Anyone but Him: The Complexity of Precluding an Alternative*

Edith Hemaspaandra
Department of Computer Science
Rochester Institute of Technology
Rochester, NY 14623, USA
eh@cs.rit.edu

Lane A. Hemaspaandra
Department of Computer Science
University of Rochester
Rochester, NY 14627, USA
lane@cs.rochester.edu

Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany
rothe@cs.uni-duesseldorf.de

Abstract

Preference aggregation in a multiagent setting is a central issue in both human and computer contexts. In this paper, we study in terms of complexity the vulnerability of preference aggregation to destructive control. That is, we study the ability of an election's chair to, through such mechanisms as voter/candidate addition/suppression/partition, ensure that a particular candidate (equivalently, alternative) does not win. And we study the extent to which election systems can make it impossible, or computationally costly (NP-complete), for the chair to execute such control. Among the systems we study—plurality, Condorcet, and approval voting—we find cases where systems immune or computationally resistant to a chair choosing the winner nonetheless are vulnerable to the chair blocking a victory. Beyond that, we see that among our studied systems no one system offers the best protection against destructive control. Rather, the choice of a preference aggregation system will depend closely on which types of control one wishes to be protected against. We also find concrete cases where the complexity of or susceptibility to control varies dramatically based on the choice among natural tie-handling rules.

Key words: preferences, computational complexity, multiagent systems.

Introduction

Voting systems provide a broad model for aggregating preferences in a multiagent setting. The literature on voting is vast and active, and spans such areas as AI, complexity, economics, operations research, and political science. As noted by Conitzer, Lang, and Sandholm (2003), voting has been proposed as a mechanism for use in decision-making in various computational settings, including planning (Ephrati & Rosenschein 1991; 1993) and collaborative filtering (Pennock, Horvitz, & Giles 2000). Voting also may be useful in many large-scale computer settings. Examples of much recent interest include the (web-page) rank aggregation problem, and related issues of reducing "spam" results in web search and improving similarity search, for which the use of voting systems has been proposed (Dwork *et al.* 2001; Fagin, Kumar, & Sivakumar 2003). In such an automated setting, it is natural to imagine decisions with thousands or millions of "voters" and "candidates."

In the seminal paper "How hard is it to control an election?" (Bartholdi, Tovey, & Trick 1992), the issue of constructive control of election systems is studied: How hard is it for a chair (who knows all voters' preferences) to—through control of the voter or candidate set or of the partition structure of an election—cause a given candidate (equivalently, alternative) to be the (unique) winner?[1] They studied plurality and Condorcet voting, and seven natural types of control: adding candidates, suppressing candidates, partition of candidates, run-off partition of candidates, adding voters, suppressing voters, and partition of voters. They found that in some cases there is *immunity* to constructive control (if his/her candidate was not already the[2] unique winner, no action of the specified type by the chair can make the candidate the unique winner), in some cases there is *(computational) resistance* to constructive control (it is NP-complete to decide whether the chair can achieve his/her desired outcome), and in some cases the system is *(computationally) vulnerable* to constructive control (there is a polynomial-time algorithm that will tell the chair how to achieve the desired outcome whenever possible[3]).

In this paper, we obtain results for each of their 14 cases (two preference aggregation systems, each under seven con-

*Supported in part by the NSF under grants NSF-CCR-0311021 and NSF-CCF-0426761 and by the DFG under grant RO 1202/9-1. This work was done in part while the authors were visiting Klaus Wagner at the University of Würzburg, and while the first author was on sabbatical at the University of Rochester.

[1]In their model, which is also adopted here, the chair has complete information on the voters' preferences. This is a natural assumption in many situations. For example, in a computer science department, after endless discussions, most people know what each person's position is on key issues. Also, since the case where complete information is available to the chair is a special subcase of the more general setting that allows information to be specified with any level of completeness, lower bounds obtained in the complete information setting are inherited by any natural incomplete information model.

[2]Really "*a* unique winner," since there may be no winner at all, but we'll usually write "*the* unique winner" when this is clear from context.

