, the inheritance behavior is not as broad, since partition control cannot necessarily be “canceled out” by setting a parameter to zero, as partition doesn’t even have a numerical parameter. Nonpartition control types do display inheritance, but for the competitive cases there is some “flipping” of the type of control and the set involved. For each constructive (respectively, destructive) control type regarding adding or deleting candidates or voters, destructive (respectively, constructive) manipulation many-one reduces to the complement of the set capturing the competitive case of the constructive (respectively, destructive) control type combined with manipulation. For example, $\mathcal{E}\text{-CUCM} \leq_m^p \overline{\mathcal{E}\text{-DCAC-CF}}$ and $\mathcal{E}\text{-DUCM} \leq_m^p \overline{\mathcal{E}\text{-CCDV-MF}}$. For the cooperative cases there is no “flipping.” For each

Problem	CF	CF-revoting	MF	MF-revoting
$\mathcal{E} - \left[\begin{array}{c} C \\ D \end{array} \right] C \left[\begin{array}{c} A \\ D \end{array} \right] \left[\begin{array}{c} C \\ V \end{array} \right]$	NP^{NP} (coDP for DV)	N/A	coNP^{NP}	N/A
$\mathcal{E} - \left[\begin{array}{c} C \\ D \end{array} \right] C \left[\begin{array}{c} PC \\ RP \\ PV \end{array} \right] - \left[\begin{array}{c} TE \\ TP \end{array} \right]$	NP^{NP}	NP^{NP}	coNP^{NP}	coNP^{NP} (TE) $\text{coNP}^{\text{NP}^{\text{NP}}}$ (TP)

Table 6.1: Upper Bounds. (N/A means not applicable.)

constructive or destructive control type regarding adding or deleting candidates or voters, manipulation many-one reduces to the cooperative case of that control type combined with manipulation. For example, $\mathcal{E}\text{-CUCM} \leq_m^p \mathcal{E}\text{-M+CCAC}$ and $\mathcal{E}\text{-DUCM} \leq_m^p \mathcal{E}\text{-M+DCAC}$.

6.4 General Upper Bounds and Matching Lower Bounds

For election systems with p -time winner problems, all the cooperative cases clearly have NP upper bounds. But the upper bounds for the competitive cases are far higher, falling in the second and third levels of the polynomial hierarchy, as described by the following theorem.

Theorem 6.4.1 *For each election system \mathcal{E} having a p -time winner problem, the bounds of Table 6.1 hold.¹*

Although the table’s upper bounds clearly follow from the structure of the problems (only for the coDP cases is this nontrivial, see Theorem 6.4.11), the bounds are very high. Can they be improved by some cleverer approach? Or are there systems with p -time winner problems that show the bounds to be tight? The following result establishes that the latter holds; each of the cells in the table is tight for at least some cases.

Theorem 6.4.2 *1. For each of the eight problems on the top line of Table 6.1, and each of the columns on that line, there exists an election system \mathcal{E} , which has a p -time winner problem, for which the named problem is complete for the named complexity class.²*

2. For each of CCPV-TP and CCPV-TE, and each of the CF, CF-revoting, and MF columns of Table 6.1, and each of the columns on that line, there exists an election

¹Where the table says N/A—not applicable—the nonrevoting bounds just to the left of the box technically still hold; we say N/A simply to be clear that revoting cannot even take place in nonpartition cases, since there is no second round.

²The CCDV-CF and DCDV-CF cases were incorrectly classified as NP^{NP} in early versions [FHH13a, FHH14b].

system \mathcal{E} , which has a p -time winner problem, for which the named problem is complete for the named complexity class.

3. There exists an election system \mathcal{E} , which has a p -time winner problem, for which CCPV-TP-MF-revoting is $\text{coNP}^{\text{NP}^{\text{NP}}}$ -complete, and there exists an election system \mathcal{E} , which has a p -time winner problem, for which CCPV-TE-MF-revoting is coNP^{NP} -complete.

The above result says that the upper bounds are not needlessly high. They are truly needed, at least for some systems. However, the constructions proving the lower bounds are artificial and the construction involving the third level of the polynomial hierarchy is lengthy and difficult. In particular, this leaves completely open the possibility that for particular, important real-world systems, even the competitive cases may be far simpler than those bounds suggest. In Section 6.5, we will see that indeed for some of the most important real-world systems, even in the presence of manipulators, the control problem is just as computationally easy as when there are no manipulators.

We now present the proof of the CCAC-CF case of Theorem 6.4.2, which illustrates the general arguments used in the proof of this theorem. For proofs of the cases of this theorem we generally reduce from Quantified Boolean Formulas (QBF) where formulas are restricted to k alternating quantifiers where each quantifier quantifies over a list of boolean variables. The problem QBF_k is the case of k alternating quantifiers beginning with \exists and similarly $\widetilde{\text{QBF}}_k$ is the case of k alternating quantifiers beginning with \forall ; QBF_2 is NP^{NP} -hard, $\widetilde{\text{QBF}}_2$ is coNP^{NP} -hard, and $\widetilde{\text{QBF}}_3$ is $\text{coNP}^{\text{NP}^{\text{NP}}}$ -hard [SM73, Wra76]. In all our proofs using QBF_k or $\widetilde{\text{QBF}}_k$ we assume without loss of generality that the same number of variables are bound to each quantifier.

Theorem 6.4.3 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCAC-CF is NP^{NP} -complete.*

Proof. Let \mathcal{E} be defined in the following way. Given an election (C, V) , if $\|V\| = 1$, $\|C\| \geq 1$ and the candidates in C listed in increasing lexicographic order are c_0, c_1, \dots, c_ℓ , and c_0 encodes a boolean formula $\psi(x_1, \dots, x_{2\ell})$, then do the following. For each i , $1 \leq i \leq \ell$, set x_i to true if the lowest-order bit of c_i is 1 and otherwise set x_i to false. For each i , $1 \leq i \leq \ell$, set $x_{\ell+i}$ to true if the voter states $c_i > c_0$ and otherwise set $x_{\ell+i}$ to false. If this is a satisfying assignment for ψ then everyone wins. In all other cases everyone loses. That completes the specification of \mathcal{E} .

Clearly \mathcal{E} has a p -time winner problem, and by Theorem 6.4.1 we know that \mathcal{E} -CCAC-CF is in NP^{NP} . So what is left is to show that \mathcal{E} -CCAC-CF is NP^{NP} -hard.

Let $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ be an instance of QBF_2 . We construct an instance of \mathcal{E} -CCAC-CF in the following way. Let the candidate set C consist of p encoding the boolean formula ψ , and let there be zero nonmanipulators and one manipulator. Let the set of unregistered candidates contain ℓ pairs where for each i , $1 \leq i \leq \ell$, there is a candidate $p \cdot i_{\text{binary}} \cdot 0$ and a candidate $p \cdot i_{\text{binary}} \cdot 1$. (where \cdot denotes concatenation and i_{binary} denotes i encoded in binary). Let the add limit $k = 2\ell$.³

If $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})] \in \text{QBF}_2$, fix an assignment to x_1, \dots, x_ℓ such that $(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ is true. For each i , $1 \leq i \leq \ell$, the chair adds the candidate, call it c_i , from the i th pair whose last bit corresponds to the value of x_i in this assignment. Note that p, c_1, \dots, c_ℓ are in increasing lexicographic order. Then no matter what assignment to $x_{\ell+1}, \dots, x_{2\ell}$ is induced by the manipulator's vote, formula ψ is satisfied and so p will win.

Conversely, if the chair makes p a winner, then the chair adds exactly ℓ candidates whose lowest-order bits give an assignment to x_1, \dots, x_ℓ such that $(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ is true. \square

The rest of this section presents the remaining proofs of cases of Theorem 6.4.2. Notable proofs include the coDP upper bound for \mathcal{E} -CCDV-CF and \mathcal{E} -DCDV-CF shown in Theorem 6.4.11 (this bound is shown to be tight in the proofs of Theorems 6.4.12 and 6.4.13), and the proof of Theorem 6.4.18, which shows completeness at the third level of the polynomial hierarchy.

Theorem 6.4.4 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCAC-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E} be as defined in the proof of Theorem 6.4.3. Then \mathcal{E} has a p -time winner problem and by Theorem 6.4.1 we know that \mathcal{E} -CCAC-MF is in coNP^{NP} . So what is left is to show that \mathcal{E} -CCAC-MF is coNP^{NP} -hard.

Let $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})]$ be an instance of $\widetilde{\text{QBF}}_2$. Our instance of \mathcal{E} -CCAC-MF is exactly the instance of \mathcal{E} -CCAC-CF from the proof of Theorem 6.4.3. The same argument as in that proof shows that $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})] \in \widetilde{\text{QBF}}_2$ if and only if the chair can always ensure that p becomes a winner. \square

Theorem 6.4.5 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCDC-CF is NP^{NP} -complete.*

³We set $k = 2\ell$ instead of the obvious choice of ℓ since then the same proof can be used for the similar cases that follow, and this also nicely handles the case of control by unlimited of adding candidates.

Proof. Let \mathcal{E} be defined as in the proof of Theorem 6.4.3. Then \mathcal{E} has a p -time winner problem and by Theorem 6.4.1 we know that \mathcal{E} -CCDC-CF is in NP^{NP} . So what is left is to show that \mathcal{E} -CCDC-CF is NP^{NP} -hard.

Let $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ be an instance of QBF_2 . Our instance of \mathcal{E} -CCDC-MF is the instance of \mathcal{E} -CCAC-CF from the proof of Theorem 6.4.3, except that we let the candidate set C consist of all $2\ell + 1$ candidates. The same argument as in the proof of Theorem 6.4.3 shows that $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})] \in \text{QBF}_2$ if and only if the chair can ensure that p always becomes a winner by deleting candidates. \square

Theorem 6.4.6 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCDC-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E} be defined as in the proof of Theorem 6.4.3. Then \mathcal{E} has a p -time winner problem and by Theorem 6.4.1 we know that \mathcal{E} -CCDC-MF is in coNP^{NP} . So what is left is to show that \mathcal{E} -CCDC-MF is coNP^{NP} -hard.

Let $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})]$ be an instance of $\widetilde{\text{QBF}}_2$. Our instance of \mathcal{E} -CCDC-MF is exactly the instance of \mathcal{E} -CCDC-CF from the proof of Theorem 6.4.5. The same argument as in that proof of Theorem 6.4.3 shows that $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})] \in \widetilde{\text{QBF}}_2$ if and only if the chair can ensure that p always becomes a winner. \square

Theorem 6.4.7 *There exists an election system, \mathcal{E}' , with a p -time winner problem, such that \mathcal{E}' -DCAC-CF is NP^{NP} -complete, \mathcal{E}' -DCAC-MF is coNP^{NP} -complete, \mathcal{E}' -DCDC-CF is NP^{NP} -complete, and \mathcal{E}' -DCDC-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E}' be defined as \mathcal{E} in Theorem 6.4.3 except replace “everyone loses” with “everyone wins” and “everyone wins” with “everyone loses.”

Note that for every election (C, V) and every candidate $p \in C$, p is an \mathcal{E} winner of (C, V) if and only if p is not an \mathcal{E}' winner of (C, V) . This immediately implies that \mathcal{E}' -DCAC-CF = \mathcal{E} -CCAC-CF, \mathcal{E}' -DCAC-MF = \mathcal{E} -CCAC-MF, \mathcal{E}' -CCDC-CF = \mathcal{E} -CCDC-CF, and \mathcal{E}' -CCDC-MF = \mathcal{E} -CCDC-MF. The result follows from Theorems 6.4.3, 6.4.4, 6.4.5, and 6.4.6. \square

Theorem 6.4.8 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCAV-CF is NP^{NP} -complete.*

Proof. Let \mathcal{E} be defined in the following way. Given an election (C, V) , if $\|C\| \geq 3$ and the candidates in C listed in increasing lexicographic order are $c_0, c_1, \dots, c_{\ell+1}$, candidate c_0 encodes a boolean formula $\psi(x_1, \dots, x_{2\ell})$, $\|V\| = 2\ell + 1$, and for each $i, 1 \leq i \leq \ell$ there are at least two voters with the same vote who rank c_i first, then do the following. For each $i, 1 \leq i \leq \ell$, set x_i to true if two voters with c_i first both state $c_{\ell+1} > c_0$ and otherwise set x_i to false. Let \hat{v} be the unique vote that occurs three times or only once in V . For each $i, 1 \leq i \leq \ell$, set $x_{\ell+i}$ to true if \hat{v} states $c_i > c_0$, else set $x_{\ell+i}$ to false. If this is a satisfying assignment for ψ then everyone wins. In all other cases everyone loses. That completes the specification of \mathcal{E} .

Clearly \mathcal{E} has a p-time winner problem, and by Theorem 6.4.1 we know that \mathcal{E} -CCAV-CF is in NP^{NP} . So what is left is to show that \mathcal{E} -CCAV-CF is NP^{NP} -hard.

Let $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ be an instance of QBF_2 . We construct an instance of \mathcal{E} -CCAV-CF in the following way. Let the candidate set C consist of p encoding ψ and $\ell + 1$ candidates all lexicographically larger than p . So, the candidates in C can be listed in increasing lexicographic order as $p, c_1, \dots, c_{\ell+1}$. Let the collection of registered voters V consist of zero nonmanipulators and one manipulator. Let the collection of unregistered voters, all nonmanipulators, consist of 2ℓ pairs where for each $i, 1 \leq i \leq \ell$, there are two voters v_i and v'_i with the same vote $(c_i > c_{\ell+1} > p > \dots)$ and two voters u_i and u'_i with the same vote $(c_i > p > c_{\ell+1} > \dots)$. Let the add limit $k = 4\ell$ and let the preferred candidate of the chair be $p \in C$.

If $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})] \in \text{QBF}_2$, fix an assignment to x_1, \dots, x_ℓ such that $(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ is true. For each $i, 1 \leq i \leq \ell$, if x_i is true in the assignment the chair adds v_i and v'_i and if x_i is false the chair adds u_i and u'_i . Note that the vote of the manipulator will be the unique vote \hat{v} that occurs three times (if the manipulator votes the same as one of the paired voters) or only once. And no matter what assignment to $x_{\ell+1}, \dots, x_{2\ell}$ is induced by \hat{v} , formula ψ is satisfied and so p will win.

