

Multiwinner Analogues of the Plurality Rule: Axiomatic and Algorithmic Perspectives

Piotr Faliszewski

AGH University
Krakow, Poland
faliszew@agh.edu.pl

Piotr Skowron

University of Oxford
Oxford, United Kingdom
p.k.skowron@gmail.com

Arkadii Slinko

University of Auckland
Auckland, New Zealand
a.slinko@auckland.ac.nz

Nimrod Talmon

TU Berlin
Berlin, Germany
nimrodtalmon77@gmail.com

Abstract

We characterize the class of committee scoring rules that satisfy the fixed-majority criterion. In some sense, the committee scoring rules in this class are multiwinner analogues of the single-winner Plurality rule, which is uniquely characterized as the only single-winner scoring rule that satisfies the simple majority criterion. We find that, for most of the rules in our new class, the complexity of winner determination is high (i.e., the problem of computing the winners is NP-hard), but we also show some examples of polynomial-time winner determination procedures, exact and approximate.

Introduction

The scoring rules in general, and the Plurality rule specifically, are among the most basic and the best studied single-winner voting rules. However, our understanding of their recently-introduced multiwinner analogues, committee scoring rules (Elkind et al. 2014), is very limited. In this paper, we attempt to rectify this situation by asking a seemingly very innocuous question: what is the committee scoring rule analogue of the Plurality rule? Using an axiomatic approach, we find a rather surprising answer. Not only is there a whole class of committee scoring rules that correspond to the Plurality rule, but also one of the most natural candidates to be the multiwinner Plurality, the single non-transferable vote rule (the SNTV rule), falls short on our criterion. On the other hand, the Bloc rule turns out to be quite a satisfying candidate, but certainly not the only one. In addition to our axiomatic study, we provide an algorithmic analysis of this new class of committee scoring rules. In particular, we show that it can be seen as a subfamily of the OWA-based rules of Skowron, Faliszewski, and Lang (2015) (also studied by Aziz et al. (2015b; 2015a); see also the work of Kilgour (2010) for a more general overview of approval-based multiwinner rules). However, the hardness results for general OWA rules do not translate directly to our case (and some indeed do not even hold). On the side, we provide an axiomatic characterization of the Bloc rule (among the committee scoring rules).

Let us now describe our setting more precisely. In a multiwinner election, each voter ranks the candidates from the

most desired one to the least desired one, and the goal is to pick a committee of a given size k that, in some sense, best matches the voters' preferences. Naturally, the exact meaning of the phrase "best matches" depends strongly on the application at hand, as well as on the societal conventions and understanding of fairness. For example, if we are to choose a size- k parliament, then it is important to guarantee proportional representation; if the goal is to pick a group of products to offer to customers, then it might be important to maintain diversity of the offer; if we are to shortlist a group of candidates for a job, then it is important to focus on the quality of the selected candidates regardless of how similar some of them might be.

In effect, there is quite a variety of multiwinner voting rules. For example, under the SNTV rule, the winning committee consists of k candidates who are ranked first more frequently than others. Under the Bloc rule, each voter gives one point to each candidate he or she ranks among his or her top k positions, and the committee consists of k candidates with the most points. Under the Chamberlin–Courant rule, the winning committee consists of k candidates such that each voter ranks his or her most preferred committee member as high as possible (for the exact definition see the original paper of Chamberlin and Courant (1983) or papers studying the rule's features and computational complexity (Procaccia, Rosenschein, and Zohar 2008; Lu and Boutilier 2011; Elkind et al. 2014; Skowron, Faliszewski, and Slinko 2015; Skowron and Faliszewski 2015)).

The three rules mentioned above are examples of committee scoring rules (a class of rules generalizing single-winner scoring rules to the multiwinner setting, recently introduced by Elkind et al. (2014); see the preliminaries for the definition).¹ Of course, there are natural multiwinner rules that cannot be expressed as committee scoring rules, such as the single transferable vote rule (the STV rule), the Monroe rule (Monroe 1995), or all the multiwinner rules based on the Condorcet principle (see, e.g., the works of Elkind et al. (2011), Fishburn (1981), and Gehrlein (1985)). Nonetheless, we believe that committee scoring rules form a very diverse class of voting rules that deserves a further study.

We ask for a committee scoring rule that can be seen as

¹Naturally, these rules were known much earlier than Elkind et al. (2014) introduced the unifying framework for them.

a “multiwinner analogue” of the Plurality rule. Intuitively, it might seem as if the SNTV rule were such a rule and the question were trivial. However, instead of following this intuition we take an axiomatic approach. We note that Plurality is the only single-winner scoring rule that has the simple majority property, i.e., that guarantees that if a candidate is ranked first by a simple majority of the voters, then he or she is the unique winner of the election. We ask for a committee scoring rule that has the fixed-majority criterion (a multiwinner analogue of the simple majority property, introduced by Debord (1993)), which requires that if there is a majority of voters each of whom ranks the same k candidates in the top k positions (perhaps in a different order), then these k candidates should form a unique winning committee.