[3]This is more like "computationally certifiably-vulnerable," see Definition 1. Vulnerability as defined in (Bartholdi, Tovey, & Trick 1992) means one can quickly decide if there *exists* a way for the chair to achieve the desired outcome.

trol schemes) in the setting of *destructive* control. In contrast with constructive control, in which a chair tries to ensure that a specified desirable candidate is the (unique) winner, in destructive control the chair tries to ensure that a specified detested candidate is not the (unique) winner. Regarding the naturalness of destructivity, the light-hearted title of this paper tries to reflect the fact that, in human terms, one often hears feelings expressed that focus strategically on precluding one candidate, and of course in other settings this also may be a goal. Regarding the reality of electoral control, from targeted "get-out-the-vote" advertisements of parties and candidates to (alleged) voter suppression efforts by independent groups, from the way a committee chair groups alternatives to any case where a faculty member hands out student course evaluations on a day some malcontent students are not in class, it is hard to doubt that the desire for electoral control—both destructive and constructive—is a real one.

Destruction has been previously studied by Conitzer, Lang, and Sandholm (2002; 2003), but in the setting of election *manipulation*—in which some voter(s) knowing all other voters' preferences are free to shift their preferences dynamically to affect the outcome. In contrast, in this paper we study destruction in the very different setting of electoral *control* (Bartholdi, Tovey, & Trick 1992)—where a chair, given fixed and unchangeable voter preferences, tries to influence the outcome via procedural/access means.

One might ask, "Why bother studying destructive control, since any rational chair would prefer to assert constructive control?" The answer is that it is plausible—and our results show it is indeed the case—that destructive control may be possible in settings in which constructive control is not. Informally put, destructive control may be easier for the chair to assert. For example, we prove formally that of the seven types of constructive control of Condorcet elections that Bartholdi, Tovey, and Trick (1992) study, the four they showed not vulnerable to constructive control are all vulnerable to destructive control. The remaining three cases regarding Condorcet voting are vulnerable to constructive control (Bartholdi, Tovey, & Trick 1992), but we show that they are immune to destructive control.[4]

Table 1 summarizes our results on the complexity of destructively controlling Condorcet, plurality, and approval elections. We also when needed obtain, for comparative purposes, new results on the complexity of constructive control, and Table 1 displays those and also constructive control results of Bartholdi, Tovey, and Trick (1992). All entries in boldface in Table 1 are new results obtained in this paper; the other results are due to Bartholdi, Tovey, and Trick (1992). For each boldface "V" in the table, "certifiably-vulnerable" is in fact also achieved by our theorems. We mention in passing that for nonboldface "V"s in the table, "certifiably-vulnerable" can be seen directly from or by modifying the algorithms of Bartholdi, Tovey, and Trick (1992).

For control-by-partition problems—which will involve subelection(s)—we distinguish between the models Ties-Eliminate (TE, for short) and Ties-Promote (TP, for short), which define what happens when there are ties in a subelection (before the final election), namely, all participating candidates are eliminated (TE), or all who tie for winner move forward (TP). Note that these models do not apply to Condorcet voting, under which when a winner exists s/he is inherently unique; so the TE/TP distinction is made only for plurality and approval voting.

The natural conclusion to draw from our results is that when selecting an election/preference aggregation system, one should at least be aware of the issue of the system's vulnerability to control—and, beyond that, one's choice of system will depend closely on which types of immunity or computational resistance one most values. Our results also show that constructive and destructive control often differ greatly: A system immune to constructive control may be vulnerable to destructive control, and vice versa. Finally, our results show that—in contrast with some comments in earlier papers—breaking ties is far from a minor issue: For both voting types where tie-handling rules are meaningful, we find cases where the complexity of or susceptibility to control varies dramatically based on the choice among natural tie-handling rules.

Preliminaries

We first define the three voting systems considered. In *approval voting*, each voter votes "Yes" or "No" for each candidate. (So, for approval voting, a voter's preferences are reflected by a 0-1 vector.) All candidates with the maximum number of "Yes" votes are winners.

Plurality and Condorcet voting are defined in terms of strict preferences. For them, an election is given by a *preference profile*, a pair (C, V) such that C is a set of candidates and V is the multiset (henceforth, we'll just say set, as a shorthand) of the voters' preference orders on C. We assume that the preference orders are irreflexive and antisymmetric

[4]Savvy readers may wonder whether the last statement is impossible. After all, to ensure that the despised candidate c is not the unique winner, we simply have to ask whether at least one of the other candidates can be ensured to unique-win-or-tie-for-winner. Formally put, destructive control polynomial-time disjunctively truth-table reduces (Ladner, Lynch, & Selman 1975) to constructive control (redefined to embrace ties), and so the destructive control problem can (within a polynomial factor) be no less hard *computationally* than the (redefined) constructive control problem (this is noted in a different setting by Conitzer and Sandholm (2002)), seemingly (though not really) contrary to the three cases just mentioned. The brief explanation of why this does not cause a paradox lies in the word "computational": Although immunity is the most desirable case in terms of security from control, the *complexity* of recognizing whether a given candidate can be precluded from winning in immune cases will most typically be in P—after all, we can *never*, when immunity holds, change a given candidate from unique winner to not the unique winner, so the related decision problem is typically easy. (Technical side remark: We say "will most typically be in P/is typically" rather than "will be in P/is" because for impractical systems that—unlike those here—have winner-testing problems that are not in P, it is in concept possible that one can have immunity and yet also have the related language problem not belong to P.)