Conversely, if the chair makes p a winner then the chair adds exactly ℓ voter pairs whose ℓ different votes give an assignment to x_1, \dots, x_ℓ such that $(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ is true. \square

Theorem 6.4.9 *There exists an election system, \mathcal{E} , with a p-time winner problem, such that \mathcal{E} -CCAV-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E} be as defined in the proof of Theorem 6.4.8. Then \mathcal{E} has a p-time winner problem and by Theorem 6.4.1 we know that \mathcal{E} -CCAV-MF is in coNP^{NP} . So what is left is to show that \mathcal{E} -CCAV-MF is coNP^{NP} -hard.

Let $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})]$ be an instance of $\widetilde{\text{QBF}}_2$. Our instance of \mathcal{E} -CCAV-MF is exactly the instance of \mathcal{E} -CCAV-CF from the proof of Theorem 6.4.8. The same argument as in that proof shows that $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})] \in \widetilde{\text{QBF}}_2$ if and only if the chair can always ensure that p becomes a winner. \square

Theorem 6.4.10 *There exists an election system, \mathcal{E}' , with a p -time winner problem, such that \mathcal{E}' -DCAV-CF is NP^{NP} -complete and \mathcal{E}' -DCAV-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E}' be defined as \mathcal{E} in Theorem 6.4.8 except replace “everyone loses” with “everyone wins” and “everyone wins” with “everyone loses.”

Note that for every election (C, V) and every candidate $p \in C$, p is an \mathcal{E} winner of (C, V) if and only if p is not an \mathcal{E}' winner of (C, V) . This immediately implies that \mathcal{E}' -DCAV-CF = \mathcal{E} -CCAV-CF and \mathcal{E}' -DCAV-MF = \mathcal{E} -CCAV-MF. The result follows from Theorems 6.4.8 and 6.4.9. \square

Unlike in the candidate cases, we can not use the same construction to show that the deleting voter cases are also hard, because the chair can delete the manipulator. In fact, we will show that for every election system \mathcal{E} with a p -time winner problem, \mathcal{E} -CCDV-CF and \mathcal{E} -DCDV-CF are in coDP (and so are not NP^{NP} -complete unless the polynomial hierarchy collapses). DP is the class of languages that are the difference of two NP languages [PY84].

Theorem 6.4.11 *For every election system \mathcal{E} with a p -time winner problem, \mathcal{E} -CCDV-CF and \mathcal{E} -DCDV-CF are in coDP .*

Proof. It is easy to see that it is always at least as good for the chair to delete a manipulator as it is to delete a nonmanipulator (though note that because the election system can be anything, deleting as many manipulators as possible may not be best; for example, if we want to make p a winner and our election system has all candidates as winners if there are four voters and no winners if there are fewer voters, we do not want to delete manipulators if there are four voters). So we have that p can be made a winner (not a winner) by deleting at most k voters if and only if there exists a $k' \leq k$ such that (letting m be the number of manipulators):

1. $k' \leq m$ and after deleting k' manipulators the remaining $m - k'$ manipulators can not preclude p from winning (not winning), or
2. $k' > m$ and after deleting all manipulators the chair can make p win (not win) by deleting at most $k' - m$ voters.

We can check if there exists a k' such that we are in case 1 in coNP and we can check if there exists a k' such that we are in case 2 in NP, and so we can write our languages as the union of a coNP set and an NP set. \square

We now show that the coDP bounds from Theorem 6.4.11 are tight.

Theorem 6.4.12 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCDV-CF is coDP-complete.*

Proof. We reduce from the coDP-complete problem $\{\langle\phi, \psi\rangle \mid \phi \in SAT \text{ or } \psi \notin SAT\}$, which is the complement of the standard DP-complete problem SAT-UNSAT [PY84]. Without loss of generality, we assume that ϕ and ψ have the same number of variables.

Let \mathcal{E} be defined in the following way. Given an election (C, V) , if $\|C\| \geq 3$ and the candidates in C listed in increasing lexicographic order are $c_0, c_1, \dots, c_{\ell+1}$, and candidate c_0 encodes the pair of boolean formulas $\langle\phi(x_1, \dots, x_\ell), \psi(x_{\ell+1}, \dots, x_{2\ell})\rangle$, then:

1. If $\|V\| = \ell$ and for each $i, 1 \leq i \leq \ell$, there is a voter who ranks c_i first, we do the following. For each $i, 1 \leq i \leq \ell$, set x_i to true if the voter with c_i first states $c_{\ell+1} > c_0$ and otherwise set x_i to false. If this is a satisfying assignment for ϕ , then everyone wins.
2. If $\|V\| = 2\ell + 1$, then if there are no voters that rank $c_{\ell+1}$ first, then everyone wins. Otherwise, if there is exactly one voter \hat{v} that ranks $c_{\ell+1}$ first then for each $i, 1 \leq i \leq \ell$, set $x_{\ell+i}$ to true if \hat{v} states $c_i > c_0$, else set $x_{\ell+i}$ to false. If this is not a satisfying assignment for ψ , then everyone wins.

In all other cases everyone loses. That completes the specification of \mathcal{E} .

Clearly \mathcal{E} has a p -time winner problem, and by Theorem 6.4.11 we know that \mathcal{E} -CCDV-CF is in coDP. So what is left is to show that \mathcal{E} -CCDV-MF is coDP-hard.

Let $\langle\phi(x_1, \dots, x_\ell), \psi(x_{\ell+1}, \dots, x_{2\ell})\rangle$ be a pair of boolean formulas. We construct an instance of \mathcal{E} -CCDV-CF in the following way. Let the candidate set C consist of p encoding $\langle\phi, \psi\rangle$ and $\ell + 1$ candidates all lexicographically larger than p . So, the candidates in C can be listed in increasing lexicographic order as $p, c_1, \dots, c_{\ell+1}$. Let the collection of voters V consist of one manipulator and 2ℓ nonmanipulators where for each $i, 1 \leq i \leq \ell$, there is a voter v_i who votes $(c_i > c_{\ell+1} > p > \dots)$ and a voter u_i who votes $(c_i > p > c_{\ell+1} > \dots)$. Let the delete limit $k = 2\ell + 1$ (any limit $\geq \ell + 1$ will do) and let the preferred candidate of the chair be $p \in C$. We need to show that $(\phi \in SAT \text{ or } \psi \notin SAT)$ if and only if control can be asserted.

Suppose $\phi \in SAT$. Fix an assignment to x_1, \dots, x_ℓ that satisfies ϕ . The chair deletes $\ell + 1$ voters. The only voters that are not deleted are for each $i, 1 \leq i \leq \ell$, v_i if x_i is true in the assignment and u_i if x_i is false in the assignment. This leaves ℓ voters that encode a satisfying assignment for ϕ and so everyone wins. Next suppose that $\psi \notin SAT$. Then we keep all voters. Since there does not exist a satisfying assignment for ψ , everyone wins.

For the converse, to have p win, we either have that $\|V\| = \ell$, in which case ϕ is satisfiable, or $\|V\| = 2\ell + 1$. In the latter case, if $\psi \in SAT$ the manipulator could induce a satisfying assignment for ψ , but then p is not a winner. It follows that $\psi \notin SAT$. \square

Theorem 6.4.13 *There exists an election system, \mathcal{E}' , with a p -time winner problem, such that \mathcal{E} -DCDV-CF is coDP-complete.*

Proof. Let \mathcal{E}' be defined as \mathcal{E} in Theorem 6.4.12 except replace “everyone loses” with “everyone wins” and “everyone wins” with “everyone loses.”

Note that for every election (C, V) and every candidate $p \in C$, p is an \mathcal{E} winner of (C, V) if and only if p is not an \mathcal{E}' winner of (C, V) . This immediately implies that \mathcal{E}' -DCDV-CF = \mathcal{E} -CCDV-CF. The result follows from Theorem 6.4.12. \square

For the CCDV-MF case, we modify the construction from Theorem 6.4.9 to basically ensure that the manipulator will not be deleted, while still making sure that p can always be made a winner for positive instances of $\widetilde{\text{QBF}}_2$.

Theorem 6.4.14 *There exists an election system, \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCDV-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E} be defined in the following way. Given an election (C, V) , if $\|C\| \geq 3$ and the candidates in C listed in increasing lexicographic order are $c_0, c_1, \dots, c_{\ell+1}$, candidate c_0 encodes a boolean formula $\psi(x_1, \dots, x_{2\ell})$, $\|V\| = \ell + 1$, and for each $i, 1 \leq i \leq \ell$, there is a voter who ranks c_i first, then do the following. For each $i, 1 \leq i \leq \ell$, set x_i to true if some voter with c_i first states $c_{\ell+1} > c_0$ and otherwise set x_i to false. If there is a voter \hat{v} that ranks $c_{\ell+1}$ first (note that there exists at most one such voter) then for each $i, 1 \leq i \leq \ell$, set $x_{\ell+i}$ to true if \hat{v} states $c_i > c_0$, else set $x_{\ell+i}$ to false. If this is a satisfying assignment for ψ then everyone wins. If there is a voter that ranks c_0 first then everyone wins. If there are two voters that rank c_i first for some $i, 1 \leq i \leq \ell$, and these voters agree on whether or not $c_{\ell+1} > c_0$ then everyone wins. In all other cases everyone loses. That completes the specification of \mathcal{E} .

Clearly \mathcal{E} has a p -time winner problem, and by Theorem 6.4.1 we know that \mathcal{E} -CCDV-MF is in coNP^{NP} . So what is left is to show that \mathcal{E} -CCDV-MF is coNP^{NP} -hard.

Let $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})]$ be an instance of $\widetilde{\text{QBF}}_2$. We construct an instance of \mathcal{E} -CCDV-MF in the following way. Let the candidate set C consist of p encoding ψ and $\ell + 1$ candidates all lexicographically larger than p . So, the candidates in C can be listed in increasing lexicographic order as $p, c_1, \dots, c_{\ell+1}$. Let the collection of voters V consist of one manipulator and 2ℓ nonmanipulators where for each $i, 1 \leq i \leq \ell$, there is a voter v_i who votes $(c_i > c_{\ell+1} > p > \dots)$ and a voter u_i who votes $(c_i > p > c_{\ell+1} > \dots)$. Let the delete limit $k = 2\ell + 1$ (any limit $\geq \ell$ will do) and let the preferred candidate of the chair be $p \in C$.

Suppose $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})] \in \widetilde{\text{QBF}}_2$. Consider a vote \hat{v} for the manipulator. If \hat{v} ranks c_0 first then the chair deletes v_i for all $i, 1 \leq i \leq \ell$ to make p a winner. If \hat{v} ranks c_i first, for some $i, 1 \leq i \leq \ell$, and states $c_{\ell+1} > c_0$, then the chair deletes $\{u_1, \dots, u_\ell\}$ to make p a winner. If \hat{v} ranks c_i first, for some $i, 1 \leq i \leq \ell$, and states $c_0 > c_{\ell+1}$, then the chair deletes $\{v_1, \dots, v_\ell\}$ to make p a winner. If \hat{v} ranks $c_{\ell+1}$ first, then consider the assignment to $x_{\ell+1}, \dots, x_{2\ell}$ induced by \hat{v} and fix an assignment to x_1, \dots, x_ℓ such that $\psi(x_1, \dots, x_{2\ell})$ is true. For each $i, 1 \leq i \leq \ell$, if x_i is true in the assignment the chair deletes u_i and if x_i is false the chair deletes v_i . This will make p a winner.

Conversely, fix an assignment to $x_{\ell+1}, \dots, x_{2\ell}$. Set the manipulator vote \hat{v} so that it induces this assignment and so that $c_{\ell+1}$ is ranked first. Consider the set of voters left after the chair has deleted voters to make p a winner. Note that this set must include \hat{v} and a set of voters that induces an assignment to x_1, \dots, x_ℓ that makes ψ true. \square

Theorem 6.4.15 *There exists an election system, \mathcal{E}' , with a p -time winner problem, such that \mathcal{E}' -DCDV-MF is coNP^{NP} -complete.*

Proof. Let \mathcal{E}' be defined as \mathcal{E} in Theorem 6.4.12 except replace “everyone loses” with “everyone wins” and “everyone wins” with “everyone loses.”

Note that for every election (C, V) and every candidate $p \in C$, p is an \mathcal{E} winner of (C, V) if and only if p is not an \mathcal{E}' winner of (C, V) . This immediately implies that \mathcal{E}' -DCDV-MF = \mathcal{E} -CCDV-MF. The result follows from Theorem 6.4.14. \square

Theorem 6.4.16 *There exists an election system \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCPV- $\begin{bmatrix} \text{TE} \\ \text{TP} \end{bmatrix}$ - $\begin{bmatrix} \emptyset \\ \text{revoting} \end{bmatrix}$ -CF are each NP^{NP} -complete.*

Proof. Let \mathcal{E} be defined in the following way. Given an election (C, V) , do the following.

If $\|C\| = 1$ then the sole candidate wins.

If $\|C\| = 2$ then the lexicographically larger candidate wins.

If $\|C\| \geq 3$, $\|V\| = 2\ell$, the candidates in C listed in increasing lexicographic order are $c_0, c_1, \dots, c_{\ell+1}$, and candidate c_0 encodes a boolean formula $\psi(x_1, \dots, x_{2\ell})$, then if for each $i, 1 \leq i \leq \ell$, there are exactly two voters with the same vote who rank c_i first no one wins, else $c_{\ell+1}$ wins.

If $\|C\| \geq 3$, $\|V\| = 2\ell + 1$, the candidates in C listed in increasing lexicographic order are $c_0, c_1, \dots, c_{\ell+1}$, candidate c_0 encodes a boolean formula $\psi(x_1, \dots, x_{2\ell})$, and for each $i, 1 \leq i \leq \ell$, there are at least two voters with the same vote who rank c_i first, then do the following. For each $i, 1 \leq i \leq \ell$, set x_i to true if two voters with c_i first both state $c_{\ell+1} > c_0$ and otherwise set x_i to false. Let \hat{v} be the unique vote that occurs three times or only once in V . For each $i, 1 \leq i \leq \ell$, set $x_{\ell+i}$ to true if \hat{v} states $c_i > c_0$, else set $x_{\ell+i}$ to false. If this is a satisfying assignment for ψ then c_0 wins.

