Sincere-Strategy Preference-Based Approval Voting Fully Resists Constructive Control and Broadly Resists Destructive Control[1]

Gábor Erdélyi, Markus Nowak, and Jörg Rothe

Abstract

We study sincere-strategy preference-based approval voting (SP-AV), a system proposed by Brams and Sanver [8], with respect to procedural control. In such control scenarios, an external agent seeks to change the outcome of an election via actions such as adding/deleting/partitioning either candidates or voters. SP-AV combines the voters' preference rankings with their approvals of candidates, and we adapt it here so as to keep its useful features with respect to approval strategies even in the presence of control actions. We prove that this system is computationally resistant (i.e., the corresponding control problems are NP-hard) to 19 out of 22 types of constructive and destructive control. Thus, SP-AV has more resistances to control, by three, than is currently known for any other natural voting system with a polynomial-time winner problem. In particular, SP-AV is (after Copeland voting, see [19]) the second natural voting system with an easy winner-determination procedure that is known to have full resistance to constructive control, and unlike Copeland voting it in addition displays broad resistance to destructive control.

1 Introduction

Voting provides a particularly useful method for preference aggregation and collective decision-making. While voting systems were originally used in political science, economics, and operations research, they are now also of central importance in various areas of computer science, such as artificial intelligence (in particular, within multiagent systems). In automated, large-scale computer settings, voting systems have been applied, e.g., for planning [11] and similarity search [15], and have also been used in the design of recommender systems [21] and ranking algorithms [10] (where they help to lessen the spam in meta-search web-page rankings). For such applications, it is crucial to explore the computational properties of voting systems and, in particular, to study the complexity of problems related to voting (see, e.g., the survey by Faliszewski et al. [17]).

The study of voting systems from a complexity-theoretic perspective was initiated by Bartholdi, Tovey, and Trick's series of seminal papers about the complexity of winner determination [2], manipulation [1], and procedural control [3] in elections. This paper contributes to the study of electoral control, where an external agent—traditionally called *the chair*—seeks to influence the outcome of an election via procedural changes to the election's structure, namely via adding/deleting/partitioning either candidates or voters (see Section 2.2 for the formal definitions of our control problems). We consider both *constructive* control (introduced by Bartholdi et al. [3]), where the chair's goal is to make a given candidate the unique winner, and *destructive* control (introduced by Hemaspaandra et al. [22]), where the chair's goal is to prevent a given candidate from being a unique winner.

[1]Supported in part by the DFG under grants RO 1202/12-1 (within the European Science Foundation's EUROCORES program LogICCC: "Computational Foundations of Social Choice") and RO 1202/11-1 and by the Alexander von Humboldt Foundation's TransCoop program. Some results of this paper were presented at the 33rd International Symposium on Mathematical Foundations of Computer Science (MFCS-08), August 2008 [13].

We investigate the same twenty types of constructive and destructive control that were studied for approval voting [22] and two additional control types introduced by Faliszewski et al. [18], and we do so for a voting system that was proposed by Brams and Sanver [8] as a combination of preference-based and approval voting. Approval voting was introduced by Brams and Fishburn ([4, 5], see also [6]) as follows: Every voter either approves or disapproves of each candidate, and every candidate with the largest number of approvals is a winner. One of the simplest preference-based voting systems is plurality: All voters report their preference rankings of the candidates, and the winners are the candidates that are ranked first-place by the largest number of voters. The purpose of this paper is to show that Brams and Sanver's combined system (adapted here so as to keep its useful features even in the presence of control actions) combines the strengths, in terms of computational resistance to control, of plurality and approval voting.

Some voting systems are *immune* to certain types of control in the sense that it is never possible for the chair to reach his or her goal via the corresponding control action. Of course, immunity to any type of control is most desirable, as it unconditionally shields the voting system against this particular control type. Unfortunately, like most voting systems approval voting is *susceptible* (i.e., not immune) to many types of control, and plurality voting is susceptible to all types of control. However, and this was Bartholdi, Tovey, and Trick's brilliant insight [3], even for systems susceptible to control, the chair's task of controlling a given election may be too hard computationally (namely, NP-hard) for him or her to succeed. The voting system is then said to be *resistant* to this control type. If a voting system is susceptible to some type of control, but the chair's task can be solved in polynomial time, the system is said to be *vulnerable* to this control type.

