Manipulation of k-Approval in Nearly Single-Peaked Electorates

Gábor Erdélyi[1], Martin Lackner[2], and Andreas Pfandler[1,2]

[1] School of Economic Disciplines, University of Siegen, Germany
erdelyi@wiwi.uni-siegen.de
[2] Institute of Information Systems, TU Wien, Austria
{lackner,pfandler}@dbai.tuwien.ac.at

Abstract. For agents it can be advantageous to vote insincerely in order to change the outcome of an election. This behavior is called manipulation. The Gibbard-Satterthwaite theorem states that in principle every non-trivial voting rule with at least three candidates is susceptible to manipulation. Since the seminal paper by Bartholdi, Tovey, and Trick in 1989, (coalitional) manipulation has been shown NP-hard for many voting rules. However, under single-peaked preferences – one of the most influential domain restrictions – the complexity of manipulation often drops from NP-hard to P.

In this paper, we investigate the complexity of manipulation for the k-approval and veto families of voting rules in nearly single-peaked elections, exploring the limits where the manipulation problem turns from P to NP-hard. Compared to the classical notion of single-peakedness, notions of nearly single-peakedness are more robust and thus more likely to appear in real-world data sets.

1 Introduction

Elections are a useful framework for preference aggregation with many applications in both human societies and in multiagent systems. Well-known examples are political elections in human societies and the design of recommender systems [16], planning [10], and machine learning [22] in multiagent systems, just to name a few.

Informally, an election is given by a set of candidates and a set of voters who have to express their preferences over the set of candidates. A voting rule describes how to aggregate the voters' preferences in order to determine the winners of a given election. In computational social choice, a central research topic is to study computational questions regarding insincere behavior in elections. A prominent example is *coalitional manipulation*. Coalitional manipulation deals with situations in which a group of voters casts their votes strategically in order to alter the outcome of an election. (If the coalition has size one, the problem is called *single manipulation*). The famous Gibbard-Satterthwaite theorem says that, in principle, every reasonable voting rule for at least three candidates is susceptible to manipulation [17, 20].

Manipulability is considered to be an undesirable property for a voting rule. In their seminal paper Bartholdi, Tovey, and Trick suggested that although voting rules are manipulable, the manipulator's task of successfully manipulating the election can still be computationally hard, i.e., NP-hard [1]. Indeed, since the paper by Bartholdi, Tovey, and Trick, (coalitional) manipulation has been shown NP-hard for many voting rules.

In contrast, under domain restrictions, the computational complexity of manipulation drops from NP-hard to P for many voting rules. One popular model of domain restriction in elections is the model of single-peaked preferences introduced by Black [2]. Unfortunately, the concept of single-peakedness is fragile and is unlikely to appear in real-world data sets. To overcome this limitation, recent research has established notions of nearly single-peaked preferences which are more robust [6, 7, 11, 14].

To the best of our knowledge, the only paper investigating (coalitional) manipulation in nearly single-peaked elections is the work by Faliszewski, Hemaspaandra, and Hemaspaandra [14] (a detailed comparison to our paper can be found in the Related Work Section). Our paper follows this new line of research and extends it with the following contributions:

- In our complexity analysis, we provide dichotomy results for constructive coalitional weighted manipulation under k-approval in the voter deletion model. The voter deletion model assumes that at most ℓ voters are not single-peaked with respect to the linear axis. Our results pinpoint the border between P membership and NP-completeness with respect to the number of approved candidates and the distance to single-peakedness.
- For veto we show how the complexity of constructive coalitional weighted manipulation behaves under seven notions of nearly single-peakedness that have been recently introduced [11, 14]. Our dichotomies show that constructive coalitional weighted manipulation in nearly single-peaked electorates is either trivial (and therefore in P) or NP-complete depending on the distance to single-peakedness.

Related Work. Our work continues the line of research on manipulation of elections. The first paper investigating manipulation in elections is the seminal paper of Bartholdi, Tovey, and Trick [1], where they studied the single manipulation problem with unweighted voters and proved the problem to be solvable in polynomial-time for all scoring rules.

Constructive coalitional weighted manipulation (CCWM, for short) was first introduced by Conitzer, Sandholm, and Lang [5]. Later, Hemaspaandra and Hemaspaandra provided a dichotomy result for the CCWM problem for scoring rules [18]. In particular, they showed that CCWM is easy for plurality, but is NP-hard for all other k-approval and k-veto rules. Procaccia and Rosenschein have extended this line of research by studying the average-case complexity of manipulation [19].

Walsh was the first who studied the complexity of manipulation in single-peaked elections, especially with a view to answer the question whether the complexity of manipulation changes under single-peaked elections [21]. In particular, he demonstrated that the complexity of CCWM under single transferable vote remains NP-hard even for single-peaked elections. Faliszewski et al. proved that for single-peaked profiles, CCWM for m-candidate 3-veto elections is NP-complete for $m = 5$ and is in P for all other m [13]. Furthermore, they showed that for single-peaked profiles, CCWM for veto is in P and they completely characterized which scoring rules have easy CCWM problems and which scoring rules have hard CCWM problems for three-candidate elections. Brandt et al. generalized the latter result for m-candidate scoring rules [3].

The present paper was mostly motivated by Faliszewski, Hemaspaandra, and Hemaspaandra [14]: Amongst others, they investigated the complexity of the CCWM problem under veto elections in nearly single-peaked societies, where they used the nearly single-peaked notion of ℓ-Voter Deletion (which is called ℓ-Maverick in their paper). We extend their results to seven common notions of nearly single-peakedness that were recently discussed in the literature [11, 14].