In all other cases, everyone loses. That completes the specification of \mathcal{E} .

Clearly \mathcal{E} has a p-time winner problem, and \mathcal{E} -CCPV- $\left[\begin{smallmatrix} \text{TE} \\ \text{TP} \end{smallmatrix}\right]$ - $\left[\begin{smallmatrix} \emptyset \\ \text{revoting} \end{smallmatrix}\right]$ -CF are each in NP^{NP} by Theorem 6.4.1. So, what is left is to show that \mathcal{E} -CCPV- $\left[\begin{smallmatrix} \text{TE} \\ \text{TP} \end{smallmatrix}\right]$ - $\left[\begin{smallmatrix} \emptyset \\ \text{revoting} \end{smallmatrix}\right]$ -CF are each NP^{NP} -hard.

Let $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ be an instance of QBF_2 . We construct an instance of \mathcal{E} -CCPV-TE-CF in the following way. Let the candidate set C consist of p encoding ψ and $\ell + 1$ candidates lexicographically larger than p . So, the candidates in C can be listed in increasing lexicographic order as $p, c_1, \dots, c_{\ell+1}$. Let there be one manipulative voter, and let the nonmanipulators consist of 2ℓ pairs where for each $i, 1 \leq i \leq \ell$, there are two voters v_i and v'_i with the same vote ($c_i > c_{\ell+1} > p > \dots$) and two voters u_i and u'_i with the same vote ($c_i > p > c_{\ell+1} > \dots$). Let the preferred candidate of the chair be $p \in C$.

If $(\exists x_1, \dots, x_\ell)(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})] \in \text{QBF}_2$, fix an assignment to x_1, \dots, x_ℓ such that $(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ is true. The chair sets V_1 to consist of the manipulator and the subcollection of the voters whose votes encode the assignment, i.e., for each $i, 1 \leq i \leq \ell$, if x_i is true in the assignment the chair adds v_i and v'_i to V_1 and if x_i is false the chair adds u_i and u'_i to V_1 . The chair puts the remaining voters from V into V_2 . Note that the vote of the manipulator will be the unique vote \hat{v} that occurs three times (if the manipulator votes for one of the paired voters) or only once in V_1 . And no matter what assignment to $x_{\ell+1}, \dots, x_{2\ell}$ is induced by \hat{v} , formula ψ is satisfied and so p is the unique winner of (C, V_1) . Since V_2 consists 2ℓ voters of the correct form, no one wins (C, V_2) . Only candidate p participates in the runoff and so p wins the runoff. Note that this argument works for the “TE” and the “TP” models with or without revoting.

Conversely, if the chair can ensure that p wins then there exists a partition such that for all manipulations p wins. It is clear that the chair must partition the voters into (V_1, V_2) such that $\|V_1\| = 2\ell + 1$ and $\|V_2\| = 2\ell$, since otherwise there are no winners. Also, for each

$i, 1 \leq i \leq \ell$, V_2 contains exactly two voters with the same vote who rank c_i first. It follows that V_1 contains the manipulator vote \hat{v} and that for each $i, 1 \leq i \leq \ell$, V_1 contains exactly two nonmanipulators with the same vote who rank c_i first. These 2ℓ nonmanipulators induce an assignment to x_1, \dots, x_ℓ . Fix this assignment. Now fix an assignment to $x_{\ell+1}, \dots, x_{2\ell}$. Set the manipulator vote \hat{v} so that it induces this assignment. Since p wins the runoff, this is a satisfying assignment for ψ . It follows that for the assignment to x_1, \dots, x_ℓ that is induced by V_1 , it holds that $(\forall x_{\ell+1}, \dots, x_{2\ell})[\psi(x_1, \dots, x_{2\ell})]$ is true \square

Theorem 6.4.17 *There exists an election system \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCPV- $\begin{bmatrix} \text{TE} \\ \text{TP} \end{bmatrix}$ -MF and \mathcal{E} -CCPV-TE-MF-revoting are each coNP^{NP} -complete.*

Proof. Let \mathcal{E} be as defined in the proof of Theorem 6.4.16. Then \mathcal{E} has a p -time winner problem and by Theorem 6.4.1 we know that \mathcal{E} -CCPV- $\begin{bmatrix} \text{TE} \\ \text{TP} \end{bmatrix}$ -MF and \mathcal{E} -CCPV-TE-MF-revoting are each in coNP^{NP} . So what is left is to show that \mathcal{E} -CCPV- $\begin{bmatrix} \text{TE} \\ \text{TP} \end{bmatrix}$ -MF and \mathcal{E} -CCPV-TE-MF-revoting are each coNP^{NP} -hard. Below we describe the reduction for the “TE” case. It is easy to see that the same reduction holds for the “TP” case. For the “TE” case with revoting observe that the same reduction also holds since in the runoff there will be at most two candidates and in election system \mathcal{E} the votes do not affect who wins in that case.

Let $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})]$ be an instance of $\widetilde{\text{QBF}}_2$. Our instance of \mathcal{E} -CCPV-TE-MF is exactly the instance of \mathcal{E} -CCPV-TE-CF from the proof of Theorem 6.4.16. Note that the vote of the manipulator will always be the unique vote \hat{v} that occurs three times or only once in V . The same argument as in the proof of Theorem 6.4.16 shows that $(\forall x_{\ell+1}, \dots, x_{2\ell})(\exists x_1, \dots, x_\ell)[\psi(x_1, \dots, x_{2\ell})] \in \widetilde{\text{QBF}}_2$ if and only if the chair can ensure that p always becomes a winner by partitioning voters. \square

When revoting is allowed after the first round in the TP case, and the manipulators go first, we find an interesting rise in complexity.

Theorem 6.4.18 *There exists an election system \mathcal{E} , with a p -time winner problem, such that \mathcal{E} -CCPV-TP-MF-revoting is $\text{coNP}^{\text{NP}^{\text{NP}}}$ -complete.*

Proof. The election system, \mathcal{E} , defined below will utilize the following special candidates.

$\langle 1, \psi \rangle$: where ψ is a boolean formula, which we refer to as a type-1 candidate.

$\langle 2, i, j \rangle$: where $i \in \mathbb{N}$ and $j \in \{0, 1\}$, which we refer to as a type-2 candidate.

$\langle 3, i, j \rangle$: where $i \in \mathbb{N}$ and $j \in \{0, 1\}$, which we refer to as a type-3 candidate.

$\langle 4, i \rangle$: where $i \in \mathbb{N}$, which we refer to as a type-4 candidate.

Let \mathcal{E} be defined in the following way.

Given an election (C, V) :

If C consists of one type-1 candidate encoding $\psi(x_1, \dots, x_{3\ell})$, 2ℓ type-2 candidates $\langle 2, 1, 0 \rangle, \langle 2, 1, 1 \rangle, \dots, \langle 2, \ell, 0 \rangle, \langle 2, \ell, 1 \rangle$, 2ℓ type-3 candidates $\langle 3, 1, 0 \rangle, \langle 3, 1, 1 \rangle, \dots, \langle 3, \ell, 0 \rangle, \langle 3, \ell, 1 \rangle$, and $\ell + 2$ type-4 candidates $\langle 4, 1 \rangle, \dots, \langle 4, \ell + 2 \rangle$, then do the following.

- If $\|V\| = 2\ell + 1$ and for each $i, 1 \leq i \leq \ell$, there are at least two voters with the same vote who rank $\langle 4, i \rangle$ first, then we have $3\ell + 2$ winners consisting of $\langle 1, \psi \rangle, \langle 4, 1 \rangle, \dots, \langle 4, \ell + 1 \rangle$, and 2ℓ candidates determined in the following way. Let \hat{v} be the unique vote that occurs three times or only once in V . For each $i, 1 \leq i \leq \ell$, $\langle 2, i, 1 \rangle$ is a winner if \hat{v} states $\langle 4, i \rangle > \langle 1, \psi \rangle$ and otherwise $\langle 2, i, 0 \rangle$ is a winner. For each $i, 1 \leq i \leq \ell$, $\langle 3, i, 1 \rangle$ is a winner if two voters who rank $\langle 4, i \rangle$ first both state $\langle 4, \ell + 1 \rangle > \langle 1, \psi \rangle$ and otherwise $\langle 3, i, 0 \rangle$ is a winner.
- If $\|V\| = 2\ell$ then if for each $i, 1 \leq i \leq \ell$, there are at least two voters with the same vote who rank $\langle 4, i \rangle$ first, no one wins, else $\langle 4, \ell + 2 \rangle$ wins.

If C consists of one type-1 candidate encoding $\psi(x_1, \dots, x_{3\ell})$, ℓ type-2 candidates of the form $\langle 2, 1, \star \rangle, \dots, \langle 2, \ell, \star \rangle$ (where $\star \in \{0, 1\}$), ℓ type-3 candidates of the form $\langle 3, 1, \star \rangle, \dots, \langle 3, \ell, \star \rangle$ (where $\star \in \{0, 1\}$), and $\ell + 1$ type-4 candidates $\langle 4, 1 \rangle, \dots, \langle 4, \ell + 1 \rangle$, $\|V\| = 4\ell + 1$, and there is a unique vote \hat{v}' that occurs three times or only once in V , then do the following. For each $i, 1 \leq i \leq \ell$, set x_i to true if $\langle 2, i, 1 \rangle$ is in C and to false if $\langle 2, i, 0 \rangle$ is in C , set $x_{\ell+i}$ to true if $\langle 3, i, 1 \rangle$ is in C and to false if $\langle 3, i, 0 \rangle$ is in C , and set $x_{2\ell+i}$ to true if \hat{v}' states $\langle 4, i \rangle > \langle 1, \psi \rangle$ and else set $x_{2\ell+i}$ to false. If this is a satisfying assignment for formula ψ , then $\langle 1, \psi \rangle$ wins. Otherwise, everyone loses.

Else, everyone loses.

That completes the specification of \mathcal{E} .

Clearly \mathcal{E} has a p-time winner problem, and \mathcal{E} -CCPV-TP-MF-revoting is in $\text{coNP}^{\text{NP}^{\text{NP}}}$ by Theorem 6.4.1. So, what is left to show is that \mathcal{E} -CCPV-TP-MF-revoting is $\text{coNP}^{\text{NP}^{\text{NP}}}$ -hard.

Let $(\forall x_1, \dots, x_\ell)(\exists x_{\ell+1}, \dots, x_{2\ell})(\forall x_{2\ell+1}, \dots, x_{3\ell})[\psi(x_1, \dots, x_{3\ell})]$ be an instance of $\widetilde{\text{QBF}}_3$. We construct an instance of \mathcal{E} -CCPV-TP-MF-revoting in the following way. Let the candidate set C consist of one type-1 candidate encoding ψ , 2ℓ type-2 candidates $\langle 2, 1, 0 \rangle, \langle 2, 1, 1 \rangle, \dots, \langle 2, \ell, 0 \rangle, \langle 2, \ell, 1 \rangle$, 2ℓ type-3 candidates $\langle 3, 1, 0 \rangle, \langle 3, 1, 1 \rangle, \dots, \langle 3, \ell, 0 \rangle, \langle 3, \ell, 1 \rangle$, and $\ell + 2$ type-4 candidates $\langle 4, 1 \rangle, \dots, \langle 4, \ell + 2 \rangle$. Let there be one manipulator and 4ℓ nonmanipulators where for each $i, 1 \leq i \leq \ell$, there are two

voters v_i and v'_i with the same vote ($\langle 4, i \rangle > \langle 4, \ell + 1 \rangle > \langle 1, \psi \rangle > \dots$) and two voters u_i and u'_i with the same vote ($\langle 4, i \rangle > \langle 1, \psi \rangle > \langle 4, \ell + 1 \rangle > \dots$). Let the preferred candidate of the chair be $\langle 1, \psi \rangle \in C$.

Suppose $(\forall x_1, \dots, x_\ell)(\exists x_{\ell+1}, \dots, x_{2\ell})(\forall x_{2\ell+1}, \dots, x_{3\ell})[\psi(x_1, \dots, x_{3\ell})] \in \widetilde{\text{QBF}}_3$. Consider a first-round vote \widehat{v} for the manipulator, and view it as an assignment to x_1, \dots, x_ℓ where for each $i, 1 \leq i \leq \ell$, if \widehat{v} states $\langle 4, i \rangle > \langle 1, \psi \rangle$ then x_i is true and otherwise x_i is false. Using this assignment, set an assignment to $x_{\ell+1}, \dots, x_{2\ell}$ such that $(\forall x_{2\ell+1}, \dots, x_{3\ell})\psi(x_1, \dots, x_{3\ell})$ is true. The chair sets V_1 to consist of the manipulator and for each $i, 1 \leq i \leq \ell$, if $x_{\ell+i}$ is true in the assignment the chair adds v_i and v'_i to V_1 and if $x_{\ell+i}$ is false the chair adds u_i and u'_i to V_1 . The chair puts the remaining voters from V into V_2 . Note that \widehat{v} will be the unique vote that occurs three times or only once in V_1 . Notice that the type-2 and type-3 candidates that proceed to the runoff “hold” the abovementioned assignments to x_1, \dots, x_ℓ and $x_{\ell+1}, \dots, x_{2\ell}$ respectively (since $\langle 2, i, 1 \rangle$ proceeds to the runoff if and only if x_i is true, $\langle 2, i, 0 \rangle$ proceeds to the runoff if and only if x_i is false, $\langle 3, i, 1 \rangle$ proceeds to the runoff if and only if $x_{\ell+i}$ is true, and $\langle 3, i, 0 \rangle$ proceeds to the runoff if and only if $x_{\ell+i}$ is false). And that no matter what assignment to $x_{2\ell+1}, \dots, x_{3\ell}$ is induced by the second-round vote \widehat{v}' of the manipulator, formula ψ is true and so $\langle 1, \psi \rangle$ wins.