The quest for a natural voting system with an easy winner-determination procedure that is universally resistant to control has lasted for more than 15 years now. Among the voting systems that have been studied with respect to control are plurality, Condorcet, approval, cumulative, Llull, and (variants of) Copeland voting [3, 22, 23, 27, 18, 19]. Among these systems, plurality and Copeland voting (denoted Copeland$^{0.5}$ in [19]) display the broadest resistance to control, yet even they are not universally control-resistant. The only system currently known to be fully resistant—to the 20 types of constructive and destructive control studied in [22, 23]—is a highly artificial system constructed via hybridization [23]. (We mention that this system was not designed for direct, real-world use as a "natural" system but rather was intended to rule out the existence of a certain impossibility theorem [23].)

While approval voting nicely distinguishes between each voter's acceptable and inacceptable candidates, it ignores the preference rankings the voters may have about their approved (or disapproved) candidates. This shortcoming motivated Brams and Sanver [8] to introduce a voting system that combines approval and preference-based voting, and they defined the related notions of sincere and admissible approval strategies, which are quite natural requirements. We adapt their sincere-strategy preference-based approval voting system in a natural way such that, for elections with at least two candidates, admissibility of approval strategies (see Definition 2.1) can be ensured even in the presence of control actions such as deleting candidates and partitioning candidates or voters. Note that, in control by partition of voters (see [14]), the run-off may have a reduced number of candidates.

The purpose of this paper is to study if, and to what extent, this hybrid system (where "hybrid" is not meant in the sense of [23] but refers to combining preference-based with approval voting in the sense of Brams and Sanver [8]) inherits the control resistances of plurality (which is perhaps the simplest preference-based system) and approval voting. Denoting this system by SP-AV, we show that SP-AV does combine all the resistances of plurality and approval voting.

More specifically, we prove that sincere-strategy preference-based approval voting is resistant to 19 and vulnerable to only three of the 22 types of control considered here. For

Number of	Condorcet	Approval	Llull	Copeland	Plurality	SP-AV
resistances	3	4	14	15	16	19
immunities	4	9	0	0	0	0
vulnerabilities	7	9	8	7	6	3
References	[3, 22]	[3, 22]	[18, 19]	[18, 19]	[3, 22, 18]	[13, 22] and this paper

Table 1: Number of resistances, immunities, and vulnerabilities to our 22 control types.

comparison, Table 1 shows the number of resistances, immunities, and vulnerabilities to our 22 control types that are known for each of Condorcet,[2] approval, Llull, plurality, and Copeland voting (see [3, 22, 18, 19]), and for SP-AV (see Theorem 3.1 and Table 2).

This paper is organized as follows. In Section 2, we define sincere-strategy preference-based approval voting, the types of control studied in this paper, and the notions of immunity, susceptibility, vulnerability, and resistance [3]. In Section 3, we prove our results on SP-AV. Finally, in Section 4 we give our conclusions and state some open problems.

2 Preliminaries

2.1 Preference-Based Approval Voting

An election $E = (C, V)$ is specified by a finite set C of candidates and a finite collection V of voters who express their preferences over the candidates in C, where distinct voters may of course have the same preferences. How the voter preferences are represented depends on the voting system used. In approval voting (AV, for short), every voter draws a line between his or her acceptable and inacceptable candidates (by specifying a 0-1 approval vector, where 0 represents disapproval and 1 represents approval), yet does not rank them. In contrast, many other important voting systems (e.g., Condorcet voting, Copeland voting, all scoring protocols including plurality, Borda count, veto, etc.) are based on voter preferences that are specified as tie-free linear orderings of the candidates. As is most common in the literature, votes will here be represented nonsuccinctly: one ballot per voter. Note that some papers (e.g., [16, 18, 19]) also consider succinct input representations for elections where multiplicities of votes are given in binary.

Brams and Sanver [8] introduced a voting system that combines approval and preference-based voting. To distinguish this system from other systems that these authors introduced with the same purpose of combining approval and preference-based voting [7], we call the variant considered here (including the conventions and rules to be explained below) *sincere-strategy preference-based approval voting* (SP-AV, for short).

Definition 2.1 ([8]) *Let (C, V) be an election, where the voters both indicate approvals/disapprovals of the candidates and provide a tie-free linear ordering of all candidates. For each voter $v \in V$, an* AV strategy *of v is a subset $S_v \subseteq C$ such that v approves of all candidates in S_v and disapproves of all candidates in $C - S_v$. The list of AV strategies*

[2] As in [22], we consider two types of control by partition of candidates (namely, with and without run-off) and one type of control by partition of voters, and for each partition case we use the rules TE ("ties eliminate") and TP ("ties promote") for handling ties that may occur in the corresponding subelections (see [14]). However, since Condorcet winners are always unique when they exist, the distinction between TE and TP is not made for the partition cases within Condorcet voting. Note further that the two additional control types in Section 2.2.1 (namely, constructive and destructive control by adding a limited number of candidates [18]) have not been considered for Condorcet voting [3, 22]. That is why Table 1 lists only 14 instead of 22 types of control for Condorcet.