The Bloc rule obviously satisfies the fixed-majority criterion. However, it turns out that Bloc is by far not the only such committee scoring rule and there is a whole family of them. We provide an (almost) full characterization of this family² and analyze the computational complexity of winner determination for rules in this family. Initially, we identify a slightly larger class of *top- k -counting* rules for which the score that a committee receives from a given voter is a function of the number of committee members that this voter ranks in the top k positions of its vote; we refer to this function as the *counting function*. We obtain the following main results:

1. For a large class of counting functions, top- k -counting rules are NP-hard to compute. There are, however, some polynomial-time computable ones (e.g., the Bloc rule and the Perfectionist rule that we introduce).
2. If the counting function is convex, then the top- k -counting rule that it defines satisfies the fixed-majority criterion (for a fairly intuitive relaxation of the convexity notion we get an “if and only if” result).
3. If the counting function is concave, then the rule it defines fails the fixed-majority criterion, but the rule seems to be easier computationally than in the convex case. We show an exact FPT algorithm for the parameterization by the number of voters and a polynomial-time $(1 - \frac{1}{e})$ -approximation algorithm.
4. Using the fixed-majority criterion and a certain monotonicity notion, we characterize the Bloc rule among the committee scoring rules.

All in all, we find that there is no single multiwinner analogue of the Plurality rule, even if we restrict ourselves to polynomial-time computable committee scoring rules. On the intuitive level SNTV is such a rule, and through our axiomatic consideration we show that Bloc and Perfectionist are also good candidates. We omit some proofs due to space constraints.

Preliminaries

An election is a pair $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ is a set of candidates and $V = (v_1, \dots, v_n)$ is a collection

²For technical reasons, we consider the case where there are at least twice as many candidates as the size of the committee.

of voters. The number $m = |C|$ will be fixed throughout the paper. Each voter v_i is associated with a preference order \succ_i in which v_i ranks the candidates from its most desirable one to its least desirable one. If X and Y are two (disjoint) subsets of C , then by $X \succ_i Y$ we mean that for each $x \in X$ and each $y \in Y$ it holds that $x \succ_i y$. For a positive integer t , we denote the set $\{1, \dots, t\}$ by $[t]$.

Single-Winner Voting Rules. A single-winner voting rule \mathcal{R} is a function that, given an election $E = (C, V)$, outputs a subset of those candidates that tie as winners. There is quite a variety of single-winner voting rules, but in this paper it suffices to consider the scoring rules. Given a voter v and a candidate c , we write $\text{pos}_v(c)$ to denote the position of c in v ’s preference order (e.g., if v ranks c first then $\text{pos}_v(c) = 1$). A scoring function for m candidates is a function $\gamma: [m] \rightarrow \mathbb{N}$ such that for each $i \in [m-1]$ we have $\gamma(i) \geq \gamma(i+1)$. Each scoring function γ defines a voting rule \mathcal{R}_γ as follows. Let $E = (C, V)$ be an election with m candidates.³ Under \mathcal{R}_γ , each candidate $c \in C$ receives $\text{score}(c) := \sum_{v \in V} \gamma(\text{pos}_v(c))$ points and the candidate with the highest number of points wins. (If there are several such candidates, then they all tie as winners; this view is known as the nonunique-winner model.) We often refer to the value $\text{score}(c)$ as the γ -score of c .

The following scoring functions are particularly interesting. The t -approval scoring function is defined as $\alpha_t(i) = 1$ for $i \leq t$ and $\alpha_t(i) = 0$ otherwise. For example, the Plurality rule is \mathcal{R}_{α_1} , the t -Approval rule is \mathcal{R}_{α_t} , and the Veto rule is $\mathcal{R}_{\alpha_{m-1}}$ (where m is the number of candidates). The Borda scoring function for m candidates is defined as $\beta_m(i) := m - i$, and \mathcal{R}_β is the Borda rule.

Multiwinner Voting Rules. A multiwinner voting rule \mathcal{R} is a function that, given an election $E = (C, V)$ and a number k representing the size of the desired committee, outputs a set of size- k subsets of C , the set of committees that tie as winners (naturally, in most practical cases we would hope to have a single winning committee).

We focus on committee scoring rules, introduced by Elkind et al. (2014). Consider an election $E = (C, V)$ and some committee S of a given size k . Let v be some voter in V . By $\text{pos}_v(S)$ we mean the sequence (i_1, \dots, i_k) that results from sorting the set $\{\text{pos}_v(c) : c \in S\}$ in increasing order. For example, if $C = \{a, b, c, d, e\}$, the preference order of v is $a \succ b \succ c \succ d \succ e$, and $S = \{a, c, d\}$, then $\text{pos}_v(S) = (1, 3, 4)$. If $I = (i_1, \dots, i_k)$ and $J = (j_1, \dots, j_k)$ are two increasing sequences of integers, then we say that I (weakly) dominates J (denoted $I \succeq J$) if $i_t \leq j_t$ for each $t \in [k]$. For positive integers m and k , $k \leq m$, by $[m]_k$ we mean the set of all increasing size- k sequences of integers from $[m]$.

