

On Coalitional Manipulation for Multiwinner Elections: Shortlisting

Robert Brederick

University of Oxford, United Kingdom;
 TU Berlin, Germany
 robert.bredereck@tu-berlin.de

Andrzej Kaczmarczyk and Rolf Niedermeier

TU Berlin, Germany
 {a.kaczmarczyk, rolf.niedermeier}@tu-berlin.de

Abstract

Shortlisting of candidates—selecting a group of “best” candidates—is a special case of multiwinner elections. We provide the first in-depth study of the computational complexity of strategic voting for shortlisting based on the most natural and simple voting rule in this scenario, ℓ -Bloc (every voter approves ℓ candidates). In particular, we investigate the influence of several tie-breaking mechanisms (e.g. pessimistic versus optimistic) and group evaluation functions (e.g. egalitarian versus utilitarian) and conclude that in an egalitarian setting strategic voting may indeed be computationally intractable regardless of the tie-breaking rule. We provide a fairly comprehensive picture of the computational complexity landscape of this neglected scenario.

1 Introduction

Multiwinner voting rules come up very naturally whenever from a large set of candidates a smaller “best” set has to be selected. Surprisingly, although at least as practically relevant as singlewinner voting rules, the multiwinner literature is much less developed than the singlewinner literature. In recent years, however, research into multiwinner voting rules, their properties and algorithmic complexity grew significantly [Aziz *et al.*, 2017a; Aziz *et al.*, 2017b; Aziz *et al.*, 2015; Barrot *et al.*, 2013; Barberà and Coelho, 2008; Barberà and Coelho, 2010; Elkind *et al.*, 2017; Faliszewski *et al.*, 2016; Meir *et al.*, 2008; Obraztsova *et al.*, 2013; Skowron, 2015; Skowron *et al.*, 2015]. When selecting a group of winning candidates different goals are interesting: proportional representation, diversity, or a short list [Elkind *et al.*, 2017]. We focus on the last scenario. Here the goal is to select the best (say highest-scoring) group of candidates.

Shortlisting comes very naturally in the context of selection committees, say for human resources departments that need to select, for a fixed number of positions, the best qualified applicants. A standard way of candidate selection in the context of shortlisting is to use scoring-based voting rules. We focus on the two most natural ones, namely SNTV (each voter may give one point to her most-liked candidate) and ℓ -Bloc (each voter may give one point to each of her ℓ most-liked candidates, so SNTV is the same as 1-Bloc). Obviously,

for such voting rules it is trivial to determine the score of each individual candidate. The main goal of our work is to model and understand coalitional manipulation in a computational sense. In this way, we complement well-known work: manipulation for singlewinner rules initiated by Bartholdi III *et al.* [1989], coalitional manipulation for singlewinner rules initiated by Conitzer *et al.* [2007], and manipulation for multiwinner rules initiated by Meir *et al.* [2008]. In coalitional manipulation scenarios, given full knowledge about other voters’ preferences, one has a set of manipulative voters who want to influence the election outcome in a favorable way by casting their votes strategically.

To come up with meaningful models for coalitional manipulation for multiwinner elections, however, we first have to identify the exact model and questions to be asked. Most natural extensions of coalitional manipulation for singlewinner elections or (non-coalitional) manipulation for multiwinner elections do not fit. Translated from the singlewinner perspective one would probably assume that the coalition agrees on making a distinguished candidate become part of the winners or that the coalition agrees on making a distinguished candidate group become part of the winners. The former is unrealistic because in multiwinner settings (especially for shortlisting) one certainly cares about more than just one winning candidate. The latter is problematic since there may be exponentially many “equally good” candidate groups for the coalition. Notably, this was not a problem in the singlewinner case where one can simply try all equally good candidates. The single-manipulator model for multiwinner rules of Meir *et al.* [2008] is a useful first step: The manipulator specifies the utility for each candidate; the utility for a candidate group is obtained by adding up the utilities of each group member. Aggregating utilities, however, becomes non-trivial for a coalition of manipulators which may have totally different utilities for single candidates but still have strong incentives to work together. Besides formalizing this either in a utilitarian or egalitarian way, our model also aims to distinguish between rather optimistic and pessimistic manipulators; leading to significant differences in the computational complexity results. Technically, this requires to study tie-breaking mechanisms and (winning) group evaluation functions.

Related Work. To the best of our knowledge, there is no previous work on coalitional manipulation in the context of multiwinner elections. We refer to recent textbooks for an

overview of the huge literature on singlewinner (coalitional) manipulation [Rothe, 2015; Brandt *et al.*, 2016]. Most relevant to this paper, Lin [2011] proved that coalitional manipulation in singlewinner elections under ℓ -Approval is tractable. Meir *et al.* [2008] introduced (non-coalitional) manipulation for multiwinner elections. While identifying several NP-hard voting rules they showed that manipulation remains polynomial-time solvable for Bloc (as multiwinner extension of ℓ -Approval). Obraztsova *et al.* [2013] extended the latter result for different tie-breaking strategies and identified further tractable special cases of multiwinner scoring rules, but conjectured manipulation to be hard in general for (other) scoring rules. Summarizing, Bloc is simple but comparably well-studied and was, hence, selected as show-case for our study of the presumably harder coalitional manipulation.