for all voters in V is called an AV strategy profile for (C, V). *(We sometimes also speak of V's* AV strategy profile for C.*) For each $c \in C$, let $score_{(C,V)}(c) = \|\{v \in V \mid c \in S_v\}\|$ denote the number of c's approvals. Every candidate c with the largest $score_{(C,V)}(c)$ is a winner of election (C, V).*

An AV strategy S_v of a voter $v \in V$ is said to be admissible *if S_v contains v's most preferred candidate and does not contain v's least preferred candidate. S_v is said to be* sincere *if for each $c \in C$, if v approves of c then v also approves of each candidate ranked higher than c (i.e., there are no gaps allowed in sincere approval strategies). An AV strategy profile for (C, V) is* admissible *(respectively,* sincere*) if the AV strategies of all voters in V are admissible (respectively, sincere).*

Admissibility and sincerity are quite natural requirements. In particular, requiring the voters to be sincere ensures that their preference rankings and their approvals/disapprovals are not contradictory. Note further that admissible AV strategies are not dominated in a game-theoretic sense [4], and that sincere strategies for at least two candidates are always admissible if voters are neither allowed to approve of everybody nor to disapprove of everybody (i.e., if we require voters v to have only AV strategies S_v with $\emptyset \neq S_v \neq C$), a convention adopted by Brams and Sanver [8] and also adopted here.[3] Henceforth, we will tacitly assume that only sincere AV strategy profiles are considered (which by the above convention, whenever there are at least two candidates,[4] necessarily are admissible), i.e., a vote with an insincere strategy will be considered void.

Preferences are represented by a left-to-right ranking (separated by a space) of the candidates (e.g., $a\ \ b\ \ c$), with the leftmost candidate being the most preferred one, and approval strategies are denoted by inserting a straight line into such a ranking, where all candidates left of this line are approved and all candidates right of this line are disapproved (e.g., "$a\ \mid\ b\ \ c$" means that a is approved, while both b and c are disapproved). In our constructions, we sometimes also insert a subset $B \subseteq C$ into such approval rankings, where we assume some arbitrary, fixed order of the candidates in B (e.g., "$a\ \mid\ B\ \ c$" means that a is approved, while all $b \in B$ and c are disapproved).

2.2 Control Problems for Preference-Based Approval Voting

The control problems considered here were introduced by Bartholdi, Tovey, and Trick [3] for constructive control and by Hemaspaandra, Hemaspaandra, and Rothe [22] for destructive control. In constructive control scenarios the chair's goal is to make a favorite candidate win, and in destructive control scenarios the chair's goal is to ensure that a despised candidate does not win. As is common, the chair is assumed to have complete knowledge of the voters' preference rankings and approval strategies (see [22] for a detailed discussion of this assumption), and as in most papers on electoral control (exceptions are, e.g., [27, 19]) we define the control problems in the unique-winner model.

To achieve his or her goal, the chair modifies the structure of a given election via adding/deleting/partitioning either candidates or voters. Such control actions—specifically,

[3]Brams and Sanver [8] actually preclude only the case $S_v = C$ for voters v. However, an AV strategy that disapproves of all candidates obviously is sincere, yet not admissible, which is why we also exclude the case of $S_v = \emptyset$.

[4]Note that an AV strategy is never admissible for less than two candidates. We mention in passing that a precursor of this paper [13] specifically required for single-candidate elections that each voter must approve of this candidate. In this version of the paper, we drop this requirement for two reasons. First, it in fact is not needed because the one candidate in a single-candidate election will always win—even with zero approvals (i.e., SP-AV is a "voiced" voting system). Second, it is very well comprehensible that a voter, when given just a single candidate (think, for example, of an "election" in the Eastern bloc before 1989), can get some satisfaction from denying this candidate his or her approval, even if he or she knows that this disapproval won't prevent the candidate from winning.

those with respect to control via deleting or partitioning candidates or via partitioning voters—may have an undesirable impact on the resulting election in that they might violate our conventions about admissible AV strategies. That is why we define the following rule that preserves (or re-enforces) our conventions under such control actions:

> Whenever during or after a control action it happens that we obtain an election (C, V) with $\|C\| \geq 2$ and for some voter $v \in V$ we have $S_v = \emptyset$ or $S_v = C$, then each such voter's AV strategy is changed to approve of his or her top candidate and to disapprove of his or her bottom candidate. This rule re-enforces $\emptyset \neq S_v \neq C$ for each $v \in V$.

