

Multiwinner Rules with Variable Number of Winners

Piotr Faliszewski¹ and Arkadii Slinko² and Nimrod Talmon³

Abstract. We consider voting rules for approval-based elections that select committees whose size is not predetermined. Unlike the study of rules that output committees with a predetermined number of winning candidates, the study of rules that select a variable number of winners has only recently been initiated. We first mention some scenarios for which such rules are applicable. Then, aiming at better understanding these rules, we study their computational properties and report on simulations regarding the sizes of their committees.

1 Introduction

We study a setting in which a group of agents (the voters) wants to select a set of candidates (a committee) based on their preferences. The agents specify which candidates they approve for inclusion into the committee and this input data needs to be aggregated. However, as opposed to the quickly growing body of work on electing committees of a fixed size (see, e.g., [14]), here we are interested in rules that derive both the size of the winning committee and its members from the voters' preferences. Recently, Kilgour [18] and Duddy et al. [12] initiated a systematic study of such voting rules (for some recent axiomatic results, we also point to the work of Brandl and Peters [8]); here we are interested in the complexity of computing the winners and in experimentally analyzing the sizes of the elected committees.

1.1 When Not To Fix the Size of the Committee?

In some applications it is not natural to fix the size of the committee in advance, and rather it is required to deduce it from the votes. Since so far committee elections with variable number of winners have not received much attention in the AI literature, below we provide examples of such settings.⁴ We will start with the following trivial but instructive example.

A Single Decision Maker. If a single voter is charged with setting up a committee, she will include all candidates she approves of. Clearly, the number of worthy candidates cannot be known in advance. With several decision makers, the solution is not so obvious.

Finding a Set of Qualifying Candidates. Finding a set of candidates that satisfy many given criteria is a common problem. Real-life examples include selecting baseball players for inclusion into a Hall of Fame and selecting students to receive an honors degree. In the former case, eligible voters (in particular, baseball writers) approve up to ten players each; then, those approved by at least 75% of the voters are chosen. In the latter case, the voting process is typically implicit: the university announces a set of criteria of excellency—which act as

voters, “approving” the students that satisfy them⁵—and set the rules such as “a student meeting at least five out of six criteria shall receive the honors degree”.

It is often desirable that the selected committee be relatively small (e.g., just a few people for the Hall of Fame and a fairly small percentage of the students for the honors degree) but this is not always the case. As an example of such case consider a different scenario in which the task is to select people for an in-depth medical check based on a number of simple criteria that jointly indicate elevated risk of a certain disease: In this example, everyone who is at risk should be checked regardless of the number of those patients.

Remark 1 *One of the first procedures formally proposed for the task of selecting a group of qualifying candidates was the majority rule (MV), suggested by Brams, Kilgour, and Sanver [7]. This rule outputs the committee that includes all the candidates approved by at least half of the voters (i.e., all the candidates that satisfy at least half of the criteria). One might consider MV with other thresholds, such as in the example of the Hall of Fame.*

Initial Screening. Consider a situation where we need to select one item—among many possible ones—that has some desirable features. The final decision is to be done by a qualified expert, but we have a number of easy to evaluate (but imperfect) criteria that the selected item should satisfy; these criteria are soft and it may be the case that the best item actually fails some of them. We view each criterion as a voter (who “approves” the items satisfying it) and we seek a committee, hopefully of a small size, of candidates from which the qualified expert will make the final choice.⁶

Remark 2 *Initial screening is related to shortlisting [2, 13]. We use a different name for it to emphasize that we do not fix the number of candidates for selection, as is often the case with shortlisting.*

Partitioning into Two Homogeneous Groups. In this case we care about partitioning the set of candidates into two *homogeneous* groups (one of them will be the committee) so that each group contains candidates that are as similar as possible. The reason might be fairness: Trying to avoid the situation when one candidate is selected for the committee and another one with similar characteristics is not.

A concrete example is partitioning students in some class into two groups, e.g., a group of beginners and a group of advanced ones (say,

¹ AGH University, Poland, email: faliszew@agh.edu.pl

² University of Auckland, New-Zealand, email: a.slinko@auckland.ac.nz

³ Ben-Gurion University, Israel, email: talmonn@bgu.ac.il

⁴ We do not mention this repeatedly, but one may wish to automate the processes in the examples below using AI techniques.

⁵ The criteria may include, e.g., never receiving a low grade from a course, taking some advanced classes, never being suspended, etc.

⁶ One of the authors was once concerned with classifying a collection of daggers for a museum. The solution was to compute several partitions of the daggers into clusters, evaluate their qualities without reference to ethnographical knowledge, and to present the best ones to the museum's experts, who chose one partition that led to an ethnographically meaningful classification.

regarding, their knowledge of a foreign language; depending on the setting, it may or may not be important to keep the sizes of these two groups close). The students are partitioned in this way to facilitate a better learning environment for everyone; in the context of voting, the issue of partitioning students was raised by Duddy et al. [12].

Finding a Representative Committee. An elected committee is *representative* if each voter approves at least one committee member (who can then represent this voter). The idea of choosing a representative committee of a fixed size received significant attention in the literature [9, 23, 13, 1] but, as pointed out by Brams and Kilgour [6], committees of fixed size cannot always provide adequate representation. Representative committees may be desired when some authorities are revising existing regulations and need to consult citizens, for which purpose they would like to select focus groups in various cities. Members of these groups do not have to represent the society proportionally (their role is to voice opinions and concerns and not to make final decisions), but should cover the spectrum of opinions in the society. Usually, small representative committees are more desirable than larger ones.

Choosing Committees with Implicit Preferences on Size. Consider the problem of choosing a group of specialists to hire. This problem is similar to that of partitioning the candidates into two non-homogeneous groups, where the first group constitutes the team of specialists to be hired (e.g., a group of software developers for a project, evaluated by their knowledge of PHP, Java, previous experience, etc.), but we additionally have implicit preferences regarding the committee size (e.g., we may wish to hire at least five developers, but from prior experience we know that a team of more than ten people is hard to manage).