Organization. We develop a model for coalitional manipulation for multiwinner elections (Section 3). Its variants respect different ways of evaluating candidate groups (utilitarian vs. egalitarian) and different kinds of manipulators behavior (optimistic vs. pessimistic). On the technical side, we first present algorithms and complexity results for computing the output of several natural tie-breaking rules modeling optimistic or pessimistic manipulators (Section 4). Then, we explore the computational complexity of coalitional manipulation, using ℓ -Bloc as show-case (Section 5). We refer to the conclusion and Table 1 for a detailed overview of our findings (Section 6). Due to space constraints, we omit many proofs.

2 Preliminaries

For a positive integer n , we define $[n] := \{1, \dots, n\}$. A problem parameterized by ρ is *fixed-parameter tractable* (in FPT) if it is solvable in time $f(\rho) \cdot |I|^{O(1)}$; $W[t]$ -hard problems are presumably not in FPT (see [Cygan *et al.*, 2015] for details).

An *election* (C, V) consists of a set C of m candidates and a multiset V of n votes. Votes are linear orders on C . For example, for $C = \{c_1, c_2, c_3\}$ we write $c_1 \succ_v c_2 \succ_v c_3$ to express that candidate c_1 is most preferred and candidate c_3 least preferred in vote v . We omit the subscripts if the corresponding vote is clear from the context.

A *multiwinner voting correspondence* is a function that, given an election (C, V) and an integer $k \in [|C|]$, outputs a family of co-winning k -egroups. We use k -egroup as an abbreviation for k -excellence-group and do not use the established term “committee” because in shortlisting applications “committee” rather refers to voters and not to candidates.

We consider *scoring rules* which assign points to candidates based on their positions in the votes. We denote by $\text{score}(c)$ the total number of points that candidate c obtains and use $\text{score}_{V'}(c)$ when restricting the election to a subset $V' \subset V$ of voters. A (multiwinner) scoring rule selects a family \mathcal{X} of co-winning k -egroups with the maximum total sum of scores. It holds that $X \in \mathcal{X}$ if and only if $\forall c \in X, c' \in C \setminus X : \text{score}(c) \geq \text{score}(c')$. We focus on ℓ -Bloc multiwinner voting correspondences which assign for each vote one point to each of the top ℓ candidates.¹

¹The case where ℓ coincides with the size k of the egroup is typically referred to as Bloc; SNTV equals 1-Bloc.

Example 1. Given a set of candidates $C = \{c_1, c_2, c_3, c_4\}$, a multiset of votes $V = \{v_1, v_2, v_3\}$, where $v_1 = v_2 = c_1 \succ c_2 \succ c_3 \succ c_4$ and $v_3 = c_2 \succ c_3 \succ c_4 \succ c_1$, and egroup size $k = 2$, SNTV selects $\{c_1, c_2\}$ whereas 3-Bloc selects $\{c_2, c_3\}$.

To select a single k -egroup from the set of co-winning k -egroups one has to consider tie-breaking rules. A *multiwinner tie-breaking rule* is a mapping that, given an election and a family of co-winning k -egroups, outputs a single k -egroup. Among them there is a set of natural rules that is of particular interest in order to model the behavior of manipulative voters. Indeed, in case of a single manipulator *pessimistic* tie-breaking as well as *optimistic* tie-breaking has been considered in addition to lexicographic and randomized tie-breaking [Meir *et al.*, 2008; Obraztsova *et al.*, 2013]. To model optimistic (resp. pessimistic) tie-breaking in a non-trivial manner², we use the model introduced by Obraztsova *et al.* [2013] in which a manipulative voter v is described not only by the preference ordering \succ_v of the candidates but also by a utility function $u : C \rightarrow \mathbb{N}$. To cover this in the tie-breaking process, *coalition-specific* tie-breaking rules get in addition (to the election containing non-manipulative and manipulative votes and the co-winning excellence-groups) the manipulators’ utility functions in the input. The formal implementations of these rules and their properties are discussed in Subsection 3.2.

3 Model for Coalitional Manipulation

In this section, we formally define and explain our model and the respective variants. To this end, we discuss how we evaluate a k -egroup in terms of utility for a coalition of manipulators and introduce tie-breaking rules which model optimistic or pessimistic viewpoints of the manipulators.

3.1 Evaluating k -egroups

As already discussed in the introduction, one should not extend the model of coalitional manipulation for singlewinner elections to multiwinner elections in the most simple way (e.g. by assuming that manipulators agree on some distinguished candidate or on some distinguished egroup). Instead, we follow Meir *et al.* [2008] and assume we are given manipulator’s utility functions over the candidates and a utility level which, if achieved, indicates a successful manipulation. Meir *et al.* [2008] compute the utility of an egroup by summing up the utility values the manipulator assigns to each member of the egroup.