We now formally define those of our control problems that are relevant for the proofs we give later; for the definition of the remaining control problems, see the full version [14]. Each problem is defined by stating the problem instance together with two questions, one for the constructive and one for the destructive case. These control problems are tailored to sincere-strategy preference-based approval voting by requiring every election occuring in these control problems (be it before, during, or after a control action—so, in particular, this also applies to the subelections in the partitioning cases) to have a sincere AV strategy profile and to satisfy the above conventions and rules. In particular, this means that when the number of candidates is reduced (due to deleting candidates or partitioning candidates or voters), approval lines may have to be moved in accordance with the above rules.

2.2.1 Control by Adding Candidates

In this control scenario, the chair seeks to reach his or her goal by adding to the election, which originally involves only "qualified" candidates, some new candidates who are chosen from a given pool of spoiler candidates. In their study of control for approval voting, Hemaspaandra et al. [22] considered only the case of adding an *unlimited* number of spoiler candidates (which is the original variant of this problem as defined by Bartholdi et al. [3]). We consider the same variant of this problem here to make our results comparable with those established in [22], but for completeness we in addition consider the case of adding a *limited* number of spoiler candidates, where the prespecified limit is part of the problem instance. This variant of this problem was introduced by Faliszewski et al. [18, 19] in analogy with the definitions of control by deleting candidates and of control by adding or deleting voters. They showed that, for the election system Copeland$^\alpha$ they investigate, the complexity of these two problems can drastically change depending on the parameter α, see [19].

Name: Control by Adding an Unlimited Number of Candidates.
Instance: An election $(C \cup D, V)$ and a designated candidate $c \in C$, where the set C of qualified candidates and the set D of spoiler candidates are disjoint.
Question (constructive): Is it possible to choose a subset $D' \subseteq D$ such that c is the unique winner of election $(C \cup D', V)$?
Question (destructive): Is it possible to choose a subset $D' \subseteq D$ such that c is not a unique winner of election $(C \cup D', V)$?

The problem Control by Adding a Limited Number of Candidates is defined analogously, with the only difference being that the chair seeks to reach his or her goal by adding at most k spoiler candidates, where k is part of the problem instance.

2.2.2 Control by Deleting Candidates

In this control scenario, the chair seeks to reach his or her goal by deleting (up to a given number of) candidates. Here it may happen that our conventions are violated by the control action, but will be re-enforced by the above rules (namely, by moving the line between some voter's acceptable and inacceptable candidates to behind the top candidate or to before the bottom candidate whenever necessary).

Name: Control by Deleting Candidates.
Instance: An election (C, V), a designated candidate $c \in C$, and a nonnegative integer k.
Question (constructive): Is it possible to delete up to k candidates from C such that c is the unique winner of the resulting election?
Question (destructive): Is it possible to delete up to k candidates (other than c) from C such that c is not a unique winner of the resulting election?

2.3 Immunity, Susceptibility, Vulnerability, and Resistance

Definition 2.2 ([**3**]) *Let $\mathcal{E}$ be an election system and let Φ be some given type of control. $\mathcal{E}$ is said to be* immune to Φ-control *if (a) Φ is a constructive control type and it is never possible for the chair to turn a designated candidate from being not a unique winner into being the unique winner via exerting Φ-control, or (b) Φ is a destructive control type and it is never possible for the chair to turn a designated candidate from being the unique winner into being not a unique winner via exerting Φ-control. $\mathcal{E}$ is said to be* susceptible to Φ-control *if it is not immune to Φ-control. $\mathcal{E}$ is said to be* vulnerable to Φ-control *if $\mathcal{E}$ is susceptible to Φ-control and the control problem associated with Φ is solvable in polynomial time. $\mathcal{E}$ is said to be* resistant to Φ-control *if $\mathcal{E}$ is susceptible to Φ-control and the control problem associated with Φ is* NP*-hard.*

For example, approval voting is known to be immune to eight of the twelve types of candidate control considered in [22]. The proofs of these results crucially employ the links between immunity/susceptibility for various control types shown in [22] and the fact that approval voting satisfies the unique version of the Weak Axiom of Revealed Preference (denoted by Unique-WARP, see [22, 3]): If a candidate c is the unique winner in a set C of candidates, then c is the unique winner in every subset of C that includes c. In contrast with approval voting, SP-AV does not satisfy Unique-WARP, and we will see later in Section 3.2 that it indeed is susceptible to each type of control considered here.

Proposition 2.3 *Sincere-strategy preference-based approval voting does not satisfy Unique-WARP.*

3 Results for SP-AV

3.1 Overview

Theorem 3.1 below (see also Table 2) shows the complexity results regarding control of elections for SP-AV. As mentioned in the introduction, with 19 resistances and only three vulnerabilities, this system has more resistances and fewer vulnerabilities to control (for our 22 control types) than is currently known for any other natural voting system with a polynomial-time winner problem.