Two recent publications have studied the complexity of computing the distance to single-peaked electorates. Erdélyi, Lackner, and Pfandler [11] have focused on the single-peaked domain whereas Bredereck, Chen, and Woeginger [4] considered distances to a larger number of domain restrictions. Both papers mostly contain NP-hardness results with a few notable exceptions such as that the candidate deletion distance is computable in polynomial time. For a practical use of nearly single-peaked preferences, it would be desirable to have efficient algorithms to compute distances. This line of research has been initiated by Elkind and Lackner [8], where several approximation and fixed-parameter algorithms have been presented.

Organization. The remainder of the paper is organized as follows. In Section 2, we recap some voting theory basics. Section 3 gives an overview on the nearly single-peakedness notions and their relations handled in this paper. Our results on manipulation are presented in Section 4. Section 5 provides some conclusions and future directions.

2 Preliminaries

Let C be a finite set of *candidates*, V be a finite set of *voters*, and let $\succ$ be a *vote* (i.e., a total order) on C. Without loss of generality let $V = \{1, \dots, n\}$. Let $\mathcal{P} = (\succ_1, \dots, \succ_n)$ be a *(preference) profile*, i.e., a collection of votes. For simplicity, we will write for each voter $i \in V$, $c_1 c_2 \dots c_m$ instead of $c_1 \succ_i c_2 \succ_i \dots \succ_i c_m$. For two preference profiles on the same set of candidates $\mathcal{P} = (\succ_1, \dots, \succ_n)$ and $\mathcal{L} = (\succ_{n+1}, \dots, \succ_s)$, let $(\mathcal{P}, \mathcal{L}) = (\succ_1, \dots, \succ_s)$ define the *union* of the two preference profiles. An *election* is defined as a triple $E = (C, V, \mathcal{P})$, where C is the set of candidates, V the set of voters, and $\mathcal{P}$ a preference profile over C. Throughout the paper let m denote the number of candidates and n the number of votes.

A *voting correspondence (or voting rule)* $\mathcal{F}$ is a mapping from a given election $E = (C, V, \mathcal{P})$ to a non-empty subset $W \subseteq C$; we call the candidates in W the *winners* of the election E. A prominent class of voting rules is the class of scoring rules, which are defined using a *scoring vector* $\alpha = (\alpha_1, \dots, \alpha_m)$, $\alpha_i \in \mathbb{N}$, $\alpha_1 \geq \dots \geq \alpha_m$. In an *$m$-candidate scoring rule* each voter has to specify a tie-free linear ordering of all candidates and gives α_i points to the candidate ranked in position i. The winners of the election are the candidates with the highest overall score. *k-approval* is an m-candidate scoring rule with $\alpha_1 = \dots = \alpha_k = 1$ and $\alpha_{k+1} = \dots = \alpha_m = 0$. *veto* is the scoring rule defined by the scoring vector $\alpha_1 = \dots = \alpha_{m-1} = 1$ and $\alpha_m = 0$.

In the case of k-approval we say that the first k candidates in a given ranking have been *approved* whereas the others have been *disapproved*. For k-approval and veto, preferences actually do not have to be full rankings but *dichotomous preferences* suffice. Dichotomous preferences only distinguish between approved and disapproved candidates. In this paper we often do not give full rankings but rather the set of approved candidates. Strictly speaking, we use this notation to describe some total order that ranks the approved candidates above the disapproved candidates; all such total orders are equivalent from the perspective of k-approval.

Definition 1. *Let an* axis A *be a total order on* C *denoted by* $<$. *Furthermore, let* $\succ$ *be a vote with* c *as its highest ranked candidate. The vote* $\succ$ *is* single-peaked with respect to A *if for any* $x, y \in C$, *if* $x < y < c$ *or* $c < y < x$ *then* $c \succ y \succ x$ *has to hold. A preference profile* $\mathcal{P}$ *is said to be* single-peaked with respect to an axis A *if each vote is single-peaked with respect to* A. *A preference profile* $\mathcal{P}$ *is said to be* single-peaked consistent *if there exists an axis* A *such that* $\mathcal{P}$ *is single-peaked with respect to* A.

Note that, given a set of approved candidates and an axis A, there exists a single-peaked total order that corresponds to these approved candidates if and only if the candidates form an interval on A. Thus, for dichotomous preferences, one could also define single-peakedness in terms of intervals on an axis. We remark that recently several other domain restrictions specifically for dichotomous preferences have been proposed and studied [9].

To establish NP-hardness results we will reduce from the well-known NP-complete problem PARTITION (see, e.g., [15]), which is defined as follows.

PARTITION	
Given:	A finite multiset $S = \{x_1, \dots, x_s\}$ of positive integers with $\sum_{i=1}^{s} x_i = 2X$ for some positive integer X.
Question:	Is there a subset $S' \subset S$ such that the sum of the elements in S' is exactly X?

3 Nearly Single-Peakedness

As we build upon the notions of nearly single-peakedness which were studied by Erdélyi, Lackner, and Pfandler [11], we briefly recapitulate the relevant definitions and results. All these notions have been previously introduced and defined in the literature [11, 12, 14].

In the following, let $E = (C, V, \mathcal{P})$ be an election and ℓ a positive integer. Also, by $\mathcal{P}[C']$ we denote the profile $\mathcal{P}$ restricted to the candidates in C'. Analogously if A is an axis over C, we denote by $A[C']$ the axis A restricted to candidates in C'.

Voter Deletion: A profile $\mathcal{P}$ is *ℓ-Voter Deletion single-peaked consistent* if by removing at most ℓ votes from $\mathcal{P}$ one can obtain a preference profile $\mathcal{P}'$ that is single-peaked consistent. (We remark that this notion is also referred to as ℓ-maverick-SP [14] and as ℓ-maverick single-peaked consistent [11].)