Definition 1 (Elkind et al. (2014)). A committee scoring function for a multiwinner election with m candidates, where we seek a committee of size k , is a function

³Technically, \mathcal{R}_γ is defined only for elections with m candidates. Typically, however, we are interested in families of scoring functions, with one function for each number of candidates.

$f: [m]_k \rightarrow \mathbb{N}$ such that for each two sequences $I, J \in [m]_k$ it holds that if $I \succeq J$ then $f(I) \geq f(J)$.

Let $E = (C, V)$ be an election with m candidates, let $k \leq m$ be the size of the desired committee, and let f be a committee scoring function (for m candidates and committees of size k). Under the committee scoring rule \mathcal{R}_f , every committee $S \subseteq C$ with $|S| = k$ receives $\text{score}(S) := \sum_{v \in V} f(\text{pos}_v(S))$ (for this notation, the election will always be clear from the context). The committee with the highest score wins. (If there are several such committees, then they all tie as winners.) It turns out that many well-known multiwinner voting rules are, in fact, committee scoring rules. Consider the following examples:

1. The SNTV, Bloc, and k -Borda rules pick k candidates with the highest Plurality, k -Approval, and Borda scores, respectively. Thus, they are defined through the scoring functions: $f_{\text{SNTV}}(i_1, \dots, i_k) := \sum_{t=1}^k \alpha_1(i_t) = \alpha_1(i_1)$, $f_{\text{Bloc}}(i_1, \dots, i_k) := \sum_{t=1}^k \alpha_k(i_t)$, and $f_{k\text{-Borda}}(i_1, \dots, i_k) := \sum_{t=1}^k \beta_m(i_t)$, respectively.
2. The Chamberlin–Courant rule is defined through the committee scoring function $f_{\text{CC}}(i_1, \dots, i_k) := \beta_m(i_1)$ (intuitively, under the Chamberlin–Courant rule each voter is represented by the committee member that this voter ranks highest; the Chamberlin–Courant rule maximizes the sum of the Borda scores that voters give to their representatives; we can—and often do—consider variants of this rule with other underlying scores).

Recently, Skowron, Faliszewski, and Lang (2015) introduced a new class of multiwinner rules based on OWA operators⁴ (a variant of this class was also studied by Aziz et al. (2015a; 2015b)). While they did not directly consider elections based on preference orders, we can implement their main ideas through committee scoring rules.

An OWA operator Λ of dimension k is a sequence $\Lambda = (\lambda_1, \dots, \lambda_k)$ of nonnegative numbers.

Definition 2. Let $\Lambda = (\lambda_1, \dots, \lambda_k)$ be an OWA operator of dimension k and let γ be a (single-winner) scoring function for elections with m candidates ($k \leq m$). Together, Λ and γ define a committee scoring function $f_{\Lambda, \gamma}$ such that for each $(i_1, \dots, i_k) \in [m]_k$ we have $f_{\Lambda, \gamma}(i_1, \dots, i_k) = \sum_{t=1}^k \lambda_t \gamma(i_t)$.

We refer to the committee scoring rules that can be obtained through Definition 2 as OWA-based. Clearly, SNTV, Bloc, k -Borda, and Chamberlin–Courant are OWA-based.

Fixed-Majority Consistent Rules

We first describe our axiomatic criterion for what it means to resemble the Plurality rule, then—in an intermediate step—we provide a class of committee scoring rules such that every committee scoring rule outside of this class fails the criterion, and finally we provide a complete characterization.

⁴OWA stands for “ordered weighted average.” OWA operators were introduced by Yager (1988) in the context of multicriteria decision making. Kacprzyk et al. (2011) describe their applications in the context of collective choice.

Initial Remarks. One of the features that distinguishes the Plurality rule among all the scoring rules is the fact that it satisfies the simple majority criterion.

Definition 3. A single-winner voting rule \mathcal{R} satisfies the simple majority criterion if, for every election $E = (C, V)$ where more than half of the voters rank some candidate c on top, it holds that $\mathcal{R}(E) = \{c\}$.

Proposition 1. Let γ be a scoring function such that \mathcal{R}_γ satisfies the simple majority criterion. Then $\gamma(1) > \gamma(2) = \dots = \gamma(m)$, that is, \mathcal{R}_γ is the Plurality rule.

There are at least two ways of generalizing the simple majority criterion to the multiwinner setting. We choose perhaps the simplest one, the *fixed-majority criterion* introduced by Debord (1993).