Another example would be a factory that can produce a number of different items. The factory should produce those items that are most popular on the market (that are “approved” by many possible customers), but it should not produce too many of them as it may involve increased costs of various kinds (e.g., adding a new product may require expensive extension of the factory, which would only pay off if the new product were very popular).

Remark 3 *Strictly speaking, we have two slightly different models for elections with a variable number of winners. In the former, there are societal preferences regarding the size of the target committee (as in the last two examples) and in the latter such preferences are not present (as in the first three examples). From this point of view, the classical model of multiwinner elections—where a committee of exactly k candidates is to be selected—is a special, extreme case of such preferences on the committee size.*

1.2 Our Contribution

In this paper we concentrate on the task of creating a committee solely from the votes, without being given hints on the desired size of the committee to be selected.

Our first goal is to study computational aspects of voting rules tailored for such elections, which we refer to as elections with variable number of winners. This research direction was pioneered by Fishburn and Pekeč [15], who introduced the class of *threshold rules* and studied their computational complexity (somewhat surprisingly, only for the case when the committee size is fixed). We are not aware of any other computational study that followed their work. We study threshold rules and a number of other rules, including those discussed by Kilgour [18]. For each, we establish whether finding a winning

committee is in P or is NP-hard, in which case we seek FPT algorithms parameterized by the two most natural parameters, namely the number of candidates and the number of voters.

Our second goal is to experimentally evaluate the average sizes of the committees elected by the rules in consideration. For this purpose we generate elections according to the approval-based variants of the Impartial Culture model and the Polya-Eggenberger urn model.

1.3 Paper Structure

After providing preliminaries (Section 2), we go on to consider several classes of voting rules for elections with variable number of winners: In Section 3 we consider Net-Approval Voting, in Section 4 we consider Capped Satisfaction Rules, and in Section 5 we consider Threshold rules. For each class of rules, we first define several rules in the class, then discuss our experimental results and analyze their computational properties. We conclude in Section 6.

2 Preliminaries

An approval-based election $E = (C, V)$ consists of a set $C = \{c_1, \dots, c_m\}$ of candidates and a collection $V = (v_1, \dots, v_n)$ of voters. Voters express their preferences by filling approval ballots. An approval ballot of a voter specifies the subset of candidates that the voter approves. To simplify notation, we denote voter v_i 's approval ballot also as v_i . A collection V of voters, interpreted as a collection of approval ballots in a certain election, is a *preference profile*. For a given subset S of the candidates from C , by \bar{S} we mean the set of candidates not in S (i.e., $\bar{S} = C \setminus S$). The *approval score* of a candidate is the number of voters that approve him.

A voting rule for elections with a variable number of winners is a function \mathcal{R} that, given an election $E = (C, V)$, returns a family of subsets of C (the set of committees which are tied as winners). The main point of difference between the type of voting rules that we study here and the voting rules typically studied in the context of multiwinner elections is that we do not fix the size of the committee to be elected and we let it be deduced by the rule⁷. For an overview of multiwinner rules using approval balloting, we point to the works of Kilgour [17, 18] (also for discussions regarding rules with a variable number of winners), Aziz et al. [1], and Lackner and Skowron [20] (for rules selecting fixed-size committees); Duddy et al. [12] and Brandl and Peters [8] discuss the Borda mean rule (geared for variable number of winners). This rule is not included in our discussion.

Computational Complexity. We assume basic knowledge of (parameterized) complexity theory. When we say that a given multiwinner rule is computable in polynomial time (in FPT time), we mean that for each set of candidates it is possible to check in polynomial-time (in FPT time) if there is a winning committee that includes all the candidates from this set (and if this set is a winning committee itself). Using such a procedure, we can compute any winning committee, by starting with an empty set and extending it with candidates one-by-one (each time checking if by adding a given candidate we still have a subset of a winning committee).

Our hardness results follow by reductions from the NP-complete problem SET COVER. An instance of SET COVER consists of a set $U = \{u_1, \dots, u_n\}$ of elements, a family $\mathcal{S} = \{S_1, \dots, S_m\}$ of subsets of U , and an integer k ; we ask whether there are at most k sets from \mathcal{S} whose union is U .

⁷ In particular, the committee may have zero candidates, i.e., it can be empty.

Experiments. We experimentally evaluate the sizes of the committees computed using the rules under consideration. In all the experiments, the elections have 20 candidates and either 20 or 100 voters (such elections are small enough to compute all the rules easily, yet they seem sufficient to demonstrate the effects we are interested in). We use two models for generating approval ballots. In the *p-Independent Approval model* (the *p-IA model*), each voter approves each candidate c independently, with probability p . In the second model, the *Polya-Eggenberger urn model* adapted to the approval setting (α -urn, where α is a parameter), we proceed as follows: At first, we have an urn with one copy of each possible ballot. To generate a vote, we draw a ballot from the urn (this is the generated vote) and return it together with $\alpha \cdot 2^m$ additional copies of it. We continue until the required number of votes is generated. E.g., for $\alpha = 1$, the first two votes are the same with probability roughly 0.5; the 0-urn model is equivalent to the 0.5-IA model. The value α is called the *parameter of contagion* and measures correlation between the generated votes. (Regarding the urn model, see, e.g., the work of Berg and Lepelley [3], as well as the works of McCabe-Dansted and Slinko [22] and Walsh [26] for examples of papers where the urn model is used).

In all our experiments, to measure a particular quantity we generate 10,000 elections and calculate the average size of a committee.

3 (Generalized) Net-Approval Voting

In the setting where the size of the target committee is fixed, one may ask for a committee of candidates whose sum of approval scores is the highest (approval voting). To adapt this idea to the variable number of winners, Brams and Kilgour [6] suggested the Net-Approval Voting (NAV) rule. This rule pays attention not only to approvals but also to disapprovals.