Candidate Deletion: A profile $\mathcal{P}$ is *ℓ-Candidate Deletion single-peaked consistent* if we can obtain a set $C' \subseteq C$ by removing at most ℓ candidates from C such that the preference profile $\mathcal{P}[C']$ is single-peaked consistent.

Local Candidate Deletion: Let A be an axis over C. A vote $\succ$ on a candidate set $C' \subset C$ is called a partial vote. A partial vote on C' is said to be *single-peaked with respect to A* if it is single-peaked with respect to $A[C']$. A profile $\mathcal{P}$ is *ℓ-Local Candidate Deletion single-peaked consistent* if there exists an axis A such that by removing at most ℓ candidates from each vote we obtain a partial profile $\mathcal{P}'$ that is single-peaked with respect to A.

Additional Axes: A profile $\mathcal{P}$ is *ℓ-Additional Axes single-peaked consistent* if there is a partition $V_1, \ldots, V_{\ell+1}$ of the voter set V such that the corresponding preference profiles $\mathcal{P}_1, \ldots, \mathcal{P}_{\ell+1}$ are single-peaked consistent.

Global Swaps: A profile $\mathcal{P}$ is *ℓ-Global Swaps single-peaked consistent* if $\mathcal{P}$ can be made single-peaked by performing at most ℓ swaps of consecutive candidates in the profile. (Note that these swaps can be performed wherever we want – we can have ℓ swaps in only one vote, or one swap each in ℓ votes.)

Local Swaps: A profile $\mathcal{P}$ is *ℓ-Local Swaps single-peaked consistent* if $\mathcal{P}$ can be made single-peaked consistent by performing no more than ℓ swaps of consecutive candidates per vote.

Candidate Partition: A profile $\mathcal{P}$ is *ℓ-Candidate Partition single-peaked consistent* if the set of candidates C can be partitioned into at most ℓ disjoint sets $C_1, \ldots, C_\ell$ with $C_1 \cup \ldots \cup C_\ell = C$ such that the profiles $\mathcal{P}[C_1], \ldots, \mathcal{P}[C_\ell]$ are single-peaked consistent.

We denote by $VD(\mathcal{P})/CD(\mathcal{P})/LCD(\mathcal{P})/AA(\mathcal{P})/GS(\mathcal{P})/LS(\mathcal{P})/CP(\mathcal{P})$ the smallest ℓ such that $\mathcal{P}$ is ℓ-Voter Deletion/ℓ-Candidate Deletion/ℓ-Local Candidate Deletion/ℓ-Additional Axes/ℓ-Global Swaps/ℓ-Local Swaps/ℓ-Candidate Partition single-peaked consistent.

Theorem 2 (cf. [11]). *Let $\mathcal{P}$ be a preference profile. Then the following inequalities hold:*

(1) $LS(\mathcal{P}) \leq GS(\mathcal{P})$.
(2) $LCD(\mathcal{P}) \leq CD(\mathcal{P})$.
(3) $CD(\mathcal{P}) \leq GS(\mathcal{P})$.
(4) $LCD(\mathcal{P}) \leq LS(\mathcal{P})$.
(5) $VD(\mathcal{P}) \leq GS(\mathcal{P})$.
(6) $AA(\mathcal{P}) \leq VD(\mathcal{P})$.
(7) $CP(\mathcal{P}) \leq CD(\mathcal{P}) + 1$.
(8) $CP(\mathcal{P}) \leq LS(\mathcal{P}) + 1$.

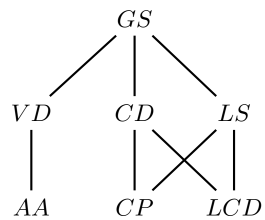

GS	...	Global Swaps
VD	...	Voter Deletion
CD	...	Candidate Deletion
LS	...	Local Swaps
AA	...	Additional Axes
CP	...	Candidate Partition
LCD	...	Local Candidate Deletion

Fig. 1. Hasse diagram of the partial order described in Theorem 2.

This list is complete in the following sense: Inequalities that are not listed here and that do not follow from transitivity do not hold in general. The resulting partial order with respect to $\leq$ is displayed in Figure 1 as a Hasse diagram.

Finally, let us summarize the complexity results concerning the detection of nearly single-peaked elections. Here, the question is whether a given election is ℓ-X single-peaked consistent. This problem is polynomial-time solvable for the Candidate Deletion distance and NP-complete for all $X \in \{$Voter Deletion, Local Candidate Deletion, Additional Axes, Global Swaps, Local Swaps$\}$ [4, 11].

4 Manipulation

In what follows we investigate the computational complexity of coalitional manipulation in scoring rules under the assumption that the underlying elections are nearly single-peaked. For this we have to formally define the coalitional weighted manipulation problem in general and for nearly single-peaked electorates. Let $\mathcal{F}$ be a voting correspondence.

$\mathcal{F}$-Constructive Coalitional Weighted Manipulation ($\mathcal{F}$-CCWM)	
Given:	An election $(C, V, \mathcal{P})$, where C is a set of candidates, V a set of nonmanipulative voters, and $\mathcal{P} = (P_1, \dots, P_h)$ a preference profile; in addition, a set of manipulative voters S with $V \cap S = \emptyset$, a weight function w from $V \cup S$ to $\mathbb{N}$, and a distinguished candidate $p \in C$.
Question:	Is there a preference profile $\mathcal{L} = (L_1, \dots, L_s)$ for the manipulative voters in S such that p is a (co-)winner in $(C, V \cup S, (\mathcal{P}, \mathcal{L}))$ with respect to the voting correspondence $\mathcal{F}$?