Definition 4. A multiwinner voting rule \mathcal{R} satisfies the fixed-majority criterion if for every election $E = (C, V)$ and every positive integer k , $1 \leq k \leq m$, if there is a committee W of size k such that more than half of the voters rank all the members of W above the non-members of W , then $\mathcal{R}(E, k) = \{W\}$.

Another way of extending the simple majority criterion to the multiwinner case would be to say that if a committee W is such that for each $c \in W$ a majority of voters ranks c among top k positions (possibly a different majority for each c), then W is a winning committee. However, consider the following votes over candidate set $\{a, b, c\}$:

$$v_1: a > b > c, \quad v_2: a > c > b, \quad v_3: b > c > a.$$

For $k = 2$, all three committees, $\{a, b\}$, $\{a, c\}$, and $\{b, c\}$, have majority support in the just-described sense. We feel that this is against the spirit of the single-winner simple majority criterion.

It is easy to verify that Bloc satisfies the fixed-majority criterion and that SNTV does not (it will also follow formally from our further discussion). This means that in a certain axiomatic sense, Bloc is closer to the Plurality rule than SNTV. This is quite interesting since one’s first idea of generalizing Plurality would likely be to think of SNTV. Yet, Bloc is certainly not the only committee scoring rule that satisfies our criterion. Let us consider the following rule.

Definition 5. Let k be the size of the committee to be elected. The Perfectionist rule is defined through the scoring function f_{Perf} such that $f_{\text{Perf}}(i_1, \dots, i_k) = 1$ if $(i_1, \dots, i_k) = (1, \dots, k)$ and $f_{\text{Perf}}(i_1, \dots, i_k) = 0$ otherwise. In other words, a voter gives score of 1 to a committee only if its members occupy the top k positions of his or her vote.

It is easy to see that the Perfectionist rule satisfies the fixed-majority property and that it closely resembles the Plurality rule. On the other hand, neither k -Borda nor Chamberlin–Courant satisfy it (for $k = 1$ they are equivalent to the Borda rule, which fails the simple majority criterion).

Top- k -Counting Rules. To characterize the committee scoring rules that satisfy the fixed-majority criterion, we introduce a class of scoring functions that depend only on the number of committee members ranked in the top k positions.

Definition 6. Let $f: [m]_k \rightarrow \mathbb{N}$ be a committee scoring function. We say that f is a top- k -counting function if there is a function $g: \{0, \dots, k\} \rightarrow \mathbb{N}$ such that for each $(i_1, \dots, i_k) \in [m]_k$ we have $f(i_1, \dots, i_k) = g(|\{t \in [k]: i_t \leq k\}|)$. (We refer to g as the counting function for f .) We say that a committee scoring rule is top- k -counting if it is defined through a top- k -counting scoring function.

Both Bloc and Perfectionist are top- k -counting rules. The former is defined through the linear counting function $g(x) = x$, while the latter is defined through the counting function g which is a step-function: $g(x) = 1$ for $x = k$ and $g(x) = 0$ otherwise.

Another example of a top- k -counting rule is the k -Approval Chamberlin–Courant rule⁵, defined through the scoring function $\alpha_{k\text{-CC}}(i_1, \dots, i_k) = \alpha_k(i_1)$ (we also refer to it as the α_k -CC rule). It is a top- k -counting rule, defined through the counting function g such that $g(0) = 0$ and $g(x) = 1$ for all $x > 0$.

Top- k -counting rules have a number of interesting features. Firstly, their counting functions have to be nondecreasing. Secondly, they are OWA-based. Thirdly, every committee scoring rule that satisfies the fixed-majority criterion is top- k -counting. We express these facts in the following three propositions. For technical reasons, in the rest of the paper we make the assumption that $m \geq 2k$.

Proposition 2. Let $f: [m]_k \rightarrow \mathbb{N}$ be a top- k -counting scoring function defined through a counting function g ($m \geq 2k$). Then, g is nondecreasing.

Proof. Consider the sequences $I_t = (1, \dots, t, k+1, \dots, k+(k-t))$ and $I_{t+1} = (1, \dots, t+1, k+1, \dots, k+(k-t-1))$ from $[m]_k$. Since $I_{t+1} \succeq I_t$, we have that $f(I_{t+1}) \geq f(I_t)$. By definition, however, we have that $f(I_{t+1}) = g(t+1)$ and that $f(I_t) = g(t)$. Hence, $g(t+1) \geq g(t)$. \square

Without the assumption that $m \geq 2k$, Proposition 2 would have to be phrased more cautiously, and would speak only of existence of nondecreasing counting functions. (For example, for $m = k$, the function g could be arbitrary.)

Proposition 3. Let $f: [m]_k \rightarrow \mathbb{N}$ be a top- k -counting scoring function defined through a counting function g . Let $\Lambda = (g(1) - g(0), g(2) - g(1), \dots, g(k) - g(k-1))$. The committee scoring rule defined through the scoring function f and the OWA-based rule defined through the k -Approval scoring function and OWA operator Λ are equivalent (i.e., they always elect the same committee of size k).