Net Approval Voting (NAV). The score of a committee S in election $E = (C, V)$ under NAV is defined to be $\sum_{v_i \in V} (|S \cap v_i| - |S \cap \bar{v}_i|)$; the committees with the highest score tie as co-winners.

Naturally, NAV is polynomial-time computable: The winning committees consist of all the candidates approved by a strict majority of the voters and a subset of candidates approved by exactly half of the voters (very similarly to the MV rule from the introduction [7]).

Experiment 1 In Figure 1a we see the probability of obtaining each committee size under the NAV rule with 20 candidates, 20 voters, and votes generated using the 0.3-IA, 0.5-IA, and 0.7-IA models, respectively. These results show that NAV is quite specific: When the probability of a voter approving a candidate is the same as the fraction of approvals needed for getting a place in the committee (i.e., when $p = 0.5$), then the typical committee sizes are nicely spread around value 9, but when p is smaller or larger, the committee ends up nearly empty or contains almost all the candidates, respectively.

Later, in Table 1, we present average sizes of committees computed by our rules (including NAV) for the 0.5-IA model (both for 20 and 100 voters). On the average, the smallest NAV committee contains a bit less than half of the candidates (e.g., 8.25 for the case of 20 candidates), which is quite intuitive, as a candidate needs a strict majority of approvals to become a committee member. We also computed the average sizes of the smallest NAV committees for the *p-IA* model with different approval-probability values p , and for the α -urn model, with different α values. The results are presented in Figures 4 and 5. In particular, Figure 4 shows that when the number of voters becomes large, the graph of the average committee size becomes very close to the step function.




Figure 1: Histogram of sizes of committees selected by the NAV rule and the MRC rule in elections with 20 candidates and 20 voters, generated according to the *p-IA* model for $p \in \{0.3, 0.5, 0.7\}$. The x -axis gives the committee size and the y -axis gives the probability that a committee of a given size is winning.

These experiments show that NAV should only be used in very specific settings (such as in the baseball Hall of Fame). In particular, NAV might be appropriate for choosing a set of qualifying candidates, as the decision of including a certain candidate c is made based on approvals for c only (e.g., whether a patient should be sent for an in-depth medical check should not depend on the health of other patients).⁸

Using the main idea behind the NAV rule, we suggest a whole family of similar rules.

Generalized NAV. Let f and g be two non-decreasing, non-negative-valued functions, $f, g: \mathbb{N} \rightarrow \mathbb{N}$, such that $f(0) = g(0) = 0$. We define the (f, g) -NAV score of a committee S in election $E = (C, V)$ to be:

$$\sum_{v_i \in V} (f(|S \cap v_i|) - g(|S \cap \bar{v}_i|)).$$

The committees with the highest score tie as co-winners.

Remark 4 These generalised rules allow us to count approvals and disapprovals differently. In fact, the lack of approval of a candidate is not identical to a disapproval; the reason for it may be simply the lack of information about this candidate.

The family of (f, g) -NAV rules is quite diverse. For example, if $f(x) = x$ and $g(x) = 2x$, then we get a rule that includes a candidate in a committee if it is approved by a fraction of at least $2/3$ of the voters (we refer to this rule also as $2/3$ -NAV). For nonlinear functions f and g , (f, g) -NAV can behave quite differently, though.

Let t_1 be a function such that $t_1(0) = 0$ and $t_1(k) = 1$ for each $k \geq 1$. The rule $(t_1, 0)$ -NAV (by 0 we mean the constant function always giving 0) seeks committees where each voter approves at least one committee member. The committee consisting of all candidates is always trivially winning for this rule, but checking if there is a winning $(t_1, 0)$ -NAV committee of at most a given size is NP-hard (as it is, in essence, the SET COVER problem). It might be more useful to seek small winning committees, thus we define the following rule.

Minimal Representing Committee rule (MRC). Under the MRC rule, we output all the committees of smallest size such that each voter (with a nonempty approval ballot) approves at least one committee member.

In essence, MRC is a variant of the approval-based Chamberlin–Courant rule (CC) [9, 25, 4]. For CC, the fact that a voter approves a

⁸ Such rules are also typical in lobbying scenarios [11, 5, 24].

candidate is typically interpreted as saying that the voter would feel represented by the candidate. For MRC, we insist that each voter is represented, and we want to keep the committee as small as possible.

For MRC, it is computationally hard to even check if a single given candidate belongs to *some* winning committee.

Theorem 1 *The problem of deciding if a given candidate belongs to some winning committee under the MRC rule is complete for Θ_2^P .*⁹

Proof. Hardness for Θ_2^P follows by a reduction from the Θ_2^P -hard Vertex Cover Member problem [16], in which we are given a graph G and a vertex x , and the goal is to decide whether x belongs to some minimum-size vertex cover. Given such an instance, for each vertex u we construct a candidate c_u and for each edge $\{u, v\}$ we construct a voter approving c_u and c_v . Then, x belongs to a minimum-size vertex cover if and only if c_x belongs to an MRC winning committee.

Containment in Θ_2^P is by a binary search on the size k of the MRC winning committees, followed by a call to an NP-oracle deciding whether c_x belongs to a committee of size k covering all voters. \square

Yet, there are FPT algorithms for computing MRC committees (parameterized either by the number of candidates or by the number of voters; omitted due to limited space).¹⁰

Experiment 2 *We repeated Experiment 1 for the case of MRC (using a brute-force search algorithm). In particular, for the case of 20 candidates, 20 voters, and the 0.5-IA model, the average size of an MRC committee is 2.68 (we compare this value later, in Table 5a). Indeed, on Figure 1b we see that in this case almost all committees contain either two or three candidates.*

We can also use the standard greedy algorithm for SET COVER to find approximate MRC committees; we view this algorithm as a voting rule in its own right.