In this paper we study $\mathcal{F}$-CCWM for nearly single-peaked preferences. For a fixed, non-negative integer ℓ and $X \in \{$Voter Deletion, Candidate Deletion, Local Candidate Deletion, Additional Axes, Global Swaps, Local Swaps, Candidate Partition$\}$, we define $\mathcal{F}$-ℓ-X-CCWM to be $\mathcal{F}$-CCWM restricted to profiles that are ℓ-X single-peaked consistent with respect to an axis A. Note that the combined election $(C, V \cup S, (\mathcal{P}, \mathcal{L}))$ has to be ℓ-X single-peaked consistent. In addition, we assume that this axis is part of the input. In the case of the additional axes distance, we assume that all axes are part of the input; in the case of the candidate partition distance, we assume that the actual partition is part of the input. To be more precise, $\mathcal{F}$-ℓ-X-CCWM is defined as follows:

$\mathcal{F}$-ℓ-X-CCWM	
Given:	An $\mathcal{F}$-CCWM instance, an axis A, additional axes $A_1, \dots, A_\ell$ if X is Additional Axes, a partition of the candidate set C if X is Candidate Partition.
Question:	Is there a preference profile $\mathcal{L} = (L_1, \dots, L_s)$ for the manipulative voters in S such that (i) p is a (co-)winner in $(C, V \cup S, (\mathcal{P}, \mathcal{L}))$ with respect to the voting correspondence $\mathcal{F}$ and (ii) $(C, V \cup S, (\mathcal{P}, \mathcal{L}))$ is ℓ-X single-peaked consistent?

Remark. As mentioned earlier, it is NP-hard to verify whether an election is ℓ-X single-peaked consistent for all notions of distance X considered in this paper except for the Candidate Deletion distance (for which this problem is in P) and for the Candidate Partition distance (for which its complexity is not known) [4, 11]. These NP-hardness results, however, do not influence the complexity of $\mathcal{F}$-ℓ-X-CCWM due to our assumption that the axis is part of the input. Given a fixed axis A and $X \in$ {Voter Deletion, Candidate Deletion, Local Candidate Deletion, Global Swaps, Local Swaps}, it requires only polynomial time to verify that an election is ℓ-X single-peaked with respect to A. The same holds for Additional Axis if all axes are given and for Candidate Partition if the partition of the candidates is given. Consequently, the complexity of $\mathcal{F}$-ℓ-X-CCWM can be studied separately from the complexity of deciding ℓ-X single-peaked consistency.

Let us start with our first result on CCWM. Following the notation of Faliszewski, Hemaspaandra, and Hemaspaandra [14], let $(\alpha_1, \alpha_2, \alpha_3)$ *elections* denote three-candidate scoring rule elections with scoring-vector $\alpha = (\alpha_1, \alpha_2, \alpha_3)$. In that paper it was proven that for each $\alpha_1 \geq \alpha_2 > \alpha_3$, $(\alpha_1, \alpha_2, \alpha_3)$-1-Voter Deletion-CCWM is NP-complete. This result implies that Veto-1-Voter Deletion-CCWM for three-candidate elections is NP-complete. In contrast, Veto-CCWM is in P for single-peaked societies. The following proposition makes use of Theorem 2 and shows that the same holds for all other notions of distance studied in this paper.

Proposition 3. *Let* $X \in$ {Candidate Deletion, Local Candidate Deletion, Additional Axes, Global Swaps, Local Swaps}. *For each* $\alpha_1 \geq \alpha_2 > \alpha_3$*, the problems* $(\alpha_1, \alpha_2, \alpha_3)$-1-$X$-CCWM *and* $(\alpha_1, \alpha_2, \alpha_3)$-2-Candidate Partition-CCWM *are* NP-*complete.*

Proof. Faliszewski, Hemaspaandra, and Hemaspaandra [14] show NP-completeness of the $(\alpha_1, \alpha_2, \alpha_3)$-1-Voter Deletion-CCWM problem. We now show that a three candidate, 1-voter deletion single-peaked consistent election is also 1-X single-peaked consistent for all X and 2-candidate partition single-peaked consistent. It is easy to see that every election E over three candidates is 1-candidate deletion, 1-local candidate deletion, and 2-candidate partition single-peaked consistent. From Theorem 2, Inequality (6), follows that E is also 1-additional axes single-peaked consistent.

Let $C = \{a, b, c\}$ be the set of candidates, and without loss of generality assume that E is 1-voter deletion single-peaked consistent along the axis $a < b < c$. Note that there are only two possible non-single-peaked votes, *acb* and *cab*. In both votes, swapping the last two candidates leaves us with single-peaked votes with respect to axis $a < b < c$. Thus, E is 1-global swaps single-peaked consistent. From Theorem 2, Inequality (1), follows that E is 1-local swaps single-peaked consistent. □

4.1 Manipulation for *k*-Approval

We continue by investigating the computational complexity of manipulation for the k-approval voting rule in ℓ-Voter Deletion single-peaked societies. CCWM for k-approval is known to be NP-complete in general [18]. This holds even for single-peaked elections in many settings [3]. We extend these results to ℓ-Voter Deletion single-peaked societies. More concretely, we show a dichotomy result: k-Approval-ℓ-Voter Deletion-CCWM is in P if and only if $\ell < \frac{2k-m}{m-k}$, and NP-complete otherwise. This gives a complete picture for k-approval with respect to ℓ-Voter Deletion single-peakedness. Both the P membership and NP-hardness result are generalizations of the results for veto elections [14] and our proofs can be seen as refinements of the corresponding proofs.