Proof. Let h be the committee scoring function defined through the k -Approval scoring function and the OWA operator Λ . Consider an arbitrary sequence $(i_1, \dots, i_k) \in [m]_k$. By definition, we know that $f(i_1, \dots, i_k) = g(s)$, where $s = |\{t \in [k]: i_t \leq k\}|$. We also have that:

$$\begin{aligned} h(i_1, \dots, i_k) &= \sum_{t=1}^k \alpha_k(i_t) \cdot (g(t) - g(t-1)) \\ &= \sum_{t=1}^s (g(t) - g(t-1)) = g(s) - g(0). \end{aligned}$$

⁵The name of the rule comes from the fact that it is defined in the same way as the Chamberlin–Courant rule, but using the k -Approval scoring function instead of the Borda scoring function.

(The second equality follows by the definition of s and α_k .) So, $f(i_1, \dots, i_k) - h(i_1, \dots, i_k) = g(0)$ is a fixed constant, thus \mathcal{R}_f and \mathcal{R}_h always elect the same committee. \square

Proposition 4. Let $m \geq 2k$ and let $f: [m]_k \rightarrow \mathbb{N}$ be a committee scoring function. If \mathcal{R}_f satisfies the fixed-majority criterion then f is a top- k -counting function.

Proof. For each positive integer t such that $0 \leq t \leq k$ we define the two following sequences from $[m]_k$:

1. $I_t = (1, \dots, t, k+1, \dots, k+k-t)$ is a sequence of positions of the candidates where the first t candidates are ranked in the top t positions and the remaining $k-t$ candidates are ranked just below the k th position.
2. $J_t = (k-t+1, \dots, k, m-(k-t)+1, \dots, m)$ is a sequence of positions where the first t candidates are ranked just above and including the k th position, whereas the remaining $k-t$ candidates are ranked at the bottom.

Among these, $I_k = (1, \dots, k)$ is the highest-scoring sequence of positions and $J_k = (m-k+1, \dots, m)$ is the lowest-scoring sequence. Further, for every t we have $I_t \succeq J_t$. Thus, $f(I_t) \geq f(J_t)$.

We claim that if there exists some $t \in \{0, \dots, k\}$ such that $f(I_t) > f(J_t)$ then \mathcal{R}_f does not have the fixed-majority property. For the sake of contradiction, assume that there is some t such that $f(I_t) > f(J_t)$. Let $E = (C, V)$ be an election with m candidates and $2n+1$ voters. The set of candidates is $C = X \cup Y \cup Z \cup D$, where $X = \{x_1, \dots, x_t\}$, $Y = \{y_{t+1}, \dots, y_k\}$, $Z = \{z_{t+1}, \dots, z_k\}$, and D is a set of sufficiently many dummy candidates so that $|C| = m$. We focus on two committees, $M = X \cup Y$ and $N = X \cup Z$.

The first $n+1$ voters have preference order $X \succ Y \succ Z \succ D$, and the following n voters have preference order $Z \succ X \succ D \succ Y$. Note that the fixed-majority criterion requires that M is the unique winning committee.

Committee M receives the total score of $(n+1)f(I_k) + nf(J_t)$, whereas committee N receives the total score of $(n+1)f(I_t) + nf(I_k)$. The difference between these values is:

$$\begin{aligned} (n+1)f(I_k) + nf(J_t) - (n+1)f(I_t) - nf(I_k) \\ &= f(I_k) + nf(J_t) - (n+1)f(I_t) \\ &= f(I_k) - f(I_t) + n(f(J_t) - f(I_t)), \end{aligned}$$

which, for a large enough value of n , is negative (by assumption, we know that $f(J_t) < f(I_t)$ and so $f(J_t) - f(I_t)$ is negative). That is, for large enough n , committee M does not win the election and \mathcal{R}_f fails the fixed-majority criterion.

In other words, if \mathcal{R}_f satisfies the fixed-majority criterion, then for every $t \in \{0, \dots, k\}$ we have that $f(I_t) = f(J_t)$. This means, however, that f is a top- k -counting rule. To see this, consider some sequence of positions $L = (\ell_1, \dots, \ell_k) \in [m]_k$ where exactly the first t entries are smaller than or equal to k . Clearly, we have that $I_t \succeq L \succeq J_t$ and so $f(I_t) = f(L) = f(J_t)$. This means that $f(i_1, \dots, i_k)$ depends only on the cardinality of the set $\{t \in [k]: i_t \leq k\}$. This completes the proof. \square

Unfortunately, the converse does not hold: α_k -CC is a top- k -counting rule that fails the fixed-majority criterion.

Criterion for Fixed-Majority Consistent Rules. We now provide a formal characterization of those top- k -counting rules that satisfy the fixed-majority criterion. Together with Proposition 4, this gives a full characterization of committee scoring rules with this property (the proof follows by applying techniques similar to those used for Proposition 4).