Theorem 3.1 *Sincere-strategy preference-based approval voting is resistant and vulnerable to our* 22 *types of control as shown in Table 2.*

Control by	SP-AV		AV	
	Constr.	Destr.	Constr.	Destr.
Adding an Unlimited Number of Candidates	R	R	I	V
Adding a Limited Number of Candidates	R	R	I	V
Deleting Candidates	R	R	V	I
Partition of Candidates	TE: R TP: R	TE: R TP: R	TE: V TP: I	TE: I TP: I
Run-off Partition of Candidates	TE: R TP: R	TE: R TP: R	TE: V TP: I	TE: I TP: I
Adding Voters	R	V	R	V
Deleting Voters	R	V	R	V
Partition of Voters	TE: R TP: R	TE: V TP: R	TE: R TP: R	TE: V TP: V

Table 2: Overview of results. Key: I means immune, R means resistant, V means vulnerable, TE means ties-eliminate, and TP means ties-promote. Results for SP-AV are new. Results for AV, stated here to allow comparison, are due to Hemaspaandra, Hemaspaandra, and Rothe [22]. (The results for control by adding a limited number of candidates for approval voting, though not stated explicitly in [22], follow immediately from the proofs of the corresponding results for the "unlimited" variant of the problem.)

3.2 Susceptibility

By definition, all resistance and vulnerability results in particular require susceptibility. The following two lemmas (the proofs of which can be found in the full version [14]) show that SP-AV is susceptible to the 22 types of control we consider.

Lemma 3.2 *SP-AV is susceptible to constructive and destructive control by adding candidates (in both the "limited" and the "unlimited" variant of the problem), by deleting candidates, and by partition of candidates (with or without run-off and for each in both tie-handling models, TE and TP).*

Lemma 3.3 *SP-AV is susceptible to constructive and destructive control by adding voters, by deleting voters, and by partition of voters in both tie-handling models, TE and TP.*

3.3 Candidate Control

Theorems 3.4, 3.5, and 3.6 below show that sincere-strategy preference-based approval voting is fully resistant to candidate control. This result should be contrasted with that of Hemaspaandra, Hemaspaandra, and Rothe [22], who proved immunity and vulnerability for all cases of candidate control within approval voting (see Table 2). In fact, SP-AV has the same resistances to candidate control as plurality, and we note that the construction presented in [22] to prove plurality resistant also works for SP-AV in all cases of candidate control mentioned in Theorem 3.4.

All resistance results in this section follow via a reduction from the NP-complete problem Hitting Set (see, e.g., Garey and Johnson [20]): Given a set $B = \{b_1, b_2, \ldots, b_m\}$, a collection $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ of subsets $S_i \subseteq B$, and a positive integer $k \leq m$, 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, $S_i \cap B' \neq \emptyset$?

Some of our proofs use constructions and arguments for SP-AV that are straightforward modifications of the constructions and arguments of the corresponding results for approval voting or plurality from [22], whereas some other of our results require new insights to make the proof work for SP-AV. For completeness, we also state the results for SP-AV that follow by straightforward modifications of known constructions for approval or plurality, attributing them to Hemaspaandra et al. [22] (such as Theorem 3.4 below). Proof sketches of these results (explicitly showing the modifications needed to make the proofs work for SP-AV) can be found in the full version [14]. Results that are not explicitly attributed to Hemaspaandra et al. [22] use novel constructions or arguments specific to SP-AV.

Theorem 3.4 ([22]) *SP-AV is resistant to all types of constructive and destructive candidate control considered here except for (a) constructive control by deleting candidates and (b) constructive and destructive control by adding a limited number of candidates.*

We will explain in more detail below why case (a) in Theorem 3.4 is missing, and we will show as Theorem 3.5 that resistance holds for this case (a). Case (b) is missing in Theorem 3.4 simply because this control type ("adding a limited number of candidates") has not been considered in [22], but we will establish resistance for this missing case (b) as Theorem 3.6 below, via modifying our reduction presented in the proof of Theorem 3.5.

As to the missing case (a) mentioned in Theorem 3.4 above: Why does the construction that works for plurality (see [22] and also the full version [14]) not work to show that SP-AV is resistant to constructive control by deleting candidates? Informally put, the reason is that candidate c is the only serious rival of candidate w in the election (C, V) defined in Construction 4.28 of [22] (see also Construction 3.5 in [14]), so by simply deleting c the chair could make w the unique SP-AV winner, regardless of whether $\mathcal{S}$ has a hitting set of size k. However, via a different construction, we can prove resistance also in this case.