Conversely, suppose that for all first-round manipulator votes there exists a partition such that for all second-round manipulator votes $\langle 1, \psi \rangle$ wins. Fix a first-round manipulator vote \widehat{v} , and let (V_1, V_2) be a partition such that $\langle 1, \psi \rangle$ wins regardless of the second-round vote of the manipulator. It is clear that $\|V_1\| = 2\ell + 1$ and $\|V_2\| = 2\ell$ (without loss of generality), and that for each $i, 1 \leq i \leq \ell$, V_2 contains exactly two voters with the same vote who rank $\langle 4, i \rangle$ first. It follows that the first-round manipulator vote \widehat{v} is the unique vote that occurs three times or only once in V_1 and that for each $i, 1 \leq i \leq \ell$, V_1 contains two voters with the same vote who rank $\langle 4, i \rangle$ first.

Fix an assignment to x_1, \dots, x_ℓ and consider the first-round manipulator vote \widehat{v} where for each $i, 1 \leq i \leq \ell$, if x_i is true then \widehat{v} states $\langle 4, i \rangle > \langle 1, \psi \rangle$ and so $\langle 2, i, 1 \rangle$ proceeds to the runoff, and if x_i is false then \widehat{v} states $\langle 1, \psi \rangle > \langle 4, i \rangle$ and so $\langle 2, i, 0 \rangle$ proceeds to the runoff. Since we know that there exists a partition (V_1, V_2) where $\langle 1, \psi \rangle$ wins the runoff, we know that for each $i, 1 \leq i \leq \ell$, $\langle 3, i, 1 \rangle$ proceeds to the runoff if v_i and v'_i are in V_1 and otherwise $\langle 3, i, 0 \rangle$ does. We can view this as an assignment to $x_{\ell+1}, \dots, x_{2\ell}$ where for each $i, 1 \leq i \leq \ell$, if $\langle 3, i, 1 \rangle$ proceeds to the runoff then x_i is true and if $\langle 3, i, 0 \rangle$ proceeds to the runoff then x_i is false. Now fix an assignment to $x_{2\ell+1}, \dots, x_{3\ell}$ and set the second-round manipulator vote \widehat{v}' so that it induces this assignment. Since $\langle 1, \psi \rangle$ wins the runoff, ψ is true for the assignment to $x_1, \dots, x_\ell, x_{\ell+1}, \dots, x_{2\ell}, x_{2\ell+1}, \dots, x_{3\ell}$. It follows that $(\forall x_1, \dots, x_\ell)(\exists x_{\ell+1}, \dots, x_{2\ell})(\forall x_{2\ell+1}, \dots, x_{3\ell})[\psi(x_1, \dots, x_{3\ell})] \in \widetilde{\text{QBF}}_3$. \square

6.5 Specific Systems

Plurality is certainly the most important of election systems, and approval is also an important system. Plurality, approval, and Condorcet elections each have easy manipulation problems, and their complexity for every standard control type is known [BTT92, HHR07]. We display these known results in Table 6.2.⁴ In this section we will show that the “M+,” “CF,” and “MF” cases whose control type is classified as P in Table 6.2, are in the with-no-manipulators case in P for each of our cooperative and competitive cases.⁵

Control by	Plurality		Condorcet		Approval	
	Constr.	Destr.	Constr.	Destr.	Constr.	Destr.
Adding Candidates	NPC	NPC	P	P	P	P
Deleting Candidates	NPC	NPC	P	P	P	P
Adding Voters	P	P	NPC	P	NPC	P
Deleting Voters	P	P	NPC	P	NPC	P
Partitioning Candidates	TE: NPC	TE: NPC	P	P	TE: P	TE: P
	TP: NPC	TP: NPC			TP: P	TP: P
Runoff Partitioning Candidates	TE: NPC	TE: NPC	P	P	TE: P	TE: P
	TP: NPC	TP: NPC			TP: P	TP: P
Partitioning Voters	TE: P	TE: P	NPC	P	TE: NPC	TE: P
	TP: NPC	TP: NPC			TP: NPC	TP: P

Table 6.2: Summary of complexity of control for plurality, Condorcet, and approval [HHR07].

Theorem 6.5.2 *Each problem contained in*

$$\bullet \left[\begin{array}{c} \text{approval} \\ \text{Condorcet} \\ \text{plurality} \end{array} \right] \text{-M+} \left[\begin{array}{c} \text{C} \\ \text{D} \end{array} \right] \text{C} \left[\begin{array}{c} \left[\begin{array}{c} \text{A} \\ \text{D} \end{array} \right] \left[\begin{array}{c} \text{C} \\ \text{V} \end{array} \right] \\ \left[\begin{array}{c} \text{PC} \\ \text{RPC} \end{array} \right] \text{-} \left[\begin{array}{c} \text{TE} \\ \text{TP} \end{array} \right] \end{array} \right],$$

⁴It should be noted that the referenced table in [HHR07] is focused on the unique-winner case, but by Observation 6.5.1 below these results carry over to the nonunique-winner model (some of the cases were previously noted in Faliszewski, Hemaspaandra, and Hemaspaandra [FHH14a] and Hemaspaandra, Hemaspaandra, and Rothe [HHR12b]). Also, note that the “AC” line of the referenced table refers to so-called unlimited adding and (as is now standard) we use “AC” to refer to (limited) adding. Additionally, in our table we use NPC instead of “R” (resistant) and P instead of “V” (vulnerable) or “I” (immune).

Observation 6.5.1 *The complexities of each of the standard control problems shown in Bartholdi, Tovey, and Trick [BTT92] and Hemaspaandra, Hemaspaandra, and Rothe [HHR07] for the unique-winner model hold also for the nonunique-winner model.*

⁵The reason we have looked at only the P cases of control for these system is that due to our inheritance results, for the NP cases, getting a P result will be impossible.

- $\left[\begin{array}{c} \text{approval} \\ \text{Condorcet} \\ \text{plurality} \end{array} \right] - \left[\begin{array}{c} \text{C} \\ \text{D} \end{array} \right] \text{C} \left[\begin{array}{c} \left[\begin{array}{c} \text{A} \\ \text{D} \end{array} \right] \left[\begin{array}{c} \text{C} \\ \text{V} \end{array} \right] \\ \left[\begin{array}{c} \text{PC} \\ \text{RPC} \\ \text{PV} \end{array} \right] - \left[\begin{array}{c} \text{TE} \\ \text{TP} \end{array} \right] \end{array} \right] - \text{CF}, \text{ or}$
- $\left[\begin{array}{c} \text{approval} \\ \text{Condorcet} \\ \text{plurality} \end{array} \right] - \left[\begin{array}{c} \text{C} \\ \text{D} \end{array} \right] \text{C} \left[\begin{array}{c} \left[\begin{array}{c} \text{A} \\ \text{D} \end{array} \right] \left[\begin{array}{c} \text{C} \\ \text{V} \end{array} \right] \\ \left[\begin{array}{c} \text{PC} \\ \text{RPC} \\ \text{PV} \end{array} \right] - \left[\begin{array}{c} \text{TE} \\ \text{TP} \end{array} \right] \end{array} \right] - \text{MF},$

whose corresponding control type is in P in Table 6.2 is in P.

The proofs of many of these cases will utilize the polynomial-time algorithms for the without-manipulators versions of the control cases. The well-known polynomial-time results from Bartholdi, Tovey, and Trick [BTT92] and Hemaspaandra, Hemaspaandra, and Rothe [HHR07] are both for the unique-winner model. Observation 6.5.1 states that each of these control cases holds for the nonunique-winner model, and we will reference this observation when referring to the polynomial-time algorithm for a given nonmanipulator control case.

As an illustration, we present the proof of plurality-M+CCPV-TE \in P here.

Proof. Note that it is not the case that the manipulators can always simply vote for p , no matter what the chair does. For example, if the chair partitions the voters such that one of the subelections contains a voter voting $p > a > b$, and the other subelection contains 100 voters voting $a > b > p$, 101 voters voting $b > a > p$, and one manipulator, the manipulator should vote for a , so that a and b are tied in the second subelection and neither goes through to the second round. Still, we will show that if a partition of the voters and a manipulation of the manipulators exist such that p wins the election, then there exists a way for p to win when all manipulators vote for p . It follows that we can check if p can be made a winner by first having all manipulators vote for p and then running the polynomial-time algorithm for plurality-CCPV-TE from [HHR07] (modified in the obvious way for the nonunique-winner case).

So, suppose that a manipulation and a partition (V_1, V_2) exist such that p is a winner of the election. Without loss of generality, suppose p is the unique winner of (C, V_1) . Then p is also the unique winner of (C, V_1) if all manipulators in V_1 vote for p , so have all manipulators in V_1 vote for p . Now consider (C, V_2) . As explained in the previous paragraph, simply changing the manipulators' votes to p could have bad effects. Instead, we do the following. While manipulators remain in V_2 whose first-choice candidate is not p , choose one of them, v , let a be v 's first-choice candidate, and do the following.

1. Change v 's vote from a to p and move v to V_1 .
2. For each candidate $b \neq a$, move a current V_2 voter for b (if any exists) from V_2 to V_1 and if it is a manipulator, change its vote to p .

Since in each iteration of the above loop we add at least one vote for p to V_1 , p will remain the unique winner of (C, V_1) . If after the loop (C, V_2) does not have a unique winner or has p as the unique winner it is immediate that p wins the runoff. The only remaining case is that after the loop (C, V_2) has a unique winner $c \neq p$. Note that in each iteration we keep the same set of winners in (C, V_2) unless V_2 becomes empty in which case all candidates become winners in (C, V_2) . This implies that c is the unique winner of (C, V_2) before the loop and thus c does not beat p in the runoff before the loop. Since the only votes that are changed in the loop are manipulator votes changed to p , after the loop p clearly is a winner of the runoff.⁶ \square

Below we state general results on election systems satisfying the Weak Axiom of Revealed Preferences (WARP) and its corresponding unique version (unique-WARP). An election system satisfies WARP if whenever a candidate is a winner among a set of candidates (under a collection of votes V ; as always, we assume that V is masked down to the candidates at hand in the given election) then that candidate is also a winner among every subset of those candidates that includes him or her (under that same collection of votes V ; as always, we assume that V is masked down to the candidates at hand in the given election). Similarly, an election system satisfies unique-WARP if whenever a candidate is a unique winner among a set of candidates then that candidate is also a unique winner among every subset of those candidates that includes him or her.⁷ It is easy to see that approval and Condorcet elections satisfy both WARP and unique-WARP [HHR07].

Theorem 6.5.3 *For every election system \mathcal{E} satisfying unique-WARP, and for each instance of the CCRPC-TE problem, it holds that control is possible if and only if the preferred candidate p is an overall winner using the partition $(C - \{p\}, \{p\})$.*

Proof. Given an election system satisfying unique-WARP, an election (C, V) , and a candidate $p \in C$, we do the following.

If p is an overall winner using partition $(C - \{p\}, \{p\})$ then clearly control is possible.

Conversely, if p is *not* an overall winner using partition $(C - \{p\}, \{p\})$ then we will show that control is not possible. There are two cases.

1. If under our collection of votes (masked down to the candidates in the election at hand in each case, of course) p does not win in the election where p is the sole candidate,

⁶There was a slight problem in the argument used in this paragraph in a previous version [FHH13a], which was fixed in a later version [FHH14b].

⁷We here and in many other places write the somewhat strange, awkward phrase “a unique winner” rather than the seemingly more natural phrase “the unique winner.” We do so to avoid giving the impression that there necessarily *is* a unique winner—as opposed for example to perhaps having no winners or perhaps having multiple winners.

then by unique-WARP p will not be a unique winner in any subelection it is part of, and so can never survive the first round, and so can never become an overall winner.

2. On the other hand, if under our collection of votes (masked down to the candidates in the election at hand in each case, of course) p wins in the election where p is the sole candidate, then p in the partition $(C - \{p\}, \{p\})$ clearly will survive the first round.

Since we are in the TE model, either zero or one candidates will survive the $C - \{p\}$ first-round subelection.

But if zero survive, then the second-round election involves just p , who we already, in our current case, have assumed wins under the votes masked down to it, so it will in fact be an overall winner (in fact, it will be the only overall winner).

On the other hand, if one candidate, call it r , survives the $C - \{p\}$ first-round subelection, note that since we assumed that p is not an overall winner, it must be the case that in the election between r and p (with the votes as always masked down to the candidates in the election), p is not a winner. So, can there be any partition, $(C - A, A)$, under the given votes, that will ensure that p is an overall winner? W.l.o.g., assume $p \in A$. If $r \in A$, then p cannot move forward, since to do that (as we are in the TE model) p would have to be a unique winner within A , and since $\{p, r\} \subseteq A$, by unique-WARP it would have been impossible for p to fail to beat r in the second-round election under partition $(C - \{p\}, \{p\})$ in our original setting, yet that is precisely what happened in our current case's assumptions. On the other hand, if $r \notin A$, then given that $C - A \subseteq C - \{p\}$, by unique-WARP we have that r wins the subelection $(C - A, V)$, and so faces p in the runoff, and we already know that in that case p will not be a winner of that contest.

By the above case analysis, we have shown that control is not possible, thus completing this second direction of the proof. □

Corollary 6.5.4 *For every election system \mathcal{E} that satisfies unique-WARP and has a p -time winner problem, \mathcal{E} -CCRPC-TE is in P.*

Theorem 6.5.3 does not hold for CCPC-TE. For example, in the election system where all candidates are winners if there are at least two candidates, and no candidates win if there is at most one candidate (note that this system vacuously satisfies unique-WARP), an election with candidates $\{a, b\}$ has no winners using partition $(\{a\}, \{p\})$, but all candidates win using partition $(\emptyset, \{a, p\})$.

Nonetheless, we have proven an analogue of Theorem 6.5.3 for the CCPC-TE case. Our analogue, however, applies to election systems that satisfy both WARP and unique-WARP.⁸

Theorem 6.5.5 *For every election system satisfying both WARP and unique-WARP, and for each instance of the CCPC-TE problem, it holds that control is possible if and only if the preferred candidate p is an overall winner using the partition $(C - \{p\}, \{p\})$.*

Proof. Given an election system satisfying both WARP and unique-WARP, an election (C, V) , and a candidate $p \in C$, we do the following.

