



Control in the presence of manipulators: cooperative and competitive cases

Zack Fitzsimmons¹  · Edith Hemaspaandra²  · Lane A. Hemaspaandra³ 

© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

Control and manipulation are two of the most studied types of attacks on elections. In this paper, we study the complexity of control attacks on elections in which there are manipulators. We study both the case where the “chair” who is seeking to control the election is allied with the manipulators, and the case where the manipulators seek to thwart the chair. In the latter case, we see that the order of play substantially influences the complexity. We prove upper bounds, holding over every election system with a polynomial-time winner problem, for all standard control cases, and some of these bounds are at the second or third level of the polynomial hierarchy, and we provide matching lower bounds to prove these tight. Nonetheless, for important natural systems the complexity can be much lower. We prove that for approval and plurality elections, the complexity of even competitive clashes between a controller and manipulators falls far below those high bounds, even as low as polynomial time. Yet for a Borda-voting case we show that such clashes raise the complexity unless $NP = coNP$.

Keywords Computational social choice · Control · Manipulation · Complexity

Earlier versions of this paper [31,34] were presented at the Twenty-Third International Joint Conference on Artificial Intelligence and the Fifth International Workshop on Computational Social Choice. This work was done in part while Zack Fitzsimmons was at RIT and while Edith Hemaspaandra and Lane A. Hemaspaandra were on sabbatical visits to ETH-Zürich and the University of Düsseldorf.

✉ Zack Fitzsimmons
zfitzsim@holycross.edu

Edith Hemaspaandra
eh@cs.rit.edu

Lane A. Hemaspaandra
<http://www.cs.rochester.edu/u/lane>

¹ Department of Mathematics and Computer Science, College of the Holy Cross, Worcester, MA 01610, USA

² Department of Computer Science, Rochester Institute of Technology, Rochester, NY 14623, USA

³ Department of Computer Science, University of Rochester, Rochester, NY 14627, USA

1 Introduction

Elections are an important tool in reaching decisions, in both human and online settings. Regarding online settings, elections have been proposed in such varied, multiagent-systems settings as planning, recommender systems/collaborative filtering, and web spam reduction [18,21,37,58]. With the growing importance of the online world and multiagent systems, the use of elections in computer-based settings will but increase. Unfortunately, given the relentless growth in the power of computers, it is natural to worry that computers will also be increasingly brought to bear in planning manipulative attacks on elections. Indeed, this is one of the central concerns of the relatively young multiagent systems subarea known as computational social choice [8,9].

Two of the most computationally studied types of attacks on elections are known as “control” and “manipulation.” Bartholdi et al. [3] defined the basic (constructive) control types and initiated their complexity-theoretic study. Manipulation has roots going back long before the work of Bartholdi, Orlin, Tovey, and Trick, but Bartholdi, Orlin, Tovey, and Trick [1,2] initiated the complexity-theoretic study of manipulation. In control, an agent, usually referred to as “the chair,” tries to make a given candidate win by adding/deleting/partitioning voters or candidates. In manipulation, some nonmanipulative voters and a coalition of manipulative voters vote under some election system, and the manipulative voters seek to make a given candidate win. Note that both of the problems of control and manipulation can also be studied in the so-called destructive case [14,42] where the goal of the chair (or the manipulators) is to ensure that a despised candidate does not win.

There is a broad literature on the computational complexity of control and manipulation (see, e.g., the surveys [15,23,28,30,45]). However, the present paper considers *control attacks against elections that contain manipulators*. We consider both the cooperative and the competitive cases.

In the cooperative case, the chair is allied with the manipulative coalition. For example, perhaps during a CS department’s hiring, the department chair, who happens to also be the senior member of the systems group, is mounting a control by partition of voters attack (in which he or she is dividing the faculty into two subcommittees, one to decide which candidates are strong enough teachers to merit further consideration, and the other to decide which candidates are strong enough researchers to merit further consideration), and also is able to directly control the votes of every one of his or her fellow members of the department’s systems faculty. The chair’s goal is to make some particular candidate, perhaps Dr. I. M. Systems, be the one chosen for hiring.

In the even more interesting competitive case, which can be thought of in a certain sense as *control versus manipulation*, we will assume that the manipulative coalition’s goal is to keep the chair from achieving the chair’s goal. For the competitive case, we will look at the case where the chair acts before the manipulators, and at the case where the manipulators act before the chair. For control attacks by so-called partition, in which there is a two-round election, we will consider the case where the manipulators can change their votes in the second round, and the case where the manipulators cannot change their votes in the second round.

Our main contributions are as follows.

- Building on the existing notions of control and manipulation, we give natural definitions that capture our cooperative and competitive notions as problems whose computational complexity can be studied, and we note how existing hardness results for control and manipulation are, or are not, inherited by our problems.

- We prove upper bounds on our problems. For the competitive case, these are as high as NP^{NP} , $coNP^{NP}$, and $coNP^{NP^{NP}}$, with the notable exception of the case of deleting voters in the “chair-first” setting, which is in $coDP$, i.e., it is the union of an NP set and a $coNP$ set.
- Despite how high those upper bounds are, we show that there are election systems (having polynomial-time winner problems) for which most of those high bounds have matching lower bounds, yielding completeness for those classes.
- For the important election systems approval, Condorcet, and plurality, we show that the complexity of control in the presence of manipulators, whether cooperative or competitive, can be much lower than those upper bounds, even falling as low as polynomial time. Many of the proofs of these cases involve interesting approaches that are quite different from merely seeking to adapt those used in the case of control without manipulators.
- We obtain results, for election systems satisfying versions—called WARP and unique-WARP—of the weak axiom of revealed preferences, on the complexity of control by runoff partitioning of candidates.

The general theme of those results is that the combinatorial explosion that causes many partition-related candidate-control problems to be NP -complete can never exist for election systems that satisfy certain nice properties, such as WARP and unique-WARP. In particular, we will show that such properties can change the challenge facing the chair (of a control problem) from that of needing to worry about every partition to just that of checking one very simple partition. From this, polynomial-time control algorithms immediately follow, as we will see.

The reason that this is interesting is that it is not applying just to one particular system, but rather is noting that some nice behavioral properties themselves ensure the simplicity of certain candidate-partition control problems for *all* systems having the properties.

- We also obtain cases, for veto (Theorem 4.12) and Borda (Theorem 4.13) elections, where competitive control-plus-manipulation is variously easier or harder than one might expect from the separate control and manipulation cases.

2 Related work

The idea of enhancing control with manipulative voters has been mentioned in the literature, namely, in a paragraph of Faliszewski et al. [24]. That paper cooperatively integrated with control, to a certain extent, a different attack type known as bribery [22]. 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 “multiprong setting,” and commanding 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 lovely line of work about “possible winners” [49] in the context of adding candidates might at first seem to be merging manipulation and control. We refer to the line of work explored in [4, 10, 11, 68]. That work considers an election with an initial set of candidates, over which all the voters have complete preferences, and a set of additional candidates over which the voters initially have no preferences, and asks whether, if the entire set of additional candidates is added, there is some way of extending the initial linear orders to now be over all the candidates, in such a way that a particular initial candidate becomes a winner of the

election. Although on its surface this might feel like a cross between manipulation and control by adding candidates, in fact, in this interesting problem there is no actual choice regarding the addition of candidates; all are simply added. Thus this problem is a generalization of manipulation (as the papers note), that happens to be done in a setting that involves adding candidates. It is not a generalization of control by adding, or even so-called unlimited adding, of candidates, as in those the chair must choose what collection of candidates to add. In short, unlike control and unlike this paper, there is no existentially quantified action by a chair. (An interesting recent paper of Baumeister et al. [5] uses the term possible winner in a new, different way, to speak of weights rather than preferences initially being partially unset. That particular paper’s question, as that paper notes, can be seen as a generalization of control by adding and deleting voters. However, their notion is not a generalization of manipulation.)