GreedyMRC. Under GreedyMRC, we output all committees obtained by the following method: we start with an empty committee and perform a sequence of iterations, where in each iteration we (a) add to the current committee a candidate c approved by the largest number of voters, and (b) remove the voters that approve c from further consideration. (After all voters with nonempty approval ballots are removed, we output the resulting committee).¹¹

Experiment 3 *By its connection to SET COVER, GreedyMRC always finds a committee that is at most a factor $O(\log m)$ larger than the MRC one (where m is the number of candidates). In our experiments for the p -IA model (with $p \in \{0.05, 0.1, \dots, 0.95\}$), and the α -urn model (with $\alpha \in \{0.05, 0.1, \dots, 0.95\}$), the average sizes of the GreedyMRC committees were no more than 8% larger than for MRC for 20 voters or no more than 11% larger for 100 voters.*

Now consider the $(0, t_1)$ -NAV rule, which elects committees that contain candidates approved by all voters. Here, as the empty set is trivially a winning committee, it is more interesting to ask about the largest winning committee.

Unanimity Voting. This rule outputs the committee containing all candidates approved by all voters.

⁹ Θ_2^P contains problems solvable using log-many calls to an NP-oracle.

¹⁰ For the number of voters, the general idea is to consider all partitions of the voters into subgroups that share a representative and test if, indeed, for each subgroup we can find such a common representative.

¹¹ Formally, we output a singleton set containing the resulting committee, because multiwinner rules are required to output sets of committees.

Both for $(t_1, 0)$ -NAV and for $(0, t_1)$ -NAV, computing *some* winning committee is easy (the set of all candidates in the former case and the empty set in the latter). However, for general (f, g) -NAV rules this is not the case.

Theorem 2 *There exists an (f, g) -NAV rule for which deciding if there exists a committee with at least a given score is NP-hard.*

Proof. We consider specific functions f and g and show that, for the corresponding (f, g) -NAV rule, it is NP-hard to decide if there exists a committee with at least a given score. The specific functions f and g we consider are as follows:

$$f(x) = \begin{cases} 0, & x = 0 \\ 4, & x \geq 1 \end{cases} \quad g(x) = \begin{cases} 0, & x = 0 \\ 1, & x = 1 \\ 2, & x \geq 2 \end{cases}$$

To show NP-hardness, we reduce from the NP-hard X3C problem: given sets $\mathcal{S} = \{S_1, \dots, S_m\}$ over elements b_1, \dots, b_n , each set containing exactly three integers and each element contained in exactly three sets, the goal is to decide whether there is a family of sets $S' \subseteq \mathcal{S}$ such that each element b_i belongs to exactly one set from S' (i.e., each element is covered by exactly one set). Without loss of generality, we assume that $n > 39$.

Given an instance of X3C we create an election: For each set S_j we create a candidate S_j . For each element b_i we create three voters: v_i^1, v_i^2 , and v_i^3 ; v_i^1 and v_i^2 are referred to as *set voters* while v_i^3 is referred to as an *antiset voter*. Both voters v_i^1 and v_i^2 approve exactly the candidates corresponding to the sets which contain b_i , while the voter v_i^3 approves exactly the candidates corresponding to the sets which do not contain b_i . With respect to the reduced election, we ask whether a committee with score at least $7n$ exists. This finishes the description of the reduction. Next we prove its correctness.

Let W be a winning committee for the reduced election and let b_i be some element of the X3C instance. First we show that W has at least six candidates in it. If it is not the case, then, since each set S_j covers exactly three elements, it follows that there are at least $n - 18$ elements not covered by W (i.e., there are at least $n - 18$ elements which are not covered by the sets the correspond to the members of W). Let b_i be an element not covered by W . Then, the voters v_i^1, v_i^2 , and v_i^3 corresponding to b_i give at most 2 points to W . If $W = \emptyset$ then each voter corresponding to b_i gives 0 points to W . Otherwise, if $W \neq \emptyset$, then each set voter (each of v_i^1 or v_i^2) gives at most -1 points to W , while the antiset voter gives at most 4 points. Thus, the voters corresponding to b_i give at most 2 points to W .

There are at most 18 elements which are covered by W . For each b_i which is covered by W , each of the voters v_i^1, v_i^2 , and v_i^3 corresponding to b_i give at most 4 points to b_i (since this is the maximum number of points any voter gives to any committee). Thus, the voters corresponding to b_i give at most 12 points to W .

Summarizing the above two paragraphs, we have that the total score of W which has at most six candidates is at most $(n - 18) \cdot 2 + 18 \cdot 12$. Since we assume, without loss of generality, that $n > 39$, we have that this quantity is strictly less than $7n$. Therefore, from now on we assume that W has at least six candidates in it.

Thus, let W be a committee with at least six candidates in it and let b_i be an element. Let $V_i = \{v_i^1, v_i^2, v_i^3\}$ and consider the following four cases depending on the number of times b_i is covered by the sets S_j corresponding to the candidates in W .

- **b_i is not covered by W :** In this case, the score given to W by V_i is at most $(-2) + (-2) + 4 = 0$. To see this, observe that each of

the set voters (v_i^1, v_i^2) gives to W exactly -2 points, since they do not approve any candidate from W but disapprove all candidates in W ; further, observe that the antiset voter (v_i^3) gives to W at most 4 points, as this is the maximum number of points any voter can give to any committee.