Theorem 4. *Let $m \geq 3, k > 1, \ell \geq 1$ be fixed integers such that $k < m$ and $\ell \geq \frac{2k-m}{m-k}$. Then k-Approval-ℓ-Voter Deletion-CCWM for elections with m candidates is NP-complete.*

Proof. Membership in NP is trivial. We reduce from Partition with a sum of $2X$. Let b be a positive integer such that $\max(1, 2k - m) \leq b \leq k - 1$. Note that such a b always exists since $1 < k < m$. Let $C = \{x, y, p, c_1, \ldots, c_{m-3}\}$, where p is the distinguished candidate. To construct the votes in $\mathcal{P}$, we split the sequence $pc_1 \cdots c_{b-1}$ into consecutive blocks of size $m - k$. (If necessary the last block is filled with candidates from $c_b, \ldots, c_{m-3}$.) Let $d = \left\lceil \frac{b}{m-k} \right\rceil$. These blocks give d sets of candidates $D_1, \ldots, D_d$. Furthermore, let $e = \frac{m-3-b}{2}$. We fix the axis A to $c_{b+\lceil e \rceil} < \cdots < c_{m-3} < x < p < c_1 < \cdots < c_{b-1} < y < c_b < \cdots < c_{b+\lceil e \rceil-1}$. The profile $\mathcal{P}$ comprises the following votes:

- $\mathcal{P}$ contains d votes of weight X: For each $i \in \{1, \ldots, d\}$,

$$V_i : C \setminus D_i \text{ is approved}, w(V_i) = X.$$

  Note that all these votes are not single-peaked.
- $\ell - d$ votes of weight 1: For $i \in \{d + 1, \ldots, \ell\}$,

$$V_i : \{x, y, p, c_1, \ldots, c_{k-4}, c_{m-4}\}, w(V_i) = 1.$$

Also these votes are not single-peaked and hence we have exactly ℓ non-single-peaked votes. Consequently, we force the manipulators to cast single-peaked votes.

Let the set $\mathcal{L}$ consist of s manipulators with weights $x_1, \ldots, x_s$.

At this point, candidates have the following scores: The candidates x and y are approved by all votes and hence they have a score of $d \cdot X + (\ell - d)$. Candidate p is approved by all votes except V_1 since $p \in D_1$ and thus has a score of $(d-1) \cdot X + (\ell - d)$. The candidates $c_1, \ldots, c_{b-1}$ have a score of at most $(d-1) \cdot X + (\ell - d)$ since they are contained in at least one D_i. The candidates $c_b, \ldots, c_{m-3}$ have a score of at most $d \cdot X + (\ell - d)$.

Since the manipulators can cast only single-peaked votes and they want to approve p but not both x and y, the manipulators approve an interval on A of length k that contains p and either x or y.

Intuitively, the manipulators can only make p a winner if they manage to give $2X$ points to p and X points to x and y such that x, y, and p are tied. Note that in a single-peaked vote where p is approved, also either x or y has to be approved. We claim that there is a subset $S' \subset S$ such that the elements in S' sum to X if and only if p can be made a winner of the election by constructive coalitional weighted manipulation.

"$\Rightarrow$": Suppose there is a subset $S' \subset S$ such that the elements in S' sum to X. Let all the manipulators whose weight is in S' approve

$$\{x, p, c_1, \ldots, c_{b-1}, c_{m-k+b-1}, \ldots, c_{m-3}\},$$

i.e., they approve a "block" of length k that starts at c_{b-1} and goes to the left on axis A. All the manipulators whose weight is in $S \setminus S'$ approve

$$\{p, y, c_1, \ldots, c_{k-2}\},$$

i.e., they approve a "block" of length k that starts at p and goes to the right. In both cases the vote is single-peaked. It is easy to see that p gains $2X$, whereas x and y only gain X points. Note that the maximum score where p ties with x and y is $(d+1) \cdot X + (\ell - d)$ as x and y obtain $d \cdot X + (\ell - d)$ points from the nonmanipulators and the manipulators can distribute a score of $2X$, but never approve x and y together. Note that none of the other candidates can surpass the score of x, y, and p. Hence, x, y, and p are among the winners tied for first place making the distinguished candidate p a winner.

"$\Leftarrow$": Suppose that p can be made a winner of the election by constructive coalitional weighted manipulation. Note that according to the scores given by the nonmanipulators, p is missing X points to be tied with x and y. The only way p can gain X points on these two candidates is if the manipulators can be divided into two groups, both weighing X points. The first group approves p, x, and suitable candidates from $\{c_1, \ldots, c_{m-3}\}$. The second group approves p, y, and suitable candidates from $\{c_1, \ldots, c_{m-3}\}$. Thereby, p gains $2X$ points, whereas x and y gain only X points each. Thus, there is a subset $S' \subset S$ such that the elements in S' sum to X. $\square$

Theorem 5. *If $\ell < \frac{2k-m}{m-k}$, then* k-Approval-ℓ-Voter Deletion-CCWM *is in* P*.*

Proof. Without loss of generality let $A : c_1 < c_2 < \cdots < c_m$. Every vote disapproves $m - k$ candidates. Consequently, the ℓ non-single-peaked voters disapprove at most $\ell \cdot (m - k) < 2k - m$ candidates. Therefore, there is at least one candidate contained in $\{c_{m-k+1}, \ldots, c_k\}$ that is approved by all non-single-peaked voters, since $|\{c_{m-k+1}, \ldots, c_k\}| = 2k - m$. (Note that $0 \le \ell < \frac{2k-m}{m-k}$ implies $2k - m > 0$.) Single-peaked voters may disapprove only candidates in $\{c_1, \ldots, c_{m-k}, c_{k+1}, \ldots, c_m\}$. Thus, candidates in $\{c_{m-k+1}, \ldots, c_k\}$ approved by all non-single-peaked voters are also approved by all single-peaked voters. Since there exists at least one candidate that is approved by all voters (including the manipulators), p is a winner if and only if it is approved by all voters (both manipulators and nonmanipulators). □

The following corollary shows how Theorem 4 carries over to Veto-ℓ-Voter Deletion-CCWM.