The present paper does combine control and manipulation, with both those playing active—and sometimes opposing—roles. Manipulation alone has been extensively studied in a huge number of papers; its complexity-theoretic study started with the seminal work of Bartholdi et al. [2] (see also [1]), which covered the constructive case. The destructive cases (i.e., those where the goal is to keep a particular candidate from winning) are due to Conitzer et al. [14]. Control alone has been extensively studied in many papers. The seminal control paper is by Bartholdi et al. [3] and covered the constructive case. The destructive cases were first studied by Hemaspaandra et al. [42]. There has been quite a bit of work on finding systems for which conducting various types of manipulation is hard, or for which conducting most types of control attacks is hard, see, e.g., [19, 20, 27, 43, 51, 52, 57] or the surveys [15, 23, 28, 30].

In the present paper, we will see that who goes first, the chair or the manipulators, is important in determining what complexity upper bounds apply. Order has also been seen to be important in the study of so-called online control attacks [46, 47], and of online manipulation attacks [44]. However, the papers just mentioned are separately about control, and about manipulation. In contrast we are mostly interested in when both are occurring, and especially when the two attacks are in conflict with each other.

Recent papers by Yin et al. [71] and Elkind et al. [17] each study control scenarios in settings where there is an “attacker” seeking to modify the election to attain his or her preferred outcome and a “defender” whose goal is to preserve the original outcome. However, this attacker–defender setting is essentially a competitive version of control in the case of Yin et al. and of bribery in the case of Elkind et al., and both cases are quite different from our model of control-plus-manipulation.

The present paper also looks at how revoting affects the complexity of elections that involve both control and manipulation. It is important to mention that, for the case of just manipulation, Narodytska and Walsh [54, 55] (see also [35]) have recently discussed revoting, and give an example that shows that revoting can sometimes be a valuable tool for the manipulator.

3 Preliminaries

An election system (a.k.a. a social choice correspondence) maps from a finite candidate set C and a finite vote collection V to a set, $W \subseteq C$, called the winner(s) [62]. Candidates each have a corresponding name, and these names play an important role in some of our results. Voters come without names, and the votes are input as a list, i.e., as ballots. For approval elections, each ballot is a length- $\|C\|$ 0-1 vector indicating whether each candidate is disapproved or

approved. The candidate getting the most approvals is the winner (or winners if candidates tie for most). For all other systems we discuss, each ballot is a tie-free linear ordering of the candidates. For plurality elections, each voter gives one point to his or her top choice and zero to the rest. For veto elections, each voter gives zero points to his or her bottom choice and one to the rest. For Borda elections, each voter gives zero points to his or her bottom choice, one point to his or her next to bottom choice, and so on through giving $\|C\| - 1$ points to his or her top choice. In the three systems just mentioned, the winner is the candidate(s) who receives the most total points. We also consider the Condorcet election system. Bartholdi et al. [3] recast the notion of a Condorcet winner [12] into an election system, and used it as one of their focus cases in their seminal control study. A candidate p is a Condorcet winner exactly if for each other candidate b it holds that strictly more than half the votes cast prefer p to b . Unlike the systems mentioned earlier in this paragraph, Condorcet elections on some inputs may have no winners.

An election system \mathcal{E} is said to have a polynomial-time winner problem if there is a polynomial-time algorithm that on input C, V , and $p \in C$, determines whether p is a winner under \mathcal{E} of the election over C with the votes being V .

We assume the reader is aware of the NP, coNP^{NP}, NP^{NP}, and coNP^{NP^{NP} levels of the polynomial hierarchy (the “exponentiation” notation denotes oracle class, informally put, having unit-cost access to a set of one’s choice from the given class) [53,63]. DP is the class of languages that are the difference of two languages in NP [56]. We assume that the reader is familiar with many-one reductions (which here always means polynomial-time many-one reductions). As is standard, we use \leq_m^p to denote many-one reductions. There are far fewer completeness results for levels of the hierarchy beyond NP, such as the abovementioned ones, than there are for NP; a collection of and discussion of such results can be found in Schaefer and Umans [60,61]. Completeness and hardness here are always with respect to many-one reductions.}

For proofs of hardness for higher levels of the polynomial hierarchy we 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 ; QBF₂ is NP^{NP}-hard, $\widetilde{\text{QBF}}_2$ is coNP^{NP}-hard, and $\widetilde{\text{QBF}}_3$ is coNP^{NP^{NP}-hard [64,67]. 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.}

Our hardness results are worst-case results. However, it is known that if there exists even one set that is hard for NP (and note that all sets hard for coNP^{NP}, NP^{NP}, or coNP^{NP^{NP} are hard for NP) and has a (deterministic) heuristic algorithm whose asymptotic error rate is subexponential, then the polynomial hierarchy collapses. See Hemaspaandra and Williams [48] for a discussion of that, and an attempt to reconcile that with the fact that in practice heuristics often do seem to do well, including for some cases related to elections, see, e.g., Walsh [66].}

The reader may wonder *why* we seek to obtain completeness results for these problems, given that the completeness results are for quite “hard” classes such as NP^{NP}, coDP, etc.—classes that are known to contain NP, which itself is already believed to be intractable. The answer is that completeness results pinpoint the nature and source of a problem’s hardness. For example, NP^{NP} is known to capture the power of computation that can be characterized by the quantifier pattern $\exists\forall$ (see [64,67]) and coDP is capturing the power of sets that can be described as the union of an NP set and a coNP set. By obtaining a completeness result for a problem, we clearly frame the problem’s nature. Also, although the just-mentioned classes each are included in P if and only if NP is included in P, that kind of stand-or-fall-

together result is not currently known to hold regarding certain other modes of computation. For example, it is an open question whether $NP \subseteq R$ (randomized polynomial time [38]) implies that $coNP \subseteq R$. And so proving that a set is $coNP$ -hard rather than merely NP -hard does provide a potentially heightened level of evidence that the set is not likely to be in R . Similarly, regarding UP (unambiguous polynomial time [65]), it is not known that $NP \subseteq UP$ implies $coNP \subseteq UP$. Thus a $coNP$ -hardness lower bound may be an improvement over an NP -hardness lower bound in terms of providing evidence that a given problem is not solvable by unambiguous nondeterminism. We refer the reader to the work of Hemaspaandra et al. [41] for more discussion—there in the context of the issue of NP -hardness versus hardness for so-called parallelized access to NP —of such cross-model issues. In addition to the above, we mention that a particularly interesting consequence of proving completeness for a quite “hard” class such as NP^{NP} is that it changes the go-to set of tools that one might use to heuristically approach the set. For example, if a set is NP -complete, SAT solvers may well be an excellent approach; but in contrast if a set is known to be NP^{NP} -complete, SAT solvers likely would not be a productive approach, and one might instead consider using ASP (answer set programming) solvers [16]. Indeed, an ASP-solver approach has recently been used on NP^{NP} -complete control problems in voting [36].

3.1 Types of electoral control

We now briefly define all standard control types. For a more formal description we refer the reader to the detailed definitions given in Faliszewski et al. [27]. Given as input an election, (C, V) , a distinguished (preferred or despised) candidate $p \in C$, and an integer $k \geq 0$, the constructive (respectively, destructive) control by deleting voters—for short CCDV (respectively, DCDV)—problem for an election system \mathcal{E} asks whether there is some choice of at most k voters such that if they are removed, p is a winner (respectively, is not a winner) of the given election under \mathcal{E} . We are in the so-called nonunique-winner model, and so we ask about making p “a winner” rather than “the one and only winner,” which is the so-called unique-winner model. Each of those problems has an adding voters (AV) analogue, in which one has a collection of registered voters, and has a collection of “unregistered” voters, and the question is whether there is some choice of at most k voters from the collection of unregistered voters such that if they are added, the goal is met. These types of control are motivated by issues ranging from voter suppression to targeted phone calls to get-out-the-vote drives. There are the natural analogous types for adding and deleting candidates, AC and DC (note: in the destructive control by deleting candidates case—DCDC—deleting p is not allowed [42]).