At first glance, summing up the utility values assigned by each manipulator to each member of an egroup seems to be

²We can not simply use ordinal preferences: Obraztsova *et al.* [2013] observed that already in case of a single manipulator one cannot simply set the fixed lexicographic order of the manipulators’ preferences (resp. the reverse of it) over candidates to model optimistic (resp. pessimistic) tie-breaking. For example, it is a strong restriction to assume that a manipulator would always prefer its first choice together with its fourth choice towards its second choice together with its third choice. It might be that only its first choice is really acceptable (in which case the assumption is reasonable) or that the first three choices are comparably good but the fourth choice absolutely unacceptable (in which case the assumption is wrong).

$u(\cdot)$	c_1	c_2	c_3	c_4
w_1	18	101	2	90
w_2	20	1	101	90

Figure 1: Example utilities two manipulators w_1, w_2 give to four candidates, c_1, c_2, c_3, c_4 .

the most natural extension for a coalition of manipulators. However, this *utilitarian variant* allows single manipulators to gain no utility. In extreme cases it could even be that some manipulator is worse off compared to voting sincerely. We consider the election from Example 1 and introduce two manipulators w_1, w_2 . We are interested in finding the best 2-egroup according to the utilitarian evaluation variant under SNTV voting correspondence. The utilities manipulators give to the candidates are depicted in Figure 1. Observe that if the manipulators vote sincerely, they give a point to c_2 and c_3 . Combining manipulators' votes with the non-manipulative ones, the winning 2-egroup consists of candidates c_1 and c_2 (there is no tie between candidates). The utilitarian value of such a group is equal to 140. Manipulator w_1 's utility is 119. However, both manipulators can do better by giving their points to candidate c_4 . Then, the winners are candidates c_1 and c_4 giving the total utility of 218. Observe that in spite of growth of the total utility, the utility value gained by w_1 , 108, is lower than in the case of sincere voting.

Despite the fact that w_1 is worse off compared to sincere voting in the above example, the utilitarian viewpoint is justified if the manipulators are able to compensate utility of an egroup e.g. by paying money to each other. For cases where manipulators cannot compensate utility of an egroup, we introduce two egalitarian evaluation variants. The (*egroup-wise*) *egalitarian* variant aims for maximizing the minimum satisfaction of the manipulators with the whole k -egroup. The *candidate-wise egalitarian* variant aims for maximizing the minimum satisfaction of the manipulators with the every single candidate. There is no "candidate-wise utilitarian" variant since this would be equivalent to the normal utilitarian variant.

We formalize the described variants of k -egroup evaluation (for w manipulators) with the definitions below.

Definition 1. Given a set of candidates C , a k -egroup $S \subseteq C$, $|S| = k$, and a family of manipulators' utility functions $U = \{u_1, u_2, \dots, u_w\}$ where $u_i: C \rightarrow \mathbb{N}$, we consider the following functions:

- $\text{util}_U(S) := \sum_{u \in U} \sum_{c \in S} u(c)$,
- $\text{egal}_U(S) := \min_{u \in U} \sum_{c \in S} u(c)$,
- $\text{candegal}_U(S) := \sum_{c \in S} \min_{u \in U} u(c)$.

Intuitively, they determine the utility of k -egroup S according to, respectively, the utilitarian and the egalitarian variant of evaluating S by a group of w manipulators (identifying manipulators with their utility functions). We omit subscripts when a family of utility functions is obvious.

By taking the sum of utilities (resp. the minimum utility) over all manipulators as utility for each candidate we obtain a helpful observation allowing us to simplify utilitarian and candidate-wise egalitarian approaches.

Observation 1. One can assume without loss of generality that every manipulator has the same utility function when using utilitarian or candidate-wise egalitarian evaluation.

3.2 Breaking Ties

Before formally defining our tie-breaking rules, we briefly discuss some necessary notation and central concepts. Consider an election (C, V) , a size k for the egroup to be chosen, and a scoring-based voting correspondence \mathcal{R} . Observe that we can partition the set of candidates C into three sets C^+, P , and C^- as follows. The set C^+ contains the *confirmed candidates*, that is, candidates who are in all co-winning k -egroups. The set P contains the *pending candidates*, that is, candidates that are only in some co-winning k -egroups. The set C^- contains the *rejected candidates*, that is, candidates that are in no co-winning k -egroup. Observe that $|C^+| \leq k$, $|C^+ \cup P| \geq k$, every candidate from $P \cup C^-$ receives fewer points than every candidate from C^+ , and all candidates in P receive the same number of points.

In order to define optimistic and pessimistic rules, we assume that in addition to C^+, P , and C^- we are given some family of utility functions.

Lexicographic \mathcal{F}_{lex} . Ties are broken lexicographically with respect to a predefined ordering $>_{\mathcal{L}}$ of the candidates from C . That is, \mathcal{F}_{lex} selects all candidates from C^+ together with the top $k - |C^+|$ candidates from P with respect to $>_{\mathcal{L}}$.

Optimistic (resp. pessimistic) $\mathcal{F}_{\text{opt}}^{\text{eval}}$ (resp. $\mathcal{F}_{\text{pess}}^{\text{eval}}$), $\text{eval} \in \{\text{util, egal, candegal}\}$. Some k -egroup S wins according to $\mathcal{F}_{\text{opt}}^{\text{eval}}$ (resp. $\mathcal{F}_{\text{pess}}^{\text{eval}}$) if and only if $C^+ \subseteq S \subseteq (C^+ \cup P)$ and there is no other k -egroup S' with $C^+ \subseteq S' \subseteq (C^+ \cup P)$ and $\text{eval}(S') > \text{eval}(S)$ (resp. $\text{eval}(S') < \text{eval}(S)$).