Theorem 5. *Let $m \geq 2k$ and let $f: [m]_k \rightarrow \mathbb{N}$ be a top- k -counting function with g as its counting function. Then \mathcal{R}_f satisfies the fixed-majority criterion if and only if g is not constant and for each pair of nonnegative integers k_1, k_2 with $k_1 + k_2 \leq k$ it holds that $g(k) - g(k - k_2) \geq g(k_1 + k_2) - g(k_1)$.*

Condition (ii) in this theorem is a relaxation of the convexity property for g and is illustrated in Figure 1.

Definition 7. *Let g be a counting function for some top- k -counting scoring function $f: [m]_k \rightarrow \mathbb{N}$. We say that g is convex if for each k' such that $2 \leq k' \leq k$, it holds that $g(k') - g(k' - 1) \geq g(k' - 1) - g(k' - 2)$. On the other hand, if for each k' with $2 \leq k' \leq k$ it holds that $g(k') - g(k' - 1) \leq g(k' - 1) - g(k' - 2)$, then we say that g is concave.*

The notions of convex and concave functions are standard, but allow us to express many features of top- k -counting rules in a very intuitive way. For example, the following corollary is an immediate consequence of Theorem 5.

Corollary 6. *Let $m \geq 2k$. Let f be a top- k -counting rule for m candidates and committee size k , and let g be its counting function. (1) If g is convex, then \mathcal{R}_f satisfies the fixed-majority criterion. (2) If g is concave but not linear (that is, \mathcal{R}_f is not Bloc) then \mathcal{R}_f fails the fixed-majority criterion.*

The counting function for Bloc is linear (and, thus, both convex and concave), and the counting function for Perfectionist is convex, so these two rules satisfy the fixed-majority condition. On the other hand, the counting function for α_k -CC is concave and, so, this rule fails the criterion.

By Proposition 3, a concave counting function g corresponds to a nonincreasing OWA operator, and a convex counting function corresponds to a nondecreasing one. Skowron et al. (2015) provided evidence that rules based on nondecreasing OWA operators are easier computationally than those based on general OWA operators (though, still tend to be NP-hard to compute). Below we show that this seems to be the case for top- k -counting rules as well, but we also provide a striking counterexample to their results.




Figure 1: Illustration of the condition from Theorem 5.

Complexity of Top- k -Counting Rules

In this section, we consider the computational complexity of winner determination for top- k -counting rules based on convex or concave counting functions. We start by considering several examples.

It is well-known that Bloc winners can be computed in polynomial time. The same holds for the Perfectionist rule.

Proposition 7. *Both the Bloc rule and the Perfectionist rule are computable in polynomial time.*

Proof. To find the winners under the Perfectionist rule, for each voter v we compute the score that the committee consisting of v 's top- k candidates has in the election. We output those committees that have the highest score. Correctness follows by noting that the committees that the algorithm considers are the only ones with nonzero scores. \square

While the result for the Perfectionist rule is very simple, it stands in sharp contrast to the results of Skowron, Faliszewski, and Lang (2015). By Proposition 3, Perfectionist is defined through the OWA operator $(0, \dots, 0, 1)$, and Skowron et al. have shown that, in general, such rules are NP-hard to compute and very difficult to approximate. However, their result relies on the fact that the voters can approve any number of candidates, while in our case they have to approve exactly k of them. Yet, this shows very clearly that even though top- k -counting rules are OWA-based, we cannot simply carry-over the hardness results of Skowron et al. (2015) or Aziz et al. (2015b).

Our discussion of the complexity of top- k -counting rules relies on the following property of the counting functions.

Definition 8. *Let g be a counting function for a top- k -counting function $f: [m]_k \rightarrow \mathbb{N}$. We define the singularity of g , denoted $s(g)$, to be $s(g) = \arg \min_{2 \leq i \leq k} (g(i) - g(i - 1)) \neq g(i - 1) - g(i - 2)$.*

Loosely speaking, $s(g)$ is the smallest integer in $\{2, \dots, k\}$ for which the differential of g changes. For the Bloc rule (which is an exception) we define $s(g)$ to be ∞ , since the differential is a constant function. For all other non-constant rules it is finite. For example, for the Perfectionist rule we have $s(g) = k$.

We generalize the polynomial-time algorithm for the Perfectionist rule to similar rules, for which the value $s(g)$ is close to k . (Every counting function for a given committee of size k can be encoded as a sequence of $k + 1$ numbers.)

Proposition 8. *Suppose a top- k -counting rule \mathcal{R}_f is defined by a counting function g of f such that $k - s(g) \leq q$. Then for each positive integer q , there exists a polynomial time algorithm that, given an election E with m candidates, computes a committee from $\mathcal{R}_f(E, k)$.*

Yet, as one might expect, not all top- k -counting rules are polynomial-time solvable and, indeed, most of them are not (under standard complexity-theoretic assumptions). For example, α_k -CC is NP-hard (notice that this results follows quite easily from Theorem 1 of Procaccia et al. (2008)).