The partition types are called runoff partition of candidates (RPC), partition of candidates (PC), and partition of voters (PV). In each of the three partition control types, the input is (C, V) and $p \in C$, and a two-stage election is performed. In RPC, the constructive (destructive) question is whether there exists a partition of C into C_1 and C_2 such that two initial elections are held, (C_1, V) and (C_2, V) with winner sets W_1 and W_2 , respectively, and p is (is not) a winner of the runoff election $(W_1 \cup W_2, V)$. (Though we write “ V ” for the voter set in each subselection, that implicitly means V masked down just to the candidates at hand in the subselection; the analogous issue holds regarding the DC case; and in the AC case, the voters’ preferences are over the set of all registered and unregistered candidates and are also similarly masked down when called upon.) Here, there are two models for what “survive” means. In the ties eliminate (TE) model, to move forward one must uniquely win a first-round election; in the ties promote (TP) model, it suffices to be a winner of a first-round

election. The PC case is similar, but the winners of the election (C_1, V) move on to a runoff with all the candidates in C_2 .¹ In PV, we instead consider a partition of the collection of voters V into V_1 and V_2 where the runoff consists of the candidates that survive at least one of the elections (C, V_1) and (C, V_2) .

3.2 Manipulation

As to manipulation, the constructive (destructive) unweighted coalitional manipulation CUCM (DUCM) problem under election system \mathcal{E} has as input (C, V) , $p \in C$, and a collection of manipulator voters (coming in as blank slates, e.g., input as a string 1^j to indicate there are j of them), and the question is whether there is some way of setting the votes of the manipulative coalition so that p is (is not) a winner of the resulting election under system \mathcal{E} with those votes and the nonmanipulative votes both being cast.

3.3 Control-plus-manipulation

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 thus to come in as blank slates. We mention that for AV, it is legal to have manipulators among the registered and/or the unregistered voters. 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 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 (and a given election system \mathcal{E}), the CF case (viewed as a set, i.e., the set of all inputs where the chair is successful) is a subset of the MF case (similarly viewed as a set), since if there exists a control action such that for all manipulations the chair is successful, then the chair is successful with this same control action when the manipulators go first.

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 $k \in \mathbb{N}$.

¹ Work by Hemaspaandra, Hemaspaandra, and Menton has shown that in the nonunique winner model some of the seemingly different standard control problems in fact are identical, as to when the given type of control can be successfully accomplished. In particular, Hemaspaandra et al. [40] have proven that for each election system \mathcal{E} , and for each given election (C, V) and each candidate $p \in C$, p can be made a nonwinner by partitioning of candidates in the TE model (respectively, the TP model) if and only if p can be made a nonwinner by runoff partitioning of candidates in the TE model (respectively, the TP model) [40].

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 $(C, (V \cup W) - V')$ the election with candidate set C and the votes² of $(V \cup W) - V'$, 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} ?

The formal definitions of the other cases are analogous, keeping in mind that for the adding voters (AV) cases it is legal to have manipulators among the registered and/or the unregistered voters.

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 $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix} \\ \begin{bmatrix} PC \\ RPC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix} \end{bmatrix}$ refers to $2(2*2 + 3*2) = 20$

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 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 for our paper. 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 (i.e., the problems have the same set of yes instances). This follows from the definitions.³ Thus this paper is 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

² In a slight abuse of notation, we here are using the W in $(C, (V \cup W) - V')$ to denote the votes that W —which came in as a collection of blank slates—was assigned by the quantification in the just-prior statement. The same holds also in the next two parts of this definition.

³ To be clear, we now give in full the definitions of \mathcal{E} -DCDV-CF-ManipulatorFocus and \mathcal{E} -DCDV-MF-ManipulatorFocus (in particular, in the nonunique-winner model that is the model of this paper). In each case, the definition is what one would naturally expect, namely, for the given problem and for the

the same as studying the complement of the *destructive* (respectively, *constructive*) version of the same problem in the model of this paper, that is, from the perspective of the chair. For example, the sets \mathcal{E} -DCDV-CF-ManipulatorFocus and $\overline{\mathcal{E}}$ -CCDV-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 this paper except in the previous sentence, since in this paper 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 (i.e., each other’s complement on syntactically legal inputs). That is, we are not claiming, for example that it holds that the sets \mathcal{E} -DCDV-CF-ManipulatorFocus and $\overline{\mathcal{E}}$ -CCDV-CF-ManipulatorFocus are the same on all syntactically legal inputs; that is not in general true. (However, we mention in passing that the discussion in footnote 5 of Hemaspaandra et al. [42] presents a partial connection between constructive and destructive cases.)

4 Results

4.1 Inheritance

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} -CCDV \leq_m^p \mathcal{E} -M+CCDV and \mathcal{E} -CCDV \leq_m^p \mathcal{E} -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 have an explicit 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} -CUCM \leq_m^p $\overline{\mathcal{E}}$ -DCAC-CF and \mathcal{E} -DUCM \leq_m^p $\overline{\mathcal{E}}$ -CCDV-MF. For the cooperative cases there is no “flipping.” For each 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} -CUCM \leq_m^p \mathcal{E} -M+CCAC and \mathcal{E} -DUCM \leq_m^p \mathcal{E} -M+DCAC.

Footnote 3 continued

given ordering of the chair and the manipulators, the definitions are capturing whether the manipulators can assure that the candidate p can be made, or blocked from being, a winner.

Name:	\mathcal{E} -CCDV-CF-ManipulatorFocus/ $\overline{\mathcal{E}}$ -CCDV-MF-ManipulatorFocus
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 (CF-ManipulatorFocus):	Does it hold that for each subcollection $V' \subseteq (V \cup W)$ satisfying $\ V'\ \leq k$, there exists a way for the manipulators to set their votes such that p is a winner of $(C, (V \cup W) - V')$ under election system \mathcal{E} ?
Question (MF-ManipulatorFocus):	Does there exist a way that the manipulators can set their votes so that for each subcollection $V' \subseteq (V \cup W)$ satisfying that $\ V'\ \leq k$, it holds that p is a winner of $(C, (V \cup W) - V')$ under election system \mathcal{E} ?

Table 1 Upper bounds

Problem	CF	CF-revoting	MF	MF-revoting
$\mathcal{E} - \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix}$	NP^{NP} (coDP for DV)	N/A	$coNP^{NP}$	N/A
$\mathcal{E} - \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix}$	NP^{NP}	NP^{NP}	$coNP^{NP}$	$coNP^{NP}$ (TE) $coNP^{NP^{NP}}$ (TP)

To make easily available within this table what the table's two “Problem” entries are speaking of, we mention that if the notations were expanded for clarity, then

(a) $\mathcal{E} - \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix}$ would be $\mathcal{E} - \begin{bmatrix} \text{Constructive} \\ \text{Destructive} \end{bmatrix} \text{Control by} \begin{bmatrix} \text{Adding} \\ \text{Deleting} \end{bmatrix} \begin{bmatrix} \text{Candidates} \\ \text{Voters} \end{bmatrix}$, and

(b) $\mathcal{E} - \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix}$ would be

$$\mathcal{E} - \begin{bmatrix} \text{Constructive} \\ \text{Destructive} \end{bmatrix} \text{Control by} \begin{bmatrix} \text{Partition of Candidates} \\ \text{Runoff Partition of Candidates} \\ \text{Partition of Voters} \end{bmatrix} - \begin{bmatrix} \text{Ties Eliminate} \\ \text{Ties Promote} \end{bmatrix}.$$

Also, in the cells where the table states N/A (not applicable), the nonrevoting bounds just to the left of the cell technically still hold; we write N/A to be clear that revoting cannot even take place in nonpartition cases, since there is no second round

4.2 General upper bounds and matching lower bounds

For election systems with polynomial-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 4.1 *For each election system \mathcal{E} having a polynomial-time winner problem, the bounds of Table 1 hold.*

The table's general upper bounds clearly follow from the structure of the problem. For the nonrevoting cases upper bounds of NP^{NP} for the “CF” cases and $coNP^{NP}$ for the “MF” cases clearly hold, and these are the bounds that were stated in early versions [31,32,34]. However, the “CF” cases of control by deleting voters exhibit a different behavior, since the chair can delete the manipulator. In fact, for every election system \mathcal{E} with a polynomial-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).