The disjunctive-truth-table connection mentioned above explains why, if P $\neq$ NP, it is impossible to have computational resistance to destructive control hold for any problem that, when redefined to embrace ties, is vulnerable to constructive control.

	Plurality		Condorcet		Approval	
Control by	Constructive	Destructive	Constructive	Destructive	Constructive	Destructive
Adding Candidates	R	**R**	I	**V**	**I**	**V**
Deleting Candidates	R	**R**	V	**I**	**V**	**I**
Partition of Candidates	TE: R TP: R	**TE: R** **TP: R**	V	**I**	**TE: V** **TP: I**	**TE: I** **TP: I**
Run-off Partition of Candidates	TE: R TP: R	**TE: R** **TP: R**	V	**I**	**TE: V** **TP: I**	**TE: I** **TP: I**
Adding Voters	V	**V**	R	**V**	**R**	**V**
Deleting Voters	V	**V**	R	**V**	**R**	**V**
Partition of Voters	**TE: V** **TP: R**	**TE: V** **TP: R**	R	**V**	**TE: R** **TP: R**	**TE: V** **TP: V**

Table 1: Summary of results. Results new to this paper are in boldface. Nonboldface results are due to Bartholdi, Tovey, and Trick (1992). Key: I = immune, R = resistant, V = vulnerable, TE = Ties-Eliminate, TP = Ties-Promote.

(i.e., every voter has strict preferences over the candidates), complete (i.e., every voter ranks each candidate), and transitive.

A *voting system* is a rule for how to determine the winner(s) of an election. Formally, any voting system is defined to be a *(social choice) function* mapping any given preference profile (or the analog with voters' 0-1 vectors for the approval voting case) to society's aggregate *choice set*, the set of candidates who have won the election.

In *plurality voting*, each candidate with a maximum number of "first preference among the candidates in the election" voters for him/her wins. In *Condorcet voting*, for each $c \in C$, c is a winner if and only if for each $d \in C$ with $d \neq c$, c defeats d by a strict majority of votes in a pairwise election between them based on the voters' preferences.

The Condorcet Paradox observes that whenever there are at least three candidates, due to cyclic aggregate preference rankings Condorcet winners may not exist (Condorcet 1785). That is, the set of winners may be empty. However, a Condorcet winner is unique whenever one does exist. In the case of plurality and approval voting, due to ties, there may exist multiple winners. Regarding ties, we—following Bartholdi, Tovey, and Trick (1992) to best allow comparison—focus in our control problems on creating a *unique* winner (constructive), and precluding a candidate from being the *unique* winner (destructive). (Ties in sub-elections, for the partition problems, are handled via the TE and TP rules described earlier.)

Results

The issue of control of an election by the authority conducting it (called the chair) can be studied under a variety of models and scenarios. For plurality and Condorcet voting, Bartholdi, Tovey, and Trick (1992)—which for the rest of this section will be referred to as "BTT92"—study constructive control by adding candidates, deleting candidates, partition of candidates, run-off partition of candidates, adding voters, deleting voters, and partition of voters. In their setting, the chair's goal is to make a given candidate uniquely win the election. Analogously, we consider in turn the corresponding seven *destructive* control problems, where the chair's goal is to preclude a given candidate from being the unique winner. For each of these seven we define the problem and present prior results and our results. (Formally, for each type of control one defines a decision problem and studies its computational complexity.) To make comparisons as easy as possible, we in stating these seven control problems whenever possible exactly follow BTT92's wording—modified to the destructive case (and when we diverge, we explain why and how).

Destructive Control by Adding Candidates As is common, we state our decision problems as "Given" instances, and a related Yes/No question. The language in each case is the set of all instances for which the answer is Yes.

Given: A set C of qualified candidates and a distinguished candidate $c \in C$, a set D of possible spoiler candidates, and a set V of voters with preferences (in the approval case, the "preferences" will, as always for that case, actually be 0-1 vectors) over $C \cup D$.