If p is an overall winner using partition $(C - \{p\}, \{p\})$ then clearly control is possible.

Conversely, if p is *not* an overall winner using partition $(C - \{p\}, \{p\})$ then we will show that control is not possible. There are two cases.

1. If under our collection of votes (masked down to the candidates in the election at hand in each case, of course) p does not win in the election where p is the sole candidate, then by WARP p will not be a winner in any larger subelection that contains him or her, and so can never be an overall winner.
2. If under our collection of votes (masked down to the candidates in the election at hand in each case, of course) p wins the election where p is the sole candidate then since p is not the overall winner using partition $(C - \{p\}, \{p\})$ and we are in the TE model, there exists a candidate $r \in C - \{p\}$ such that r is the unique winner of the subelection $(C - \{p\}, V)$ and p does not win the runoff election $(\{p, r\}, V)$. Since the given election system satisfies unique-WARP and r is the unique winner of $(C - \{p\}, V)$, r will be the unique winner of every subelection that does not involve p . And since the given election satisfies WARP and p does not win $(\{p, r\}, V)$, p is not a winner in any subelection involving r . Notice that p participates in the runoff only if r also participates in the runoff. So it is clear to see that control is not possible.

□

⁸Is it going unnaturally far to study systems that satisfy both WARP and unique-WARP? We do not think so. Indeed, to put our use of two properties in context, we mention that even combined they are a weaker assumption about the election system than is even a certain different version of WARP that is sometimes used. The version of WARP that we are using here is precisely that found for example in Baumeister and Rothe’s survey of preference aggregation [BR16]. This version focuses on the individual candidate and what happens when other candidates are removed, namely, that winning does not turn into not winning for any unremoved candidate. The other version, and to avoid confusion let us refer to it as WARP’, focuses on whether when one removes candidates the winner set is always *exactly* the previous winner set intersected with the remaining set of candidates. WARP’ clearly implies both WARP and unique-WARP. And so Theorem 6.5.5 would certainly remain true if in it one were to replace the phrase “both WARP and unique-WARP” with simply “WARP’.”

Corollary 6.5.6 *For every election system \mathcal{E} that satisfies both WARP and unique-WARP and has a p -time winner problem, \mathcal{E} -CCPC-TE is in P.*

Corollary 6.5.7 *For every election system \mathcal{E} satisfying both WARP and unique-WARP, \mathcal{E} -CCPC-TE = \mathcal{E} -CCRPC-TE.*

We now present proofs of the remaining cases of Theorem 6.5.2. In some of these proofs, we use the notation $score_{(C,V)}(a)$ to denote the score of candidate a in election (C, V) . When it is clear from context, we may leave out C, V , or both.

6.5.1 Plurality

Theorem 6.5.8 *For plurality elections, the following hold.*

1. $M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} V$ are each in P.
2. $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} V - \begin{bmatrix} CF \\ MF \end{bmatrix}$ are each in P.

Proof. For the constructive cooperative and the destructive competitive cases it is clear that the manipulators should all vote for p .

For the destructive cooperative and the constructive competitive cases the optimal action for the manipulators is to all vote for the same highest-scoring candidate in $C - \{p\}$.

In all cases we can determine if the chair can be successful by assuming the manipulators vote as above and using the corresponding p -time algorithm for control from Bartholdi, Tovey, and Trick [BTT92] (for the constructive cases) or from Hemaspaandra, Hemaspaandra, and Rothe [HHR07] (for the destructive cases), modified in the obvious way for the nonunique-winner case (see Observation 6.5.1). \square

For the remaining proofs in this section, given an election (C, V) containing k manipulators, we say that a candidate r is a *rival* of p if r can beat p pairwise, i.e., if $score_{\{p,r\}}(r) + k > score_{\{p,r\}}(p)$.

Lemma 6.5.9 *If there exists a partition such that p is an overall winner in the “TE” model when all manipulators vote for the same highest-scoring rival r and put p last, then there exists a partition such that p is always an overall winner.*

Proof. Given an election (C, V) where V contains k manipulators, a candidate $p \in C$, and a candidate $r \in C - \{p\}$ such that $score_{\{p,r\}}(r) + k > score_{\{p,r\}}(p)$, we do the following.

Let (V_1, V_2) be a partition such that p is an overall winner when all manipulators vote for r and put p last. Let k_1 be the number of manipulators in V_1 , let k_2 be the number of

manipulators in V_2 , let ℓ_1 be the number of nonmanipulator votes for r in V_1 , and let ℓ_2 be the number of nonmanipulator votes for r in V_2 . Without loss of generality assume that p is the unique winner of (C, V_1) when all manipulators vote for r .

Now we will construct a new partition $(\widehat{V}_1, \widehat{V}_2)$ that will work regardless of how the manipulators vote. Let \widehat{V}_2 consist of ℓ_2 nonmanipulator votes for r , $score_{V_2}(p)$ nonmanipulator votes for p , for every rival $\widehat{r} \neq r$, $\min(\ell_2, score(\widehat{r}))$ votes for \widehat{r} , for every nonrival $c \neq p$ all the nonmanipulator votes for c , and k_2 manipulators. Let $\widehat{V}_1 = V - \widehat{V}_2$.

We first show that p is always the unique winner of (C, \widehat{V}_1) . We know that $score_{\widehat{V}_1}(r) + k_1 = \ell_1 + k_1 < score_{V_1}(p) = score_{\widehat{V}_1}(p)$. For every nonrival $c \neq p$, $score_{\widehat{V}_1}(c) + k_1 = k_1 < score_{\widehat{V}_1}(p)$. Finally, for every rival $\widehat{r} \neq r$, $score(\widehat{r}) \leq score(r) = \ell_1 + \ell_2$, and so $score_{\widehat{V}_1}(\widehat{r}) \leq \ell_1$, which implies that $score_{\widehat{V}_1}(\widehat{r}) + k_1 \leq \ell_1 + k_1 < score_{\widehat{V}_1}(p)$. It follows that p is always the unique winner of (C, \widehat{V}_1) .

So the only way in which p can be precluded from winning the runoff is if there exist a manipulation and a rival \widehat{r} of p such that \widehat{r} is the unique winner of (C, \widehat{V}_2) . Then $\ell_2 + k_2 > score(c)$ for every nonrival $c \neq p$, and $\ell_2 + k_2 > score_{\widehat{V}_2}(p) = score_{V_2}(p)$. Now consider (C, V_2) and let all manipulators vote for r . Then the score of r in (C, V_2) (after the manipulation) is $\ell_2 + k_2$, and r is the unique winner of (C, V_2) . Then p is not an overall winner of (C, V) when all manipulators vote for r , which contradicts our assumption.

It follows that p is always a winner of $(\widehat{V}_1, \widehat{V}_2)$. □

Theorem 6.5.10 plurality-CCPV-TE-CF *is in* P.

Proof. Given an election (C, V) and a preferred candidate of the chair $p \in C$, p can be made a winner if and only if there exists a partition (V_1, V_2) such p is always an overall winner.

If no rivals of p exist, then clearly control is possible if and only if $C = \{p\}$ or there is at least one vote for p (in the latter case, let V_1 consist of one voter for p).

Otherwise, let r be a highest-scoring rival of p . It is immediate from Lemma 6.5.9 that control is possible if and only if there exists a partition such that p wins when all manipulators vote for r and put p last. This can be determined by running the polynomial-time algorithm for plurality-CCPV-TE from [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1). □

Theorem 6.5.11 plurality-CCPV-TE-CF = plurality-CCPV-TE-MF.

Proof. It immediately follows from the definition that plurality-CCPV-TE-CF \subseteq plurality-CCPV-TE-MF.

Now suppose that “MF” control is possible. Then for all manipulations there exists a partition such that the preferred candidate p wins. Then either no rival to p exists, in which case “CF” control is possible since either p is the only candidate or there exists at least one vote for p . When a rival r to p exists, control is certainly possible when all the manipulators vote for r and put p last. By Lemma 6.5.9 we know that then there exists a partition where p is always a winner, so “CF” control is possible. \square

Corollary 6.5.12 *plurality-CCPV-TE-MF is in P.*

Theorem 6.5.13 *plurality-M+DCPV-TE is in P.*

Proof. Given an election (C, V) and a despised candidate of the chair $p \in C$, we can determine in polynomial time if p can be precluded from winning by partitioning voters as follows. If there are no manipulators, run the polynomial-time algorithm for plurality-DCPV-TE from [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1).

So, let $k > 0$ denote the number of manipulators in V . If there exists a rival r to p (i.e., a candidate that can beat p pairwise, i.e., a candidate for which $score_{\{p,r\}}(p) < score_{\{p,r\}}(r) + k$), then control is possible: Let V_2 consist of one manipulator and let all manipulators vote for r .

If there are no rivals, we must ensure that p doesn’t make it to the runoff. It is easy to see that this can be done if and only if we are in one of the following two cases.

1. There are at least two candidates, c is a highest-scoring candidate in $C - \{p\}$, and $score(p) \leq score(c) + k$. (Have all manipulators vote for c and use partition (V, \emptyset) .)
2. There are at least three candidates, c and d are two highest-scoring candidates in $C - \{p\}$, and $score(p) \leq score(c) + score(d) + k$. (Have V_1 consist of $\min(score(p), score(c))$ votes for p and all votes for c . The remaining votes, including all manipulators, who will vote for d , will be in V_2 .)

\square

Lemma 6.5.14 *If there exists a partition of voters such that p is not a plurality winner in the “TE” model when all manipulators vote for p , then there exists a partition such that p can never be made a plurality winner by the manipulators.*

Proof. Given an election (C, V) and a candidate $p \in C$, we do the following.

Let (V_1, V_2) be a partition such that p is not a winner when all manipulators vote for p . If p can never be made a winner by the manipulators in this partition then we are done. So, suppose there exists a manipulation such that p is an overall winner (with the partition (V_1, V_2)). Without loss of generality assume that p is the unique winner of (C, V_1) . Then p is also the unique winner in (C, V_1) if all manipulators vote for p . However, since p is not an overall winner if all manipulators vote for p there is a candidate $c \neq p$ such that if all manipulators vote for p , c is the unique winner of (C, V_2) and c is the unique winner of the runoff $(\{p, c\}, V)$.

Now move all manipulators from V_2 to V_1 . Note that c remains the unique winner of (C, V_2) and that c is always the unique winner of $(\{p, c\}, V)$. It follows that in this new partition, p is never a winner, no matter what the manipulators do. \square

Lemma 6.5.14 implies that plurality-DCPV-TE-CF is in P, since control is possible if and only if control is possible when all manipulators vote for p . This can be checked using the polynomial-time algorithm for plurality-DCPV-TE from Hemaspaandra, Hemaspaandra, and Rothe [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1).

Theorem 6.5.15 *plurality-DCPV-TE-CF is in P.*

We will now show that Lemma 6.5.14 also implies that plurality-DCPV-TE-MF is in P.

Theorem 6.5.16 *plurality-DCPV-TE-MF is in P.*

Proof. Given an election (C, V) and a despised candidate of the chair $p \in C$, we will show that we can determine in polynomial time if p can be precluded from winning by partitioning voters.

As in the “CF” case we will use Lemma 6.5.14 to show that control is possible if and only if there exists a partition such that p is precluded from winning when all manipulators vote for p . This also implies that plurality-DCPV-TE-CF = plurality-DCPV-TE-MF.

It immediately follows from the definition that if the instance of plurality-DCPV-TE-MF is positive, then there exists a partition such that p is not a winner when all manipulators vote for p .

For the other direction, by Lemma 6.5.14 if there exists a partition such that p is not a winner when all the manipulators vote for p , then there exists a partition (V_1, V_2) such that p can never be made a winner by the manipulators. This implies that no matter what the manipulators do, there exists a partition (in fact, always the same partition) such that p is not a winner. This then implies that the instance of plurality-DCPV-TE-MF is positive. \square

6.5.2 Condorcet

Theorem 6.5.17 *For Condorcet elections, the following hold.*

1. $M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} C$ are each in P .
2. $M + DC \begin{bmatrix} A \\ D \end{bmatrix} V$ are both in P .
3. $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} C - \begin{bmatrix} CF \\ MF \end{bmatrix}$ are each in P .
4. $DC \begin{bmatrix} A \\ D \end{bmatrix} V - \begin{bmatrix} CF \\ MF \end{bmatrix}$ are both in P .

Proof. For the constructive cooperative and the destructive competitive cases it is clear that the manipulators should all vote for p .

For the destructive cooperative and the constructive competitive cases the optimal action for the manipulators is to rank p last.

In all cases we can determine if the chair can be successful by assuming the manipulators vote as above and using the corresponding p -time algorithm for control from Bartholdi, Tovey, and Trick [BTT92] (for the constructive cases) or from Hemaspaandra, Hemaspaandra, and Rothe [HHR07] (for the destructive cases), modified in the obvious way for the nonunique-winner case (see Observation 6.5.1). \square

We now prove the Condorcet partition cases. Since Condorcet winners are always unique, the “TE” and “TP” cases coincide and so we will leave out this notation, following [HHR07].

Theorem 6.5.18 *Condorcet- $M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \end{bmatrix}$ are each in P .*

Proof. Given an election (C, V) and a preferred candidate of the chair $p \in C$, we can determine in polynomial time if p can be made a winner by partitioning of candidates and by runoff partitioning of candidates as follows.

For the constructive cases we do the following. Since Condorcet elections satisfy both WARP and unique-WARP, we know from Theorems 6.5.3 and 6.5.5 that control is possible if and only if control is possible using partition $(C - \{p\}, \{p\})$. Set all manipulators to rank p first. Rank the candidates that do not beat p pairwise next in all manipulator votes (in any order). Then, as long as there exists an unranked candidate c that can never be a Condorcet winner in $(C - \{p\}, V)$, rank c next in all manipulator votes.

Let \widehat{C} be the set of candidates not yet ranked by the manipulators. Notice that every $c \in \widehat{C}$ beats p pairwise, and every $c \in \widehat{C}$ can become a Condorcet winner in (\widehat{C}, V) (and thus also in (C, V)).

