
The Shield that Never Was: Societies with Single-Peaked Preferences are More Open to Manipulation and Control

Piotr Faliszewski
Dept. of Computer Science
AGH Univ. of Sci. & Techn.
30-059 Kraków, Poland

Edith Hemaspaandra
Dept. of Computer Science
Rochester Inst. of Technology
Rochester, NY 14623, USA

Lane A. Hemaspaandra
Dept. of Computer Science
University of Rochester
Rochester, NY 14627, USA

Jörg Rothe
Institut für Informatik
Universität Düsseldorf
40225 Düsseldorf, Germany

Abstract

Much work has been devoted, during the past twenty years, to using complexity to protect elections from manipulation and control. Many results have been obtained showing NP-hardness shields, and recently there has been much focus on whether such worst-case hardness protections can be bypassed by frequently correct heuristics or by approximations. This paper takes a very different approach: We argue that when electorates follow the canonical political science model of societal preferences the complexity shield never existed in the first place. In particular, we show that for electorates having single-peaked preferences, many existing NP-hardness results on manipulation and control evaporate.

1 Introduction

Elections are a broadly used preference aggregation model in both human societies and multiagent systems. For example, elections have been proposed as a mechanism for collaborative decision-making in such multiagent system contexts as recommender systems/collaborative filtering [Pennock *et al.*, 2000] and planning [Ephrati and Rosenschein, 1991; 1997]. The importance of elections explains why elections are studied intensely over a wide range of fields, including political science, mathematics, social choice, artificial intelligence, economics, and operations research. Formally, an election system takes as input a set of candidates (or alternatives) and a set of votes (usually each a 0-1 approval vector over the candidates or a linear order over the candidates) and outputs a subset of the candidates as the winner(s).

Control refers to attempts to make a given candidate win [Bartholdi *et al.*, 1992] (or not win [Hemaspaandra *et al.*, 2007a]) an election by such participation-change actions as adding or deleting voters or candidates. Manipulation refers to attempts to make a given candidate

win [Bartholdi *et al.*, 1989; Bartholdi and Orlin, 1991] (or not win [Conitzer *et al.*, 2007]) by some coalition of voters who strategically change their votes. There is a large literature, started by the insightful contributions of Bartholdi, Orlin, Tovey, and Trick [Bartholdi and Orlin, 1991; Bartholdi *et al.*, 1989; 1992], on choosing election systems that make control and manipulation NP-hard, and by doing so seek to make control and manipulation computationally prohibitive (see the survey [Faliszewski *et al.*, 2009]). Recently, there has been a flurry of work seeking to bypass worst-case manipulation hardness results by frequently correct heuristics or approximation algorithms (as a few pointers into that literature, see, e.g., [Friedgut *et al.*, 2008; Conitzer and Sandholm, 2006; Procaccia and Rosenschein, 2007; Brelsford *et al.*, 2008]).

The present paper takes a radically different approach. We study elections where the vote set must be “single-peaked.” We’ll discuss single-peakedness in more detail after stating our main contributions. But, briefly put, single-peakedness means that there is some linear ordering of the candidates relative to which (in the model in which votes are linear orders) each voter’s preferences always increase, always decrease, or first increase and then decrease, or (in the model in which votes are approval vectors) each voter’s approved candidates are contiguous within the linear order. Single-peaked preferences, introduced by Black [1948] (see also [Black, 1958]), model societies that are heavily focused on one issue (taxes, war, etc.), and the single-peaked framework is so central to political science that it has been described as “*the* canonical setting for models of political institutions” [Gailmard *et al.*, to appear], and indeed is typically the model of societal voting first covered in an introductory course on positive (i.e., theoretical) political science. This paper’s main contributions are the following:

- We introduce the study of single-peakedness for approval voting. We as Theorem 2.1 provide a polynomial-time linear programming algorithm to test single-peakedness of a collection of approval vectors and to find a linear order witnessing the single-peakedness.

Copyright is held by the author/owner(s).
TARK '09, July 6-8, 2009, California
ISBN: 978-1-60558-560-4...\$10.00

- In Section 3 we show that for both voting by approval vectors and for voting by linear orders, many election control problems known to be NP-hard in the general case have polynomial-time algorithms in the single-peaked case.
 - In Section 4 we show that many election manipulation problems known to be NP-hard in the general case have polynomial-time algorithms in the single-peaked case.
- However—and in this we are inspired by the path-setting work of Walsh [2007], who showed that Single Transferable Voting (for at least 3 candidates and weighted votes) remains NP-hard to manipulate even in the single-peaked case—we in Section 4 also show that many manipulation problems remain NP-hard even when restricted to the single-peaked case.
- We show, contrary to intuition and the tacit assumptions of some papers, that even for natural systems there are cases where increasing the number of candidates decreases the complexity. In particular, we as Theorem 4.2 show that in 3-Veto elections (i.e., one votes against three candidates and for all others) manipulation is in P for up to four candidates, is NP-hard for five candidates, and is in P for six or more candidates.

We mostly defer discussion of related work until after our results, as the related work will then have more context and definitions to draw on.

2 Preliminaries

Elections and Preferences An *election* consists of a set C of candidates and a collection V of votes. We will consider two different models for votes. One is that each vote is a vector (an *approval vector*) from $\{0, 1\}^{|C|}$, denoting approval (1) or disapproval (0) for each candidate. The other is that each vote is a linear order (by which we mean a strict, linear order—a complete, transitive, antisymmetric relation) over the candidates, e.g., Bob > Alice > David > Carol. An *election system* is a mapping that takes as input a candidate set C and a set V of votes over that candidate set, and outputs an element of 2^C , i.e., outputs which candidates are *winners* of the election. (We, like Bartholdi, Tovey, and Trick [1992], do not expressly forbid elections with no winners, although all of the many natural election systems discussed in this paper have the property that there is always at least one winner when there is at least one candidate.) Except where we explicitly state otherwise, V is a list of votes (ballots), so if three votes are the same, they will appear three times in the list. We will use *succinct input* [Faliszewski *et al.*, 2006] to describe the quite different input model in which each preference that is held by one

or more voters appears just once in the list and is accompanied by a binary number stating how many voters have that preference.

The election systems of most interest to us in this paper are the following ones. In *approval voting*, voters vote by approval vectors, and whichever candidate(s) get the most approvals are the winner(s). A *scoring protocol* election, which is always defined for a specific number m of candidates, is specified by a *scoring vector*, $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$, satisfying $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$. Votes are linear orders. Each vote contributes α_1 points to that vote's most preferred candidate, α_2 points to that vote's second most preferred candidate, and so on. And whichever candidate(s) get the most total points are the winner(s). Among the most important scoring protocols for m candi-

dates are plurality, with $\alpha = (1, \overbrace{0, \dots, 0}^{m-1})$, j -veto ($j \leq m$), with $\alpha = (\overbrace{1, \dots, 1}^{m-j}, \overbrace{0, \dots, 0}^j)$, and Borda, with $\alpha = (m-1, m-2, \dots, 0)$. For m candidates, j -approval and $(m-j)$ -veto are the same system. We'll also speak of the veto system for an unbounded number of candidates, by which we mean that on inputs with m candidates, each voter gives 1 point to all candidates other than her least favorite candidate and 0 points to her least favorite candidate. We'll similarly also speak of plurality for an unbounded number of candidates, by which we mean that on inputs with m candidates, each voter gives 0 points to all candidates other than her favorite candidate and 1 point to her favorite candidate. And the general case of approval voting is always for an unbounded number of candidates.

Single-Peaked Preferences A collection V of votes, each vote v_i being a linear order P_i over C , is said to be *single-peaked* exactly if there exists a linear order over C , call it L , such that for each triple of candidates c , d , and e , it holds that:

$$(cLdLe \vee eLdLc) \implies (\forall i) [cP_i d \implies dP_i e].$$

That is just a formal way of saying that with respect to L , each voter's degree of preference rises to a peak and then falls (or just rises or just falls). The loose intuition behind this is captured in Figure 1. If we imagine an electorate completely focused on one issue (say taxes), with each person having a single-peaked (in the natural analogous sense of that notion applied to curves) utility curve (but potentially different utility curves for different voters), one gets precisely this notion (give or take the issue of ties). In Figure 1's example, the preferences of v_1 would be $c_1 > c_2 > c_3 > c_4 > c_5$, of v_2 would be $c_3 > c_4 > c_2 > c_1 > c_5$, and of v_3 would be $c_4 > c_3 > c_2 > c_1 > c_5$.

The seminal study of single-peaked preferences was done by Black [1948], and that work and many subsequent studies (e.g., [Niemi and Wright, 1987; Davis *et al.*, 1970;

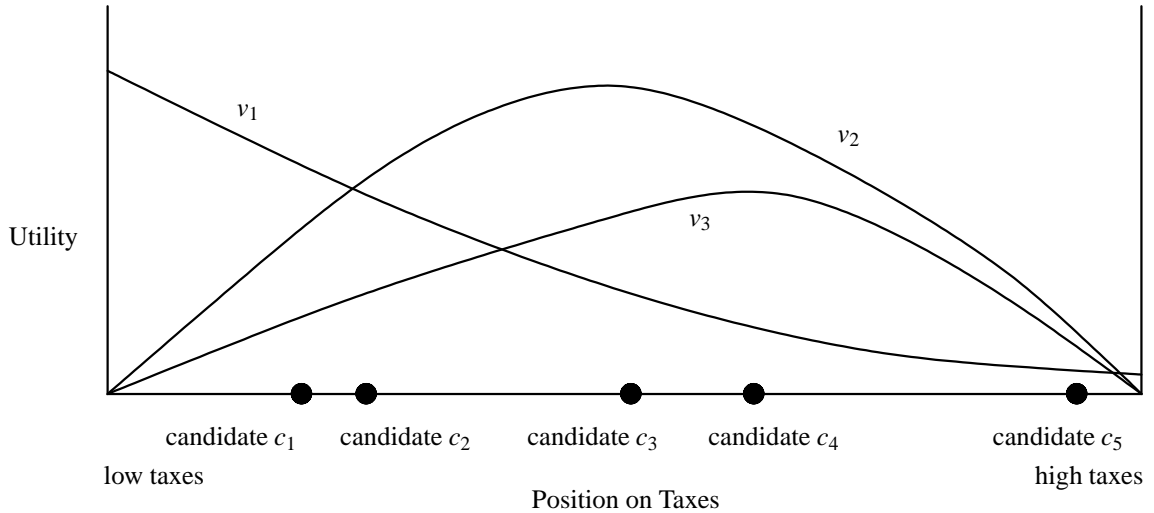


Figure 1: Single-Peaked Preference Motivation: Utility Curves for Three Voters.

Poole and Rosenthal, 1997; Krehbiel, 1998; Black, 1958]) argued that single-peaked preferences (in the unidimensional spatial model) are a broadly useful model of electoral preferences that captures many important settings.¹ Of course, issues that are multidimensional—as many issues are—are typically not captured by single-peaked preferences (by which we always mean the unidimensional case). And even in a society that is completely focused on one issue, some maverick individuals may focus on other issues, so one should keep in mind that the single-peaked case is a widely studied but extreme model.