The way we define the optimistic and pessimistic tie-breaking rules allows for co-winning k -egroups, but every co-winning k -egroup is evaluated equally by the coalition. For simplicity, one may assume that an arbitrary co-winning k -egroup is selected.

3.3 Limits of Lexicographic Tie-Breaking

From the above discussion, we can conclude that lexicographic tie-breaking is very straight-forward in the case of scoring-based multiwinner voting correspondences. Basically any subset of the desired cardinality from the set of pending candidates can be chosen. In particular, the best candidates with respect to the given ordering can be chosen. We remark that this is not necessarily the case for general multiwinner voting correspondences.

It remains to be clarified whether one can find a reasonable ordering of the pending candidates in order to model optimistic or pessimistic tie-breaking rules in a simple way. Using Observation 1 we show that this is possible for every $\mathcal{F}_{\text{bhav}}^{\text{eval}}$, $\text{eval} \in \{\text{util, candegal}\}$, $\text{bhav} \in \{\text{opt, pess}\}$. We have a counterexample for $\text{eval} = \text{egal}$.

We say tie-breaking rule \mathcal{F} *simulates* tie-breaking rule \mathcal{F}' with respect to a set C of candidates, a family U of utility functions, and an evaluation function eval if and only if $\text{eval}_U(S) = \text{eval}_U(S')$ for every integer k , $0 < k < |C|$, and for every partition of C into C^+, P , and C^- (as described above) where S is the k -egroup selected by \mathcal{F} and S' is the k -egroup selected by \mathcal{F}' .

Proposition 1. Let C be a set of candidates, U be a family of utility functions, and $\text{bhav} \in \{\text{opt}, \text{pess}\}$.

One can compute in $O(m \cdot (w + \log m))$ time a lexicographic tie-breaking rule that simulates $\mathcal{F}_{\text{bhav}}^{\text{eval}}$ with respect to C , U , and $\text{eval} \in \{\text{util}, \text{candegal}\}$.

There is no lexicographic tie-breaking rule that simulates $\mathcal{F}_{\text{bhav}}^{\text{eval}}$ with respect to C , U , and $\text{eval} = \text{egal}$.

Proof. If $\text{eval} \in \{\text{util}, \text{candegal}\}$, then we can assume that every manipulator has the same utility function (Observation 1). We compute such utility function u^* which is “equivalent” to U in $O(m \cdot w)$ time. We say an order $>_{\mathcal{L}}$ of the candidates is *consistent* with some utility function u if $c >_{\mathcal{L}} c'$ implies $u(c) \geq u(c')$ for optimistic tie-breaking and $c >_{\mathcal{L}} c'$ implies $u(c) \leq u(c')$ for pessimistic tie-breaking. Any lexicographic tie-breaking rule defined by an order $>_{\mathcal{L}}$ that is consistent with the utility function u^* simulates $\mathcal{F}_{\text{bhav}}^{\text{eval}}$. We compute a consistent ordering by sorting the candidates according to u^* in $O(m \cdot \log m)$ time.

We build a counterexample to show impossibility of simulation optimistic or pessimistic tie-breaking by a lexicographic tie-breaking rule for egalitarian evaluation. Consider four candidates, c_1 to c_4 , two manipulators, w_1 and w_2 , with utility functions as depicted in Figure 1, and 2-egroup. When pending set P contains all candidates then the optimal optimistic (resp. pessimistic) 2-egroup consists of candidates c_1 and c_4 (resp. c_3). This induces a lexicographic order preferring candidates c_1 and c_4 (resp. c_3) to the others. However, if confirmed set P contains only candidate c_2 (resp. c_4) and the rest are pending ones then it is optimal to choose candidate c_3 (resp. c_2) to the solution; A contradiction to the order induced by the first case. \square

4 Complexity of Tie-Breaking

It is natural to ask whether the proposed tie-breaking rules are practical in terms of computational complexity. If not, then there is no hope for coalitional manipulation because tie-breaking is a subtask to be solved by the manipulators.

By definition, we can apply every lexicographic tie-breaking rule that is defined through some predefined ordering of the candidates in linear time. Hence, we focus on the rules that model optimistic or pessimistic manipulators. To this end, we analyze the following computational problem.

$\mathcal{F}_{\text{bhav}}^{\text{eval}}$ -TIE-BREAKING ($\mathcal{F}_{\text{bhav}}^{\text{eval}}$ -TB),
($\text{eval} \in \{\text{util}, \text{egal}, \text{candegal}\}$, $\text{bhav} \in \{\text{opt}, \text{pess}\}$)

Input: A set of candidates C partitioned into a set P of pending candidates and a set C^+ of confirmed candidates, the size $k > |C^+|$ of the excellence-group, a family of manipulators’ utility functions $U = \{u_1, u_2, \dots, u_w\}$ where $u_i : C \rightarrow \mathbb{N}$, and a non-negative integer q .