So, to determine if control is possible, we must determine if the manipulators can vote in such a way that there is no Condorcet winner in (\widehat{C}, V) , i.e., $\forall c \in \widehat{C} \exists c' \in \widehat{C}$ such that c' ties-or-beats c pairwise.

For $\|V\|$ even, assume that there are at least two candidates in \widehat{C} and for $\|V\|$ odd, assume there are at least three candidates in \widehat{C} (otherwise there will always be Condorcet winners). We have the following cases, depending on whether or not there is a Condorcet winner in (\widehat{C}, V) before the manipulators vote and depending on the parity of $\|V\|$. Let $k \geq 1$ denote the number of manipulators in V .

1. If there exists a Condorcet winner and $\|V\|$ is even, then let c be the Condorcet winner, and let $d \in \widehat{C} - \{c\}$. It is easy to see that each of the manipulators can vote $(c > d > \widehat{C} - \{c, d\})$ or $(d > c > \widehat{C} - \{c, d\})$ in such a way that c ties d pairwise. So, c is no longer a Condorcet winner and no other candidate becomes a Condorcet winner, since c ties-or-beats every other candidate pairwise.
2. If there exists a Condorcet winner and $\|V\|$ is odd, then let c be the Condorcet winner, and let $a, b \in \widehat{C} - \{c\}$ be such that a ties-or-beats b pairwise. Have $\lfloor k/2 \rfloor$ manipulators vote $(a > b > c > \widehat{C} - \{a, b, c\})$ and $\lfloor k/2 \rfloor$ manipulators vote $(b > c > a > \widehat{C} - \{a, b, c\})$. After this manipulation, b beats c pairwise, a beats b pairwise, and c beats every candidate in $\widehat{C} - \{b, c\}$ pairwise.
3. If there is no Condorcet winner and $\|V\|$ is even, then have $\lfloor k/2 \rfloor$ manipulators vote \widehat{C} (i.e., the candidates in \widehat{C} in some fixed order) and $\lfloor k/2 \rfloor$ manipulators vote $\overleftarrow{\widehat{C}}$ (i.e., the candidates in \widehat{C} in reverse order). When k is odd, let the remaining manipulator vote arbitrarily. It is clear that no Condorcet winners are created by the manipulators.
4. If there is no Condorcet winner and $\|V\|$ is odd, then we have the following cases.
 - (a) If k is even, then have $k/2$ manipulators vote \widehat{C} and the remaining $k/2$ manipulators vote $\overleftarrow{\widehat{C}}$.
 - (b) If k is odd and there is no weak Condorcet winner (a weak Condorcet winner is a candidate that ties-or-beats every other candidate pairwise), then have $\lfloor k/2 \rfloor$ manipulators vote \widehat{C} and $\lfloor k/2 \rfloor$ manipulators vote $\overleftarrow{\widehat{C}}$. Let the remaining manipulator vote arbitrarily. It is clear that no Condorcet winner is created by the manipulators.
 - (c) If k is odd and there exists a weak Condorcet winner, then let c be a weak Condorcet winner and let a be a candidate such that a ties c pairwise. We have the following two cases.

- i. If for all $b \in \widehat{C} - \{a, c\}$, a beats b pairwise and c beats b pairwise, then have $\lfloor k/2 \rfloor$ manipulators vote $(\widehat{C} - \{a, c\} > a > c)$ and have the remaining $\lfloor k/2 \rfloor$ manipulators vote $(c > \widehat{C} - \{a, c\} > a)$. So, now a beats c pairwise, and for all $b \in \widehat{C} - \{a, c\}$, c beats b pairwise and b beats a pairwise, and thus there is still no Condorcet winner.
- ii. Otherwise, there exists a candidate $b \in C - \{a, c\}$ such that it is not the case that a and c both beat b pairwise. Suppose there are at least three manipulators, and set their votes in the following way. (If there is only one manipulator, then since each candidate in \widehat{C} can become a Condorcet winner, all candidates in \widehat{C} tie pairwise. And so there is always a Condorcet winner after manipulation.)
 - A. If a does not beat b pairwise, then let $\lfloor k/3 \rfloor$ manipulators vote $(c > b > a > \widehat{C} - \{a, b, c\})$, $\lfloor k/3 \rfloor$ manipulators vote $(b > a > c > \widehat{C} - \{a, b, c\})$, and $\lfloor k/3 \rfloor$ manipulators vote $(a > c > b > \widehat{C} - \{a, b, c\})$. Note that a beats c pairwise, b beats a pairwise, and c beats every candidate in $\widehat{C} - \{a, c\}$ pairwise, so there is no Condorcet winner. If two manipulators remain, then have one vote \widehat{C} and the other vote $\overleftarrow{\widehat{C}}$. Otherwise, if a single manipulator remains, since a beats c pairwise after the manipulators act as above, when the one remaining manipulator votes $(c > \dots)$, no Condorcet winner is created.
 - B. If a beats b pairwise, then c does not beat b pairwise. It follows that c ties b pairwise. Now switch candidates a and b , and we are in the previous case.

For the destructive cases, since Condorcet elections satisfy unique-WARP, the chair cannot, by partitioning of candidates or by runoff partitioning of candidates, cause a candidate that is a unique winner to no longer be a unique winner [HHR07]. This implies that control is possible if and only if the manipulators can vote so that p is not a winner in (C, V) . It is immediate that the optimal action for the manipulators is to put p last. \square

Theorem 6.5.19 Condorcet- $[\frac{C}{D}]C[\frac{PC}{RPC}]$ - $[\frac{CF}{MF}]$ are each in P.

Proof. Given an election (C, V) and a preferred candidate of the chair $p \in C$, we can determine in polynomial time if p can be made a winner by partitioning candidates and by runoff partitioning of candidates as follows.

For the constructive cases, since Condorcet elections satisfy both WARP and unique-WARP, we know from Theorems 6.5.3 and 6.5.5 (which each apply only to the TE model,

but since the Condorcet election system never has more than one winner, for Condorcet elections TE and TP are in effect identical) that control is possible if and only if control is possible using partition $(C - \{p\}, \{p\})$. The manipulators can preclude p from winning if and only if there is a candidate $c \neq p$ that can be made to uniquely win $(C - \{p\}, V)$ and ties-or-beats p pairwise. This can easily be checked by having all manipulators vote for c .

For the destructive cases, since Condorcet elections satisfy unique-WARP, the chair cannot, by partitioning of candidates or by runoff partitioning of candidates, cause a candidate that is a unique winner to no longer be a unique winner [HHR07]. This implies that control is possible if and only if the manipulators cannot vote so that p becomes a winner in (C, V) . It is immediate that the optimal action for the manipulators is to vote for p . \square

Theorem 6.5.20 Condorcet-M+DCPV is in P.

Proof. Given an election (C, V) and a despised candidate of the chair $p \in C$, we can determine in polynomial time if p can be precluded from winning by partitioning voters as follows.

If there exists a candidate $r \in C - \{p\}$ such that when all manipulators rank p last, r ties-or-beats p pairwise, then control is possible by having all manipulators rank p last and using partition (V, \emptyset) .

If no such candidate exists, the only way to ensure that p is not a winner is to ensure that p does not participate in the runoff. Suppose there exists a partition and a manipulation such that p is not a unique winner of either subelection. If in this partition we set all manipulators to rank p last, p still does not win either subelection. So, we can check whether we are in this case by having all manipulators rank p last, and then use the polynomial-time algorithm for Condorcet-DCPV from [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1). \square

Below we state a lemma analogous to Lemma 6.5.14, but for Condorcet elections.

Lemma 6.5.21 *If there exists a partition of voters such that p is not a Condorcet winner when all manipulators vote for p , then there exists a partition such that p can never be made a winner by the manipulators.*

Proof. Given an election (C, V) and a candidate $p \in C$, we do the following.

Let (V_1, V_2) be a partition such that p is not a winner when all manipulators vote for p . So, either there exists a candidate $r \in C - \{p\}$ such that r ties-or-beats p pairwise when all manipulators vote for p , or p is not a unique winner of either subelection.

In the former case the partition (V, \emptyset) will always work, and in the latter case it is clear to see that there is no way for the manipulators to make p a unique winner of either subelection, so we are done. \square

Lemma 6.5.21 implies that Condorcet-DCPV-CF is in P, since control is possible if and only if control is possible when all manipulators vote for p . This can be checked using the polynomial-time algorithm from [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1).

Theorem 6.5.22 Condorcet-DCPV-CF *is in P.*

A similar argument as in the proof of Theorem 6.5.16 shows that Lemma 6.5.21 above also implies that the corresponding “MF” case is also in P.

Theorem 6.5.23 Condorcet-DCPV-MF *is in P.*

6.5.3 Approval

Theorem 6.5.24 *For approval elections, the following hold.*

1. $M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} C$ are each in P.
2. $M + DC \begin{bmatrix} A \\ D \end{bmatrix} V$ are both in P.
3. $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} C - \begin{bmatrix} CF \\ MF \end{bmatrix}$ are each in P.
4. $DC \begin{bmatrix} A \\ D \end{bmatrix} V - \begin{bmatrix} CF \\ MF \end{bmatrix}$ are each in P.

Proof. For the constructive cooperative and the destructive competitive cases it is clear that the manipulators should all approve of only p .

For the destructive cooperative and the constructive competitive cases the optimal action for the manipulators is approve of all candidates except p .

In all cases we can determine if the chair can be successful by assuming the manipulators vote as above and using the corresponding p-time algorithm for control from Bartholdi, Tovey, and Trick [BTT92] (for the constructive cases) or from Hemaspaandra, Hemaspaandra, and Rothe [HHR07] (for the destructive cases), modified in the obvious way for the nonunique-winner case (see Observation 6.5.1). \square

Theorem 6.5.25 approval- $M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix}$ are each in P.

Proof. Given an election (C, V) and a preferred candidate of the chair $p \in C$, we can determine in polynomial time if p can be made a winner by partitioning of candidates and by runoff partitioning of candidates as follows. Let k denote the number of manipulators in V .

For the constructive “TE” cases we do the following. Since approval elections satisfy both WARP and unique-WARP, we know from Theorems 6.5.3 and 6.5.5 that control is possible if and only if control is possible using partition $(C - \{p\}, \{p\})$. Set all manipulators to approve of p . If that makes p an overall winner of the election, we are done. If not, let c be the unique winner of subelection $(C - \{p\}, V)$ (since p will participate in the runoff, the only way p can fail to then win overall is if there is a unique winner of $(C - \{p\}, V)$ who beats p in the runoff). As just mentioned parenthetically, note that after manipulation, c ’s score in this case must be greater than p ’s score. If for all $d \in C - \{p, c\}$, $score(c) > score(d) + k$, c will always be the unique winner of $(C - \{p\}, V)$ and so p will never be an overall winner. If there exists a candidate d in $C - \{p, c\}$ such that $score(c) \leq score(d) + k$, let $score(c) - score(d)$ voters approve of d (in addition to p). In this case, $(C - \{p\}, V)$ does not have a unique winner and so p is an overall winner.

For the constructive “TP” cases, note that control is possible if and only if the manipulators can vote so that p becomes a winner in (C, V) . So the optimal action for the manipulators is to approve of only p . Similarly, for the destructive cases, control is possible if and only if the manipulators can vote so that p does not win (for the “TP” cases) or does not uniquely win (for the “TE” cases) in (C, V) . So the optimal action for the manipulators is to approve of all candidates except p . \square

Theorem 6.5.26 approval- $\left[\begin{smallmatrix} C \\ D \end{smallmatrix} \right] C \left[\begin{smallmatrix} PC \\ RPC \end{smallmatrix} \right] - \left[\begin{smallmatrix} TE \\ TP \end{smallmatrix} \right] - \left[\begin{smallmatrix} CF \\ MF \end{smallmatrix} \right]$ are each in P.

Proof. Given an election (C, V) and a preferred candidate of the chair $p \in C$, we can determine in polynomial time if p can be made a winner by partitioning candidates and by runoff partitioning of candidates as follows.

For the constructive “TE” cases, since approval elections satisfy both WARP and unique-WARP, we know from Theorems 6.5.3 and 6.5.5 that control is possible if and only if control is possible using partition $(C - \{p\}, \{p\})$. The manipulators can preclude p from winning if and only if there is a candidate $c \neq p$ that can be made to uniquely win using partition $(C - \{p\}, \{p\})$. This can easily be checked by having all manipulators approve of only c .

For the constructive “TP” cases, note that control is possible if and only if the manipulators cannot vote so that p does not become a winner in (C, V) . So the optimal action for the manipulators, regardless of who goes first, is to approve of all candidates except p .

Similarly, for the destructive cases, control is possible if and only if the manipulators cannot vote so that p becomes a winner (for the “TP” cases) or a unique winner (for the “TE” cases) in (C, V) . So the optimal action for the manipulators, regardless of who goes first, is to approve of only p . \square

Theorem 6.5.27 approval-M+DCPV- $\left[\begin{smallmatrix} \text{TE} \\ \text{TP} \end{smallmatrix}\right]$ is in P.

Proof. Given an election (C, V) and a despised candidate of the chair $p \in C$, we can determine in polynomial time if p can be precluded from winning by partitioning voters for the “TE” case as follows.

1. If there is a candidate, $c \neq p$ such that $\text{score}(p) \leq \text{score}(c) + k$, then control is possible by having all manipulators disapprove of only p and using partition (V, \emptyset) .
2. If we are not in Case 1, the only way to preclude p from being a winner is if p doesn't make it to the runoff, i.e., if there exist a partition and a manipulation such that p is not a unique winner of either subelection. If in this partition we make all manipulators vote to disapprove of only p , p is still not a unique winner of either subelection. So, we can check whether we are in this case by having all manipulators vote to disapprove of only p , and then using the polynomial-time algorithm for approval-DCPV-TE from [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1).

For the “TP” case, replace “ \leq ” by “ $<$ ” in Case 1, and “unique winner” by “winner” and “approval-DCPV-TE” by “approval-DCPV-TP” in Case 2. \square

Below we state a lemma analogous to Lemma 6.5.14, but for approval elections.