Theorem 4.2 *For every election system \mathcal{E} with a polynomial-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 cannot 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

The bounds presented in Table 1 are very high. It is natural to ask: Can they be improved by some cleverer approach? Or are there systems with polynomial-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 (and sometimes for all) cases.

Theorem 4.3 1. *For each of the eight problems on the top line of Table 1, and each of the columns on that line, there exists an election system \mathcal{E} , which has a polynomial-time winner problem, for which the named problem is complete for the named complexity class.*

2. *For each of CCPV-TP and CCPV-TE, and each of the CF, CF-revoting, and MF columns of Table 1, and each of the columns on that line, there exists an election system \mathcal{E} , which has a polynomial-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 polynomial-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 polynomial-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.⁴ 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 the coming section, 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 4.3, which illustrates the general arguments used in the proof of this theorem. The proofs of the other cases use similar approaches and can be found in the technical-report version of this paper [33].

Theorem 4.4 *There exists an election system, \mathcal{E} , with a polynomial-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

⁴ The third-level case has to overcome the specific, and as far as we know now, worry that in the second round, the first-round vote of the manipulators is no longer available. Yet in a “ $\forall \exists \forall$ ” context (which is the quantifier structure that models $\text{coNP}^{\text{NP}^{\text{NP}}}$), a particular existential choice has to handle only a particular value of the first \forall . So to make the construction work, we need to in some sense have the first-round votes, which are no longer available, still cast a clear and usable shadow forward into the second round, at least in certain cases in the image of the reduction. We achieve this, in particular by shaping the election system itself carefully to help realize this unusual effect. Otherwise, we would not be capturing the right quantifier structure.

encodes⁵ 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 polynomial-time winner problem, and by Theorem 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 QBF₂. 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 $k = 2\ell$.⁶

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

4.3 Specific systems

This section looks at control in the presence of manipulators, for the case of certain concrete systems.

Section 4.3.1 is in the model we have been using so far in this paper, namely, the voters are not weighted.

Section 4.3.2 looks at a different case, namely, the case of weighted voters, a model that better captures cases such as stockholder elections (in which each voter's vote is weighted by the number of shares the voter owns of that stock) and the Electoral College that is used in the United States to select the president. Starting in 1866, Sweden even started weighting voter's votes based on wealth, with some voters having their votes amplified by factors of up to 5000, though the system was discontinued [13, chapter 14]. Weighted manipulation (not in the presence of control) was introduced by Conitzer et al. [14], and weighted control (not in the presence of manipulators) was extensively studied by Faliszewski et al. [26].

⁵ We are using here the fact that a boolean formula can in a straightforward way be represented as a string and encoded in the name of candidate c_0 .

⁶ We set $k = 2\ell$ instead of the obvious choice of ℓ since then the same proof can be used for the similar cases that appear in the technical report [33], and this also nicely handles the case of “control by unlimited adding of candidates.”

4.3.1 Unweighted voters

Plurality is certainly one of the most important 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 [3,42]. For each of these election systems, we will show that all the “M+,” “CF,” and “MF” cases whose control type is classified as P (i.e., as a “V” or “I” in the original notation of the 2007 table that we are about to mention) in the with-no-manipulators table of complexities for that election system in [42, p. 258] are in P for all our cooperative and competitive cases.

Theorem 4.5 *Each problem contained in*

$$\begin{aligned}
 & - \begin{bmatrix} \text{approval} \\ \text{Condorcet} \\ \text{plurality} \end{bmatrix} - M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix} \\ \begin{bmatrix} PC \\ RPC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix} \end{bmatrix}, \\
 & - \begin{bmatrix} \text{approval} \\ \text{Condorcet} \\ \text{plurality} \end{bmatrix} - \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix} \\ \begin{bmatrix} PC \\ RPC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix} \end{bmatrix} - \text{CF, or} \\
 & - \begin{bmatrix} \text{approval} \\ \text{Condorcet} \\ \text{plurality} \end{bmatrix} - \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} \begin{bmatrix} A \\ D \end{bmatrix} \begin{bmatrix} C \\ V \end{bmatrix} \\ \begin{bmatrix} PC \\ RPC \\ PV \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix} \end{bmatrix} - \text{MF,}
 \end{aligned}$$

whose corresponding control type is in P in Table 1 of page 258 of [42],⁷ 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

⁷ So that the reader has it easily at hand while viewing Theorem 4.5, we include here the central content of that table (each of whose results is from Bartholdi et al. [3] or Hemaspaandra et al. [42]):

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 TP: NPC	TE: NPC TP: NPC	TE: P TE: P	TE: P TE: P	TE: P TP: P	TE: P TP: P
Runoff partitioning candidates	TE: NPC TP: NPC	TE: NPC TP: NPC	TE: P TE: P	TE: P TE: P	TE: P TP: P	TE: P TP: P
Partitioning voters	TE: P TP: NPC	TE: P TP: NPC	TE: NPC TE: NPC	TE: P TE: P	TE: NPC TP: NPC	TE: P TP: P

We mention that our theorem’s claim holds despite the fact that this table of [42] is about the unique-winner case, and despite the fact that the AC/Adding Candidates case of the table refers to so-called unlimited adding and (as is now standard) we use AC/Adding Candidates to refer to (limited) adding.

Bartholdi et al. [3] and Hemaspaandra et al. [42] are both for the unique-winner model. The following observation, Observation 4.6, states that the results from those papers mentioned in [42, Table 1]—which is essentially reproduced in footnote 7—carry over to the nonunique-winner model (some of the cases were previously noted in Faliszewski et al. [25] and Hemaspaandra et al. [47]).

Observation 4.6 *The complexities of each of the standard control problems shown in Bartholdi et al. [3] and Hemaspaandra et al. [42] for the unique-winner model hold also for the nonunique-winner model.*

One might wonder why Theorem 4.5 speaks only of showing that the P boxes from the table remain P cases in the corresponding case within our theorem. The reason is that, due to our inheritance results, for the NP-completeness cases of the table getting a P result in our corresponding case is not possible (unless P = NP). Regarding this claim, one might worry as to the two NP-completeness cases regarding the “Adding Candidates” table line, since (see footnote 7) the corresponding case in our theorem is about limited adding of candidates yet the table is about the unlimited adding of candidates model. However, that in fact is not a problem, since if the limited adding of candidates case is in P for a given election system, then the unlimited adding of candidates case for that same election system is also in P.

As an illustration as to proving Theorem 4.5, we present the proof of plurality-M+CCPV-TE $\in P$ here. The proofs of the remaining cases of Theorem 4.5 can be found in Appendix A.

Proof In some cases, there are possible partitions by the chair such that if one simply has all the manipulators vote for p we don’t achieve our goal yet there is a different action by the manipulators under that same partition that would achieve our goal. For example, if the candidates are a , b , and p , and the chair partitions the voters such that one of the subselections contains a voter voting $p > a > b$ (the “ $>$ ” denotes strict preference), and the other subselection 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 subselection and neither goes through to the second round. In particular, 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 partition such that p wins 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 Hemaspaandra et al. [42] (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. \square

We now will seem to change directions, and will briefly study “standard” control problems, i.e., ones not in the presence of manipulators. However, we do so in service of the goals of this paper. The results we will obtain below will be crucially used to prove parts of Theorem 4.5, though the proofs that do so are found not in the body of the paper but in four proofs in Appendix A that draw on the results below.