Theorem 3.5 *SP-AV is resistant to constructive control by deleting candidates.*

Proof. Susceptibility holds by Lemma 3.2. To prove resistance, we provide a reduction from Hitting Set. Let $(B, \mathcal{S}, k)$ be a given instance of Hitting Set, where $B = \{b_1, b_2, \ldots, b_m\}$ is a set, $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ is a collection of subsets $S_i \subseteq B$, and $k < m$ is a positive integer.[5]

Define the election (C, V), where $C = B \cup \{w\}$ is the candidate set and where V consists of the following $4n(k+1) + 4m - 2k + 3$ voters:

1. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form: $S_i \mid (B - S_i)\ w$.
2. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form: $(B - S_i)\ w \mid S_i$.
3. For each j, $1 \leq j \leq m$, there are two voters of the form: $b_j \mid w\ (B - \{b_j\})$.
4. There are $2(m-k)$ voters of the form: $B \mid w$.
5. There are three voters of the form: $w \mid B$.

Since for each $b_j \in B$, the difference

$$score_{(C,V)}(w) - score_{(C,V)}(b_j) = 2n(k+1) + 3 - (2n(k+1) + 2 + 2(m-k)) = 1 - 2(m-k)$$

is negative (due to $k < m$), w loses to each member of B and so does not win election (C, V).

We claim that $\mathcal{S}$ has a hitting set B' of size k if and only if w can be made the unique SP-AV winner by deleting at most $m - k$ candidates.

[5]Note that if $k = m$ then B is always a hitting set of size at most k (provided that $\mathcal{S}$ contains only nonempty sets—a requirement that doesn't affect the NP-completeness of the problem), and we thus may require that $k < m$.

From left to right: Suppose $\mathcal{S}$ has a hitting set B' of size k. Then, for each $b_j \in B'$,

$$\begin{aligned} & score_{(B' \cup \{w\},V)}(w) - score_{(B' \cup \{w\},V)}(b_j) \\ = \;& 2n(k+1) + 2(m-k) + 3 - (2n(k+1) + 2 + 2(m-k)) \;=\; 1, \end{aligned}$$

since the approval line is moved for $2(m-k)$ voters of the third group, thus transferring their approvals from members of $B - B'$ to w. So w is the unique SP-AV winner of election $(B' \cup \{w\}, V)$. Since $B' \cup \{w\} = C - (B - B')$, it follows from $\|B\| = m$ and $\|B'\| = k$ that deleting $m-k$ candidates from C makes w the unique SP-AV winner.

From right to left: Let $D \subseteq B$ be any set such that $\|D\| \leq m-k$ and w is the unique SP-AV winner of election $(C-D, V)$. Let $B' = (C-D) - \{w\}$. Note that $B' \subseteq B$ and that we have the following scores in $(B' \cup \{w\}, V)$:

$$\begin{aligned} score_{(B' \cup \{w\},V)}(w) &= 2(n-\ell)(k+1) + 2(m - \|B'\|) + 3, \\ score_{(B' \cup \{w\},V)}(b_j) &\leq 2n(k+1) + 2(k+1)\ell + 2 + 2(m-k) \quad \text{for each } b_j \in B', \end{aligned}$$

where ℓ is the number of sets $S_i \in \mathcal{S}$ that are not hit by B', i.e., $B' \cap S_i = \emptyset$. Since w is the unique SP-AV winner of $(B' \cup \{w\}, V)$, w has more approvals than any candidate b_j in B':

$$\begin{aligned} & score_{(B' \cup \{w\},V)}(w) - score_{(B' \cup \{w\},V)}(b_j) \\ \geq \;& 2(n-\ell)(k+1) + 2(m - \|B'\|) + 3 - 2n(k+1) - 2\ell(k+1) - 2 - 2(m-k) \\ = \;& 1 + 2(k - \|B'\|) - 4\ell(k+1) \;>\; 0. \end{aligned}$$

Solving this inequality for ℓ, we obtain $0 \leq \ell < \frac{1+2(k-\|B'\|)}{4(k+1)} < \frac{4+4k}{4(k+1)} = 1$. Thus $\ell = 0$. It follows that $1 + 2(k - \|B'\|) > 0$, which implies $\|B'\| \leq k$. Thus, B' is a hitting set of size at most k. ❑

Now consider the missing case (b) in Theorem 3.4. Control by adding a limited number of candidates has not been considered in [22], but we now prove resistance in this case by modifying the construction of the previous proof.

Theorem 3.6 *SP-AV is resistant to constructive and destructive control by adding a limited number of candidates.*

Proof. Susceptibility holds by Lemma 3.2 in both the constructive and destructive case.