Corollary 6 (also shown in [14]). *Let $m, \ell \in \mathbb{N}$ be fixed such that $\ell > m - 3$. Then* Veto-ℓ-Voter Deletion-CCWM *is* NP*-complete. Otherwise,* Veto-ℓ-Voter Deletion-CCWM *is in* P*.*

4.2 Manipulation for Veto

In this section we study the complexity of constructive coalitional weighted manipulation in nearly single-peaked societies under the veto rule. For veto, CCWM is NP-complete in general [18], whereas the problem is in P for single-peaked elections [13]. In contrast to the previous section, we study here only a single rule, veto, but consider a variety of notions for nearly single-peakedness. Table 1 summarizes the complexity results regarding Veto-ℓ-X-CCWM under several notions of nearly single-peakedness. Note that all results in this table yield dichotomies.

In the following we will prove each entry of this table. We assume throughout this section that there are at least three candidates, since for less than three candidates Veto-CCWM is in P [5, 18].

All P membership proofs in this section follow the same reasoning as the proof of Theorem 5. More specifically, we show that there is at least one candidate that is never vetoed. As a consequence, a candidate p can only be amongst the winners if p is never vetoed (both by the nonmanipulators and the manipulators). Clearly, it is possible in polynomial time to determine whether a candidate is approved by all nonmanipulative voters and to construct the manipulator's votes that approve p. In the following P membership proofs we only argue that there is indeed at least one candidate that is never vetoed and omit the remainder of the argument.

Theorem 7. *Let $m \ge 3$. For each $\ell \ge 0$, if $\ell \le m - 3$* Veto-ℓ-Candidate Deletion-CCWM *is in* P *and* NP*-complete otherwise.*

X	in P	NP-complete	Reference
Voter Deletion	$\ell \leq m-3$	$\ell > m-3$	[14] & Cor. 6
Candidate Deletion	$\ell \leq m-3$	$\ell > m-3$	Thm. 7
Local Candidate Del.	$\ell = 0$	$\ell \geq 1$	Prop. 8
Global Swaps	$m = 2k$: $\ell \leq k^2 - k - 1$	$\ell > k^2 - k - 1$	Thm. 10
	$m = 2k-1$: $\ell \leq k^2 - 2k$	$\ell > k^2 - 2k$	Thm. 10
Local Swaps	$\ell < \lfloor \frac{m-1}{2} \rfloor$	$\ell \geq \lfloor \frac{m-1}{2} \rfloor$	Thm. 11
Candidate Partition	$\ell < \frac{m}{2}$	$\ell \geq \frac{m}{2}$	Thm. 12
Additional Axes	$\ell < \frac{m}{2} - 1$	$\ell \geq \frac{m}{2} - 1$	Thm. 13

Table 1. Complexity results regarding Veto-ℓ-X-CCWM under several notions of nearly single-peakedness, assuming $m \geq 3$.

Proof. We are first handling the $\ell \leq m-3$ case. Let A be the axis along which the election is nearly single-peaked and let c_l and c_r be the leftmost and rightmost candidates in A, respectively. Note that in a veto election over a single-peaked society, only c_l and c_r can be vetoed. In an ℓ-Candidate Deletion single-peaked society there are at most ℓ additional candidates vetoed in those votes not consistent with the axis A. Thus, there are at most $\ell + 2 \leq m - 1$ candidates that are vetoed. Consequently, there has to be at least one candidate who never got vetoed.

We now turn to the case where $\ell > m - 3$. In this case, hardness follows immediately from the fact that every profile is $(m-2)$-candidate deletion single-peaked consistent and Veto-CCWM is an NP-hard problem [5, 18, 19]. □

In the following proposition we require that $\ell \geq 1$. The $\ell = 0$ case would mean that the election is single-peaked, for which Brandt et al. [3] proved that constructive coalitional weighted manipulation under the veto rule is in P.

Proposition 8. *For each $m \geq 3$ and $\ell \geq 1$,* Veto-ℓ-Local Candidate Deletion-CCWM *is* NP-*complete.*

Proof. The crucial observation here is that with $\ell \geq 1$ every candidate can be vetoed, since the vetoed candidate can be the one that is locally deleted. Thus, this problem is equivalent to Veto-CCWM, which is NP-complete for each $m \geq 3$ [5, 18, 19]. □

In the following, for any two candidates $c_1, c_2 \in C$ let $d_A(c_1, c_2)$ be the *distance of two candidates on the axis A*. For example, for the axis $A = c_1 < c_3 < c_5 < c_4 < c_2 < c_6$ the distance $d_A(c_1, c_2) = 4$.

Lemma 9. *Let $E = (C, V, \mathcal{P})$ be a single-peaked election along the axis A, where c_l and c_r are the leftmost and rightmost candidates, respectively. The number of swaps required to make a candidate $c \in C$ the lowest-ranked candidate in a vote $v \in V$ is at least $min(d_A(c, c_r), d_A(c, c_l))$.*

Proof. Without loss of generality assume that c is closer to c_r in A than to c_l (i.e., $d_A(c, c_r) < d_A(c, c_l)$). If a vote v coincides with the axis A then clearly exactly $min(d_A(c, c_r), d_A(c, c_l)) = d_A(c, c_r)$ swaps are needed to make c the candidate who gets vetoed in v.