Proposition 9. *Deciding the existence of a committee with at least a given score is NP-hard for α_k -CC.*

We generalize this NP-hardness result to the case of convex top- k -counting rules for which there is some constant c such that $k - s(g) \geq k/c$ (that is, for convex counting functions where the differential changes “early”). An analogous result for concave counting functions follows from the works of Skowron, Faliszewski, and Lang (2015) and Aziz et al. (2015b).

Theorem 10. *Let c be some constant and let \mathcal{R}_f be a top- k -counting rule whose counting function g satisfies the following conditions: (1) it does not depend on the number of candidates in the election (but may depend on k), (2) it is computable in polynomial time for each committee size k and, for each committee size k , its highest value is polynomially bounded in k , (3) for each committee size k (greater than some fixed constant) the counting function g is convex, and $s(g) \geq k/c$. Then, deciding if there is a committee with at least a given score is NP-hard for \mathcal{R} .*

Perhaps the only truly controversial assumption in this theorem is the requirement that for a given committee size k the highest value of the counting function is polynomially bounded in k . The reason for having it is that if the highest value were extremely large (say, exponentially large with respect to k) then for sufficiently few voters (e.g., polynomially many) the rule might degenerate to a polynomial-time computable rule (e.g., it might resemble the Perfectionist rule for this case). Indeed, to avoid such problems, in our proof we use a number of voters that depends on $g(k)$.

While top- k -counting functions tend to be NP-hard, for concave top- k -counting rules we can obtain constant-factor approximation algorithms (and FPT results; see later).

Theorem 11. *Let \mathcal{R}_f be a top- k -counting rule whose counting function g satisfies the following conditions: (1) it is computable in polynomial time for each number m of candidates and each committee size k , and (2) for each committee size k it is concave. Then, there is a polynomial-time $(1 - \frac{1}{e})$ -approximation algorithm for computing the score of a winning committee under \mathcal{R}_f .*

Such a general result for convex counting functions seems impossible. Let us consider a convex counting function $g_2(x) = \max(x - 1, 0)$ that is nearly identical to the one used by Bloc. If we had a polynomial-time constant-factor approximation algorithm for a rule defined by g_2 , we would have a constant-factor approximation algorithm for the densest at most K subgraph problem (DAMKS). By the results of Khuller and Saha (2009), Raghavendra and Steurer (2010), and Alon et al. (2011), this seems very unlikely.

Theorem 12. *There is no polynomial-time constant-factor approximation algorithm for the problem of computing the score of a winning committee under the rule defined by the counting function g_2 unless such an algorithm exists for the DAMKS problem.*

On the other hand, for top- k -counting rules that are not too far from α_k -CC, we have a polynomial-time approximation scheme (PTAS), i.e., an algorithm that can achieve any desired approximation ratio, provided the number of candidates is not too large. This result holds even for rules that are not concave (provided they satisfy the conditions of the

theorem); the result follows by noting that our voters have non-finicky utilities (Skowron, Faliszewski, and Lang 2015).

Theorem 13. *Let \mathcal{R} be a top- k -counting rule with counting function which is constant for arguments greater than ℓ . Further, assume that $m = o(k^2)$. Then, there is a PTAS for computing the score of a winning committee under f .*

We finish with the following fixed-parameter tractability results (for a more detailed description of parametrized complexity we point the readers to the books of Downey and Fellows (1999) or Niedermeier (2006)).

Proposition 14. *There is an algorithm that, given a counting function g and an election E , computes a winning committee for E under the top- k -counting rule \mathcal{R}_f defined by g in FPT time with respect to the number of candidates.*

Proof. The algorithm simply computes the score of every possible committee and outputs the one with the highest score. With m candidates and committee size k , the algorithm has to check $\binom{m}{k} = O(m^m)$ committees, and checking each committee takes a polynomial number of steps. \square

Theorem 15. *There is an algorithm that, given a concave counting function g and an election E , outputs a winning committee for E under the top- k -counting function defined by g , in FPT time with respect to the number n of voters.*

Our proof is based on the idea of solving a mixed integer linear program (MILP) in FPT time with respect to the number of integral variables. We use Bredebeck et al.’s. (2015) technique, where we enforce that non-integral variables take integral values in the optimal solution.

Characterization of the Bloc rule

Using our observations regarding the fixed-majority criterion and the notion of non-crossing monotonicity (Elkind et al. 2014), we provide a characterization of the Bloc rule.

Definition 9 (Elkind et al. (2014)). *A multiwinner rule \mathcal{R} is said to be non-crossing monotone if for each election $E = (C, V)$ and each $k \in [|C|]$ the following holds: if $c \in W$ for some $W \in \mathcal{R}(E, k)$, then, for each E' obtained from E by shifting c forward by one position in some vote without passing another member of W , we have that $W \in \mathcal{R}(E', k)$.*

For the committees that contain at most half of the candidates, Bloc is the only committee scoring rule that is fixed-majority consistent and non-crossing monotone.