To prove resistance in the constructive case, we slightly modify the construction presented in the proof of Theorem 3.5 that provides a reduction from Hitting Set. Let $(B, \mathcal{S}, k)$ be a given instance of Hitting Set, where $B = \{b_1, b_2, \ldots, b_m\}$ is a set, $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ is a collection of subsets $S_i \subseteq B$, and $k < m$ is a positive integer (see Footnote 5 for why $k < m$ may be assumed).

We define the election (C, V) as in the proof of Theorem 3.5, except that we introduce a new candidate, a, and insert a into the preference lists of the voters. So $C = B \cup \{a, w\}$ is now the candidate set and V consists now of the following $4n(k+1) + 4m - 2k + 3$ voters:

1. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form: $S_i \mid (B - S_i) \;\; a \;\; w$.
2. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form: $(B - S_i) \;\; a \;\; w \mid S_i$.
3. For each j, $1 \leq j \leq m$, there are two voters of the form: $b_j \mid w \;\; (B - \{b_j\}) \;\; a$.
4. There are $2(m-k)$ voters of the form: $B \mid a \;\; w$.
5. There are three voters of the form: $w \mid a \;\; B$.

Our reduction maps the Hitting Set instance $(B, \mathcal{S}, k)$ to the instance $((C' \cup B, V), w, k)$ of Constructive Control by Adding a Limited Number of Candidates, where $C' = \{a, w\}$ is the set of qualified candidates, B is the set of spoiler candidates, w is the distinguished candidate, and k is the limit on the number of candidates that may be added. Note that $score_{(C',V)}(w) = 2m + 3$ and $score_{(C',V)}(a) = 4n(k+1) + 2(m-k)$. So $score_{(C',V)}(a) > score_{(C',V)}(w)$, and thus a is the unique winner of the election (C', V). However, by an argument analogous to that given in the proof of Theorem 3.5, we can show that $\mathcal{S}$ has a hitting set of size k if and only if w can be made the unique SP-AV winner by adding at most k candidates from B.

To prove resistance in the destructive case, we modify the above construction as follows: We map our Hitting Set instance $(B, \mathcal{S}, k)$ as above to the instance $((C' \cup B, V), w, k)$ of Destructive Control by Adding a Limited Number of Candidates, where the election (C, V) has the same candidate set $C = B \cup \{a, w\}$ as above, but V now consists of the following $4n(k+1) + 4m - 2k + 2$ voters:

1. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form: $S_i \mid w \ (B - S_i) \ a$.
2. For each i, $1 \leq i \leq n$, there are $2(k+1)$ voters of the form: $(B - S_i) \ w \ a \mid S_i$.
3. For each j, $1 \leq j \leq m$, there are two voters of the form: $b_j \mid w \ (B - \{b_j\}) \ a$.
4. There are $2(m-k)$ voters of the form: $B \mid a \ w$.
5. There are two voters of the form: $w \mid a \ B$.

That is, the destructive case differs from the constructive case as follows: a and w have switched their positions in the voters of the first two groups and there are only two instead of three voters of the form $w \mid a \ B$ in the fifth group. In particular, w is now the unique winner of election (C', V), where $C' = \{a, w\}$, since $score_{(C',V)}(w) = 4n(k+1) + 2m + 2$ and $score_{(C',V)}(a) = 2(m-k)$. Again, by an argument analogous to that given in the proof of Theorem 3.5, we can show that $\mathcal{S}$ has a hitting set of size k if and only if it can be ensured by adding at most k candidates from B that w is not the unique SP-AV winner of the resulting election. This completes the proof. ❑

3.4 Voter Control

Turning now to voter control, most of the proofs for SP-AV follow from modifications of the corresponding constructions for approval voting given in [22], except for destructive control by partition of voters in model TE. Let us first state the former results.

Theorem 3.7 ([22]) *SP-AV is resistant to constructive control by adding voters and by deleting voters, to constructive and destructive control by partition of voters in model TP, and to constructive control by partition of voters in model TE. SP-AV is vulnerable to destructive control by adding voters and by deleting voters.*

Now, we turn to destructive control by partition of voters in model TE. While our polynomial-time algorithm showing vulnerability for SP-AV in this case is based on the corresponding polynomial-time algorithm for approval voting in [22], it extends their algorithm in a nontrivial way. The proof of Theorem 3.8 can be found in the full version [14].

Theorem 3.8 *SP-AV is vulnerable to destructive control by partition of voters in model TE.*

4 Conclusions and Open Questions

We have shown that Brams and Sanver's sincere-strategy preference-based approval voting system [8] combines the resistances of approval and plurality voting to procedural control: SP-AV is resistant to 19 of the 22 previously studied types of control. On the one hand, like Copeland voting [19], SP-AV is fully resistant to constructive control, yet unlike Copeland it additionally is broadly resistant to destructive control. On the other hand, like plurality [3, 22], SP-AV is fully resistant to candidate control, yet unlike plurality it additionally is broadly resistant to voter control. Thus, for these 22 types of control, SP-AV has more resistances, by three, and fewer vulnerabilities to control than is currently known for any other natural voting system with a polynomial-time winner problem.