Looking again at Figure 1, let us imagine that each voter has a utility threshold at which she starts approving of candidates. Then in the approval-vector-as-method-of-voting we'll have that for some linear order L , each voter either disapproves of everyone or approves of precisely a set of candidates that are contiguous with respect to L . Formally, we say a collection V , made up of approval vectors $v_1, v_2, \dots, v_{\|C\|}$ over the candidate set C , with Approves_i being the set of candidates that v_i approves, is *single-peaked* exactly if there is a linear order over C , call it L , such that for each triple of candidates c, d , and e , it holds that:

$$cLdLe \implies (\forall i) [\{c, e\} \subseteq \text{Approves}_i \implies d \in \text{Approves}_i].$$

The reasons that single-peaked approval voting is compellingly natural to study (and the reasons why one should not go overboard and claim that it is a universally appropriate notion) are essentially the same as those, touched on

¹We mention in passing that since with single-peaked preferences there is always a Condorcet winner when the number of voters is odd, every Condorcet-consistent voting rule has an efficient winner-determination algorithm for odd numbers of voters. Of course, many interesting voting rules are not Condorcet-consistent. (Also, see Peleg and Sudhölter [1999] for discussion on the relationship between single-peakedness and strategy-proofness for certain voting contexts.)

above, involving single-peaked voting with respect to preference orders. In many cases it is natural to assume that there is some unidimensional issue steering society and that each person's range of comfort on that issue is some contiguous segment along that issue's dimension. In such a case, each person's set of approved candidates will form a contiguous block among the candidates when they are ordered by their positions on that issue. To the best of our knowledge, this is the first paper to study single-peaked approval vectors—although the literature is so vast that we would not be at all surprised if this notion had been previously defined.

In our control and manipulation problems, we'll for the single-peaked case follow Walsh [2007] and take as part of the input a particular linear order of the candidates relative to which the votes are single-peaked. This is arguably natural—as we may view candidates' positions on the issue that defines the entire election as being openly known. However, one may wonder how hard it is to, given a set of voters, tell whether it is single-peaked. For the case of votes being linear orders, Bartholdi and Trick [1986] (see also [Escoffier *et al.*, 2008]; a somewhat related paper is [Trick, 1989]) show by a path-based graph algorithm that the problem is in P, and Escoffier, Lang, and Öztürk [2008] show how to produce in polynomial time a linear order witnessing the single-peakedness when such an ordering exists. For the case of votes being approval vectors, by a very different approach—linear programming—we have the following result, which provides the analogue of both of the results just mentioned.

Theorem 2.1 *Given a collection V of approval vectors over C , in polynomial time we can produce a linear order L witnessing V 's single-peakedness or can determine that V is not single-peaked.*

We give a very brief, telegraphic sketch of the easy, linear-programming-based proof of Theorem 2.1.

Proof Sketch for Theorem 2.1. Recall that in this theorem we are in the approval-vector model of voting (the following statement actually can be seen not to hold in the vote-by-linear-order model). It is not hard to see that V is single-peaked if and only if the candidates in C can be assigned, one each in some potentially permuted order, to the points $1, 2, \dots, \|C\|$ on the real line and each voter v_i can be assigned to an integer point or an “integer-plus-1/2” point and given a threshold t_i such that v_i approves of those candidates whose distance (absolute value) from v_i ’s point is less than or equal to t_i . (As to the “only if” direction, given any single-peaked ordering, put the candidates in that order at real-line positions $1, 2, \dots, \|C\|$ and put each voter who approves of, say, exactly the candidates located from position j to position k , $j \leq k$, at the point $\frac{j+k}{2}$, and put voters who disapprove of all the candidates at, say, location 0 on the real line.) We mentioned this mostly to give some intuition about single-peaked approval. Relatedly, and more to the point for this proof, note that clearly V is single-peaked if and only if there exists some geometric embedding of this sort (in which we no longer require that the candidate/voter points be integer and “integer-plus-1/2” points).

Now let us in light of the geometric embedding observation of the second of the above “if and only if” statements see how to, given the votes, find whether there is a single-peaked societal ordering and, if so, how to produce one. We build a linear program constraint satisfaction problem that will for each v_i and each pair (c, d) where v_i approves of c and disapproves of d have a constraint of the form (where the “ Loc ” are location-on-the-real-line variables) $|Loc_{v_i} - Loc_c| \leq |Loc_{v_i} - Loc_d| - 0.01$. The “0.01” is a trick to speak of “ $<$ ” in effect; using a standard linear programming trick our set of equations can be transformed in polynomial time to a linear programming optimization problem without absolute-value operations; and with some discussion one can see that if the linear program assigns multiple candidates to the same location we can arbitrarily break that tie in stating our order. This linear program will give locations for the candidates that in effect give us a single-peaked ordering when one exists. \square

Note that Theorem 2.1 does not seem to naturally follow from the analogous results for linear orders, as approval vectors contain much less information. More to the point, a collection of non-single-peaked linear orders may, in some sense, correspond to a collection of single-peaked approval vectors. To see this, we point the reader to Lemma 1 of [Escoffier *et al.*, 2008], which says that in a set of single-peaked linear preferences at most two candidates are ranked last (that is, the set of candidates that appear on the last position of some voter’s preference or-

der contains at most two elements). However, a collection of single-peaked approval vectors might contain more than two candidates that are not approved by some voter. Thus, a trivial completion of approval vectors to form linear orders would not necessarily work.

Control and Manipulation Problems Due to space we describe these very briefly. For an election system \mathcal{E} , the *Constructive Control by Adding Candidates* problem is the set of all (C, V, p, k, C') , where V consists of votes over $C \cup C'$, $p \in C$, and $C \cap C' = \emptyset$, such that there is a set $C'' \subseteq C'$ with $\|C''\| \leq k$ for which candidate p is the unique winner under election system \mathcal{E} when the voters in V vote over $C \cup C''$ (i.e., we restrict each voter down to her induced preferences over $C \cup C''$). The *Destructive Control by Adding Candidates* problem is the same except now the question is whether there is such a C'' ensuring that p is *not* a unique winner.