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 vote set 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 vote set 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.⁸ It is easy to see that approval and Condorcet elections satisfy both WARP and unique-WARP [42].

⁸ We 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.

We mention in passing that WARP itself is very closely connected to immunity to destructive control by deleting candidates (DCDC); in particular, they are the same. To see this, we need to discuss a notion from the literature: immunity. An election system is said to be immune to destructive control by deleting candidates if for every election instance (C, V) and every candidate $c \in C$ it holds that: If c is a winner in that election instance, then for every candidate set C' satisfying $\{c\} \subseteq C' \subseteq C$ it holds that c is a winner in the election with candidate set C and vote set V (masked down to the candidates in C'). This notion, destructive control by deleting candidates, is due to the seminal control paper of Bartholdi et al. [3], except their paper is in the unique-winner model and our paper is in the nonunique-winner model. Yang [69] has observed that WARP implies, in the nonunique-winner model, immunity to destructive control by deleting candidates. We here add the observation that the converse also holds, since the definitions of the two concepts are in fact the same. Thus the following holds.

An election system \mathcal{E} satisfies WARP if and only if \mathcal{E} is immune to DCDC (destructive control by deleting candidates).

Again, like all the results in this paper, the above if and only if statement is with respect to the nonunique-winner model. We mention, for context, that in the unique-winner model (which is not the model we are using in this paper), the analogous result holds if one looks instead at unique-WARP, namely, we have the following result.

An election system \mathcal{E} satisfies unique-WARP if and only if \mathcal{E} is, in the unique-winner model, immune to DCDC (destructive control by deleting candidates).

This result’s “only if” direction is stated in [42] and this result’s “if” direction clearly also holds, again as the definitions of the two notions in fact are the same. Finally, Yang [69] (respectively, Hemaspaandra et al. [42]) states that in the nonunique-winner model (respectively, unique-winner model), that WARP (respectively, unique-WARP) implies immunity to constructive control by adding candidates. We observe that the converse directions for each of those claims hold, for the same reasons as mentioned above for the DCDC cases, thus yielding two additional if and only if results.

Though as mentioned above these results are rather crucially used as tools within our proofs about control in the presence of manipulators, we feel they are of interest in their own right. Let us take as an example the coming Theorem 4.7, which loosely put says that for every election system satisfying unique-WARP, and for each instance of CCRPC-TE, it holds that the partition whose parts are “all candidates other than p ” and “ p ” will cause p to win if and only if the chair has *any* partition choice that will cause p to win.

The result is interesting because it is directly attacking what is the heart of the complexity of partition problems: combinatorial explosion, i.e., the fact that there are an enormous number of partitions and the chair must determine whether any one of them makes p a winner. This is precisely why such problems so often turn out to be NP-hard. However, Theorem 4.7 says that for systems obeying the unique-WARP axiom, that potential complexity is completely side-stepped: There is a single partition that is the only one that needs to be examined. This immediately shows that the control type is of polynomial-time complexity for systems satisfying unique-WARP.

Viewed more broadly, by linking the complexity of control to social-choice properties, this part of our work is trying to take a step away from analyzing systems one at a time, and is trying to more generally determine what it is that can yield computational simplicity. Work having that goal is most typically done by studying the class of so-called scoring systems, each of which is defined by a so-called scoring vector, and finding some simple property of the scoring vector that determines the complexity of various manipulative attacks on elections. To give as an example just one family of such results, we mention the line doing this regarding manipulation of elections in the general case and in the so-called single-peaked case [7, 14, 29, 39, 59]. However, that work focuses on the direct definitions of the election systems, and our work in contrast is focusing on how possession of an axiomatic property can itself force simplicity. We mention in passing that Yang and Guo [70], in the context of parameterized complexity, presented fixed-parameter tractability results for control by adding/deleting voters for systems satisfying the Smith-IIA property.

Let us now turn to our results of this type.

We in the following and elsewhere will say that “control is possible” as a shorthand way of saying that there exists a control action of the given type that achieves the goal (e.g., if it is a constructive control type, then making p be a winner, and if it is a destructive control type, making p not be a winner).

Theorem 4.7 *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 set 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 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 set 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 candidates or one candidate 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? Without loss of generality, 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. \square

Corollary 4.8 *For every election system \mathcal{E} that satisfies unique-WARP and has a polynomial-time winner problem, \mathcal{E} -CCRPC-TE is in P.*

Theorem 4.7 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 4.7 for the CCPC-TE case. Our analogue, however, applies to election systems that satisfy both WARP and unique-WARP.⁹

Theorem 4.9 *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.

⁹ Is 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 [6]. 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 4.9 would certainly remain true if in it one were to replace the phrase "both WARP and unique-WARP" with simply "WARP".

1. If under our set 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 set 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 that r participates in. 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. \square

Corollary 4.10 For every election system \mathcal{E} that satisfies both WARP and unique-WARP and has a polynomial-time winner problem, \mathcal{E} -CCPC-TE is in P.

Corollary 4.11 For every election system \mathcal{E} satisfying both WARP and unique-WARP, \mathcal{E} -CCPC-TE = \mathcal{E} -CCRPC-TE.

What equality means, in Corollary 4.11 and elsewhere when used on election problems, is as follows, using the above as an example. Each of \mathcal{E} -CCPC-TE and \mathcal{E} -CCRPC-TE is a set, and by speaking of them being equal, we mean the sets are equal, i.e., no string is in exactly one of the two sets.

4.3.2 Weighted voters

We now give results for veto and Borda, including, for the latter, an interesting increase in complexity.

In weighted elections every voter has a positive integer weight, and a voter with weight w counts as w 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 [14], destructive coalitional manipulation is in P [14], and CCAV and CCDV are both in P [26]. The following result shows that for this system $CC\begin{bmatrix} A \\ D \end{bmatrix}V\begin{bmatrix} CF \\ MF \end{bmatrix}$ are all in P—not NP-complete.

Theorem 4.12 For 3-candidate weighted veto elections, the following hold.

1. $M+CC\begin{bmatrix} A \\ D \end{bmatrix}V$ are both NP-complete.
2. $CC\begin{bmatrix} A \\ D \end{bmatrix}V\begin{bmatrix} CF \\ MF \end{bmatrix}$ are each in P.

Proof The first case follows directly from the fact that constructive coalitional manipulation for weighted veto elections is NP-complete [14] and the inheritance observations from Sect. 4.1 (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 Faliszewski et al. [26]. \square

An increase in complexity is shown by 3-candidate weighted Borda elections. The known results for this system are that constructive coalitional manipulation is NP-complete [14], destructive coalitional manipulation is in P [14], and CCAV and CCDV are both NP-complete [26] 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 $\text{NP} \cap \text{coNP}$.

Theorem 4.13 *For 3-candidate weighted Borda elections, the following hold.*

1. $\text{M+CC} \left[\begin{matrix} \text{A} \\ \text{D} \end{matrix} \right] \text{V}$ are both NP-complete.
2. $\text{CC} \left[\begin{matrix} \text{AV-CF} \\ \text{DV-} \left[\begin{matrix} \text{CF} \\ \text{MF} \end{matrix} \right] \end{matrix} \right]$ are each NP-hard.
3. CCAV-MF is NP-hard and coNP-hard.
4. $\text{CC} \left[\begin{matrix} \text{A} \\ \text{D} \end{matrix} \right] \text{V-CF}$ are both NP-complete.

Proof The first case follows directly from the fact that manipulation is NP-complete [14] and the inheritance observations from Sect. 4.1 (as the relevant result there holds even for the weighted case).

The remaining NP-hardness results follow from the NP-completeness of CCAV and CCDV and the inheritance observations from Sect. 4.1.

Let a denote one of the candidates other than p (either one, but which one it is will now be fixed for the rest of this proof), and let b denote the remaining candidate (the one that is neither a nor p). 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. Partition is the set of all nonempty sequences of positive integers k_1, \dots, k_t such that the sequence sums to $2K$ for some integer K and there exists a partition (i.e., a subsequence of k_1, \dots, k_t) that sums to K ; this set is well-known to be NP-complete (see [50]; from the NP-completeness of its version of Partition, the above version is easily seen to be NP-complete, since empty lists are not a substantive case and 0s do not change the total of a sum). 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 $3K$, respectively. 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, i.e., there is no subsequence of k_1, \dots, k_t that sums to K . 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

5 Conclusions and open directions

We have established general inheritance results and complexity upper bounds for control in the presence of manipulators, for both cooperative and competitive settings. For the upper bounds we provided matching lower bounds, but also showed that for many natural systems the complexity is far lower than the general upper bounds.

Many open directions remain. For example, regarding 3-candidate weighted Borda elections, we have shown that CCAV-MF is NP-hard and coNP-hard, and although our upper-bound theorem is not explicitly about weighted cases, clearly this problem, for exactly the same reason as in our upper-bound theorem, is in coNP^{NP}. But precisely where within that range does it fall? Also, what happens for real-world election systems that themselves are complex to manipulate and/or control, such as Llull, Copeland, fallback, sincere-preference approval, and Schulze elections? Do some of these systems themselves provide natural systems that might for our competitive cases be complete for some of the high complexity classes given in Table 1?

Acknowledgements Earlier versions of this paper appeared in IJCAI-2013 [31] and COMSOC-2014 [34]. We are extremely grateful to the anonymous conference and journal referees, whose comments and suggestions much improved this paper. This work was supported in part by NSF Graduate Research Fellowship DGE-1102937 and by NSF Grants CCF-0915792, CCF-1101452, CCF-1101479, DUE-1819546, and CCF-2006496.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

A Specific systems

In some of the proofs in this section, 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.

A.1 Plurality

In this subsection we prove the remaining plurality cases of Theorem 4.5.

Theorem A.1 *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 polynomial-time algorithm for control from Bartholdi et al. [3] (for the constructive cases) or from Hemaspaandra et al. [42] (for the destructive cases), modified in the obvious way for the nonunique-winner case (see Observation 4.6). \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 $score_{\{p,r\}}(r) + k > score_{\{p,r\}}(p)$. Recall that the manipulators come in without votes (i.e., as blank slates), and so $score_{\{p,r\}}(r)$ and $score_{\{p,r\}}(p)$ denote the scores of r and p in the election comprised of the candidates p and r , and the nonmanipulators in V .

Lemma A.2 *If there exists a partition of voters such that p is an overall winner in the “TE” model when all manipulators vote for the same highest-scoring rival r , then there exists a partition such that p is always (i.e., regardless of the votes of the manipulators) an overall winner.*

Proof Given an election (C, V) where V contains k manipulators, a candidate $p \in C$, and a highest-scoring rival r , 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 . 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 . Since we can exchange a manipulator in V_2 that votes for r with a nonmanipulator in V_1 that votes for r without changing the outcome of the election, we can assume that there are no manipulators in V_2 (i.e., $k_2 = 0$) or that there are no nonmanipulators voting for r in V_1 (i.e., $\ell_1 = 0$).

We first consider the case where $k_2 = 0$. 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} , and for every nonrival $c \neq p$ all the nonmanipulator votes for c . 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 $score_{V_1}(r) + k_1 = \ell_1 + k_1 < score_{V_1}(p) = score_{\widehat{V}_1}(p)$. Since there are k_1 manipulators in \widehat{V}_1 , it suffices to show that for all $c \neq p$, $score_{\widehat{V}_1}(c) \leq \ell_1$. This is immediate from the construction: $score_{\widehat{V}_1}(r) = \ell_1$, for every nonrival $c \neq p$, $score_{\widehat{V}_1}(c) = 0$, and for every rival $\widehat{r} \neq r$, $score(\widehat{r}) \leq score(r) = \ell_1 + \ell_2$, and so $score_{\widehat{V}_1}(\widehat{r}) \leq \ell_1$.

So the only way in which p can be precluded from winning the runoff is if there is a rival \widehat{r} of p such that \widehat{r} is the unique winner of (C, \widehat{V}_2) (recall that there are no manipulators in \widehat{V}_2). Since $score_{\widehat{V}_2}(r) = \ell_2$ and $score_{\widehat{V}_2}(\widehat{r}) \leq \ell_2$ for every rival \widehat{r} , it then follows that r is the unique winner of (C, \widehat{V}_2) . But then $score(\widehat{r}) < \ell_2$ for every rival $\widehat{r} \neq r$ and $score(c) < \ell_2$

for every nonrival $c \neq p$, and $score_{V_2}(p) = score_{\widehat{V}_2}(p) < \ell_2$. It follows that r is the unique winner of (C, V_2) . Then p is never an overall winner of (C, V) using the partition (V_1, V_2) , which contradicts our assumption. It follows that p is always a winner of (C, V) using the partition $(\widehat{V}_1, \widehat{V}_2)$.

It remains to show the claim for $k_2 > 0$ and $\ell_1 = 0$. We will again construct a new partition $(\widehat{V}_1, \widehat{V}_2)$ that will work regardless of how the manipulators vote. Let \widehat{V}_1 consist of $score_{V_1}(p)$ nonmanipulator votes for p and k_1 manipulators. Then \widehat{V}_2 consists of $score_{V_2}(p)$ nonmanipulator votes for p , all nonmanipulator votes for $c, c \neq p$, and k_2 manipulators.

Since p is the unique winner of (C, V_1) when all manipulators vote for r , it follows that $k_1 < score_{V_1}(p) = score_{\widehat{V}_1}(p)$, and so 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 is manipulation and a rival \widehat{r} of p such that \widehat{r} is the unique winner of (C, \widehat{V}_2) . Then $score(\widehat{r}) + k_2 = score_{\widehat{V}_2}(\widehat{r}) > score_{\widehat{V}_2}(c)$ for all $c \neq \widehat{r}$. In particular, $score(\widehat{r}) + k_2 > score(c)$ for every nonrival $c \neq p$ and $score(\widehat{r}) + k_2 > score_{V_2}(p)$. Since $score(\widehat{r}) \leq score(r) = \ell_2$, it follows that r is the unique winner of (C, V_2) when all manipulators vote for r . This contradicts our assumption. It follows that p is always a winner of (C, V) using the partition $(\widehat{V}_1, \widehat{V}_2)$. \square

Theorem A.3 *For plurality elections, 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 that 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 A.2 that control is possible if and only if there exists a partition such that p wins when all manipulators vote for r and rank the remaining candidates in lexicographic order. This can be determined by running the polynomial-time algorithm for plurality-CCPV-TE from [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6). \square

Theorem A.4 *It holds that plurality- CCPV-TE-CF = plurality- CCPV-TE-MF.*

Proof It immediately follows from the definition that $\text{plurality-CCPV-TE-CF} \subseteq \text{plurality-CCPV-TE-MF}$.

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 . By Lemma A.2 we know that then there exists a partition where p is always a winner, so “CF” control is possible. \square

Corollary A.5 *For plurality elections, CCPV-TE-MF is in P.*

Theorem A.6 *For plurality elections, 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 [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6).

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, i.e., we need to ensure there is a candidate $c \neq p$ such that c ties-or-beats p in both subselections or there are two distinct candidates $c, d \neq p$ such that c ties-or-beats p in the first subselection and d ties-or-beats p in the second subselection. 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 A.7 *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 A.7 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 et al. [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6).

Theorem A.8 *For plurality elections, DCPV-TE-CF is in P.*

We will now show that Lemma A.7 also implies that plurality-DCPV-TE-MF is in P.

Theorem A.9 *For plurality elections, 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 A.7 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 A.7 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

A.2 Condorcet

In this subsection we prove the Condorcet cases of Theorem 4.5.

Theorem A.10 *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 rank p first. The remaining candidates can be ranked in any order since this does not affect whether or not p is a Condorcet winner.

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 polynomial-time algorithm for control from Bartholdi et al. [3] (for the constructive cases) or from Hemaspaandra et al. [42] (for the destructive cases), modified in the obvious way for the nonunique-winner case (see Observation 4.6). \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 [42].

Theorem A.11 *For Condorcet elections, $M + \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \end{bmatrix}$ 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 (in the presence of manipulators) by partitioning of candidates and by runoff partitioning of candidates as follows.

The use we’re about to make of the notion of control being possible is as described immediately before Theorem 4.7. For the constructive cases we do the following. Since Condorcet elections satisfy both WARP and unique-WARP, we know from Theorems 4.7 and 4.9 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 (\widehat{C}, V) has no Condorcet winner, i.e., for each $c \in \widehat{C}$ there exists a $c' \in \widehat{C} - \{c\}$ 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 $\lceil k/2 \rceil$ 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 $\lceil k/2 \rceil$ 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 $\lceil k/2 \rceil$ 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 $\lceil k/2 \rceil$ 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 $c > b > a > \widehat{C} - \{a, b, c\}$, $\lceil k/3 \rceil$ 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 [42]. 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 A.12 *For Condorcet elections, $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \end{bmatrix} - \begin{bmatrix} CF \\ MF \end{bmatrix}$ 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 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 4.7 and 4.9 (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 [42]. 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 A.13 *For Condorcet elections, 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 (and ranking the remaining candidates in any order), 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 [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6). \square

Below we state a lemma analogous to Lemma A.7, but for Condorcet elections.

Lemma A.14 *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 A.14 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 [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6).

Theorem A.15 *For Condorcet elections, DCPV-CF is in P.*

A similar argument as in the proof of Theorem A.9 shows that Lemma A.14 above also implies that the corresponding “MF” case is also in P.

Theorem A.16 *For Condorcet elections, DCPV-MF is in P.*

A.3 Approval

In this subsection we prove the Approval cases of Theorem 4.5.

Theorem A.17 *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 polynomial-time algorithm for control from Bartholdi et al. [3] (for the constructive cases) or from Hemaspaandra et al. [42] (for the destructive cases), modified in the obvious way for the nonunique-winner case (see Observation 4.6). \square

Theorem A.18 For approval elections, $M+ \begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix}$ 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 (in the presence of manipulators) 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 4.7 and 4.9 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 A.19 For approval elections, $\begin{bmatrix} C \\ D \end{bmatrix} C \begin{bmatrix} PC \\ RPC \end{bmatrix} - \begin{bmatrix} TE \\ TP \end{bmatrix} - \begin{bmatrix} CF \\ MF \end{bmatrix}$ 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 (in the presence of manipulators) 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 4.7 and 4.9 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 A.20 For approval elections, $M+DCPV - \begin{bmatrix} TE \\ TP \end{bmatrix}$ are both in P.

Proof Given an election (C, V) containing k manipulators, 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. Let k denote the number of manipulators in V .

1. If there is a candidate, $c \neq p$ such that $score(p) \leq score(c) + k$, 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 [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6).

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 A.7, but for approval elections.

Lemma A.21 *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 A.7, 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 A.21 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 et al. [42], modified in the obvious way for the nonunique-winner case (see Observation 4.6).

Theorem A.22 *For approval elections, DCPV- $\begin{bmatrix} \text{TE} \\ \text{TP} \end{bmatrix}$ -CF are both in P.*

Lemma A.21 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 A.9.

Theorem A.23 *For approval elections, DCPV- $\begin{bmatrix} \text{TE} \\ \text{TP} \end{bmatrix}$ -MF are both in P.*

References

1. Bartholdi III, J., & Orlin, J. (1991). Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4), 341–354.
2. Bartholdi III, J., Tovey, C., & Trick, M. (1989). The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3), 227–241.
3. Bartholdi III, J., Tovey, C., & Trick, M. (1992). How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8–9), 27–40.
4. Baumeister, D., Roos, M., & Rothe, J. (2011). Computational complexity of two variants of the possible winner problem. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems* (pp. 853–860). International Foundation for Autonomous Agents and Multiagent Systems.
5. Baumeister, D., Roos, M., Rothe, J., Schend, L., & Xia, L. (2012). The possible winner problem with uncertain weights. In *Proceedings of the 20th European Conference on Artificial Intelligence* (pp. 133–138). IOS Press.
6. Baumeister, D., & Rothe, J. (2016). Preference aggregation by voting. In J. Rothe (Ed.), *Economics and computation: An introduction to algorithmic game theory, computational social choice, and fair division* (pp. 197–325). Berlin: Springer.
7. Brandt, F., Brill, M., Hemaspaandra, E., & Hemaspaandra, L. (2015). Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research*, 53, 439–496.
8. Brandt, F., Conitzer, V., & Endriss, U. (2013). Computational social choice. In G. Weiss (Ed.), *Multiagent systems* (2nd ed., pp. 213–284). Cambridge: MIT Press.
9. Chevaleyre, Y., Endriss, U., Lang, J., & Maudet, N. (2007). A short introduction to computational social choice. In *Proceedings of the 33rd International Conference on Current Trends in Theory and Practice of Computer Science*, Lecture Notes in Computer Science #4362 (pp. 51–69). Springer.
10. Chevaleyre, Y., Lang, J., Maudet, N., & Monnot, J. (2010). Possible winners when new candidates are added: The case of scoring rules. In *Proceedings of the 24th AAAI conference on artificial intelligence* (pp. 762–767). AAAI Press.
11. Chevaleyre, Y., Lang, J., Maudet, N., Monnot, J., & Xia, L. (2012). New candidates welcome! Possible winners with respect to the addition of new candidates. *Mathematical Social Sciences*, 64(1), 74–88.
12. Condorcet, J. *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.
13. Congleton, R. (2010). *Perfecting parliament: Constitutional reform, liberalism, and the rise of Western democracy*. Cambridge: Cambridge University Press.
14. Conitzer, V., Sandholm, T., & Lang, J. (2007). When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3), 14.
15. Conitzer, V., & Walsh, T. (2016). Barriers to manipulation in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, & A. Procaccia (Eds.), *Handbook of computational social choice* (pp. 127–145). Cambridge: Cambridge University Press.
16. Eiter, T., Gottlob, G., & Mannila, H. (1997). Disjunctive datalog. *ACM Transactions on Database Systems*, 22(3), 364–418.
17. Elkind, E., Gan, J., Obraztsova, S., Rabinovich, Z., & Voudouris, A. (2019). Protecting elections by recounting ballots. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (pp. 259–265).
18. Ephrati, E., & Rosenschein, J. (1997). A heuristic technique for multi-agent planning. *Annals of Mathematics and Artificial Intelligence*, 20(1–4), 13–67.
19. Erdélyi, G., Fellows, M., Rothe, J., & Schend, L. (2015). Control complexity in Bucklin and fallback voting: A theoretical analysis. *Journal of Computer and System Sciences*, 81(4), 632–660.
20. Erdélyi, G., Nowak, M., & Rothe, J. (2009). Sincere-strategy preference-based approval voting fully resists constructive control and broadly resists destructive control. *Mathematical Logic Quarterly*, 55(4), 425–443.
21. Fagin, R., Kumar, R., & Sivakumar, D. (2003). Efficient similarity search and classification via rank aggregation. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data* (pp. 301–312). ACM Press.
22. Faliszewski, P., Hemaspaandra, E., & Hemaspaandra, L. (2009). How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35, 485–532.
23. Faliszewski, P., Hemaspaandra, E., & Hemaspaandra, L. (2010). Using complexity to protect elections. *Communications of the ACM*, 53(11), 74–82.

24. Faliszewski, P., Hemaspaandra, E., & Hemaspaandra, L. (2011). Multimode control attacks on elections. *Journal of Artificial Intelligence Research*, 40, 305–351.
25. Faliszewski, P., Hemaspaandra, E., & Hemaspaandra, L. (2014). The complexity of manipulative attacks in nearly single-peaked electorates. *Journal of Artificial Intelligence Research*, 207, 69–99.
26. Faliszewski, P., Hemaspaandra, E., & Hemaspaandra, L. (2015). Weighted electoral control. *Journal of Artificial Intelligence Research*, 52, 507–542.
27. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2009). Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35, 275–341.
28. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2009). A richer understanding of the complexity of election systems. In S. Ravi & S. Shukla (Eds.), *Fundamental problems in computing: Essays in honor of Professor Daniel J. Rosenkrantz* (pp. 375–406). Berlin: Springer.
29. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2011). The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2), 89–107.
30. Faliszewski, P., & Rothe, J. (2016). Control and bribery in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, & A. Procaccia (Eds.), *Handbook of computational social choice* (pp. 146–168). Cambridge: Cambridge University Press.
31. Fitzsimmons, Z., Hemaspaandra, E., & Hemaspaandra, L. (2013). Control in the presence of manipulators: Cooperative and competitive cases. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence* (pp. 113–119). AAAI Press.
32. Fitzsimmons, Z., Hemaspaandra, E., & Hemaspaandra, L. (2013). *Control in the presence of manipulators: Cooperative and competitive cases*. Tech. Rep. [arXiv:1308.0544v1](https://arxiv.org/abs/1308.0544v1) [cs.GT], Computing Research Repository, [arXiv.org/corr/](https://arxiv.org/corr/).
33. Fitzsimmons, Z., Hemaspaandra, E., & Hemaspaandra, L. (2013). *Control in the presence of manipulators: Cooperative and competitive cases*. Tech. Rep. [arXiv:1308.0544](https://arxiv.org/abs/1308.0544) [cs.GT], Computing Research Repository, [arXiv.org/corr/](https://arxiv.org/corr/). Revised, June 2017.
34. Fitzsimmons, Z., Hemaspaandra, E., & Hemaspaandra, L. (2014). Control in the presence of manipulators: Cooperative and competitive cases. In *Proceedings (Workshop Notes) of the 5th International Workshop on Computational Social Choice*.
35. Fitzsimmons, Z., Hemaspaandra, E., & Hemaspaandra, L. (2016). Manipulation complexity of same-system runoff elections. *Annals of Mathematics and Artificial Intelligence*, 77(3–4), 159–189.
36. Fitzsimmons, Z., Hemaspaandra, E., Hoover, A., & Narváez, D. (2019). Very hard electoral control problems. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence* (pp. 1933–1940).
37. Ghosh, S., Mundhe, M., Hernandez, K., & Sen, S. (1999). Voting for movies: The anatomy of recommender systems. In *Proceedings of the 3rd Annual Conference on Autonomous Agents* (pp. 434–435). ACM Press.
38. Gill, J. (1977). Computational complexity of probabilistic Turing machines. *SIAM Journal on Computing*, 6(4), 675–695.
39. Hemaspaandra, E., & Hemaspaandra, L. (2007). Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1), 73–83.
40. Hemaspaandra, E., Hemaspaandra, L., & Menton, C. (2020). Search versus decision for election manipulation problems. *ACM Transactions on Computation*, 12(1), Article 3, 1–42.
41. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (1997). Raising NP lower bounds to parallel NP lower bounds. *SIGACT News*, 28(2), 2–13.
42. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2007). Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6), 255–285.
43. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2009). Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly*, 55(4), 397–424.
44. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2014). The complexity of online manipulation of sequential elections. *Journal of Computer and System Sciences*, 80(4), 697–710.
45. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2016). The complexity of manipulative actions in single-peaked societies. In J. Rothe (Ed.), *Economics and computation* (pp. 327–360). Berlin: Springer.
46. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2017). The complexity of controlling candidate-sequential elections. *Theoretical Computer Science*, 678, 14–21.
47. Hemaspaandra, E., Hemaspaandra, L., & Rothe, J. (2017). The complexity of online voter control in sequential elections. *Autonomous Agents and Multi-Agent Systems*, 31(5), 1055–1076.
48. Hemaspaandra, L., & Williams, R. (2012). An atypical survey of typical-case heuristic algorithms. *SIGACT News*, 43(4), 71–89.

49. Konczak, K., & Lang, J. (2005). Voting procedures with incomplete preferences. In *Proceedings (Workshop Notes) of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling* (pp. 124–129).
50. Lewis, H., & Papadimitriou, C. (1981). *Elements of the theory of computation*. Englewood Cliffs: Prentice-Hall.
51. Menton, C. (2013). Normalized range voting broadly resists control. *Theory of Computing Systems*, 53(4), 507–531.
52. Menton, C., & Singh, P. (2013). Control complexity of Schulze voting. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence* (pp. 286–292). AAAI Press.
53. Meyer, A., & Stockmeyer, L. (1972). The equivalence problem for regular expressions with squaring requires exponential space. In *Proceedings of the 13th IEEE Symposium on Switching and Automata Theory* (pp. 125–129). IEEE Press.
54. Narodytska, N., & Walsh, T. (2012). Manipulating two stage voting rules. In *Proceedings (Workshop Notes) of the 4th International Workshop on Computational Social Choice* (pp. 323–334).
55. Narodytska, N., & Walsh, T. (2013). Manipulating two stage voting rules. In *Proceedings of the 10th International Conference of Autonomous Agents and Multiagent Systems* (pp. 423–430). International Foundation for Autonomous Agents and Multiagent Systems.
56. Papadimitriou, C., & Yannakakis, M. (1984). The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2), 244–259.
57. Parkes, D., & Xia, L. (2012). A complexity-of-strategic-behavior comparison between Schulze's rule and ranked pairs. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence* (pp. 1429–1435). AAAI Press.
58. Pennock, D., Horvitz, E., & Giles, C. (2000). Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the 17th National Conference on Artificial Intelligence* (pp. 729–734). AAAI Press.
59. Procaccia, A., & Rosenschein, J. (2007). Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28, 157–181.
60. Schaefer, M., & Umans, C. (2002). Completeness in the polynomial-time hierarchy: Part I: A compendium. *SIGACT News*, 33(3), 32–49.
61. Schaefer, M., & Umans, C. (2002). Completeness in the polynomial-time hierarchy: Part II. *SIGACT News*, 33(4), 22–36.
62. Shoham, Y., & Leyton-Brown, K. (2009). *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge: Cambridge University Press.
63. Stockmeyer, L. (1976). The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1), 1–22.
64. Stockmeyer, L., & Meyer, A. (1973). Word problems requiring exponential time. In *Proceedings of the 5th ACM Symposium on Theory of Computing* (pp. 1–9). ACM Press.
65. Valiant, L. (1976). The relative complexity of checking and evaluating. *Information Processing Letters*, 5(1), 20–23.
66. Walsh, T. (2009). Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence* (pp. 324–329). AAAI Press.
67. Wrathall, C. (1976). Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3(1), 23–33.
68. Xia, L., Lang, J., & Monnot, J. (2011). Possible winners when new alternatives join: New results coming up! In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems* (pp. 829–836). International Foundation for Autonomous Agents and Multiagent Systems.
69. Yang, Y. (2017). The complexity of control and bribery in majority judgment. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems* (pp. 1169–1177). International Foundation for Autonomous Agents and Multiagent Systems.
70. Yang, Y., & Guo, J. (2014). Controlling elections with bounded single-peaked width. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems* (pp. 629–636). International Foundation for Autonomous Agents and Multiagent Systems.
71. Yin, Y., Vorobeychik, Y., An, B., & Hazon, N. (2018). Optimal defense against election control by deleting voter groups. *Artificial Intelligence*, 259, 32–51.