- **b_i is covered exactly once by W :** In this case, the score given to W by V_i is $2 + 2 + 4 - 1 = 7$. To see this, observe that each of the set voters (v_i^1, v_i^2) gives to W exactly 2 points, since they approve one candidate from W (the one candidate corresponding to the one set covering b_i) and disapprove all other candidates in W ; further, observe that the antiset voter (v_i^3) gives to W exactly 3 points, since it approves at least one candidate in W and disapprove exactly one candidate in W (the one candidate corresponding to the one set covering b_i).
- **b_i is covered more than once by W :** In this case, the score given to W by V_i is $2+2+4-2 = 6$. To see this, observe that each of the set voters (v_i^1, v_i^2) gives to W exactly 2 points, since they approve more than one candidate from W (the two or three candidates corresponding to the two or three sets covering b_i) and disapprove all other candidates in W ; further, observe that the antiset voter (v_i^3) gives to W exactly 2 points, since it approves at least one candidate in W and disapprove two or three candidates in W (the two or three candidates corresponding to the two or three sets covering b_i).

As there are exactly n elements, it follows from the case analysis above that a committee W with score at least $7n$ corresponds to an exact cover of the elements b_1, \dots, b_n by sets from S . \square

4 (Net-)Capped Satisfaction Rules

Kilgour and Marshall [19] introduced the following rule in the context of electing committees of fixed size, and Kilgour [18] recalled it in the context of elections with a variable number of winners, suggesting its net version:

Capped Satisfaction Approval (CSA). The Capped Satisfaction Approval (CSA) score of a committee S is $\sum_{v_i \in V} \frac{|S \cap v_i|}{|S|}$. The committees with the highest score tie as co-winners.

Net Capped Satisfaction Approval (NCSA). The NCSA rule uses the “net” variant of CSA score; specifically, the score of a committee S is $\sum_{v_i \in V} \left[\frac{|S \cap v_i|}{|S|} - \frac{|S \cap \bar{v}_i|}{|S|} \right]$ and the committees with the highest score tie as co-winners.

In the definitions above, the idea behind dividing the scores by the size of the committee is to ensure that the rule is biased towards smaller committees. Unfortunately, for the rules as defined by Kilgour [18], this effect is too strong, leading mostly to committees containing only the candidate(s) with the highest approval score. We explain why this is the case and suggest a modification.

Consider an election $E = (C, V)$ with candidate set $C = \{c_1, \dots, c_m\}$ and preference profile $V = (v_1, \dots, v_n)$. Let $s(c_1), \dots, s(c_m)$ be the approval scores of the candidates, and, without loss of generality, assume that $s(c_1) \geq s(c_2) \geq \dots \geq s(c_m)$. Note that, if there are no ties regarding the approval scores, then for each k , the highest-scoring CSA committee of size k is simply $S_k = \{c_1, \dots, c_k\}$ and its score is

$$\sum_{v_i \in V} \frac{|S_k \cap v_i|}{|S_k|} = \frac{1}{k} \sum_{v_i \in V} |S_k \cap v_i| = \frac{1}{k} (s(c_1) + \dots + s(c_k)).$$

This value never increases with k and, so, typically CSA outputs very small committees (which contain only the candidates with the highest approval score; the same reasoning applies to NCSA). Thus, we introduce the q -CSA and q -NCSA rules, where q is a real number, $0 \leq q \leq 1$, and (a) the q -CSA score of a committee S in election $E = (C, V)$ is $\sum_{v_i \in V} \frac{|S \cap v_i|^q}{|S|^q}$, and (b) the q -NCSA score of this committee is

$$\sum_{v_i \in V} \frac{|S \cap v_i|^q}{|S|^q} - \frac{|S \cap \bar{v}_i|^q}{|S|^q}.$$

Note that for $q = 1$ these rules are, simply, CSA and NCSA, whereas 0-NCSA is NAV and 0-CSA is a rule that outputs the committee that includes all the candidates that receive at least one approval.

By the reasoning above, for each rational value of q , both q -CSA and q -NCSA are polynomial-time computable.




Figure 2: Average committee sizes (y -axis) under q -CSA and q -NCSA rules for different values of q (x -axis); elections with 20 candidates, votes generated using the 0.5-IA model.




Figure 3: Percentages of committees with a given size for (a) 0.9-CSA and (b) 0.9-NCSA (committee sizes are on the x -axis; elections have 20 candidates and 20 voters, and are generated according to the p -IA model for $p \in \{0.3, 0.5, 0.7\}$). The x -axis gives the committee size and the y -axis gives the probability that a committee of a given size is winning.

Experiment 4 To obtain a better understanding of the influence of the parameter q on the size of the committees elected according to q -CSA and q -NCSA, we have computed the average sizes of their winning committees under the 0.5-IA model, for the case of both 20 voters and 100 voters, for q values between 0 and 1 with step 0.01. The average sizes of the committees we obtained are presented in Figure 2. While the average committee size for q -NCSA does not depend very strongly on the number of voters (and its dependence on q is appealing), the results for q -CSA are worrying. Not only does the rule elect (nearly) all candidates for most values of q , but also for the values where it is more selective (e.g., $q = 0.9$), the average size of its committees depends very strongly on the number of voters. In Figure 3 we show how frequently 0.9-CSA and 0.9-NCSA choose

committees of particular sizes (for the case of 20 voters and the p -IA model with $p \in \{0.3, 0.5, 0.7\}$). Notice that there is much more variance in the behavior of 0.9-CSA than in that of 0.9-NCSA.

In Figure 4 we show average sizes of 0.9-CSA and of 0.9-NCSA committees for the p -IA model, depending on probability p . These figures confirm our worries regarding the q -CSA rules: While the dependence of the average committee size on the candidate approval probability for 0.9-NCSA has the same nature irrespective of the number of voters (it is, roughly speaking, convex both for 20 and 100 voters), the same dependence for 0.9-CSA changes its nature (from roughly convex for 20 voters to roughly concave for 100 voters). On the other hand, Figure 5 shows how the average committee sizes under 0.9-CSA and 0.9-NCSA change under the α -urn model, depending on the contagion parameter. Apparently, as soon as there is some vote correlation ($\alpha > 0.1$, say), 0.9-CSA stabilizes.