Question: Is there a choice of candidates from D whose entry into the election would assure that c is not the unique winner?

The above type of destructive control captures the idea that the chair tries to dethrone the despised candidate c by introducing new "spoiler" candidates.

With this first problem stated, now is a good time to define our notions of control. Our terminology will closely follow the notions in BTT92, to allow comparison.

Definition 1 *We say that a voting system is* immune to destructive control *in a given model (of control) if it is never possible for the chair to change a given candidate from the unique winner to being not the unique winner by using his/her allowed model of control. If a system is not immune to a type of control, it is said to be* susceptible *to that type of control.*

A voting system is said to be (computationally) vulnerable to control *if it is susceptible to control and the corresponding language problem is computationally easy (i.e., solvable in polynomial time). If a system is not just vulnerable but one can, given a control problem, produce in polynomial time* *the actual action of the chair to execute control the "best" way (namely, by adding or deleting the smallest number of candidates or voters for add/delete problems; for partition problems, any legal partition that works is acceptable), we say the system is* (computationally) certifiably-vulnerable to control.[5]

A voting system is said to be resistant to control *if it is susceptible to control but the corresponding language problem is computationally hard (i.e.,* NP-*complete).*[6]

For the theory of NP-completeness, see, e.g., (Garey & Johnson 1979; Hopcroft & Ullman 1979). For space reasons we do not explicitly define the corresponding notions for constructive control (except briefly in the introduction) and the analogous seven constructive control decision problems from BTT92, but except when noted below they are exactly analogous. For example, for the above control scenario, the analogous constructive problem pairs the same "Given" with the question "Is there a choice of candidates from D whose entry into the election would assure that c is the unique winner?"

As to what is known about Control by Adding Candidates, BTT92 showed that plurality is resistant and Condorcet is immune. Our results are:

Theorem 2 *Approval (voting) is immune to constructive control by adding candidates, and plurality, Condorcet, and approval (voting) are respectively resistant, vulnerable/certifiably-vulnerable, and vulnerable/certifiably-vulnerable to destructive control by adding candidates.*

So, though Condorcet and approval are immune to constructive control of this sort, they both are vulnerable to destructive control. This reverses itself for:

Destructive Control by Deleting Candidates

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

Question: Are there k or fewer candidates other than c in C whose disqualification would assure that c is not the unique winner?

[5]For the seven problems studied here, certifiably-vulnerable implies vulnerable (but we list both, since if one studied add/delete problems stated not in terms of "is there some subset" or "by adding/deleting $\leq k$" but rather in terms of "by adding/deleting exactly k," then for certain systems the implication need not hold).

[6]It would be more natural to define resistance as meaning the corresponding language is (many-one) NP-hard. However, in this paper, we define resistance in terms of NP-completeness. One reason is that this matches the way the term is used by BTT92. More importantly, all the problems discussed in this paper have obvious NP upper bounds since testing whether a given candidate has won a given election for the systems considered here is obviously in P. So for the problems in this paper, NP-completeness and NP-hardness stand or fall together. We mention in passing that there are natural election systems whose complexity seems beyond NP. The first such case established was for the election system defined by Lewis Carroll in 1876 (Dodgson 1876), where even the complexity of determining whether a given candidate has won is now known to be hard for parallel access to NP (Hemaspaandra, Hemaspaandra, & Rothe 1997).

In this type of control, the chair seeks to influence the outcome of the election by suppressing certain candidates (other than c), in hopes that their voters now support another candidate to ensure stopping c. Note that this formalization is not a perfect analog of the constructive case of BTT92 in that we explicitly prevent deleting c, since otherwise any voting system in which the winners can efficiently be determined would be trivially vulnerable to this type of control.

Here, BTT92 establish for constructive control resistance for plurality and immunity for Condorcet. Our results are:

Theorem 3 *Approval is vulnerable/certifiably-vulnerable to constructive control by deleting candidates. Plurality, Condorcet, and approval voting are respectively resistant, immune, and immune to destructive control by deleting candidates.*

We now handle jointly the two types of partition of candidates, since they yield identical results.

Destructive Control by Partition of Candidates

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

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

Destructive Control by Run-Off Partition of Candidates

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

Question: Is there a partition of C into C_1 and C_2 such that c is not the unique winner of the election in which those candidates surviving (with respect to the tie-handling rule) subelections (C_1, V) and (C_2, V) have a run-off with voter set V.

These two types of control express settings in which the chair tries to, overall, partition the candidates in such a clever way that the hated candidate c fails to be the unique winner—one via a cascading setup, and one via a run-off setup. Here, BTT92 show that for constructive control plurality is resistant (and their results hold in both our TE and TP models) and Condorcet is vulnerable. Our results are:

Theorem 4 *Approval is vulnerable/certifiably-vulnerable to constructive control by partition of candidates and run-off partition of candidates in model TE and immune to constructive control by partition of candidates and run-off partition of candidates in model TP. Plurality, Condorcet, and approval voting are, in models TE and TP, respectively resistant, immune, and immune to destructive control by partition of candidates and by run-off partition of candidates.*

So Condorcet, though vulnerable to constructive control, is immune to destructive control here. And, perhaps more interesting, for constructive control, approval changes from vulnerable to immune depending on the tie-handling rule.

We now turn to control of the voter set. The intuition behind seeking destructive control by adding or deleting voters is clear, e.g., getting out the vote and vote suppression. We handle these two cases together as their results are identical.

Destructive Control by Adding Voters

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

Question: Are there k or fewer voters from W whose registration would assure that c is not the unique winner?

Destructive Control by Deleting Voters

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

Question: Are there k or fewer voters in V whose disenfranchisement would assure that c is not the unique winner?

Here, BTT92 show that for constructive control plurality is vulnerable and Condorcet is resistant. Our results are:

Theorem 5 *Approval is resistant to constructive control by adding voters and by deleting voters. Plurality, Condorcet, and approval voting are all vulnerable/certifiably-vulnerable to destructive control by adding voters and by deleting voters.*

So Condorcet and approval, though resistant to constructive control, are vulnerable to destructive control here.

The final problem here results in a surprise.

Destructive Control by Partition of Voters

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

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

In this last type of control, the voter set is partitioned into two "subcommittees" that both separately select their "nominees," who run against each other in the final decision stage. Unlike BTT92, we again distinguish between the two models Ties-Eliminate and Ties-Promote defined above. That is, in the Ties-Eliminate model, if two or more candidates tie for winning in a subcommittee's election, no candidate is nominated by that subcommittee. In contrast, in the Ties-Promote model, all the candidates who tie for winning in a subcommittee's election are nominated to run in the final decision stage. If both subcommittees nominate the same candidate (and no one else), we by convention declare this candidate the run-off winner—one cannot eliminate oneself.

We mention that both of our two tie-handling models, TE and TP, differ from the model adopted in BTT92, where they for vulnerability results about this problem adopt a third model in which ties are handled not by a tie-handling rule but rather by changing the decision problem itself to require the chair to find a partition that completely avoids ties in any subcommittee. We find our model the more natural, but we mention that they obtained for this case, in their tie model, a constructive-control vulnerability result for plurality. For Condorcet and constructive control, they proved that resistance holds. Our results are:

Theorem 6 *Approval is resistant to constructive control by partition of voters in models TE and TP, vulnerable/certifiably-vulnerable to destructive control by partition of voters in models TE and TP. Plurality is vulnerable/certifiably-vulnerable to both constructive and destructive control by partition of voters in model TE, and is resistant to both constructive and destructive control by partition of voters in model TP. Condorcet is vulnerable/certifiably-vulnerable to destructive control by partition of voters.*

The most striking behavior here is that plurality voting varies between being vulnerable and being resistant, depending on the tie-handling rule. The loose intuition for this is that in TE, at most one candidate wins each subcommittee and in polynomial time we can explore every way this can happen. In contrast, under TP potentially any subset of candidates may move forward, and in this particular setting, that flexibility is enough to support NP-completeness. Also interesting is that both Condorcet and approval, while resistant to constructive control, are vulnerable to destructive control.

Proof Comments

This section tries to give some general feeling for the proofs. Complete proofs will be included in the full version of this paper.

Immunity results are generally clear from the definitions. For example, note that any system satisfying, as does approval voting, the "unique" version of the Weak Axiom of Revealed Preference—a unique winner among a collection of candidates always remains a unique winner among every subcollection of candidates that includes him/her—is immune to destructive control by deleting candidates, by partition of candidates, and by run-off partition of candidates.

The certifiably-vulnerable results (which here imply the vulnerable results) range from clear greedy algorithms to trickier algorithms based on characterizing the ways in which a candidate can be precluded from winning (or for the constructive case, made to win). The more surprising of these have to do with the tie-handling cases of partition problems—where the chair can at times do shrewd things (e.g., shift voters counterintuitively to induce ties that kill off stronger candidates).

For illustration, we now prove that approval voting is vulnerable/certifiably-vulnerable to destructive control by partition of voters in the TP model. An easy example shows that approval voting is susceptible to this type of control: Let $C = \{a, b, c\}$, and define V to consist of the following ten voters (specified by 0-1 vectors): $v_1 = v_2 = v_3 = v_4 = 100$, $v_5 = v_6 = v_7 = 010$, and $v_8 = v_9 = v_{10} = 001$. In (C, V), a is the unique winner. But if V is partitioned into $V_1 = \{v_1, v_2, v_5, v_6, v_7\}$ and $V_2 = V - V_1$, b and c are nominated by the subcommittees (C, V_1) and (C, V_2), respectively, and tie for winner in the run-off.

Now, given a set of candidates C, a distinguished candidate $c \in C$, and a voter set V, our polynomial-time algorithm for this problem works as follows. If $C = \{c\}$, output "control impossible," as c must win; else if c already is not the unique winner, output $(V, \emptyset)$ as a successful partition.

Otherwise, if $||C|| = 2$, output "control impossible," since in the current case c *is* the unique winner, so c will win at least one subcommittee and also the run-off. If none of the above cases applies, for each $a, b \in C$ with $||\{a, b, c\}|| = 3$, we test whether we can make a strictly beat c in (C, V_1) and make b strictly beat c in (C, V_2). For each voter in V, we focus just on his/her approval of a, b, and c, represented as the vector $abc \in \{0, 1\}^3$. Denote the number of voters with preference 000 or 111 by N, with 001 by W_c, with 110 by L_c, with 100 by S_a, with 010 by S_b, with 101 by S_{ac}, and with 011 by S_{bc}. Since c is the approval winner, $W_c + S_{bc} - (L_c + S_a) > 0$ and $W_c + S_{ac} - (L_c + S_b) > 0$. If $W_c - L_c > S_{ac} + S_a + S_{bc} + S_b - 2$ or $S_{ac} + S_a = 0$ or $S_{bc} + S_b = 0$, then this a and b are hopeless, so move on to consider the next a and b in the loop. Otherwise, we have $W_c - L_c \leq S_{ac} + S_a + S_{bc} + S_b - 2$ and $S_{ac} + S_a > 0$ and $S_{bc} + S_b > 0$, and output (V_1, V_2) as a successful partition, where V_1 contains all voters contributing to S_{ac} and S_a, and also $\min(W_c, S_{ac} + S_a - 1)$ voters contributing to W_c, and $V_2 = V - V_1$. In (C, V_1), a (strictly) beats c, since a gets $S_{ac} + S_a - \min(W_c, S_{ac} + S_a - 1)$ more Yes votes than c. And in (C, V_2), b (strictly) beats c, since b has $S_{bc} + S_b + L_c - (W_c - \min(W_c, S_{ac} + S_a - 1))$ more Yes votes than c. So, for the construction to work, we must argue that $S_{bc} + S_b + L_c + \min(W_c, S_{ac} + S_a - 1) - W_c > 0$. That is, we need $W_c - L_c < \min(W_c, S_{ac} + S_a - 1) + S_{bc} + S_b$. If $W_c \leq S_{ac} + S_a - 1$, this reduces to $0 < L_c + S_{bc} + S_b$, which follows from the fact that in the current case $S_{bc} + S_b > 0$. And if $W_c > S_{ac} + S_a - 1$, the desired inequality follows from the known fact that $W_c - L_c \leq S_{ac} + S_a + S_{bc} + S_b - 2$. To conclude the polynomial algorithm, if in no loop iteration did we find an a and b that allowed us to output a partition of voters dethroning c, then output "control impossible."

Finally, the resistance results are based on clear containments in NP, plus (polynomial-time, many-one) reductions establishing NP-hardness. We whenever possible try to achieve multiple resistance results via a single proof.

For example, with a single proof we establish all seven resistance results for plurality voting: destructive control by adding, deleting, partition (TE and TP), and run-off partition (TE and TP) of candidates,[7] and by partition of voters (TP). We now sketch this proof, which is achieved via one general construction that yields the reductions, each from the NP-complete problem Hitting Set: Given a set $B = \{b_1, b_2, \ldots, b_m\}$, a family $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ of subsets S_i of B, and a positive integer k, does $\mathcal{S}$ have a hitting set of size at most k, i.e., is there a set $B' \subseteq B$ with $||B'|| \leq k$ such that for each i, $S_i \cap B' \neq \emptyset$? Given a triple $(B, \mathcal{S}, k)$, we construct the following election. The candidate set is $C = B \cup \{c, w\}$. The voter set V is defined as follows. There are $2(m-k)+2n(k+1)+4$ voters of the form $w > c > \cdots$, where "$\cdots$" means that the remaining candidates follow in some arbitrary order. There are $2n(k+1)+5$ voters of the form $c > w > \cdots$. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form $S_i > w > \cdots$, where "S_i" denotes the elements of S_i in some arbitrary order. Finally, for each j, $1 \leq j \leq m$, there are two voters of the form $b_j > c > \cdots$.