If v does not coincide with A, we have to distinguish three cases. First, let c be the peak of v. In this case it is clear that c has to be swapped with all other candidates to get vetoed and thus we need exactly $d_A(c, c_r) + d_A(c, c_l) \geq min(d_A(c, c_r), d_A(c, c_l))$ swaps. Second, let c be left from v's peak on axis A. This means that, according to the definition of single-peakedness, all the candidates left from c on axis A are ranked lower than c in v. To swap c through to the last position, we will have to make at least $d_A(c, c_l) > min(d_A(c, c_r), d_A(c, c_l))$ swaps. Finally, let c be right from v's peak on axis A. This means that all the candidates on axis A right from c are ranked lower than c in v. To swap c through to the last position, we will have to make at least $d_A(c, c_r) = min(d_A(c, c_r), d_A(c, c_l))$ swaps. □

Using Lemma 9, the following two theorems can be shown.

Theorem 10. *Let $k \geq 2$ be a positive integer.*

1. *Let the number of candidates be $m = 2k$. For each $\ell \geq 0$,* VETO-ℓ-GLOBAL SWAPS-CCWM *is in* P *if $\ell \leq k^2 - k - 1$ and* NP*-complete otherwise.*
2. *Let the number of candidates be $m = 2k-1$. For each $\ell \geq 0$,* VETO-ℓ-GLOBAL SWAPS-CCWM *is in* P *if $\ell \leq k^2 - 2k$ and* NP*-complete otherwise.*

Proof. Without loss of generality, let $A : c_1 < \cdots < c_m$ be the axis for which the election is nearly single-peaked. Let us consider case (1) first, i.e., $m = 2k$. We count the number of candidates who can be vetoed. These are the two candidates c_1 and c_m, and those candidates that can be "swapped" to the last position with at most ℓ swaps. Observe that it requires at least one swap each for swapping c_2 and c_{m-1} to the lowest position in a vote (cf. Lemma 9). For c_3 and c_{m-2} at least three swaps are required, etc. We can make at most $k^2 - k - 1 = -1 + 2\sum_{i=0}^{k-1} i$ swaps and consequently less than m different candidates can be swapped to a last position in some vote (cf. Lemma 9). Thus, there is at least one candidate who is never vetoed. In case (2), i.e., $m = 2k - 1$, note that $k^2 - 2k = -1 + (k-1) + 2\sum_{i=0}^{k-2} i$ and hence less than m can be vetoed.

To show hardness we reduce from PARTITION. Let $m = 2k$ and $\ell \geq k^2 - k$. (The case that $m = 2k-1$ works analogously.) Given a multiset $S = \{x_1, \ldots, x_s\}$ of s integers that sum to $2X$, define the following instance of VETO-ℓ-GLOBAL SWAPS-CCWM. Let $C = \{p, c_l, c_r, c_1, \ldots, c_{m-3}\}$ be the set of candidates and let p be the distinguished candidate. Let A be the axis for which the election is nearly single-peaked and let candidates c_l and c_r be the leftmost and rightmost candidates in A. Let the nonmanipulative voters consist of $m-2$ voters, each with weight X such that for every candidate $c \in C\setminus\{c_l, c_r\}$ there is a nonmanipulative voter who ranks c last but otherwise the votes are identical with the axis A or its reverse $\overline{A}$ (if c is closer to c_l on A, then we choose the axis as vote which ranks c_l last). Note that in this case we need $2\sum_{i=0}^{k-1} i = k^2 - k$ global swaps to make the profile single-peaked which is still less or equal ℓ. Let $\mathcal{L}$ consist of s manipulators with weights $x_1, \dots, x_s$.

We claim that there is a subset $S' \subset S$ such that the elements in S' sum to X if and only if p can be made a winner of the election by constructive coalitional weighted manipulation.

"$\Rightarrow$": Suppose there is a subset $S' \subset S$ such that the elements in S' sum to X. Let all the manipulators whose weight is in S' vote identically to the axis A, and all the manipulators whose weight is in $S \setminus S'$ vote reverse. It is easy to see that now all candidates tie for first place and, thus, the distinguished candidate p is a winner.

"$\Leftarrow$": Suppose that p can be made a winner of the election by constructive coalitional weighted manipulation. Note that p ties (with $c_1, \dots, c_{m-3}$) for the third place X points behind both candidates c_l and c_r. The only way p can gain X points on these two candidates is if the manipulators can be divided into two groups, both weighing X points and vetoing candidates c_l and c_r, respectively. Thus, there is a subset $S' \subset S$ such that the elements in S' sum to X. □

Theorem 11. *Let $m \geq 3$ denote the number of candidates. For each $\ell \geq 0$,* Veto-ℓ-Local Swaps-CCWM *is in* P *if $\ell < \lfloor \frac{m-1}{2} \rfloor$ and* NP*-complete otherwise.*

Proof. Let A be the axis along the election is nearly single-peaked, and let c_l and c_r be the leftmost and rightmost candidates in A, respectively. Observe that there is a candidate on A with distance at least $\lfloor \frac{m-1}{2} \rfloor$ to both c_l or c_r. Thus, for $\ell < \lfloor \frac{m-1}{2} \rfloor$, there is a candidate that is never vetoed.

For showing hardness, note that when we start with the single-peaked votes $c_1 \succ c_2 \succ \dots \succ c_m$ or $c_m \succ \dots \succ c_2 \succ c_1$, $\lfloor \frac{m-1}{2} \rfloor$ swaps suffice to make any candidate rank last. Thus, every candidate can be vetoed and Veto-ℓ-Local Swaps-CCWM for $\ell \geq \lfloor \frac{m-1}{2} \rfloor$ is equivalent to Veto-CCWM, which is NP-complete for $m \geq 3$ [5, 18, 19]. □

Theorem 12. *Let $m \geq 3$ be the number of candidates in an election E. For each $\ell \geq 1$,* Veto-ℓ-Candidate Partition-CCWM *is in* P *if $\ell < \frac{m}{2}$ and* NP*-complete otherwise.*

Proof. In the $\ell < \frac{m}{2}$ case we again count the number of candidates who can be vetoed. As we can veto at most two candidates per partition and we have ℓ single-peaked partitions, there can be at most 2ℓ different candidates being vetoed. Since $2\ell < m$, there has to be at least one candidate who is never vetoed.

For the other case, $\ell \geq \frac{m}{2}$, note that since there are at least $\frac{m}{2}$ partitions, all candidates can be vetoed while preserving candidate partition single-peakedness. Hardness for this case follows from the result for the general case [5, 18, 19]. □

Finally, we turn to Veto-ℓ-Additional Axes-CCWM.

Theorem 13. *Let $m \geq 3$. For each $\ell \geq 0$,* Veto-ℓ-Additional Axes-CCWM *is in* P *if $\ell < \frac{m}{2} - 1$ and* NP*-complete otherwise.*

Proof. The proof is similar to the candidate partition case (Theorem 12). The important observation is that for this purpose candidate partition and alternative axes provide the same freedom: In both cases at most two candidates per axis or partition can be vetoed. Note that the -1 in the bound on ℓ comes from the fact that ℓ *additional* axes give us $\ell + 1$ axes in total. □

5 Conclusions and Open Questions

We have investigated the computational complexity of manipulation in nearly single-peaked elections, where we focused on k-approval and veto. For veto we have studied seven notions of nearly single-peakedness that were recently studied in the literature [11, 14]. In contrast, for k-approval, we have explored how k influences the complexity if we consider voter deletion as notion for nearly single-peakedness. In both cases we proved dichotomies that exactly pinpoint the border of tractability. These results give insight into the sources of hardness and reveal in which settings we can hope for computationally hard instances.

There are several ways to continue with this direction of research. Extending our results to k-approval (or even arbitrary scoring rules) for all notions of nearly single-peakedness is certainly an important direction to go. Another way is to consider other notions of strategic behavior such as bribery and control in the light of nearly single-peakedness.

Acknowledgements. We thank the anonymous ADT-2015 referees for their very helpful comments and suggestions. This work was supported by the Austrian Science Fund (FWF): P25518, Y698, and the German Research Foundation (DFG): ER 738/2-1.

References

1. Bartholdi, J., Tovey, C., Trick, M.: The computational difficulty of manipulating an election. Social Choice and Welfare 6(3), 227–241 (1989)
2. Black, D.: On the rationale of group decision making. Journal of Political Economy 56(1), 23–34 (1948)
3. Brandt, F., Brill, M., Hemaspaandra, E., Hemaspaandra, L.A.: Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In: Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI 2010). pp. 715–722. AAAI Press (2010)
4. Bredereck, R., Chen, J., Woeginger, G.J.: Are there any nicely structured preference profiles nearby? In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013). pp. 62–68. AAAI Press (2013)
5. Conitzer, V., Sandholm, T., Lang, J.: When are elections with few candidates hard to manipulate? Journal of the ACM 54(3), Article 14 (2007)
6. Cornaz, D., Galand, L., Spanjaard, O.: Bounded single-peaked width and proportional representation. In: Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012). Frontiers in Artificial Intelligence and Applications, vol. 242, pp. 270–275. IOS Press (2012)

7. Cornaz, D., Galand, L., Spanjaard, O.: Kemeny elections with bounded single-peaked or single-crossing width. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013). pp. 76–82. AAAI Press (2013)
8. Elkind, E., Lackner, M.: On detecting nearly structured preference profiles. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014). pp. 661–667. AAAI Press (2014)
9. Elkind, E., Lackner, M.: Structure in dichotomous preferences. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015). AAAI Press (2015), to appear
10. Ephrati, E., Rosenschein, J.: A heuristic technique for multi-agent planning. Annals of Mathematics and Artificial Intelligence 20(1–4), 13–67 (1997)
11. Erdélyi, G., Lackner, M., Pfandler, A.: Computational aspects of nearly single-peaked electorates. In: Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013). pp. 283–289. AAAI Press (July 2013)
12. Escoffier, B., Lang, J., Öztürk, M.: Single-peaked consistency and its complexity. In: Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008). Frontiers in Artificial Intelligence and Applications, vol. 178, pp. 366–370. IOS Press (2008)
13. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L.A., Rothe, J.: The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. Information and Computation 209(2), 89–107 (2011)
14. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L.A.: The complexity of manipulative attacks in nearly single-peaked electorates. Artificial Intelligence 207, 69–99 (2014)
15. Garey, M., Johnson, D.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman and Company (1979)
16. Ghosh, S., Mundhe, M., Hernandez, K., Sen, S.: Voting for movies: The anatomy of recommender systems. In: Proceedings of the 3rd Annual Conference on Autonomous Agents (AGENTS '99). pp. 434–435. ACM Press (1999)
17. Gibbard, A.: Manipulation of voting schemes. Econometrica 41(4), 587–601 (1973)
18. Hemaspaandra, E., Hemaspaandra, L.: Dichotomy for voting systems. Journal of Computer and System Sciences 73(1), 73–83 (2007)
19. Procaccia, A.D., Rosenschein, J.S.: Junta distributions and the average-case complexity of manipulating elections. Journal of Artificial Intelligence Research (JAIR) 28, 157–181 (2007)
20. Satterthwaite, M.: Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. Journal of Economic Theory 10(2), 187–217 (1975)
21. Walsh, T.: Uncertainty in preference elicitation and aggregation. In: Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI 2007). pp. 3–8. AAAI Press (2007)
22. Xia, L.: Designing social choice mechanisms using machine learning. In: Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013). pp. 471–474. International Foundation for Autonomous Agents and Multiagent Systems (2013)