Given the above experiments, we believe that for practical applications, where we may have limited control on the number of candidates, the number of voters, and the types of the votes cast, choosing an appropriate value of the parameter q for q -CSA rules (e.g., to promote committees close to a particular size) would be very difficult. On the other hand, q -NCSA might be robust enough to be practical.

We conclude this section with a different rule of Kilgour [18], which is not an (N)CSA rule, but is somewhat similar as it also chooses a certain number of candidates with the highest approval scores.

FirstMajority. This rule outputs all the committees of the smallest possible size such that (a) each candidate in the committee has at least as high approval score as each candidate outside, and (b) the sum of the approval scores of the candidates in the committee is higher than the sum of the scores of the candidates outside.

The definition of FirstMajority gives a polynomial-time algorithm for computing its winning committees: Sort the candidates in decreasing order of their approval scores and iteratively add them until constraint (b) is satisfied.

Experiment 5 In all our experiments (for the p -IA model and for the α -urn model), FirstMajority always outputs committees whose average size is slightly below half of the number of candidates (see Figures 4 and 5 and Table 1).

5 Threshold Rules

We conclude our discussion by considering the threshold rules of Fishburn and Pekeč [15]. Let $t: \mathbb{N} \rightarrow \mathbb{N}$ be some function referred to as the *threshold function*. The t -Threshold rule is defined as follows.

t -Threshold. Consider an election $E = (C, V)$. Under the t -Threshold rule, we say that a voter $v_i \in V$ approves a committee S if $|S \cap v_i| \geq t(|S|)$. The t -Threshold rule outputs those committees that are approved by the largest number of voters.

We consider the following (in some sense extreme) threshold functions: (a) the unit function $t_{\text{unit}} = t_1$ (defined earlier); (b) the majority function, $t_{\text{maj}}(k) = k/2$; and (c) the full function, $t_{\text{full}}(k) = k$.

The t_{unit} -Threshold rule is the same as $(t_1, 0)$ -NAV since a voter approves a committee exactly if it includes at least one candidate that this voter approves; the rule outputs all committees S such that each voter with a nonempty approval ballot approves some member of S . On the other hand, the t_{full} -Threshold rule outputs exactly such

committees S that (a) each candidate in S has the highest approval score and (b) all the candidates in S are approved by the same group of voters. Thus the rule does not seem to be very useful.

In spite of this discouraging beginning, the t_{maj} -Threshold rule, introduced and studied by Fishburn and Pekeč [15], is intriguing: t_{maj} -Threshold winning committees receive broad voter support, and—as suggested by Fishburn and Pekeč—should be “of moderate size”. Computing them is NP-hard, but there are FPT algorithms.

Theorem 3 Deciding if there is a nonempty committee that satisfies all the voters under the t_{maj} -threshold rule is NP-hard.

Proof. We describe a reduction from SET COVER: let our input instance be $I = (U, \mathcal{S}, k)$, where $U = \{u_1, \dots, u_n\}$ is a set of elements, $\mathcal{S} = \{S_1, \dots, S_m\}$ is a family of subsets of U , and k is a positive integer. Without loss of generality, we assume that $m > k$ (otherwise there would be a trivial solution for our input instance).

We form an election with the candidate set $C = F \cup \mathcal{S}$, where $F = \{f_1, \dots, f_k\}$ is a set of *filler candidates* and \mathcal{S} is a set of candidates corresponding to the sets from the SET COVER instance (by a small abuse of notation, we use the same symbols for \mathcal{S} and its contents irrespective if we interpret it as part of the SET COVER instance or as candidates in our elections). We introduce $kn + 2$ voters:

1. The first voter approves all the filler candidates and the second voter approves all the set candidates. We refer to these voters as the *balancing voters*.
2. For each element $u_i \in U$, we have a group of k voters, so that the j th voter in this group ($j \in [k]$) approves all filler candidates except f_j , and also those set candidates that correspond to sets containing u_i .

We claim that there is a nonempty committee S such that every voter approves at least half of the members of S (i.e., every voter is satisfied) if and only if I is a *yes*-instance.

Assume that S is a committee that satisfies all the voters. We note that S must contain the same number of filler and set candidates. If it contained more set candidates than filler candidates then the first balancing voter would not be satisfied, and if it contained more filler candidates than set candidates, then the second balancing voter would not be satisfied. Thus, there is a number k' such that $|S| = 2k'$, $k' \leq k$, and S contains exactly k' filler and k' set candidates.

We claim that these k' set candidates correspond to a cover of U . Consider some arbitrary element u_i and some filler candidate f_j such that f_j does not belong to S (since $m > k \geq k'$ such candidates must exist). There is a voter that approves all the filler candidates except f_j and all the set candidates that contain u_i . Thus, the committee contains exactly $k' - 1$ filler candidates that this voter approves and—to satisfy this voter—must contain at least one set candidate that contains u_i . Since u_i was chosen arbitrarily, we conclude that the set candidates from S form a cover of U . There are at most k of them, so I is a *yes*-instance.

On the other hand, if there is a family of $k' \leq k$ sets that jointly cover U , then a committee that consists of arbitrarily chosen k' filler candidates and the set candidates corresponding to the cover satisfies all the voters. \square

Theorem 4 Let t be a linear function (i.e., $t(k) = \alpha k$, $\alpha \in [0, 1]$). The t -Threshold rule is computable in FPT time for parameterizations by the number of candidates and by the number of voters.

Proof. For parameterization by the number of candidates it suffices to try all possible committees. For parameterization by the number of voters, we combine the candidate-type technique of Chen et al. [10] and an integer linear programming (ILP) approach. The *type* of candidate c is the subset of voters that approve c . For an election with n voters, each candidate has one of at most 2^n types. We describe an algorithm for computing a committee approved by at least N voters (where N is part of the input; it suffices to try all values of $N \in [n]$ to find a committee with the highest score). We focus on computing the largest winning committee.