We show that $\mathcal{S}$ has a hitting set of size less than or equal to k if and only if destructive control by adding candidates can be executed for the election with qualified candidates $\{c, w\}$, spoiler candidates B, distinguished candidate w, and voters V. For every candidate d, let $sc(d)$ denote the number of voters who rank d first. If B' is a hitting set of $\mathcal{S}$ of size k, then in the election $(B' \cup \{c, w\}, V)$, $sc(w) = 2(m-k) + 2n(k+1) + 4$ and $sc(c) = 2n(k+1) + 5 + 2(m-k)$. It follows that w is not the unique winner of $(B' \cup \{c, w\}, V)$. For the converse, suppose that w is not the unique winner of an election $(B' \cup \{c, w\}, V)$, where $B' \subseteq B$. We show that then B' is a hitting set of $\mathcal{S}$ of size at most k. First note that for all $B' \subseteq B$ and for all $b \in B'$, $sc(b) < sc(w)$ in $(B' \cup \{c, w\}, V)$. So, if w is not the unique winner of $(B' \cup \{c, w\}, V)$, then $sc(c) \geq sc(w)$ in $(B' \cup \{c, w\}, V)$.

In $(B' \cup \{c, w\}, V)$, $sc(c) = 2n(k+1)+5+2(m-||B'||)$ and $sc(w) = 2(m-k) + 2n(k+1) + 4 + 2(k+1)\ell$, where ℓ is the number of sets in $\mathcal{S}$ that are not hit by B' (i.e., that have an empty intersection with B'). Since $sc(w) \leq sc(c)$, it follows that $2(m-k) + 2(k+1)\ell \leq 1 + 2(m - ||B'||)$, which implies $(k+1)\ell + ||B'|| - k \leq 0$. So $\ell = 0$. Thus, B' is a hitting set of $\mathcal{S}$ of size at most k. It follows that plurality is resistant to destructive control by adding candidates.

A similar argument can be used to show that $\mathcal{S}$ has a hitting set of size at most k if and only if the election with candidates C, distinguished candidate w, and voters V can be destructively controlled by deleting at most $m-k$ candidates. (The proof is omitted.) Thus, plurality is resistant to destructive control by deleting candidates.

Furthermore, $\mathcal{S}$ has a hitting set of size at most k if and only if the election with candidates C, distinguished candidate w, and voters V can be destructively controlled by partitioning C into $C_1 = B' \cup \{c, w\}$ and $C_2 = B - B'$, where B' is a hitting set of $\mathcal{S}$ of size $\leq k$. (The proof is omitted. Since w never ties for winner in a subelection, the proof works for both TE and TP.) A similar argument works for destructive control for run-off partition of candidates. Thus, plurality is resistant to destructive control by partition and run-off partition of candidates (both in TE and TP).

Finally, we show that plurality is resistant to destructive control by partition of voters in the TP model. To this end, we reduce from the Hitting Set problem restricted to instances where $n(k+1)+1 \leq m-k$. (The proof that this restriction of Hitting Set is still NP-complete is omitted.)

We first show that, in the election (C, V) constructed above, for every partition of V into V_1 and V_2, w is a winner of (C, V_1) or of (C, V_2). For a contradiction, suppose that w is a winner of neither (C, V_1) nor (C, V_2). Let $x \in B \cup \{c\}$ be a winner of (C, V_1) with score $sc_{V_1}(x)$, and let $y \in B \cup \{c\}$ be a winner of (C, V_2) with score $sc_{V_2}(y)$. Then $sc_{V_1}(x) + sc_{V_2}(y) \geq sc_V(w) + 2$. Since w's score in (C, V) is greater than that of any other candidate, $x \neq y$. It follows that $sc_{V_1}(x) + sc_{V_2}(y) \leq sc_V(c) + sc_V(b_i) \leq 2n(k+1)+5+2n(k+1)+2 \leq 2n(k+1)+5+2(m-k) = sc_V(w)+1$, a contradiction. We can now show that, assum-