Theorem 16. *Let $m \geq 2k$ and let \mathcal{R}_f be a committee scoring rule based on the scoring function $f: [m]_k \rightarrow \mathbb{N}$. If \mathcal{R}_f is fixed-majority consistent and non-crossing monotone, then \mathcal{R}_f is the Bloc rule.*

If we had a characterization of committee scoring rules, as we do for single-winner scoring rules (Young 1975), we would have a full characterization of Bloc.

Outlook

We aimed at finding a multiwinner analogue of the single-winner Plurality rule and we have shown that the answer

is quite involved. While intuitively SNTV is a natural analogue of the Plurality rule, it fails the fixed-majority criterion (which the Plurality rule satisfies in the single-winner setting). We have found that among committee scoring rules, only the top- k -counting rules—a class of rules we have defined in this paper—have a chance of satisfying our criterion, and we have characterized exactly when this happens. Since the research on multiwinner voting is still in an early, exploratory stage, we believe that it is important and valuable to identify such interesting classes of multiwinner rules.

Our work leads to a number of open questions. On the axiomatic front, we believe that it would be interesting to provide a characterization of committee scoring rules along the lines of Young's (1975) characterization for their single-winner counterparts. On the computational front, it would be interesting to find more powerful algorithms for computing top- k -counting rules.

Acknowledgments

We are grateful to AAAI reviewers for their helpful comments. Piotr Faliszewski and Piotr Skowron were supported in part by NCN grant DEC-2012/06/M/ST1/00358. Piotr Faliszewski was supported in part by AGH University grant 11.11.230.124 (statutory research). Arkadii Slinko was supported by the Royal Society of NZ Marsden Fund UOA-254. Nimrod Talmon was supported by the DFG Research Training Group MDS (GRK 1408). The research was also partially supported through the COST action IC1205 (Piotr Faliszewski's visit to TU Berlin).

References

- Alon, N.; Arora, S.; Manokaran, R.; Moshkovitz, D.; and Weinstein, O. 2011. Inapproximability of densest k -subgraph from average case hardness. Technical report.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2015a. Justified representation in approval-based committee voting. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 784–790.
- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015b. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, 107–115.
- Bredereck, R.; Faliszewski, P.; Niedermeier, R.; Skowron, P.; and Talmon, N. 2015. Elections with few candidates: Prices, weights, and covering problems. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory*, 414–431.
- Chamberlin, B., and Courant, P. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77(3):718–733.
- Debord, B. 1993. Prudent k -choice functions: Properties and algorithms. *Mathematical Social Sciences* 26:63–77.
- Downey, R., and Fellows, M. 1999. *Parameterized Complexity*. Springer-Verlag.
- Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2014. Properties of multiwinner voting rules. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, 53–60.
- Elkind, E.; Lang, J.; and Saffidine, A. 2011. Choosing collectively optimal sets of alternatives based on the condorcet criterion. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 186–191.
- Fishburn, P. 1981. An analysis of simple voting systems for electing committees. *SIAM Journal on Applied Mathematics* 41(3):499–502.
- Gehrlein, W. 1985. The condorcet criterion and committee selection. *Mathematical Social Sciences* 10(3):199–209.
- Kacprzyk, J.; Nurmi, H.; and Zadrozny, S. 2011. The role of the OWA operators as a unification tool for the representation of collective choice sets. In *Recent Developments in the Ordered Weighted Averaging Operators*, 149–166.
- Khuller, S., and Saha, B. 2009. On finding dense subgraphs. In *Proceedings of the 36th International Colloquium on Automata, Languages, and Programming*, 597–608.
- Kilgour, M. 2010. Approval balloting for multi-winner elections. In *Handbook on Approval Voting*. Springer. Chapter 6.
- Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 280–286.
- Monroe, B. 1995. Fully proportional representation. *American Political Science Review* 89(4):925–940.
- Niedermeier, R. 2006. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press.
- Procaccia, A.; Rosenschein, J.; and Zohar, A. 2008. On the complexity of achieving proportional representation. *Social Choice and Welfare* 30(3):353–362.
- Raghavendra, P., and Steurer, D. 2010. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, 755–764.
- Skowron, P., and Faliszewski, P. 2015. Fully proportional representation with approval ballots: Approximating the MaxCover problem with bounded frequencies in FPT time. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2124–2130.
- Skowron, P.; Faliszewski, P.; and Lang, J. 2015. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*.
- Skowron, P.; Faliszewski, P.; and Slinko, A. 2015. Achieving fully proportional representation: Approximability result. *Artificial Intelligence* 222:67–103.
- Yager, R. 1988. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on Systems, Man and Cybernetics* 18(1):183–190.
- Young, H. 1975. Social choice scoring functions. *SIAM Journal on Applied Mathematics* 28(4):824–838.