Question: Is there a size- k set $S \subseteq C$ such that S wins according to $\mathcal{F}_{\text{bhav}}^{\text{eval}}$, $C^+ \subseteq S$, and $\text{eval}(S) \geq q$.

4.1 Utilitarian and Candidate-wise Egalitarian: Tie-Breaking is Easy

As a warm-up, we show that tie-breaking can be applied and evaluated efficiently if k -egroups are evaluated according to the utilitarian or candidate-wise egalitarian variant. The corresponding result follows almost directly from Proposition 1.

Corollary 1. One can solve $\mathcal{F}_{\text{bhav}}^{\text{eval}}$ -TIE-BREAKING in $O(m \cdot (w + \log m))$ time for $\text{eval} \in \{\text{util}, \text{candegal}\}$, $\text{bhav} \in \{\text{opt}, \text{pess}\}$.

Proof. The algorithm works in two steps. First, compute a lexicographic tie-breaking rule \mathcal{F}_{lex} that simulates $\mathcal{F}_{\text{bhav}}^{\text{eval}}$ in $O(m \cdot (w + \log m))$ time as described in Proposition 1. Second, apply \mathcal{F}_{lex} in $O(|P|)$ time and evaluate the resulting k -egroup in $O(k \cdot w)$ time. \square

4.2 Egalitarian: Being Optimistic is Hard

In the remainder of this section, we consider the defined tie-breaking rules when applied for searching a k -egroup evaluated in terms of the egalitarian variant. First, we show that applying and evaluating egalitarian tie-breaking is computationally easy for pessimistic manipulators but computationally intractable for optimistic manipulators even if the size of the egroup is small. Being pessimistic, the main idea is to “guess” the manipulator that is least satisfied and select the candidates appropriately. We show the computational hardness of the optimistic case via a reduction from SET COVER.

Theorem 1. One can solve $\mathcal{F}_{\text{pess}}^{\text{egal}}$ -TIE-BREAKING in $O(w \cdot m \log m)$ time, but $\mathcal{F}_{\text{opt}}^{\text{egal}}$ -TIE-BREAKING is NP-hard and W[2]-hard when parameterized by k even if $u_{\max} = 1$.

A small egroup size (alone) obviously does not make the problem computationally tractable even for very simple utility functions. Next, using parameterized reduction from MULTICOLORED CLIQUE, we show that there is still no hope for fixed-parameter tractability even for the combined parameter $w + k$, which covers situations where few manipulators are going to influence an election for a small egroup.

Theorem 2. Parameterized by $w + k$, $\mathcal{F}_{\text{opt}}^{\text{egal}}$ -TIE-BREAKING is W[1]-hard.

Finally, we show that our problem becomes tractable when parameterized by $w + u_{\max}$, covering situations with few manipulators that have simple utility functions e.g. if few voters have 0/1 utility functions. We use an Integer Linear Program with at most $O(u_{\max}^w)$ integer variables to solve our problem and apply the famous result of Lenstra [1983] to obtain FPT.

Theorem 3. Parameterized by $w + u_{\max}$, $\mathcal{F}_{\text{opt}}^{\text{egal}}$ -TIE-BREAKING is fixed-parameter tractable.

5 Complexity of (Coal.) Manipulation

In this section, we analyze the computational difficulty of voting strategically for a coalition of manipulators. To this end, we formalize our central computational problem. Let \mathcal{R} be a multiwinner voting correspondence and let \mathcal{F} be a multiwinner tie-breaking rule.

\mathcal{R} - \mathcal{F} -eval-COALITIONAL MANIPULATION

(\mathcal{R} - \mathcal{F} -eval-CM), ($\text{eval} \in \{\text{util}, \text{egal}, \text{candegal}\}$)

Input: An election (C, V) , a committee size k , a coalition size w , for each manipulator $x \in [w]$ a utility function $u_x : C \rightarrow \mathbb{N}$, and a non-negative integer q .

Question: Is there a multiset W of manipulative votes over C such that $S \subseteq C$ wins the election $(C, V \cup W)$ under \mathcal{R} and \mathcal{F} , $|S| = k$, $|W| = w$, and $\text{eval}(S) \geq q$?

5.1 Utilitarian and Candidate-wise Egalitarian: Manipulation is Tractable

We show that ℓ -Bloc- \mathcal{F} -eval-CM can indeed be solved in polynomial time for any constant $\ell \in \mathbb{N}$, any eval $\in \{\text{util}, \text{candegal}\}$, and any $\mathcal{F} \in \{\mathcal{F}_{\text{lex}}, \mathcal{F}_{\text{opt}}^{\text{eval}}, \mathcal{F}_{\text{pess}}^{\text{eval}}\}$.

For Bloc, that is, $\ell = k$, we give a quadratic time algorithm. We start with an algorithm for a version of problem ℓ -Bloc- \mathcal{F} -eval-CM where all manipulators have to cast exactly the same vote and show that this algorithm can also be used to solve Bloc- \mathcal{F} -eval-CM. The main idea is to “guess” the minimum score which a member of the egroup can achieve and then carefully (with respect to tie-breaking) select the best candidates that can reach this score.

Theorem 4. *One can solve ℓ -Bloc- \mathcal{F} -eval-COALITIONAL MANIPULATION with consistent manipulators in $O((n + m)n)$ time for any eval $\in \{\text{util}, \text{candegal}\}$ and $\mathcal{F} \in \{\mathcal{F}_{\text{opt}}^{\text{eval}}, \mathcal{F}_{\text{pess}}^{\text{eval}}\}$.*

For Bloc, we can show that manipulators can always vote identically to achieve an optimal k -egroup. In a nutshell, for every excellence-group the manipulators can only do better by voting exactly for its members, because Bloc is monotonic. This leads us to the following corollary.

Corollary 2. *One can solve Bloc- \mathcal{F} -eval-COALITIONAL MANIPULATION in $O((n + m)n)$ time for any eval $\in \{\text{util}, \text{candegal}\}$ and $\mathcal{F} \in \{\mathcal{F}_{\text{opt}}^{\text{eval}}, \mathcal{F}_{\text{pess}}^{\text{eval}}\}$.*

To complete our analysis we provide a polynomial-time algorithm which also works for a general case where the number of approvals may be different to the size of the excellence-group and manipulators can vote differently from each other.

Theorem 5. *One can solve ℓ -Bloc- \mathcal{F} -util-COALITIONAL MANIPULATION in $O(kn^2m^3)$ time for any $\mathcal{F} \in \{\mathcal{F}_{\text{lex}}, \mathcal{F}_{\text{opt}}^{\text{util}}, \mathcal{F}_{\text{pess}}^{\text{util}}\}$.*

Proof. The basic idea of our algorithm is to fix certain parameters of the solution and then reduce the resulting subproblem to a KNAPSACK variant with polynomial-sized weights.

In the first phase of the algorithm, we iterate through all possible value combinations of the following two parameters: (1) the lowest final score $z < |V \cup W|$ of any member of the k -egroup and (2) the candidate \hat{c} which is the least-preferred member of the k -egroup with final score z with respect to the tie-breaking rule \mathcal{F} .

Having fixed z , let C^+ denote the set of candidates which get at least $z + 1$ approvals from the non-manipulative winners or which are preferred to \hat{c} with respect to \mathcal{F} and get exactly z approvals from the non-manipulative winners. Assuming that the combination of parameter values is correct, all candidates from $C^+ \cup \{\hat{c}\}$ must belong to the k -egroup. Set $k^+ := |C^+|$. For sanity, check whether $k^+ < k$ and discard the corresponding combination of solution parameter values if the check fails. Next, we ensure that \hat{c} obtains final score exactly z . If \hat{c} receives less than $z - w$ or more than z approvals from non-manipulative voters, then we discard this combination of solution parameter values. Otherwise, let $\hat{s} := z - \text{score}_V(\hat{c})$ denote number of additional approvals candidate \hat{c} needs in order to get final score z . Let $k^* := k - k^+ - 1$ be the number of remaining (not yet fixed)

members of the k -egroup. Let $s^* := w \cdot \ell - \hat{s}$ be the number of approvals to be distributed to candidates in $C \setminus \{\hat{c}\}$.

We analyze the remaining subproblem to be solved in the second phase. The manipulators influence further k^* candidates to join the k -egroup (so far only consisting of $C^+ \cup \{\hat{c}\}$) and they distribute exactly s^* approvals in total to candidates in $C \setminus \{\hat{c}\}$ and at most w approvals per candidate. To this end, let C^* denote the set of candidates which can possibly join the k -egroup. For each candidate $c \in C \setminus (C^+ \cup \{\hat{c}\})$ it holds that $c \in C^*$ if and only if $z - w \leq \text{score}_V(c) \leq z - 1$ if c is preferred to \hat{c} with respect to \mathcal{F} , or $z - w + 1 \leq \text{score}_V(c) \leq z$ if \hat{c} is preferred to c with respect to \mathcal{F} .

The simplest idea is to select the k^* elements from C^* which have the highest value (that is, utility) for the coalition. However, there can be two issues. First, s^* might be too small, that is, there are too few approvals to ensure that each of the k^* most-valued candidates gets final score at least z (resp. at least $z + 1$). Second, s^* might be too large, that is, there are too many approvals to be distributed so that there is no way to do this without causing unwanted candidates to get a final score of at least z (resp. at least $z + 1$). Fortunately, we can show that at most one of both issues is possible for a given combination of solution parameter values and that we can easily see which one. In either case, we can reduce the remaining problem to an instance of EXACT k -ITEM KNAPSACK. In both cases, the value of each $c^* \in C^*$ is set to its utility for the coalition (this is well-defined: due to Observation 1 we can assume that each manipulator has the same utility function) and the weight is set to $z - \text{score}_V(c^*)$ if c^* is preferred to \hat{c} with respect to \mathcal{F} and otherwise to $(z + 1) - \text{score}_V(c^*)$. The capacity is set to s^* if $s^* \leq w \cdot k^*$ and to $k^* \cdot (w - 1) - (s^* - s^+ + k^*)$ if $s^* > w \cdot k^*$, where s^+ denotes the total number of approvals which can be safely distributed to candidates in $C \setminus \{\hat{c}\}$ without causing one of the candidates from C^* to reach score at least z (resp. at least $z + 1$). We omit technical details (especially for the case where $s^* \leq w \cdot k^*$).

As for the running time, we iterate through at most $n + w \leq 2n$ values for z , at most m values for c^+ , and we can solve EXACT k -ITEM KNAPSACK (the variant of KNAPSACK where the solution needs to have exactly k items) in $O(knm^2)$ time by using dynamic programming via weights [Kellerer *et al.*, 2004, Chapter 9.7.3] (note that the overall sum of weights is trivially bounded by $n \cdot m$). \square

5.2 Egalitarian: Hard for Simple Tie-Breaking

In Subsection 4.2, we showed that already breaking ties might be computationally intractable. The intractability results only hold with respect to egalitarian evaluation and optimistic manipulators. We now show that this intractability transfers to coalitional manipulation for any tie-breaking rules and egalitarian evaluation. This includes the pessimistic egalitarian case which we consider to be highly relevant as it models searching for a “safe” strategic voting strategy in a very natural way.

Proposition 2. *There is a polynomial-time many-one reduction from $\mathcal{F}_{\text{opt}}^{\text{egal}}$ -TIE-BREAKING to ℓ -Bloc- \mathcal{F} -egal-COALITIONAL MANIPULATION for any tie-breaking rule \mathcal{F} .*

$\mathcal{F}_{\text{behav}}^{\text{eval}}$ -TIE-BREAKING, easy cases (Cor. 1):		
settings (evaluation, behavior)	complexity	
utilitarian/cand.wise egalitarian, optimistic/pessimistic	$O(m \cdot (w + \log m))$	
egalitarian, pessimistic	$O(w \cdot m \log m)$	
$\mathcal{F}_{\text{opt}}^{\text{egal}}$ -TIE-BREAKING (egalitarian, optimistic):		
parameters, restrictions	complexity	ref.
general	NP-complete	Thm. 1
$k, 0/1$ utilities	W[2]-hard	Thm. 1
$w + k$	W[1]-hard	Thm. 2
$w + u_{\max}$	ILP-FPT	Thm. 3

ℓ -Bloc- \mathcal{F} -eval-COALITIONAL MANIPULATION utilitarian/cand.wise egalitarian, optimistic/pessimistic:		
restrictions	complexity	ref.
general	$O(kn^2m^3)$	Thm. 5
consistent manipulators	$O((n+m)n)$	Thm. 4
$\ell = k$	$O((n+m)n)$	Cor. 2
ℓ -Bloc- \mathcal{F} -eval-COALITIONAL MANIPULATION egalitarian, optimistic/pessimistic:		
parameters, restrictions	complexity	ref.
general	NP-complete	Cor. 3
$k, 0/1$ utilities	W[2]-hard	Cor. 3
$w + k$	W[1]-hard	Cor. 3
$w + u_{\max}$	ILP-FPT	Thm. 6

Table 1: Computational complexity of tie-breaking and coalitional manipulation. Our results for ℓ -Bloc hold for any $\ell \geq 1$, and thus cover SNTV. The parameters are the size k of the excellence-group, the number w of manipulators, and the maximum utility value $u_{\max} := \max_{i \in [w], c \in C} u_i(c)$. Furthermore, m is the number of candidates and n is the number of voters.

The reduction keeps the egroup size and the number of manipulators in the ℓ -Bloc- \mathcal{F} -egal-CM is upper-bounded by $w + k$. This leads us to the following corollary.

Corollary 3. *Parameterized by $w + k$, ℓ -Bloc- \mathcal{F} -egal-COALITIONAL MANIPULATION is W[1]-hard. Parameterized by k , ℓ -Bloc- \mathcal{F} -egal-COALITIONAL MANIPULATION is W[2]-hard even if $u_{\max} = 1$.*

Finally by using ideas from Theorem 5 and an adaption of the ILP from Theorem 3 as subroutine, we can show that fixed-parameter tractability for the combined parameter $w + u_{\max}$ transfers to coalitional manipulation.

Theorem 6. *Parameterized by $w + u_{\max}$, ℓ -Bloc- \mathcal{F} -egal-COALITIONAL MANIPULATION is fixed-parameter tractable.*

6 Conclusion

We developed a new model for and started the first systematic study of coalitional manipulation for multiwinner elections. Our analysis shows that multiwinner coalitional manipulation requires models which are significantly more complex than those for singlewinner coalitional manipulation or multiwinner (non-coalitional) manipulation. On the one hand, we generalize tractability results for coalitional manipulation of ℓ -Approval by Conitzer *et al.* [2007] and Lin [2011] and for (non-coalitional) manipulation of Bloc by Meir *et al.* [2008] and Obraztsova *et al.* [2013] to tractability of coalitional manipulation of ℓ -Bloc in case of utilitarian or candidate-wise egalitarian evaluation of egroups. On the other hand, we show that coalitional manipulation becomes computationally intractable in case of egalitarian evaluation of egroups.