[7] Our constructions ensure that the distinguished candidate is never tied for winner in any subelection in the image of the NP-hardness reduction. Thus, these results hold both in the Ties-Eliminate and Ties-Promote models.

ing $n(k+1)+1 \leq m-k$, $\mathcal{S}$ has a hitting set of size at most k if and only if V can be partitioned such that w is not the unique winner of (C, V) in the TP model. If B' is a hitting set of size at most k, let V_1 consist of one voter each of the form $c > w > \cdots$ and $b_j > c > \cdots$, for each j, $1 \leq j \leq m$, and let $V_2 = V - V_1$. Then the winners of (C, V_1) are $B' \cup \{c\}$, and the winner of (C, V_2) is w. It is easy to see that c is the unique winner in $(B' \cup \{c, w\}, V)$. For the converse, suppose there is a partition of V such that w is not the unique winner of the election. However, w is a winner of a subcommittee's election, as shown above. It follows that w is not the unique winner of a run-off election involving w, i.e., w is not the unique winner in $(D \cup \{w\}, V)$, where $D \subseteq B \cup \{c\}$. An argument similar to that in the "destructive control by adding candidates" case shows that $\mathcal{S}$ has a hitting set of size at most k. This completes the proof sketch.

Conclusions

In this paper, we studied the computational resistance and vulnerability of three voting systems—plurality, Condorcet, and approval voting—to destructive control by an election's chair in each of seven control scenarios: candidate addition, suppression, partition, and run-off partition, and voter addition, suppression, and partition. We classified each case as immune, vulnerable, or computationally resistant. We also studied the analogous constructive control cases and fully resolved those that were not considered by Bartholdi, Tovey, and Trick.

We identified cases where a system immune to constructive control still can be vulnerable to destructive control (e.g., Condorcet voting for control by adding candidates), and vice versa (e.g., approval voting for control by deleting candidates). We saw that, among the systems studied, none is globally superior to the others. Rather, when choosing a voting system, one's choice will depend on the types of control against which protection is most desired. Finally, we saw that—in contrast to some comments in earlier papers—tie-breaking is a far from minor issue: For those control types that involve partitions of the candidate or voter set, we studied two natural tie-handling rules, and we found specific cases in which the complexity of the corresponding control problem varies crucially depending on which tie-handling rule is adopted.

Acknowledgments

We thank Jeroen Snippe for helpful comments, and we thank Klaus Wagner and his group for hosting a visit during which this work was done in part.

References

Bartholdi, III, J.; Tovey, C.; and Trick, M. 1992. How hard is it to control an election? *Mathematical and Computer Modeling* 16(8/9):27–40.

Black, D. 1958. *Theory of Committees and Elections*. Cambridge University Press.

Condorcet, M. 1785. *Essai sur l'Application de L'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale.

Conitzer, V., and Sandholm, T. 2002. Complexity of manipulating elections with few candidates. In *Proceedings of the 18th National Conference on Artificial Intelligence*, 314–319. AAAI Press.

Conitzer, V.; Lang, J.; and Sandholm, T. 2003. How many candidates are needed to make elections hard to manipulate? In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge*, 201–214. ACM Press.

Dodgson, C. 1876. A method of taking votes on more than two issues. Pamphlet printed by the Clarendon Press, Oxford, and headed "not yet published" (see the discussions in (McLean & Urken 1995; Black 1958), both of which reprint this paper).

Dwork, C.; Kumar, S.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, 613–622.

Ephrati, E., and Rosenschein, J. 1991. The Clarke tax as a consensus mechanism among automated agents. In *Proceedings of the 9th National Conference on Artificial Intelligence*, 173–178. AAAI Press.

Ephrati, E., and Rosenschein, J. 1993. Multi-agent planning as a dynamic search for social consensus. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, 423–429.

Fagin, R.; Kumar, R.; and Sivakumar, D. 2003. Efficient similarity search and classification via rank aggregation. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*, 301–312. ACM Press.

Garey, M., and Johnson, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company.

Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 1997. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM* 44(6):806–825.

Hopcroft, J., and Ullman, J. 1979. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley.

Ladner, R.; Lynch, N.; and Selman, A. 1975. A comparison of polynomial time reducibilities. *Theoretical Computer Science* 1(2):103–124.

McLean, I., and Urken, A. 1995. *Classics of Social Choice*. University of Michigan Press.

Pennock, D.; Horvitz, E.; and 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*, 729–734. AAAI Press.