Let $E = (C, V)$ be the input election with n voters. We form an instance of ILP as follows. For each candidate type i , $i \in [2^n]$, we introduce integer variable x_i (intuitively x_i is the number of candidates of type i that are included in the winning committee). For each $i \in [2^n]$, we form constraint $0 \leq x_i \leq n_i$, where n_i is the number of candidates of type i in election E . We also add constraint $\sum_{i \in [2^n]} x_i \geq 1$ as the winning committee must be nonempty.

For each voter $j \in [n]$, we define an integer variable v_j (the intention is that v_j is 1 if the j th voter approves of the committee specified by variables x_0, \dots, x_{2^n-1} and it is 0 otherwise; see also comments below). For each $j \in [n]$, we introduce constraints $0 \leq v_j \leq 1$, and:

$$\left(\sum_{i \in \text{types}(v_j)} x_i \right) - t \left(\sum_{i \in [2^n]} x_i \right) \geq -(1 - v_j)n, \quad (1)$$

where $\text{types}(v_j)$ is the set of all candidate types approved by voter v_j . To understand these constraints, note that $\sum_{i \in [2^n]} x_i$ is the size of the selected committee, $\sum_{i \in \text{types}(v_j)} x_i$ is the number of committee members approved by the j th voter, and, thus, Eq. (1) is satisfied either if $v_j = 0$ or if $v_j = 1$ and there is an integer k such that the j th voter approves at least $t(k)$ members of the selected size- k committee. We add constraint $v_1 + \dots + v_n \geq N$ (i.e., we require that at least N voters are satisfied with the selected committee; this also prevents satisfying Eq. (1) by setting $v_j = 0$ for all $j \in [n]$).

To compute the largest committee approved by at least N voters, we find a feasible solution (for the above-described integer linear program) that maximizes $\sum_{i \in [2^n]} x_i$ (we use the famous FPT-time algorithm of Lenstra [21]). \square

Table 1: Average committee sizes (20 candidates and either 20 or 100 voters, the 0.5-IA model of generating approvals). Rules are sorted with respect to the average committee size for 20 voters (results in **bold** are those that would change their position if we sorted for the average committee size with 100 voters). $t_{\text{maj}}\text{-Thr. (min)}$ and $t_{\text{maj}}\text{-Thr. (max)}$ refer to the smallest and largest committees under the $t_{\text{maj}}\text{-Threshold}$ rule.

rule	avg. committee size ± its std. deviation		complexity
	(20 voters)	(100 voters)	
0.5-Independent Approval model			
2/3-NAV	1.02 ± 1.01	0.01 ± 0.09	P
0.9-NCSA	1.52 ± 0.89	1.50 ± 0.78	P
MRC	2.63 ± 0.48	4.08 ± 0.26	NP-hard
GreedyMRC	2.75 ± 0.46	4.55 ± 0.53	P
$t_{\text{maj}}\text{-Thr. (min)}$	2.75 ± 1.33	2.05 ± 0.34	NP-hard
0.5-NCSA	5.57 ± 2.14	5.57 ± 2.18	P
0.9-CSA	5.63 ± 3.02	14.67 ± 2.75	P
$t_{\text{maj}}\text{-Thr. (max)}$	7.68 ± 3.27	2.20 ± 0.78	NP-hard
NAV	8.25 ± 2.19	9.19 ± 2.23	P
FirstMajority	9.51 ± 0.43	9.50 ± 0.25	P
0.5-CSA	19.74 ± 0.52	20.00 ± 0.00	P

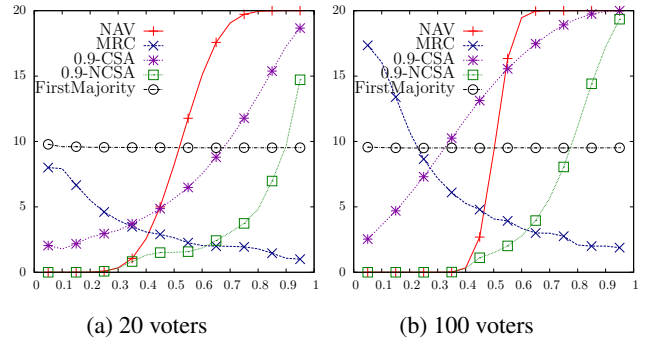


Figure 4: Average committee sizes for some of our rules (20 candidates and either 20 or 100 voters; p -IA model, p is on the x -axis).

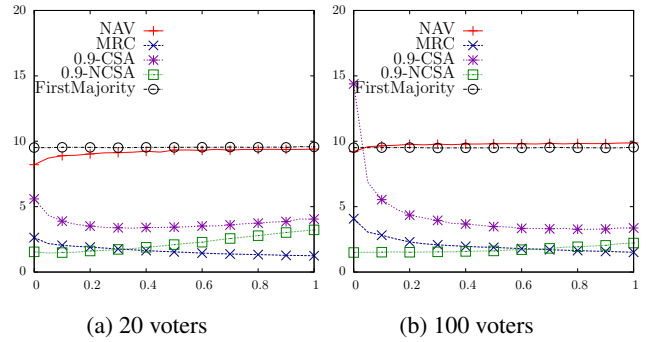


Figure 5: Average committee sizes for some of our rules (20 candidates and either 20 or 100 voters; α -urn model, α is on the x -axis).

Experiment 6 For the 0.5-IA model and 20 voters, the average size of the smallest $t_{\text{maj}}\text{-Threshold}$ committee was 2.84. On the other hand, the largest committee contained, on average, 7.52 candidates. Yet, for the case of 100 voters the difference between the sizes of the largest committee and the smallest committee are much more modest (see Table 1).