Let us discuss particular findings in more details (Table 1 surveys our results in full detail). We studied lexicographic, optimistic, and pessimistic tie-breaking and showed that, with the exception of egalitarian group evaluation, winner groups can be determined very efficiently. The intractability (NP-hardness, parameterized hardness (W-hardness)) for the egalitarian case, however, turns out to hold even for

quite restricted scenarios (including two parameter combinations). We also demonstrate that most tie-breaking rules can be “simulated” by (carefully chosen) lexicographic tie-breaking, again except for the egalitarian case. Interestingly, however, the hardness of egalitarian tie-breaking holds only for the optimistic case while for the pessimistic case it is efficiently computable. Hardness for the egalitarian optimistic scenario, however, translates into hardness results for coalitional manipulation *regardless* of the specific tie-breaking rule. On the contrary, coalitional manipulation becomes tractable for all other evaluation functions (different from the “global” egalitarian one, that is, it is tractable for a “candidate-wise” egalitarian one).

In our study, we entirely focus on shortlisting as one of the simplest tasks for multiwinner elections to design our evaluation function. It seems interesting and non-trivial to develop models for multiwinner rules that aim for proportional representation or diversity. For shortlisting, extending our studies to non-approval-like scoring-based voting rules would be a natural next step. In this context, already seeing what happens if one extends the set of individual scores from being only 0 or 1 to more (but few) numbers might be of interest. Moreover, we focused on deterministic tie-breaking mechanisms, ignoring randomized tie-breaking.

Beyond the above, further research on manipulators behavior directing towards game-theoretic aspects seems promising as well. Intuitively, the utility for every voter that is a part of the manipulating coalition should not be below the utility the voter receives when voting sincerely. This is of course only a necessary condition to ensure the stability of a coalition. A more sophisticated analysis of stability needs to consider game-theoretic aspects such as Nash or core stability.

Acknowledgments

RB was from 09/2016 to 09/2017 at the University of Oxford, supported by the DFG fellowship BR 5207/2. AK was supported by the DFG project AFFA (BR 5207/1 and NI 369/15).

References

- [Aziz *et al.*, 2015] Haris Aziz, Serge Gaspers, Joachim Gudmundsson, Simon Mackenzie, Nicholas Mattei, and Toby Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*, pages 107–115, 2015.
- [Aziz *et al.*, 2017a] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- [Aziz *et al.*, 2017b] Haris Aziz, Edith Elkind, Piotr Faliszewski, Martin Lackner, and Piotr Skowron. The Condorcet principle for multiwinner elections: From shortlisting to proportionality. In *Proceedings of the 26th International Conference on Artificial Intelligence, IJCAI'17*, 2017. This proceedings.
- [Barberà and Coelho, 2008] Salvador Barberà and Danilo Coelho. How to choose a non-controversial list with k names. *Social Choice and Welfare*, 31(1):79–96, 2008.
- [Barberà and Coelho, 2010] Salvador Barberà and Danilo Coelho. On the rule of k names. *Games and Economic Behavior*, 70(1):44 – 61, 2010.
- [Barrot *et al.*, 2013] Nathanaël Barrot, Laurent Gourvès, Jérôme Lang, Jérôme Monnot, and Bernard Ries. Possible winners in approval voting. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory, ADT '13*, pages 57–70, 2013.
- [Bartholdi III *et al.*, 1989] John J. Bartholdi III, Craig A. Tovey, and Michael A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [Brandt *et al.*, 2016] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Conitzer *et al.*, 2007] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.
- [Cygan *et al.*, 2015] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [Elkind *et al.*, 2017] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3):599–632, 2017.
- [Faliszewski *et al.*, 2016] Piotr Faliszewski, Piotr Skowron, Arkadii M. Slinko, and Nimrod Talmon. Multiwinner analogues of the plurality rule: Axiomatic and algorithmic perspectives. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence, AAAI'16*, pages 482–488, 2016.
- [Kellerer *et al.*, 2004] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Knapsack Problems*. Springer, 2004.
- [Lenstra, 1983] Hendrik W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.
- [Lin, 2011] Andrew Lin. The complexity of manipulating k -approval elections. In *Proceedings of the 3rd International Conference on Agents and Artificial Intelligence, ICAART '11*, pages 212–218, 2011.
- [Meir *et al.*, 2008] Reshef Meir, Ariel D. Procaccia, Jeffrey S. Rosenschein, and Aviv Zohar. Complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research*, 33(1):149–178, 2008.
- [Obraztsova *et al.*, 2013] Svetlana Obraztsova, Yair Zick, and Edith Elkind. On manipulation in multi-winner elections based on scoring rules. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '13*, pages 359–366, 2013.
- [Rothe, 2015] J. Rothe, editor. *Economics and Computation*. Springer, 2015.
- [Skowron *et al.*, 2015] Piotr Skowron, Piotr Faliszewski, and Arkadii M. Slinko. Achieving fully proportional representation: Approximability results. *Artificial Intelligence*, 222:67–103, 2015.
- [Skowron, 2015] Piotr Skowron. What do we elect committees for? A voting committee model for multi-winner rules. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, pages 1141–1147, 2015.