Lemma 6.5.28 *If there exists a partition of voters such that p is not an approval winner in the “TE” (“TP”) model when all manipulators approve of only p , then there exists a partition such that p can never be made an approval winner by the manipulators in the same tie-breaking model.*

Proof. The proof for the “TE” case follows similarly to the proof of Lemma 6.5.14, so we just provide the proof of the “TP” case.

Given an election (C, V) and a candidate $p \in C$, we do the following.

Let (V_1, V_2) be a partition such that p is not a winner when all manipulators approve of only p . If p can never be made a winner by the manipulators in this partition then we are

done. So, suppose there exists a manipulation such that p is an overall winner (with the partition (V_1, V_2)). Without loss of generality p is a winner of the subelection (C, V_1) . Then if all manipulators in V_1 approve of only p , we know that p remains a winner of (C, V_1) . Note we don't get any new winners in (C, V_1) . Since p is not an overall winner if all manipulators approve of only p there is a candidate $c \neq p$ such that if all manipulators vote for p , c is a winner of (C, V_2) and $score(c) > score(p)$.

Now move all manipulators from V_2 to V_1 . Note that c remains a winner of (C, V_2) and that c will always beat p in the runoff. It follows that in this new partition, p is never a winner, no matter what the manipulators do. \square

Lemma 6.5.28 implies that approval-DCPV-TE-CF and approval-DCPV-TP-CF are both in P, since control is possible if and only if (nonmanipulator) control is possible when all manipulators approve of only p . This can be checked using the corresponding polynomial-time algorithms from Hemaspaandra, Hemaspaandra, and Rothe [HHR07], modified in the obvious way for the nonunique-winner case (see Observation 6.5.1).

Theorem 6.5.29 approval-DCPV- $\begin{bmatrix} TE \\ TP \end{bmatrix}$ -CF are both in P.

Lemma 6.5.28 above also implies that the corresponding manipulators-first cases are both in P. The proof of the following theorem follows from a similar argument as the proof of Theorem 6.5.16.

Theorem 6.5.30 approval-DCPV- $\begin{bmatrix} TE \\ TP \end{bmatrix}$ -MF are both in P.

6.6 Weighted Voters

We now give results for veto and Borda, including, for the latter, an interesting increase in complexity, by considering weighted elections.

Recall that in a weighted election every voter has a positive integer weight, and a voter with weight ω counts as ω voters. In weighted voter control cases, the addition/deletion limit still pertains to the number of voters that can be added or deleted. Consider the case of 3-candidate weighted veto elections. The known results on this are that constructive coalitional manipulation is NP-complete [CSL07], destructive coalitional manipulation is in P [CSL07], and CCAV and CCDV are both in P [FHH15]. The following result, whose second part may be surprising, shows that for this system $CC\begin{bmatrix} A \\ D \end{bmatrix}V\text{-}\begin{bmatrix} CF \\ MF \end{bmatrix}$ are all in P—not NP-complete.

Theorem 6.6.1 For 3-candidate weighted veto elections, the following hold.

1. $M+CC\left[\begin{smallmatrix} A \\ D \end{smallmatrix}\right]V$ are both NP-complete.
2. $CC\left[\begin{smallmatrix} A \\ D \end{smallmatrix}\right]V-\left[\begin{smallmatrix} CF \\ MF \end{smallmatrix}\right]$ are each in P.

Proof. The first case follows directly from the fact that constructive manipulation is NP-complete [CSL07] and the inheritance observations from Section 6.3 (as the relevant result there holds even for the weighted case).

For the competitive cases, note that the only action that makes sense for the manipulators is to veto p . This holds regardless of whether the manipulators or the chair goes first. So, we let the manipulators veto p and then run the polynomial-time algorithm for CCAV and CCDV from [FHH15]. \square

3-candidate weighted Borda elections show a true increase in complexity. The known results for this system are that constructive coalitional manipulation is NP-complete [CSL07], destructive coalitional manipulation is in P [CSL07], and CCAV and CCDV are both NP-complete [FHH15] and thus all these problems are in NP. Yet we show that CCAV-MF is coNP-hard, and so cannot be in NP unless the polynomial hierarchy collapses to $NP \cap coNP$.

Theorem 6.6.2 *For 3-candidate weighted Borda elections, the following hold.*

1. $M+CC\left[\begin{smallmatrix} A \\ D \end{smallmatrix}\right]V$ are both NP-complete.
2. $CC\left[\begin{smallmatrix} AV-CF \\ DV-\left[\begin{smallmatrix} CF \\ MF \end{smallmatrix}\right] \end{smallmatrix}\right]$ are each NP-hard.
3. CCAV-MF is NP-hard and coNP-hard.
4. $CC\left[\begin{smallmatrix} A \\ D \end{smallmatrix}\right]V-CF$ is NP-complete.

Proof. The first case follows directly from the fact that manipulation is NP-complete [CSL07] and the inheritance observations from Section 6.3.

The remaining NP-hardness results follow from the NP-completeness of CCAV and CCDV and the inheritance observations from Section 6.3.

To show that CCAV-CF is in NP, guess a set of voters to add, and then check that the manipulators can't make p not win. We do this by setting all manipulators to $(a > b > p)$, checking that p is a winner, and then setting all manipulators to $(b > a > p)$, and checking that p is a winner. A similar argument shows that CCDV-CF is in NP.

It remains to show that CCAV-MF is coNP-hard, i.e., that the complement of CCAV-MF is NP-hard. We will reduce from Partition. Given a nonempty sequence of positive integers k_1, \dots, k_t that sums to $2K$, we will construct an election such that there is a partition (i.e.,

a subsequence of k_1, \dots, k_t that sums to K) if and only if the manipulators can vote in such a way that the chair won't be able to make p a winner.

We construct the following election: We have manipulators with weights k_1, \dots, k_t . The manipulators are registered voters. We have two unregistered voters, both with weight $3K - 1$. One of these voters votes ($p > a > b$) and one votes ($p > b > a$). We have addition limit one, i.e., the chair can add at most one voter.

If there is a partition, then the manipulators vote so that a total of K vote weight casts the vote ($a > b > p$) and a total of K vote weight casts the vote ($b > a > p$). So, the scores of p , a , and b are 0, $3K$, and $3K$. There is no way for the chair to make p a winner by adding at most one voter. If the chair adds the weight $3K - 1$ voter voting ($p > a > b$), the score of p is $6K - 2$ and the score of a is $3K + (3K - 1) = 6K - 1$ and so p is not a winner. Adding the other voter gives a score of $6K - 2$ for p and a score of $6K - 1$ for b and again p is not a winner.

Now consider the case that there is no partition. Look at the scores of the candidates after the manipulators have voted. Without loss of generality, assume that $score(a) \leq score(b)$. Then $score(a) \leq 3K - 1$ (since there is no partition) and $score(b) \leq 4K$. Now the chair adds the weight $3K - 1$ voter voting ($p > a > b$). After adding that voter, p 's score is $6K - 2$, a 's score is at most $(3K - 1) + (3K - 1)$ and b 's score is at most $4K$. It follows that p is a winner. \square

6.7 Related Work

Many variants of control and manipulation have been studied, with some combining different aspects of each type of attack.

The model of bribery introduced by Faliszewski, Hemaspaandra, and Hemaspaandra [FHH09] can be viewed as control by deleting voters followed by manipulation, but in that case the manipulators would only take the place of the deleted voters (and it would only be the case where the chair and the manipulators cooperate). The idea of enhancing control with manipulative voters has been mentioned in the literature, namely, in a paragraph of [FHH11]. That paper cooperatively integrated with control, to a certain extent, bribery [FHH09]. In that paper's conclusions and open directions, there is a paragraph suggesting that manipulation could and should also be integrated into that paper's "multi-prong setting," and commending such future study to interested readers. That paragraph was certainly influential in our choice of this direction. However, it is speaking just of the cooperative case, and provides no results on this since it is suggesting a direction for study.

The possible winner problem in the setting where new candidates are added to an election, first introduced by Konczak and Lang [KL05], and expanded on by later work (see, e.g., [CLMM10, BRR11, XLM11, CLM⁺12]), considers an election where there is an initial set of candidates (that the voters all have total-order preferences over), a specified candidate, and a set of additional candidates (that the voters have no specified preferences over), and asks if there exists a way to extend the preferences of the voters to total orders that include the additional candidates such that the specified candidate wins. This may seem like it is combining manipulation and control by adding candidates, but in the possible winner model there is no choice over which candidates are added to the election, which there is in the case of control. So it is better thought of as a generalization of the manipulation problem (as the abovementioned papers note). We also mention the work by Baumeister et al. [BRR⁺12], which uses the term possible winner in a different way, since instead of considering the votes initially unset over the new candidates, weights are initially unset. So as that paper notes, the model can be seen as a generalization of control by adding and deleting voters. However, the model is not a generalization of manipulation.

When we consider control by partitioning voters/candidates in our setting we consider both the case where the manipulators can change their votes after the first round and the case when they cannot. We mention here the recent work by Narodytska and Walsh [NW12, NW13] and Fitzsimmons, Hemaspaandra, and Hemaspaandra [FHH16] that explore how revoting affects the complexity of manipulation.

Our results show that the order of who goes first (the chair or the manipulators) can affect the hardness of our problems. Order was also seen to play an important role in recent work in the settings of online control [HHR12a, HHR12b], and online manipulation [HHR14].

6.8 Conclusions

We introduced the model of control in the presence of manipulators both for the case where the election chair and the manipulators share a common goal and the case where they have conflicting goals.

We have established general inheritance results and complexity upper bounds for control in the presence of manipulators, for both cooperative and competitive settings. We for the upper bounds provided matching lower bounds, but also showed that for many natural systems the complexity is far lower than the general upper bounds.

It would be interesting to explore the complexity of control in the presence of manipulators for election systems whose unweighted manipulation and/or control problems are computationally difficult, and to expand on the results we have for weighted elections. Ad-

ditionally, this model seems to be able to reach quite high levels of complexity, and it would be interesting to see what complexity results could be realized for natural systems with computationally difficult winner problems, such as Kemeny elections.

Chapter 7

Conclusions

Elections are a way for a group of agents to aggregate their preferences to reach a decision, whether they are citizens of a country, members on a committee, or agents in an electronic multiagent system. There will often be incentives for agents to try to change the outcome of an election and so it is important that we understand the complexity of these attacks and how they should be modeled.

In this thesis we examined how different natural models for the votes, the electorate, the representation, and the structure of the attack affect the computational complexity of election-attack problems.

In Chapter 3 we examined how allowing voters to state ties in their preferences affects the computational complexity of the three standard election attacks of manipulation, bribery, and control. We found the first natural case where the complexity of manipulation increases when moving from tie-free votes to votes with ties, found cases where the complexity of bribery can increase and when it can decrease, and stated a general observation about the complexity of control.

In Chapter 4 we considered four important models of single-peakedness for votes with ties, and showed that for each model it is in P to determine if a given electorate satisfies that model for votes with ties. We also expanded on our work on manipulation for votes with ties by considering the complexity of single-peaked votes with ties, and we found that the complexity can depend on the model used.

We considered how the succinct representation of the voters affects the complexity of different election problems in Chapter 5, and even though this representation can be exponentially smaller, we found that the complexity of election attacks rarely increases (in the length of the input), and explained this behavior by showing that several techniques used to show polynomial-time results for nonsuccinct voters can be adapted for the succinct case.

We presented a new election-attack model in Chapter 6 that models the natural case

of control in the setting where there are manipulators. We considered both the case where the chair and the manipulators have the same goal, and the case where they have directly conflicting goals. We presented general upper bounds for the complexity of control in the presence of manipulators, and we showed most of these bounds to be tight for the second and third levels of the polynomial hierarchy. However, for important natural systems we found that the complexity can be much lower than these bounds.

Our examination of the behavior of the complexity of election-attack problems in different natural settings gives new insights into the behavior of election systems as well as of complexity classes. It will be interesting to see which ideas from our work can be applied to the study of related combinatorial problems, and also what new insights can be gained by considering other measures of hardness.

Bibliography

- [AMO93] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, 1993.
- [Arr50] K. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.
- [Bar01] S. Barberà. An introduction to strategy-proof social choice functions. *Social Choice and Welfare*, 18(4):619–653, 2001.
- [Bar07] S. Barberà. Indifferences and domain restrictions. *Analyse & Kritik*, 29(2):146–162, 2007.
- [BBHH15] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53:439–496, 2015.
- [BCE⁺16] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [BCW13] R. Bredereck, J. Chen, and G. Woeginger. A characterization of the single-crossing domain. *Social Choice and Welfare*, 41(4):989–998, 2013.
- [BD09] N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections based on scoring rules. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science*, pages 124–136. Springer-Verlag *Lecture Notes in Computer Science #5734*, August 2009.
- [BF83] S. Brams and P. Fishburn. *Approval Voting*. Birkhäuser, Boston, 1983.
- [BFLR12] D. Baumeister, P. Faliszewski, J. Lang, and J. Rothe. Campaigns for lazy voters: Truncated ballots. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 577–584, June 2012.

- [BH08] H. Buhrman and J. Hitchcock. NP-hard sets are exponentially dense unless $\text{coNP} \subseteq \text{NP}/\text{poly}$. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity*, pages 1–7, June 2008.
- [BHKS13] F. Brandt, P. Harrenstein, K. Kardel, and H. Seedig. It only takes a few: On the hardness of voting with a constant number of agents. In *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 375–382, May 2013.
- [BL76] K. Booth and G. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *Journal of Computer and System Sciences*, 13(3):335–379, 1976.
- [Bla48] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
- [Bla58] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- [BNW11] N. Betzler, R. Niedermeier, and G. Woeginger. Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 55–60, July 2011.
- [BO91] J. Bartholdi, III and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.
- [BR16] D. Baumeister and J. Rothe. Preference aggregation by voting. In J. Rothe, editor, *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, pages 197–325. Springer, 2016.
- [BRR11] D. Baumeister, M. Roos, and J. Rothe. Computational complexity of two variants of the possible winner problem. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 853–860. International Foundation for Autonomous Agents and Multiagent Systems, May 2011.
- [BRR⁺12] D. Baumeister, M. Roos, J. Rothe, L. Schend, and L. Xia. The possible winner problem with uncertain weights. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 133–138, August 2012.

- [BT76] I. Borosh and L. Treybig. Bounds on positive integral solutions of linear Diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976.
- [BT86] J. Bartholdi, III and M. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.
- [BTT89a] J. Bartholdi, III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [BTT89b] J. Bartholdi, III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.
- [BTT92] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.
- [Can04] D. Cantala. Choosing the level of a public good when agents have an outside option. *Social Choice and Welfare*, 22(3):491–514, 2004.
- [CCHO05] J. Cai, V. Chakaravarthy, L. Hemaspaandra, and M. Ogihara. Competing provers yield improved Karp–Lipton collapse results. *Information and Computation*, 198(1):1–23, 2005.
- [CLM⁺12] Y. Chevaleyre, J. Lang, N. Maudet, J. Monnot, and L. Xia. New candidates welcome! Possible winners with respect to the addition of new candidates. *Mathematical Social Sciences*, 64(1):74–88, 2012.
- [CLMM10] Y. Chevaleyre, J. Lang, N. Maudet, and J. Monnot. Possible winners when new candidates are added: The case of scoring rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 762–767. AAAI Press, July 2010.
- [Con85] J.-A.-N. de Caritat, Marquis de Condorcet. *Essai sur l’Application de L’Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. 1785. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale.
- [Cop51] A. Copeland. A “reasonable” social welfare function. Mimeographed notes from a Seminar on Applications of Mathematics to the Social Sciences, University of Michigan, 1951.

- [CSL07] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.
- [CT07] T. Coleman and V. Teague. On the complexity of manipulating elections. In *Proceedings of the Thirteenth Computing: The Australasian Theory Symposium*, pages 25–33. Australian Computer Society Inc., January/February 2007.
- [DF94] J.-P. Doignon and J.-C. Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16(2):218–233, 1994.
- [DKNS01] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, pages 613–622. ACM Press, March 2001.
- [DKNW11] J. Davies, G. Katsirelos, N. Narodytska, and T. Walsh. Complexity and algorithms for Borda manipulation. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pages 657–662. AAAI Press, August 2011.
- [Dod76] C. Dodgson. A method of taking votes on more than two issues. Pamphlet printed by the Clarendon Press, Oxford, and headed “not yet published” (see the discussions in [MU95, Bla58], both of which reprint this paper), 1876.
- [DS00] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard–Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.
- [EFS12] E. Elkind, P. Faliszewski, and A. Slinko. Clone structures in voters’ preferences. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 496–513, June 2012.
- [ELÖ08] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence*, pages 366–370, July 2008.
- [Eme13] P. Emerson. The original Borda count and partial voting. *Social Choice and Welfare*, 40(2):352–358, 2013.
- [FBC14] R. Freeman, M. Brill, and V. Conitzer. On the axiomatic characterization of runoff voting rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 675–681, July 2014.

- [FG65] D. Fulkerson and O. Gross. Incidence matrices and interval graphs. *Pacific Journal of Math*, 15(3):835–855, 1965.
- [FH15] Z. Fitzsimmons and E. Hemaspaandra. Complexity of manipulative actions when voting with ties. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, pages 103–119, September 2015.
- [FH16a] Z. Fitzsimmons and E. Hemaspaandra. Complexity of manipulative actions when voting with ties. In *Proceedings (Workshop Notes) of the 6th International Workshop on Computational Social Choice*, June 2016.
- [FH16b] Z. Fitzsimmons and E. Hemaspaandra. The complexity of succinct elections. Technical Report arXiv:1611.08927 [cs.GT], arXiv.org, November 2016.
- [FH16c] Z. Fitzsimmons and E. Hemaspaandra. Modeling single-peakedness for votes with ties. In *8th European Starting AI Researcher Symposium*, pages 63–74, August 2016.
- [FH17] Z. Fitzsimmons and E. Hemaspaandra. The complexity of succinct elections. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 4921–4922, February 2017.
- [FHH09] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.
- [FHH10] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.
- [FHH11] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Multimode attacks on elections. *Journal of Artificial Intelligence Research*, 40:305–351, 2011.
- [FHH13a] Z. Fitzsimmons, E. Hemaspaandra, and L. Hemaspaandra. Control in the presence of manipulators: Cooperative and competitive cases. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 113–119. AAAI Press, August 2013.
- [FHH13b] Z. Fitzsimmons, E. Hemaspaandra, and L. Hemaspaandra. Control in the presence of manipulators: Cooperative and competitive cases. Technical Report arXiv:1308.0544 [cs.GT], Computing Research Repository, arXiv.org/corr/, August 2013. Revised, June 2017.

- [FHH14a] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *Journal of Artificial Intelligence Research*, 207:69–99, 2014.
- [FHH14b] Z. Fitzsimmons, E. Hemaspaandra, and L. Hemaspaandra. Control in the presence of manipulators: Cooperative and competitive cases. In *Proceedings (Workshop Notes) of the 5th International Workshop on Computational Social Choice*, June 2014.
- [FHH15] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Weighted electoral control. *Journal of Artificial Intelligence Research*, 52:507–542, 2015.
- [FHH16] Z. Fitzsimmons, E. Hemaspaandra, and L. Hemaspaandra. Manipulation complexity of same-system runoff elections. *Annals of Mathematics and Artificial Intelligence*, 77(3–4):159–189, 2016.
- [FHHR09] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.
- [FHHR11] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107, 2011.
- [FHS08] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: Ties matter. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 983–990, May 2008.
- [FHS10] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Manipulation of Copeland elections. In *Proceedings of the 9th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 367–374, May 2010.
- [FHS12] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Weighted manipulation for four-candidate Llull is easy. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 318–323, August 2012.
- [Fis73] P. Fishburn. *The Theory of Social Choice*, volume 264. Princeton University Press, 1973.
- [Fit15] Z. Fitzsimmons. Single-peaked consistency for weak orders is easy. In *Proceedings of Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge*, pages 127–140, June 2015.

- [FP10] P. Faliszewski and A. Procaccia. AI’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
- [Gär73] P. Gärdenfors. Positionalist voting functions. *Theory and Decision*, 4(1):1–24, 1973.
- [Gib73] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.
- [GJ79] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
- [GK68] M. Garman and M. Kamien. The paradox of voting: Probability calculations. *Behavioral Science*, 13(4):306–316, 1968.
- [GMHS99] S. Ghosh, M. Mundhe, K. Hernandez, and S. Sen. Voting for movies: The anatomy of recommender systems. In *Proceedings of the 3rd Annual Conference on Autonomous Agents*, pages 434–435. ACM Press, 1999.
- [Hem89] L. Hemachandra. The strong exponential hierarchy collapses. *Journal of Computer and System Sciences*, 39(3):299–322, 1989.
- [HH07] E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.
- [HHM13] E. Hemaspaandra, L. Hemaspaandra, and C. Menton. Search versus decision for election manipulation problems. In *Proceedings of the 30th Annual Symposium on Theoretical Aspects of Computer Science*, pages 377–388. Leibniz International Proceedings in Informatics (LIPIcs), February/March 2013.
- [HHR] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The complexity of online voter control in sequential elections. *Autonomous Agents and Multi-Agent Systems*. doi:10.1007/s10458-016-9349-1, To appear.
- [HHR97] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll’s 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.
- [HHR07] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.

- [HHR09] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly*, 55(4):397–424, 2009.
- [HHR12a] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Controlling candidate-sequential elections. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 905–906. IOS Press, August 2012. Journal version [HHR17].
- [HHR12b] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Online voter control in sequential elections. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pages 396–401. IOS Press, August 2012. Journal version [HHR].
- [HHR14] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The complexity of online manipulation of sequential elections. *Journal of Computer and System Sciences*, 80(4):697–710, 2014.
- [HHR17] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The complexity of controlling candidate-sequential elections. *Theoretical Computer Science*, 678:14–21, 2017.
- [HHS14a] E. Hemaspaandra, L. Hemaspaandra, and H. Schnoor. A control dichotomy for pure scoring rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 712–720, July 2014.
- [HHS14b] E. Hemaspaandra, L. Hemaspaandra, and H. Schnoor. A control dichotomy for pure scoring rules. Technical Report arXiv:1404.4560 [cs.GT], arXiv.org, April 2014.
- [HP01] G. Hägele and F. Pukelsheim. The electoral writings of Ramon Llull. *Studia Lulliana*, 41(97):3–38, 2001.
- [HS16] E. Hemaspaandra and H. Schnoor. Dichotomy for pure scoring rules under manipulative electoral actions. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1071–1079, Aug/Sep 2016.
- [HSV05] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.
- [HW12] L. Hemaspaandra and R. Williams. An atypical survey of typical-case heuristic algorithms. *SIGACT News*, 43(4):71–89, December 2012.

- [Kar72] R. Karp. Reducibilities among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103, 1972.
- [Kem59] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.
- [KL05] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of the 1st Multidisciplinary Workshop on Advances in Preference Handling (M-PREF 2005)*, pages 124–129, August 2005.
- [KS60] J. Kemeny and L. Snell. *Mathematical Models in the Social Sciences*. Ginn, 1960.
- [Lac14] M. Lackner. Incomplete preferences in single-peaked electorates. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 742–748, July 2014.
- [Len83] H. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.
- [Lev86] L. Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.
- [Lin12] A. Lin. *Solving Hard Problems in Election Systems*. PhD thesis, Rochester Institute of Technology, Rochester, NY, 2012.
- [McG53] D. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.
- [MFG12] N. Mattei, J. Forshee, and J. Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Proceedings (Workshop Notes) of the 4th International Workshop on Computational Social Choice*, September 2012.
- [Mir71] J. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971.
- [ML15] V. Menon and K. Larson. Complexity of manipulation in elections with top-truncated ballots. Technical Report arXiv:1505.05900 [cs.GT], arXiv.org, July 2015.

- [ML16] V. Menon and K. Larson. Reinstating combinatorial protections for manipulation and bribery in single-peaked and nearly single-peaked electorates. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, February 2016.
- [Mou80] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [MPC13] A. Mao, A. Procaccia, and Y. Chen. Better human computation through principled voting. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, July 2013.
- [MPR13] E. Mossel, A. Procaccia, and M. Rácz. A smooth transition from powerlessness to absolute power. *Journal of Artificial Intelligence Research*, 48:923–951, 2013.
- [MS72] A. Meyer and L. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *Proceedings of the 13th IEEE Symposium on Switching and Automata Theory*, pages 125–129, October 1972.
- [MU95] I. McLean and A. Urken. *Classics of Social Choice*. University of Michigan Press, 1995.
- [MW13] N. Mattei and T. Walsh. PREFLIB: A library for preferences. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 259–270, 2013.
- [NW12] N. Narodytska and T. Walsh. Manipulating two stage voting rules. In *Proceedings (Workshop Notes) of the 4th International Workshop on Computational Social Choice*, pages 323–334, September 2012.
- [NW13] N. Narodytska and T. Walsh. Manipulating two stage voting rules. In *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 423–430, May 2013.
- [NW14] N. Narodytska and T. Walsh. The computational impact of partial votes on strategic voting. In *Proceedings of the 21st European Conference on Artificial Intelligence*, August 2014.
- [Pel81] B. Peleg. Monotonicity properties of social choice correspondences. *Game Theory and Mathematical Economics*, pages 97–101, 1981.

- [PRVW11] M. Pini, F. Rossi, K. Venable, and T. Walsh. Incompleteness and incomparability in preference aggregation: Complexity results. *Artificial Intelligence*, 175(7):1272–1289, 2011.
- [PY84] C. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2):244–259, 1984.
- [PZ83] C. Papadimitriou and S. Zachos. Two remarks on the power of counting. In *Proceedings 6th GI Conference on Theoretical Computer Science*, pages 269–276. Springer-Verlag *Lecture Notes in Computer Science #145*, January 1983.
- [RO73] W. Riker and P. Ordeshook. *An Introduction to Positive Political Theory*. Prentice-Hall, 1973.
- [RSV03] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.
- [Rus07] N. Russell. Complexity of control of Borda count elections. Master’s thesis, Rochester Institute of Technology, 2007.
- [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(2):187–217, 1975.
- [Sch03] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag, 2003.
- [Sch11] M. Schulze. A new monotonic and clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36(2):267–303, 2011.
- [SM73] L. Stockmeyer and A. Meyer. Word problems requiring exponential time. In *Proceedings of the 5th ACM Symposium on Theory of Computing*, pages 1–9. ACM Press, April/May 1973.
- [Smi73] J. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.
- [Sto76] L. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

- [SU02a] M. Schaefer and C. Umans. Completeness in the polynomial-time hierarchy: Part I: A compendium. *SIGACT News*, 33(3):32–49, 2002.
- [SU02b] M. Schaefer and C. Umans. Completeness in the polynomial-time hierarchy: Part II. *SIGACT News*, 33(4), 2002.
- [SZ12] M. Sanver and W. Zwicker. Monotonicity properties and their adaptation to irresolute social choice rules. *Social Choice and Welfare*, 39(2-3):371–398, 2012.
- [Tid87] T. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- [Wag87] K. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51(1-2):53–80, 1987.
- [Wal07] T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 3–8, July 2007.
- [Wal11] T. Walsh. Where are the hard manipulation problems? *Journal of Artificial Intelligence Research*, 42(1):1–29, 2011.
- [Wra76] C. Wrathall. Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):23–33, 1976.
- [XC11] L. Xia and V. Conitzer. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 446–451, July 2011.
- [XLM11] L. Xia, J. Lang, and J. Monnot. Possible winners when new alternatives join: New results coming up! In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 829–836. International Foundation for Autonomous Agents and Multiagent Systems, May 2011.
- [Yan15] Y. Yang. Manipulation with bounded single-peaked width: A parameterized study. In *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 77–85, 2015.
- [YL78] H. Young and A. Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.
- [You75] H. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, 1975.

- [You77] H. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16(2):335–353, 1977.
- [Zwi16] W. Zwicker. The theory of voting: An introduction. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*, pages 23–56. Cambridge University Press, 2016.