6 Conclusions and Further Research

We demonstrated the usefulness of voting rules that output a variable number of winners and analyzed a number of such rules, including generalizations of rules that have been already presented in the literature [17, 6, 15]. We found polynomial algorithms in most cases, but also identified interesting NP-hard rules and proposed ways to break their computational intractability. We have also performed a number of experiments; interestingly, we found that for the p -Independent Approval model (where each voter approves each candidate independently, with probability p), the average sizes of the committees that our rules output vary strongly depending on p (see Figure 4), but the dependence of the committee sizes on the contagion parameter (for the α -urn model, Figure 5) is much more modest. We believe that the reason for this is that in the urn model, on the average, each voter approves of 50% of the candidates and this parameter has stronger influence on the result than the correlation between the votes (measured by the contagion parameter).

Further axiomatic and computational analysis, as well as more extensive simulations (including experiments on real-world data), are the most pressing directions for future research. In particular, studying approximation algorithms, further multivariate analysis, also with respect to certain domain restrictions, as well as studying the complexity of performing manipulative actions (e.g., control and

bribery) to change the size of the winning committee are interesting directions.

Finally, we note that in many practical cases there is a societal preference on the size of the committee to be elected, which is usually single-peaked. Incorporating this preference into the voting rules is an interesting direction of research.

Acknowledgments

Piotr Faliszewski was partially supported by the funds of Polish Ministry of Science and Higher Education assigned to AGH University of Science and Technology. Nimrod Talmon was supported by the Israel Science Foundation (grant No. 630/19).

REFERENCES

- [1] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh, ‘Justified representation in approval-based committee voting’, *Social Choice and Welfare*, **48**(2), 461–485, (2017).
- [2] S. Barberà and D. Coelho, ‘How to choose a non-controversial list with k names’, *Social Choice and Welfare*, **31**(1), 79–96, (2008).
- [3] S. Berg and D. Lepelley, ‘On probability models in voting theory’, *Statistica Neerlandica*, **48**(2), 133–146, (1994).
- [4] N. Betzler, A. Slinko, and J. Uhlmann, ‘On the computation of fully proportional representation’, *Journal of Artificial Intelligence Research*, **47**, 475–519, (2013).
- [5] D. Binkele-Raible, G. Erdélyi, H. Fernau, J. Goldsmith, N. Mattei, and J. Rothe, ‘The complexity of probabilistic lobbying’, *Discrete Optimization*, **11**(2), 1–21, (2014).
- [6] S. Brams and M. Kilgour, ‘Satisfaction approval voting’, in *Voting Power and Procedures*, 323–346, Springer, (2014).
- [7] S. Brams, M. Kilgour, and R. Sanver, ‘A minimax procedure for negotiating multilateral treaties’, in *Diplomacy games*, 265–282, Springer, (2007).
- [8] Florian Brandl and Dominik Peters, ‘An axiomatic characterization of the borda mean rule’, *Social choice and welfare*, **52**(4), 685–707, (2019).
- [9] B. Chamberlin and P. Courant, ‘Representative deliberations and representative decisions: Proportional representation and the Borda rule’, *American Political Science Review*, **77**(3), 718–733, (1983).
- [10] J. Chen, P. Faliszewski, R. Niedermeier, and N. Talmon, ‘Elections with few voters: candidate control can be easy’, *Journal of Artificial Intelligence Research*, **60**, 937–1002, (2017).
- [11] R. Christian, M. Fellows, F. Rosamond, and A. Slinko, ‘On complexity of lobbying in multiple referenda’, *Review of Economic Design*, **11**(3), 217–224, (2007).
- [12] C. Duddy, A. Piggins, and W. Zwicker, ‘Aggregation of binary evaluations: a Borda-like approach’, *Social Choice and Welfare*, **46**(2), 301–333, (2016).
- [13] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko, ‘Properties of multiwinner voting rules’, *Social Choice and Welfare*, **48**(3), 599–632, (2017).
- [14] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon, ‘Multiwinner voting: A new challenge for social choice theory’, in *Trends in Computational Social Choice*, ed., U. Endriss, AI Access Foundation, (2017).
- [15] P. Fishburn and A. Pekeč. Approval voting for committees: Threshold approaches. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.727.465&rep=rep1&type=pdf>, 2004. Retrieved: February 3rd, 2017.
- [16] E. Hemaspaandra, H. Spakowski, and J. Vogel, ‘The complexity of Kemeny elections’, *Theoretical Computer Science*, **349**(3), 382–391, (2005).
- [17] M. Kilgour, ‘Approval balloting for multi-winner elections’, in *Handbook on Approval Voting*, pp. 105–124. Springer, (2010). Chapter 6.
- [18] M. Kilgour, ‘Approval elections with a variable number of winners’, *Theory and Decision*, **81**(2), 1–13, (2016).
- [19] M. Kilgour and E. Marshall, ‘Approval balloting for fixed-size committees’, in *Electoral Systems: Paradoxes, Assumptions, and Procedures*, ed., M. Machover D. Felsenthal, 305–326, Springer, (2012).
- [20] M. Lackner and P. Skowron, ‘Consistent approval-based multi-winner rules’, in *Proceedings of EC ’18*, pp. 47–48, (2018).
- [21] H. Lenstra, Jr., ‘Integer programming with a fixed number of variables’, *Mathematics of Operations Research*, **8**(4), 538–548, (1983).
- [22] J. McCabe-Dansted and A. Slinko, ‘Exploratory analysis of similarities between social choice rules’, *Group Decision and Negotiation*, **15**(1), 77–107, (2006).
- [23] B. Monroe, ‘Fully proportional representation’, *American Political Science Review*, **89**(4), 925–940, (1995).
- [24] I. Nehama, ‘Complexity of optimal lobbying in threshold aggregation’, in *Proceedings of ADT ’15*, pp. 379–395, (2015).
- [25] A. Procaccia, J. Rosenschein, and A. Zohar, ‘On the complexity of achieving proportional representation’, *Social Choice and Welfare*, **30**(3), 353–362, (2008).
- [26] T. Walsh, ‘Where are the hard manipulation problems’, *Journal of Artificial Intelligence Research*, **42**(1), 1–29, (2011).