The Constructive/Destructive Control by Deleting Candidates problems are analogous, with inputs of the form (C, V, p, k) , where k is a limit on how many candidates from C can be deleted, and it is forbidden to delete p . The Constructive/Destructive Control by Adding/Deleting Voters problems are analogous, with inputs respectively (C, V, p, k, V') and (C, V, p, k) , where V' , $V' \cap V = \emptyset$, is a pool of “unregistered” voters, and k denotes the bound on how many voters from V' we can add (Adding case) or how many voters from V we can delete (Deleting case).

Constructive/Destructive Control by Unlimited Adding Candidates is the same as Constructive/Destructive Control by Adding Candidates except there is no “ k ”; one may add any or all of the members of C' .

Most of these problems were introduced in the seminal control paper of Bartholdi, Tovey, and Trick [1992], and the remaining ones were introduced in [Hemaspaandra *et al.*, 2007a; Faliszewski *et al.*, 2007]. These problems model such real-world problems as introducing a candidate to run to split off another candidate’s support; urging an independent candidate to withdraw for the good of the country; spreading rumors that people with outstanding warrants who try to vote will be arrested; and sending vans to retirement homes to drive car-less members of one’s party to the voting place.

In each case, the inputs we specified are for the “general” case. For the single-peaked case there will be an additional input, a linear order L over the candidates such that relative to L the election is single-peaked (with respect to all voters and candidates—even those in C' and in V' in problems having those extra sets). (If L is not such a linear order, the input is immediately rejected. L ’s goodness can be easily tested in polynomial time, simply by looking at each vote.) In assuming that L is given, we are following Walsh’s [2007] natural model— L is the broadly known po-

sitioning of the candidates.²

The (constructive coalition weighted) manipulation problem was introduced by Conitzer, Sandholm, and Lang ([Conitzer *et al.*, 2007] and its conference precursors, building on the seminal manipulation papers [Bartholdi and Orlin, 1991; Bartholdi *et al.*, 1989]) and takes as input a list of candidates C , a list of nonmanipulative voters each specified by preferences (as a linear order or as an approval vector, depending on our election system) over C and an integer weight, a list of the weights of the voters in our manipulating coalition, and the candidate p that our coalition seeks to make a winner. The set of all such inputs for which there is some assignment of preferences to the manipulators that makes p a winner is the *Constructive Coalition Weighted Manipulation* problem for that election system.

We follow the convention from [Bartholdi *et al.*, 1989; Bartholdi and Orlin, 1991; Bartholdi *et al.*, 1992] of focusing on the unique-winner case for control and the winner case (i.e., asking whether the candidate can be made a winner or be prevented from being a winner; following the literature, we'll sometimes refer to that as the “nonunique-winner model/case”) for manipulation. In many cases we've shown a given result in both models and will sometimes mention that in passing. All of the results of Conitzer, Sandholm, and Lang of relevance to us hold in both models (see Footnote 7 of [Conitzer *et al.*, 2007]).

For the single-peaked case, a linear order L is given as part of the input and the manipulating coalition's votes must all be single-peaked with respect to L (as must all other voters) for the input to be accepted.

Complexity Notions for Control and Manipulation Problems If by a given type of constructive control action we can never change p from not unique winning (not winning) to unique winning (winning) we say the problem is *immune* to that control type, and otherwise we call it *susceptible* to that control type. The destructive case is analogous. If a susceptible problem is in P we call it

²Both for linear-order votes ([Escoffier *et al.*, 2008], see also [Bartholdi and Trick, 1986]) and for approval-vector votes (Theorem 2.1) it holds that even if L is not given one can find a good L in polynomial time if one exists. This fact may be comforting to those who would prefer that L not be given. But we caution that a given vote set V may have many valid L 's. And for some problems, which valid L one uses may affect the problem and its complexity. Indeed, not having L as part of the input might even open the door to time-sequence issues, e.g., in deletion of candidates/voters, should we ask instead whether single-peakedness should only have to hold after the deletion? (We'd say, ideally, “no.”) These issues are reasonable for further control/manipulation-related study (and the issue of whether L is known *a priori* is much studied in the political economy literature already, see, e.g., [Austen-Smith and Banks, 2004, Section 2.4]), but we find the [Walsh, 2007] model to be most natural and compelling. In many cases our proofs work fine if L is not part of the input, and at times we'll mention that.

vulnerable, and if it is NP-hard we call it *resistant*. A manipulation problem (such as the constructive coalition weighted manipulation problem defined above) is said to be *vulnerable* if it is in P and *resistant* if it is NP-hard. In each control/manipulation case in this paper where we assert vulnerability (or membership in P), not only is the decision problem in P but also in polynomial time we can produce a successful control/manipulation action if one exists (i.e., the problem is what [Hemaspaandra *et al.*, 2007a] calls *certifiably vulnerable*). Most of these notions are in detail from or in the general spirit of the work of Bartholdi, Orlin, Tovey, and Trick [Bartholdi *et al.*, 1992; 1989; Bartholdi and Orlin, 1991], as in some cases naturally modified or extended in [Hemaspaandra *et al.*, 2007b; 2007a].

3 Control

We now turn to our results on control of elections. The theme of this paper is that electorates limited to being single-peaked often are simpler to control and manipulate. Intuitively speaking, we will show that the more limited range of vote collections allowed by single-peaked voting (as opposed to general voting) is so restrictive that the reductions showing NP-hardness fall apart—that those reductions centrally need complex collections of votes to work. Of course, for all we know, perhaps $P = NP$, and so attempting to show that no reduction at all can exist from SAT to our single-peaked control problems would be a fool's errand (or a Turing Award-scale quest). Rather, we'll show our single-peaked control problems easy the old-fashioned way: We'll prove that they are in P .

The techniques we use to prove them in P vary from easy observations to smart greedy schemes to dynamic programming. For reasons of space, in this extended abstract we (in addition to the proof sketch of Theorem 2.1) present only the proof of one result (Theorem 4.2); complete proofs of all our results will be made available in the full version.

The control complexity of approval voting is studied in detail by Hemaspaandra, Hemaspaandra, and Rothe [2007a], and see regarding limited adding of candidates [Erdélyi *et al.*, 2008] (see also the survey [Baumeister *et al.*, to appear]). Among all the adding/deleting control cases defined in this paper (ten in total), precisely two are resistant, and the other eight (all five destructive cases and the three constructive candidate cases) are immune or vulnerable. The two resistant cases are Constructive Control by Adding Voters and Constructive Control by Deleting Voters.

The following theorem shows that both of these complexity shields evaporate for societies with single-peaked preferences.

Theorem 3.1 *For the single-peaked case, approval vot-*

ing is vulnerable to constructive control by adding voters and to constructive control by deleting voters, in both the unique-winner model and the nonunique-winner model, for both the standard input model and the succinct input model.³

Briefly put, the challenge here is that the set V' of voters to add may be filled with “incomparable” voter pairs—pairs such that regarding their interval with respect to the linear order defining single-peakedness, neither is a subset of the other—and so it is not immediately obvious what voters to add. We solve this by a “smart greedy” approach, breaking votes first into broad groups based on where their intervals’ right edges fall with respect to just a certain “dangerous” subset of the candidates, and then re-sort those based on their left edges, and we argue that if any strategy will reach the control goal then this one will.

We turn from approval voting to plurality voting. Plurality voting’s constructive control complexity was studied in detail in the seminal control paper of Bartholdi, Tovey, and Trick [1992], and the destructive cases were added by [Hemaspaandra *et al.*, 2007a] (see [Faliszewski *et al.*, 2007] for the limited adding of candidates cases). For plurality, the situation is close to backwards from approval. Regarding all the adding/deleting control cases defined in this paper, the four voter cases are vulnerable but all six candidate cases are resistant. However, our results here again are in keeping with our paper’s theme: All six of these cases become vulnerable for single-peaked societies.

Theorem 3.2 *For the single-peaked case, plurality voting is vulnerable to constructive and destructive control by adding candidates, by adding unlimited candidates, and by deleting candidates, and these results hold in both the unique-winner model and the nonunique-winner model.*

Our proofs of the different parts of this theorem vary greatly in approach. One approach that is particularly useful is dynamic programming.

For two very important voting systems—plurality and approval—we have seen that in every single adding/deleting case where they are known to have NP-hardness complexity shields, the complexity shield evaporates for societies with single-peaked preferences. In the coming manipulation section we will also see a number of cases where NP-hardness shields melt away for single-peaked societies, but we will also see some cases—and earlier such a case was found by Walsh [2007]—where

³This result holds in our settled model in which the linear order specifying the society’s order on the candidates is part of the input, and also holds in the model—which due to space we won’t really cover here—in which the linear order is not part of the input but rather the question is whether there exists any valid linear order relative to which there is a way of achieving our control goal.

existing NP-hardness shields remain in place even if one adds the restriction to a single-peaked society.

The previous paragraph raises the following natural question: Given that restricting to single-peaked preferences can sometimes remove complexity shields, e.g., Theorem 3.1, and can sometimes leave them in place, e.g., Theorem 4.3, can restricting to single-peaked preferences ever erect a complexity shield? It would be very tempting to hastily state that such behavior is impossible. After all, the single-peaked case if anything thins out the flood of possible control-actions/manipulations. But we’re talking about complexity here, and fewer options doesn’t always mean a less complex problem. In particular, one of those manipulation/control-action options that single-peakedness takes off the table might have been a single, sure-fire path to victory under some election system. In fact, using precisely that approach, and a few other tricks, we have built an artificial election system for which unweighted constructive manipulation by size-3 coalitions is in P in the general case but is NP-complete for the single-peaked case. The system’s votes are approval vectors. The system is highly unnatural,⁴ highly unsatisfying, and would never be considered for real-world use. However, its goal is just to show that (in a stomach-turningly contrived way) restricting to single-peakedness in concept can, perhaps surprisingly, raise complexity. We conjecture that that strange behavior will never be seen for any existing, natural election system. We summarize this paragraph in the following (manipulation) theorem.

Theorem 3.3 *There exists an election system \mathcal{E} , whose votes are approval vectors, for which constructive size-3-coalition unweighted manipulation is in P for the general case but is NP-complete in the single-peaked model.*

4 Manipulation

In this section we study constructive coalition weighted manipulation. Recall that in the single-peaked case the manipulators must cast votes that are consistent with the linear ordering of the candidates (which is part of the input in the single-peaked case) that defines the society’s single-peakedness. However, all our “single-peaked case is in P” results in this section also hold in a model in which the linear order of the candidates is not part of the input.

This paper’s theme is that single-peakedness removes many NP-hardness shields, and for manipulation we show that

⁴The election system first uses the P algorithm from Theorem 2.1 to see if there is a linear order making the votes single-peaked and if not the system has all candidates win. This feature of the system allows a manipulating coalition to in the general (single-peakedness not required) case cast votes precluding single-peakedness, an attack the single-peaked case does not allow, and by other tricks we ensure that the single-peakedness case will be NP-hard.

via the following theorems. (Note that in the general case, the election systems of parts 1 and 3 of this theorem and the “ $k_1 \geq 2 \wedge k_0 \geq 1$ ” cases of part 2 of this theorem are known to be NP-complete [Hemaspaandra and Hemaspaandra, 2007; Procaccia and Rosenschein, 2007; Conitzer *et al.*, 2007], and the remaining part 2 cases are well known and easily seen to be in P.)

Theorem 4.1 *For the single-peaked case, the constructive coalition weighted manipulation problem (in both the nonunique-winner model and the unique-winner model) for each of the following election systems is in P:*

1. *The scoring protocol $\alpha = (2, 1, 0)$, i.e., 3-candidate Borda elections.*

2. *Each scoring protocol $\alpha = (\overbrace{1, \dots, 1}^{k_1}, \overbrace{0, \dots, 0}^{k_0})$, $k_1 \geq k_0$. (This includes a variety of ℓ -veto and ℓ -approval protocols, e.g., the 3-veto cases for $m \geq 6$ candidates in Theorem 4.2.)*

3. *Veto.*

We now come to an unusual case. For general votes (not limited to single-peaked societies), 3-veto is in P for three or four candidates, and 3-veto is NP-complete (and so resistant) for five or more candidates (see [Hemaspaandra and Hemaspaandra, 2007]). (3-veto is not meaningfully defined for two or fewer candidates.) However, for single-peaked votes, 3-veto shows a remarkable behavior: Moving from five to six candidates *lowers* the complexity. In some papers, authors state that if one knows the number of candidates (if any) at which a system switches from easy to hard, it is easy by adding dummy candidates to see that for all larger number of candidates the system remains hard. The following theorem should stand as a caution to take that view only if one has carefully built and checked a “dummy candidates” construction for one’s specific case.

We present the proof of Theorem 4.2 and point out that the same proof ideas can be used and generalized to prove some of the other theorems of this paper.

Theorem 4.2 *For the single-peaked case, the constructive coalition weighted manipulation problem for m -candidate 3-veto elections is in P for $m \in \{3, 4, 6, 7, 8, \dots\}$ and is resistant (indeed, NP-complete) for $m = 5$.*

Proof. Consider the following five cases.

$m = 3$. In this case, we are looking at the scoring protocol $(0, 0, 0)$. It is immediate that in this scoring protocol all candidates are always tied for winner, and so p can always be made a winner.

$m = 4$. In this case, we are looking at the scoring protocol $(1, 0, 0, 0)$. It is immediate that p can be made a winner

if and only if p is a winner in the election where every manipulator ranks p first.

$m = 5$. In this case, we are looking at the scoring protocol $(1, 1, 0, 0, 0)$. It is immediate that the constructive coalition weighted manipulation problem is in NP (simply guess single-peaked votes for the manipulators and verify that p is a winner of the election). To show NP-hardness, we reduce from the well-known NP-complete problem PARTITION: Given a set $\{k_1, \dots, k_n\}$ of n distinct positive integers that sum to $2K$, does there exist a subset that sums to K ? This problem was also used to prove general (i.e., not required to be single-peaked) constructive coalition weighted manipulation problems NP-hard (see, e.g., [Conitzer *et al.*, 2007]).

Given an instance of PARTITION, i.e., a set of n distinct positive integers $\{k_1, \dots, k_n\}$ that sums to $2K$, construct the following instance of constructive coalition weighted manipulation: The set of nonmanipulators S consists of two voters, each of weight K . One of the voters votes $c > a > p > b > d$ and the other voter votes $d > b > p > a > c$. (Note that these two voters fix society’s order.) We also have a set T of n manipulators. The weights of the manipulators are k_1, k_2, \dots, k_n . We claim that there is a partition if and only if the manipulators can cast single-peaked votes that make p a winner.

First suppose that there exists a subset S' of $\{k_1, \dots, k_n\}$ that sums to K . For every $i \in \{1, \dots, n\}$, we set the weight k_i manipulator to $p > a > b > c > d$ if k_i in S' and to $p > b > a > c > d$ otherwise. In the resulting election, a, b , and p each score $2K$ points and c and d score K points. It follows that p is a winner of the election.

For the converse, suppose the voters in T vote single-peaked such that p is a winner of the election. For \hat{c} a candidate, we will write $score_T(\hat{c})$ for the number of points \hat{c} gets from the voters in T . Since for every single-peaked vote, if p gets a point then a or b gets a point, it follows that $score_T(p) \leq score_T(a) + score_T(b)$. Since p is a winner of the resulting election, $score_T(p) \geq K + score_T(a)$ and $score_T(p) \geq K + score_T(b)$. It follows that $score_T(p) = 2K$ and that $score_T(a) = score_T(b) = K$. But then the weights of the voters in T that give a point to a sum to K and so we have found a partition.

$m = 6$. In this case, we are looking at the scoring protocol $(1, 1, 1, 0, 0, 0)$. Consider society’s order. Without loss of generality, suppose that p is right of the middle in this order. Let A be the set of all candidates to the right of p . Note that $\|A\| \leq 2$. For every set of single-peaked votes consistent with society’s preference, and for every $a \in A$, it holds that $score(p) \geq score(a)$ in

every election. It follows that p can be made a winner if and only if p is a winner if every manipulator ranks the candidates from right to left.

$m \geq 7$. Consider society's order. Let c be an arbitrary candidate such that there are at least 3 candidates to the left of c and at least three candidates to the right of c . Note that for every single-peaked vote, candidate c scores a point. This implies that p can be made a winner if and only if p is doesn't lose any points in the set of nonmanipulators.

This completes the proof of Theorem 4.2. \square

Finally, we present some cases that are known (see [Hemaspaandra and Hemaspaandra, 2007]) to be NP-hard in the general case and that we can prove remain hard even in the single-peaked case.

Theorem 4.3 *For the single-peaked case, the constructive coalition weighted manipulation problem is resistant (indeed, NP-complete) for the following scoring protocols:*

1. $\alpha = (3, 1, 0)$.
2. $\alpha = (3, 2, 1, 0)$, i.e., 4-candidate Borda elections.

In fact, we can extend the idea behind the proofs of part 1 of Theorem 4.1 and part 1 of Theorem 4.3 to obtain the following dichotomy result for the 3-candidate case.

Theorem 4.4 *Consider a 3-candidate scoring protocol, namely, $\alpha = (\alpha_1, \alpha_2, \alpha_3)$, $\alpha_1 \geq \alpha_2 \geq \alpha_3$, $\alpha_1 \in \mathbb{N}$, $\alpha_2 \in \mathbb{N}$, $\alpha_3 \in \mathbb{N}$. For the single-peaked case, the constructive coalition weighted manipulation problem is resistant (indeed, NP-complete) when $\alpha_1 - \alpha_3 > 2(\alpha_2 - \alpha_3) > 0$ and is in P otherwise.*

Finally, we mention that despite the NP-hardness result of part 2 of Theorem 4.3, for the single-peaked 4-candidate Borda case there is a polynomial-time manipulation algorithm for the case when the candidate the coalition wants to win is either the top or the bottom candidate in society's input linear order on the candidates.

5 Related Work

The paper that inspired our work is Walsh's "Uncertainty in Preference Elicitation and Aggregation" [Walsh, 2007]. Among other things, in that paper he raises the issue of manipulation in single-peaked societies. Our paper follows his model of assuming society's linear ordering of the candidates is given and that manipulative voters must be single-peaked with respect to that ordering. However, our theme and his differ. His manipulation results present

cases where single-peakedness leaves an NP-completeness shield intact. In particular, for both the constructive and the destructive cases, he shows that the coalition weighted manipulation problem for the single transferable vote election rule for three or more candidates remains NP-hard in the single-peaked case. Although our Theorem 4.3 follows this path of seeing where shields remain intact for single-peaked preferences, the central focus of our paper is that single-peaked preferences often remove complexity shields on manipulation and control. Walsh's paper for a different issue—looking at incomplete profiles and asking whether some/all the completions make a candidate a winner—proves both P results and NP-completeness results. We're greatly indebted to his paper for raising and exploring the issue of manipulation for single-peaked electorates.

As mentioned in the main text, Bartholdi and Trick [1986] and Escoffier, Lang, and Öztürk [2008] have provided efficient algorithms for testing single-peakedness and producing a valid candidate linear ordering, for the case when votes are linear orders.

Other work is more distant from our work but worth mentioning. Conitzer [to appear] has done an interesting, detailed study showing that (in the model where votes are linear orders) preferences in single-peaked societies can be quickly elicited via comparison queries ("Do you prefer candidate i to candidate j ?"). He studies the case when the linear order of society is known and the case when it is not. We mention in passing that we have looked at the issue of preference elicitation in single-peaked societies (where the linear order is given) of approval vectors via approval queries ("Do you approve of candidate i ?"). It is immediately obvious that single-peakedness gives no improvement for approval vectors and 1-approval vectors (approval vectors with exactly one 1; this is a vote type and should not be confused with "1-approval" as it would be used in scoring systems, where the actual vote is a linear order). But for j -candidate k -approval vectors (approval vectors with exactly k 1's), $k \geq 1$, it is easy to see that the general-case elicitation query complexity is exactly 0 when $j = k$ and is exactly $j - 1$ when $j > k$, but for the single-peaked case, the elicitation query complexity for ik -candidate k -approval vectors is at most $i - 1 + \lceil \log_2 k \rceil$. That is, we get a savings of a multiplicative factor of about k from single-peakedness.

Single-peaked preferences of course have been studied extensively in political science. We in particular mention that Ballester and Haeringer [2007] provide a precise mathematical characterization of single-peakedness, that Lepelley [1996] shows that single-peakedness removes some negative results about the relationship between scoring protocols and Condorcet-type criteria, and Gailmard, Patty, and Penn [to appear] discuss Arrow's Theorem on single-peaked domains. We refer the interested reader also to the coverage of single-peaked preferences in the excellent text-

book by Austen-Smith and Banks [2004].

6 Conclusions and Future Directions

The central point of this paper is that single-peaked preferences remove many complexity-theoretic shields against control and manipulation. That is, we showed that those shields, already under frequency-of-hardness and approximation attacks from other quarters, for single-peaked preferences didn't even exist in the first place. It follows that when choosing election systems for electorates one suspects will be single-peaked, one must not rely on results for those systems that were proven in the standard, unrestricted preference model.

This paper's work suggests many directions for future efforts. For single-peaked manipulation of scoring protocols, we gave some NP-complete cases and some P cases. But for manipulation of scoring protocols (in the general model where society is not required to be single-peaked), Hemaspaandra and Hemaspaandra [2007] (see also [Procaccia and Rosenschein, 2007; Conitzer *et al.*, 2007]) have provided a dichotomy theorem clearly classifying each case as NP-complete or in P. Can we obtain a dichotomy theorem for manipulation of scoring protocols in single-peaked societies? Theorem 4.4 achieves this for the case of three candidates.

Throughout this paper, single-peaked has meant the uni-dimensional case. Do the shield removals of this paper hold in, for example, an appropriate two-dimensional (or k -dimensional) analogue? (We mention in passing that every profile of n voters voting by linear orders can be embedded into \mathbb{R}^n in such a way that each voter and candidate is a point in \mathbb{R}^n and each voter prefers c_i to c_j if her Euclidean distance to c_i is less than to c_j .)

Finally, in a human society with a large number of voters, even if one issue, e.g., the economy, is almost completely polarizing the society, there are bound to be a few voters whose preferences are shaped by quite different issues, e.g., a given candidate's religion. So it would be natural to ask whether the shield-evaporation results of this paper can be extended even to societies that are "very nearly" single-peaked (see [Escoffier *et al.*, 2008, Section 6] and [Conitzer, to appear, Section 6] for discussion of nearness to single-peakedness in other contexts).

Acknowledgements

We thank the referees for their many helpful comments. This work was supported in part by DFG grants RO-1202/{11-1,12-1}, NSF grants CCF-0426761 and IIS-0713061, AGH University of Science and Technology grant 11.11.120.777, the European Science Foundation's EUROCORES program LogICCC, and Friedrich Wilhelm Bessel Research Awards to Edith Hemaspaandra and Lane

A. Hemaspaandra. This work was done in part during visits by the first three authors to Heinrich-Heine-Universität Düsseldorf and by the fourth author to the University of Rochester.

References

- [Austen-Smith and Banks, 2004] D. Austen-Smith and J. Banks. *Positive Political Theory II: Strategy and Structure*. University of Michigan Press, 2004.
- [Ballester and Haeringer, 2007] M. Ballester and G. Haeringer. A characterization of the single-peaked domain, 2007. Manuscript.
- [Bartholdi and Orlin, 1991] J. Bartholdi, III and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.
- [Bartholdi and Trick, 1986] J. Bartholdi, III and M. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.
- [Bartholdi *et al.*, 1989] J. Bartholdi, III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [Bartholdi *et al.*, 1992] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.
- [Baumeister *et al.*, to appear] D. Baumeister, G. Erdélyi, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Computational aspects of approval voting. In J. Laslier and R. Sanver, editors, *Handbook of Approval Voting*. Springer, to appear. Available as University of Rochester Computer Science Department Technical Report TR-2009-944, May 2009.
- [Black, 1948] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
- [Black, 1958] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- [Brelsford *et al.*, 2008] E. Brelsford, P. Faliszewski, E. Hemaspaandra, H. Schnoor, and I. Schnoor. Approximability of manipulating elections. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, pages 44–49. AAAI Press, July 2008.
- [Conitzer and Sandholm, 2006] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 627–634. AAAI Press, July 2006.
- [Conitzer *et al.*, 2007] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.
- [Conitzer, to appear] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, to appear.
- [Davis *et al.*, 1970] O. Davis, M. Hinich, and P. Ordeshook. An expository development of a mathematical model of the electoral process. *American Political Science Review*, 54(2):426–448, 1970.

- [Ephrati and Rosenschein, 1991] E. Ephrati and J. Rosenschein. The Clarke Tax as a consensus mechanism among automated agents. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 173–178. AAAI Press, July 1991.
- [Ephrati and Rosenschein, 1997] E. Ephrati and J. Rosenschein. A heuristic technique for multi-agent planning. *Annals of Mathematics and Artificial Intelligence*, 20(1–4):13–67, 1997.
- [Erdélyi *et al.*, 2008] G. Erdélyi, M. Nowak, and J. Rothe. Sincere-strategy preference-based approval voting fully resists constructive control and broadly resists destructive control. Technical Report arXiv:0806.0535 [cs.GT], arXiv.org, September 2008. To appear in *Mathematical Logic Quarterly*, 55(4):425–443, 2009.
- [Escoffier *et al.*, 2008] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence*, pages 366–370. IOS Press, July 2008.
- [Faliszewski *et al.*, 2006] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of bribery in elections. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 641–646. AAAI Press, July 2006.
- [Faliszewski *et al.*, 2007] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting broadly resist bribery and control. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 724–730. AAAI Press, July 2007.
- [Faliszewski *et al.*, 2009] 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*, pages 375–406. Springer, 2009.
- [Friedgut *et al.*, 2008] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science*, pages 243–249. IEEE Computer Society, October 2008.
- [Gailmard *et al.*, to appear] S. Gailmard, J. Patty, and E. Penn. Arrow’s theorem on single-peaked domains. In E. Aragonés, C. Beviá, H. Llavador, and N. Schofield, editors, *The Political Economy of Democracy*. To appear.
- [Hemaspaandra and Hemaspaandra, 2007] E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.
- [Hemaspaandra *et al.*, 2007a] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5-6):255–285, 2007.
- [Hemaspaandra *et al.*, 2007b] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1308–1314. AAAI Press, January 2007. To appear in *Mathematical Logic Quarterly*, 55(4):397–424, 2009.
- [Krehbiel, 1998] K. Krehbiel. *Pivotal Politics: A Theory of U.S. Lawmaking*. University of Chicago Press, 1998.
- [Lepelley, 1996] D. Lepelley. Constant scoring rules, Condorcet criteria, and single-peaked preferences. *Economic Theory*, 7(3):491–500, 1996.
- [Niemi and Wright, 1987] R. Niemi and J. Wright. Voting cycles and the structure of individual preferences. *Social Choice and Welfare*, 4(3):173–183, 1987.
- [Peleg and Sudhölter, 1999] B. Peleg and P. Sudhölter. Single-peakedness and coalition-proofness. *Review of Economic Design*, 4(4):381–387, 1999.
- [Pennock *et al.*, 2000] D. Pennock, E. Horvitz, and C. Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 729–734. AAAI Press, July/August 2000.
- [Poole and Rosenthal, 1997] K. Poole and H. Rosenthal. *Congress: A Political-Economic History of Roll-Call Voting*. Oxford University Press, 1997.
- [Procaccia and Rosenschein, 2007] A. Procaccia and J. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.
- [Trick, 1989] M. Trick. Recognizing single-peaked preferences on a tree. *Mathematical Social Sciences*, 17(3):329–334, 1989.
- [Walsh, 2007] T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 3–8. AAAI Press, July 2007.