As a work in progress, we are currently expanding our study of SP-AV's behavior with respect to procedural control towards other areas of computational social choice. In addition, we propose as an interesting and extremely ambitious task for future work the study of SP-AV (and other voting systems as well) beyond the worst-case—as we have done here—and towards an appropriate typical-case complexity model; see, e.g., [25, 26, 9, 24, 12] for interesting results and discussion in this direction.

Acknowledgments: We thank the anonymous MFCS-2008 and COMSOC-2008 referees for their helpful comments on preliminary versions of this paper.

References

[1] J. Bartholdi III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[2] J. Bartholdi III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.

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

[4] S. Brams and P. Fishburn. Approval voting. *American Political Science Review*, 72(3):831–847, 1978.

[5] S. Brams and P. Fishburn. *Approval Voting*. Birkhäuser, Boston, 1983.

[6] S. Brams and P. Fishburn. Voting procedures. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 173–236. North-Holland, 2002.

[7] S. Brams and R. Sanver. Voting systems that combine approval and preference. In S. Brams, W. Gehrlein, and F. Roberts, editors, *The Mathematics of Preference, Choice, and Order: Essays in Honor of Peter C. Fishburn*. Springer. To appear.

[8] S. Brams and R. Sanver. Critical strategies under approval voting: Who gets ruled in and ruled out. *Electoral Studies*, 25(2):287–305, 2006.

[9] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. AAAI'06*, pages 627–634. AAAI Press, July 2006.

[10] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proc. WWW'01*, pages 613–622. ACM Press, 2001.

[11] E. Ephrati and J. Rosenschein. Multi-agent planning as a dynamic search for social consensus. In *Proc. IJCAI'93*, pages 423–429, 1993.

[12] G. Erdélyi, L. Hemaspaandra, J. Rothe, and H. Spakowski. On approximating optimal weighted lobbying, and frequency of correctness versus average-case polynomial time. In *Proc. FCT'07*, pages 300–311. Springer-Verlag LNCS #4639, August 2007.

[13] G. Erdélyi, M. Nowak, and J. Rothe. Sincere-strategy preference-based approval voting broadly resists control. In *Proc. MFCS'08*. Springer-Verlag LNCS, August 2008. To appear.

[14] G. Erdélyi, M. Nowak, and J. Rothe. Sincere-strategy preference-based approval voting fully resists constructive control and broadly resists destructive control. Technical Report cs.GT/0806.0535, ACM Computing Research Repository (CoRR), June 2008.

[15] R. Fagin, R. Kumar, and D. Sivakumar. Efficient similarity search and classification via rank aggregation. In *Proc. ACM SIGMOD Intern. Conf. on Management of Data*, pages 301–312. ACM Press, 2003.

[16] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of bribery in elections. In *Proc. AAAI'06*, pages 641–646. AAAI Press, 2006.

[17] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*. Springer. To appear.

[18] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting broadly resist bribery and control. In *Proc. AAAI'07*, pages 724–730. AAAI Press, 2007.

[19] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Copeland voting fully resists constructive control. In *Proc. AAIM'08*, pages 165–176. Springer-Verlag LNCS #5034, June 2008.

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

[21] S. Ghosh, M. Mundhe, K. Hernandez, and S. Sen. Voting for movies: The anatomy of recommender systems. In *Proc. 3rd Annual Conference on Autonomous Agents*, pages 434–435. ACM Press, 1999.

[22] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.

[23] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. In *Proc. IJCAI'07*, pages 1308–1314. AAAI Press, 2007.

[24] C. Homan and L. Hemaspaandra. Guarantees for the success frequency of an algorithm for finding Dodgson-election winners. *Journal of Heuristics.* To appear. Full version available as Department of Computer Science, University of Rochester Technical Report TR-881, September 2005, revised June 2007.

[25] J. McCabe-Dansted, G. Pritchard, and A. Slinko. Approximability of Dodgson's rule. *Social Choice and Welfare*, 2008. DOI 10.1007/s00355-007-0282-8.

[26] A. Procaccia and J. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.

[27] A. Procaccia, J. Rosenschein, and A. Zohar. Multi-winner elections: Complexity of manipulation, control, and winner-determination. In *Proc. IJCAI'07*, pages 1476–1481. AAAI Press, 2007.

Gábor Erdélyi, Markus Nowak, and Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany