

What Do Multiwinner Voting Rules Do?

An Experiment Over the Two-Dimensional Euclidean Domain

Edith Elkind
University of Oxford
Oxford, UK
elkind@cs.ox.ac.uk

Piotr Skowron
University of Oxford
Oxford, UK
piotr.skowron@cs.ox.ac.uk

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

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

Jean-François Laslier
Paris School of Economics
Paris, France
jean-francois.laslier@ens.fr

Nimrod Talmon
Weizmann Institute of Science
Rehovot, Israel
nimrod.talmon@weizmann.ac.il

Abstract

We visualize aggregate outputs of popular multiwinner voting rules—SNTV, STV, Bloc, k -Borda, Monroe, Chamberlin–Courant, and PAV—for elections generated according to the two-dimensional Euclidean model. We consider three applications of multiwinner voting, namely, parliamentary elections, portfolio/movie selection, and shortlisting, and use our results to understand which of our rules seem to be best suited for each application. In particular, we show that STV (one of the few nontrivial rules used in real high-stake elections) exhibits excellent performance, whereas the Bloc rule (also often used in practice) performs poorly.

Introduction

The goal of this paper is to develop a better understanding of a number of well-known multiwinner voting rules, by analyzing their behavior in elections where voters’ preferences are generated according to a two-dimensional spatial model. By focusing on this preference domain, we can visualize the election results and check if they agree with the intuition and motivation behind these rules. Our study can be seen as an experimental counterpart of the work of Elkind et al. (2014; 2017), who analyze multiwinner rules axiomatically.

In a multiwinner election, the goal is to select a size- k committee (i.e., a set of k candidates, where $k \in \mathbb{N}$ is part of the input) based on the voters’ preferences. Usually, voters can express their preferences by listing the candidates from best to worst or by indicating which candidates they approve; we focus on the former setting, as it fits the spatial preference model better.

Applications of multiwinner voting range from choosing a parliament through preparing a portfolio of company’s products (Lu and Boutilier 2011; 2015) or choosing movies to offer to passengers on a long flight (Elkind et al. 2014; 2017; Skowron, Faliszewski, and Lang 2016) to shortlisting runners-up for an award (Barberà and Coelho 2008; Elkind et al. 2014; 2017). As a consequence, there is also

quite a variety of different multiwinner voting rules. For instance, for parliamentary elections an important desideratum is proportional representation of the voters, and there are voting rules such as STV or the Monroe rule (we define all rules considered in this paper in the next section) that have been designed with this idea in mind. On the other hand, in the context of portfolio or movie selection we primarily care about the diversity of the selected committee, and it has been argued that the Chamberlin–Courant rule is good for this purpose (Lu and Boutilier 2011; Skowron, Faliszewski, and Lang 2016). For shortlisting, our primary concern is fairness: if there are two similar candidates, we want to select both or neither, and increasing the target committee size should not result in any of the selected candidates being dropped; these requirements are satisfied by k -Borda. Naturally, there are other scenarios which require other normative properties.

The examples above indicate that choosing a good multiwinner rule is not a trivial task. It is therefore natural to ask how we can facilitate the decision-making process of a user who is facing this choice. There are several good answers to this question. First, some rules are specifically designed for certain tasks. For example, STV and the Monroe rule have explicit built-in mechanisms ensuring that every sufficiently large group of like-minded voters is represented. Second, we can analyze axiomatic properties of the rules. This line of work, was extensively pursued for single-winner rules; for the case of multiple winners it was initiated by Felsenthal and Maoz (1992) and Debord (1992), with recent contributions including the work of Elkind et al. (2014; 2017) and Aziz et al. (2015a; 2017). Finally, one can use empirical analysis to compare different rules under particular conditions. For example, Diss and Doghmi (2016) consider a few multiwinner voting rules and experimentally investigate how frequently they pick Condorcet committees.¹ All these approaches are useful, and the choice of a voting rule should take all of them into account.

Nonetheless, a non-expert user may still feel ill at ease when deciding which rule to choose for his or her particu-

¹In a Condorcet committee, every committee member is preferred to every non-member by a majority of the voters.

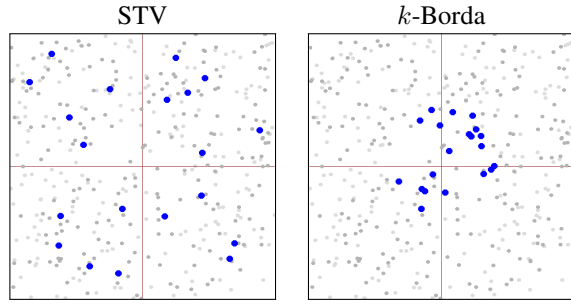


Figure 1: Results of an election (generated using the 2D Euclidean model) according to STV (left) and k -Borda (right). Voters are depicted as dark gray dots, candidates as light gray dots, and the winners as larger blue dots.

lar application. In this case, a picture may be worth a thousand words: a simple graph that clearly explains differences between rules can be very informative. The contribution of this paper is to propose a novel approach to selecting a suitable multiwinner rule, which is based on graphical information. That is, we provide images that we expect to be helpful in discussions of multiwinner voting rules. Naturally, reality is too complicated for a single picture to constitute a definite argument, but we believe that, on the one hand, our results provide good illustrations confirming intuitions regarding various multiwinner rules and, on the other hand, they highlight some faults of the rules that otherwise would not be easily visible.

Our Methodology. The outcome of an election depends both on the voting rule and on the set of candidates. In this work, we focus on the former aspect and ask what multiwinner rules do when choosing from a set of candidates that is representative of the electorate, i.e., under what one may call the *representative candidacy* assumption. We evaluate a number of multiwinner voting rules (SNTV, STV, Bloc, k -Borda, Chamberlin–Courant, Monroe, and PAV) on elections generated using the two-dimensional Euclidean model of preferences. In this model each candidate and each voter is represented by a point on a plane, and voters form their preference orders by ranking the candidates that are closer to them above the ones that are further away.

This model is very appealing and extensively studied (Enelow and Hinich 1984; 1990) because of its natural interpretations: A point representing a candidate or a voter simply specifies his or her position regarding two given issues. In the world of politics, these two issues could be, for example, the preferred levels of taxation and immigration, or the extent to which the individual believes in personal and economic freedom. While in some settings more dimensions may be necessary, the popularity of the Nolan Chart, which is used to represent the spectrum of political opinions, indicates that two dimensions are often sufficient to provide a good approximation of voters’ preferences.

In Figure 1 we show a sample election (the points for candidates and voters are generated using uniform distribution over a square) and the committees selected by STV (left)

and k -Borda (right). It is quite evident that the committee on the left would form a far more representative parliament than the one on the right, whereas the one on the right would probably be a better choice for the set of candidates that are shortlisted for a position, since they are similar to each other and receive broad support among the voters (in particular, no voter ranks them close to the bottom of their list).

Our main contributions are as follows:

1. For each of our voting rules and four distributions of candidates and voters (Gaussian, uniform on a disc, uniform on a square, and a mix of four Gaussians), we have generated 10 000 elections and built histograms (Figure 3) indicating how likely it is that a candidate from a given position will be selected.
2. We consider three applications of multiwinner voting, and, for each application, we identify the voting rules in our collection that are most appropriate for it. We make these recommendations based on our histograms and certain statistical properties of the elected committees. E.g., we confirm that STV is an excellent rule for parliamentary elections, even superior to the Monroe rule; PAV can also be seen as an interesting rule that chooses fairly representative committees, ignoring candidates with extreme opinions. We also provide evidence that Bloc should be treated very carefully since it may not perform as well as one might expect (this is particularly important because Bloc is among the most popular multiwinner rules).

Due to space restrictions, we omit some of our results (in particular, the analysis of approximation algorithms for the Monroe and Chamberlin–Courant rules); these results are available in the full version of this paper.

Preliminaries

For every positive integer n , we write $[n]$ to denote the set $\{1, \dots, n\}$.

Elections. An *election* $E = (C, V)$ consists of a set $C = \{c_1, \dots, c_m\}$ of *candidates* and a list $V = (v_1, \dots, v_n)$ of *voters*. Each voter v_i has a *preference order* \succ_i , i.e., a ranking of the candidates from the most to the least favored one (according to this voter). For a voter v and a candidate c , we write $\text{pos}_v(c)$ to denote the position of c in v ’s preference order (where the top-ranked candidate has position 1). A *committee* is a subset of C .

A *multiwinner voting rule* is a function \mathcal{R} that, given an election $E = (C, V)$ and a target committee size k ($1 \leq k \leq |C|$), outputs a nonempty set of size- k committees; these committees are said to *tie as election winners*. In practice, one has to use some tie-breaking mechanism. For our experiments, whenever we need to break a tie (possibly at an intermediate stage in the execution of the rule), we make a random choice with a uniform distribution over all possibilities.

(Single-Winner) Scoring Functions. For an election with m candidates, a *scoring function* γ_m associates each position j , $j \in [m]$, with a *score* $\gamma_m(j)$. The γ_m -score that

candidate c receives from voter v is $\gamma_m(\text{pos}_v(c))$. The γ_m -score of candidate c in election E is the sum of the γ_m -scores that c receives from the voters in E . We consider the following two prominent families of scoring functions:

1. The Borda scoring function, β_m , is defined as $\beta_m(j) = m - j$.
2. For each $t \in [m]$, the t -Approval scoring function, α_t , is defined as $\alpha_t(j) = 1$ if $j \leq t$ and $\alpha_t(j) = 0$ otherwise. The candidate's 1-Approval score is known as her Plurality score.

Multiwinner Rules. We focus on the following multiwinner rules (in the description below we consider an election $E = (C, V)$ and committee size k):

SNTV. The Single Nontransferable Vote rule (SNTV) outputs k candidates with the highest Plurality scores.

STV. The Single Transferable Vote rule (STV) executes a series of iterations, until it finds k winners. A single iteration operates as follows: If there is at least one candidate with Plurality score at least $q = \lfloor \frac{n}{k+1} \rfloor + 1$, then a candidate with the highest Plurality score is added to the committee; then q voters that rank him or her first are removed from the election (our randomized tie-breaking plays an important role here), and the selected candidate is removed from all voters' preference orders. If there is no such candidate, then a candidate with the lowest Plurality score is removed from the election (again, ties are broken uniformly at random). The Plurality scores are then recomputed.

Bloc. Under the Bloc rule we output k candidates with the highest k -Approval scores (intuitively, each voter is asked to name his or her k favorite committee members, and those mentioned most frequently are elected).

k -Borda. Under the k -Borda rule we output k candidates with the highest Borda score.

β -CC. The (classical) Chamberlin–Courant rule (β -CC) is defined as follows (Chamberlin and Courant 1983). A k -CC-assignment function is a function $\Phi: V \rightarrow C$ such that $|\Phi(V)| \leq k$ (i.e., Φ associates each voter with a candidate in a set $W \subseteq C$, $|W| \leq k$; for a voter v , candidate $\Phi(v)$ is referred to as v 's representative). The β -CC score of an assignment Φ is defined as $\beta(\Phi) = \sum_{v \in V} \beta_m(\text{pos}_v(\Phi(v)))$ (i.e., it is the sum of the Borda scores of voters' representatives). β -CC finds a k -CC-assignment Φ that maximizes $\beta(\Phi)$ and outputs the committee $\Phi(V)$ (if it happens that $|\Phi(V)| < k$ —a situation that occurs, e.g., when all the voters have identical preference orders—then β -CC supplements $\Phi(V)$ with $k - |\Phi(V)|$ candidates selected at random).

β -Monroe. The (classical) Monroe rule (Monroe 1995) is similar to β -CC, except that it is restricted to k -Monroe-assignments. A k -Monroe-assignment is a k -CC-assignment that satisfies the following constraints: (a) $|\Phi(V)| = k$, and (b) for each candidate c such that $\Phi^{-1}(c) \neq \emptyset$ (i.e., for each selected representative) it holds

that $\lfloor \frac{n}{k} \rfloor \leq |\Phi^{-1}(c)| \leq \lceil \frac{n}{k} \rceil$. Intuitively, under the Monroe rule each selected candidate represents, roughly, the same number of voters.

α_k -PAV. Consider a scoring function γ . For a voter v and a committee W such that v ranks the members of W on positions $p_1 < \dots < p_k$, the γ -PAV score that v assigns to W is $\gamma\text{-PAV}(W, v) = \sum_{t=1}^k \frac{1}{t} \gamma(p_t)$. For an election $E = (C, V)$, the γ -PAV score of a committee W is defined as $\gamma\text{-PAV}(W, V) = \sum_{v \in V} \gamma\text{-PAV}(W, v)$. The rule outputs a committee with the highest γ -PAV score. In this paper, we consider α_k -PAV (originally the rule was defined for approval ballots, see e.g., the overview of Kilgour (2010); as we work with preference orders, we modify the definition accordingly).

With our tie-breaking, STV, SNTV, Bloc, and k -Borda are computable in polynomial time using straightforward algorithms. Unfortunately, the Chamberlin–Courant and Monroe rules are NP-hard to compute (Procaccia et al. (2008) show this for variants of these rules that use t -Approval scores α_t instead of β ; for the Borda-based variants defined here, the results for the Chamberlin–Courant rule and the Monroe rule are due to Lu and Boutilier (2011) and Betzler et al. (2013), respectively). We compute these rules by solving their integer linear programming (ILP) formulations (suggested by Lu and Boutilier (2011) for the case of Chamberlin–Courant, and by Skowron et al. (2015) for the case of Monroe). PAV is also NP-hard to compute (Aziz et al. 2015b; Skowron, Faliszewski, and Lang 2016)², and we use a simplified version of the ILP formulation proposed by Skowron et al. (2016); see the full version of the paper.

Euclidean Preferences. Given two points on the plane, $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$, we write $d(p_1, p_2)$ to denote the distance $\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$ between them.

In a two-dimensional Euclidean election $E = (C, V)$, each entity e (i.e., either a candidate or a voter) is associated with a point $p(e) = (x(e), y(e))$. Given a pair of candidates $c_i, c_j \in C$, a voter $v \in V$ prefers c_i to c_j if $d(p(v), p(c_i)) < d(p(v), p(c_j))$. Note that this condition does not constrain voter's preferences over two equidistant candidates. In our case, since we draw our elections at random, such situations are unlikely to happen. When they do, we break the tie arbitrarily.

Euclidean preferences are very useful to realistically model political preferences and, in many cases, to model preferences in shortlisting tasks. Unfortunately, they are not nearly as useful for modeling preferences over movies. The reason is that people often do not have a single most favorite type of a movie, but rather like various genres for different reasons. Nonetheless, investigating rules meant for the movie selection application (i.e., for selecting diverse committees) in our framework is still important. On the one hand, movie selection is not the only application where diverse committees are needed, and, on the other hand, if a rule behaves badly on the Euclidean domain, then it is unlikely that it would behave well for richer preference models.

²the hardness proofs for this rule are in the approval model, but can be easily adapted to the preference-order based one.

Main Results and Analysis

Experimental Setup. We assume that both the candidates and the voters have ideal positions in a two-dimensional Euclidean issue space that are drawn from the same distributions. For each voting rule and each distribution, we generated 10 000 elections, each with $m = 200$ candidates and $n = 200$ voters, and for each of them we computed a winning committee of size $k = 20$.

We consider four distributions of the ideal positions:

Gaussian. Ideal points are generated using symmetric Gaussian distribution with mean $(0, 0)$ and standard deviation 1.

Uniform Square. Ideal points are distributed uniformly on the square $[-3, 3] \times [-3, 3]$.

Uniform Disc. Ideal points are distributed uniformly on the disc with center $(0, 0)$ and radius 3.

4-Gaussian. Ideal points are generated using four symmetric Gaussian distributions with standard deviation 0.5, but different mean values, namely, $(-1, 0)$, $(1, 0)$, $(0, -1)$ and $(0, 1)$; each mean is used to generate 25% of the points.

We use the Gaussian distribution to model a society with one dominant idea (e.g., where being moderate is the most popular position, or where a single dominant party exists). Since the boundary plays a significant role in the case of uniform distributions (we will discuss this effect below), we have chosen the Gaussian distribution, as its density vanishes close to the boundary.

The 4-Gaussian distribution models a structured society, with four well-established positions (for the movie selection scenario, these might correspond to, e.g., a combination of two genres and two typical budget values; in the world of politics, these could be four political parties).

We also use the uniform distributions, on a square and on a disc, as intermediate cases, and in order to study specific behavior of voting rules at the border and, in case of the square, at the corners of the support of the distribution.

Raw Results. For each rule and each distribution, we have computed a histogram, showing how frequently winners from a given location were selected. These histograms, together with examples of elections and their winning committees, are presented in Figure 3 (the first row presents the distributions themselves).

The histograms were generated as follows. For each rule and distribution, all the winners were always within the $[-3, 3] \times [-3, 3]$ square. We have partitioned this square into 120×120 cells (each cell is a 0.05×0.05 square), and—for each given distribution and rule—counted how many times a member of the winning committee fell into a given cell (we refer to this value as the *frequency* of this cell). Then we have transformed the frequencies into color intensities (the more winners fall into a particular cell, the darker it is in Figure 3). Since there are big differences among frequencies of cells across various rules and distributions (e.g., the highest frequency of a cell for k -Borda with the Gaussian distribution is over 27 times larger than the highest frequency of a

cell for SNTV under the uniform square distribution), we took the following approach. Given a cell of frequency x , we compute its color intensity y ($0 \leq y \leq 1$; the closer is y to 1 the darker is the cell) using the following formula:

$$y = \frac{1}{\pi/2} \arctan\left(\frac{x}{\varepsilon T}\right), \quad (1)$$

where T is the sum of the frequencies of all the cells (so in our case $T = 20 \cdot 10000$) and ε is a parameter. We used $\varepsilon = 0.0004$, so for the highest frequency of a cell in all our experiments (found for k -Borda with the Gaussian distribution) we have $x/(\varepsilon T) = 10.9$; for most other rules and distributions this value is below 1.5 and thus falls into the part where our function behaves fairly linearly (see Figure 2). To present the distributions themselves, we computed histograms of the ideal points generated using our distributions (on the technical side, to generate these histograms, we used candidate positions from 10 000 generated elections for each distribution; since formula (1) is normalized, the pictures in the first row of Figure 3 are comparable to those in the other rows).

Analysis. We now consider the three applications of multiwinner rules that we mentioned in the introduction and analyze which of our rules are most suitable for each application.

Parliamentary Elections. We start with the case of parliamentary elections. Intuitively, in this application we value proportional representation, which requires that the distribution of the winners (as seen through the histograms) should be as close as possible to the underlying distribution of the voters. Thus, at first sight, among our rules SNTV would be the champion in this category. In addition, SNTV satisfies a number of axioms studied by Elkind et al. (2014; 2017), especially those geared towards proportional representation. However, at the same time, it is intuitively clear that SNTV is not a very good rule because it only takes the voters' top choices into account, thus ignoring most of the information in voters' preferences. A look at the sample elections for SNTV (Figure 3) shows that this intuition is correct: The reason why SNTV has such an appealing histogram is that it selects committee members in areas that, by random chance, have above-average density of voters and below-average density of candidates. Over all 10 000 elec-




Figure 2: Plot of the function $y = \frac{1}{\pi/2} \arctan\left(\frac{x}{\varepsilon T}\right)$ that we use for converting cell frequencies to color intensities.

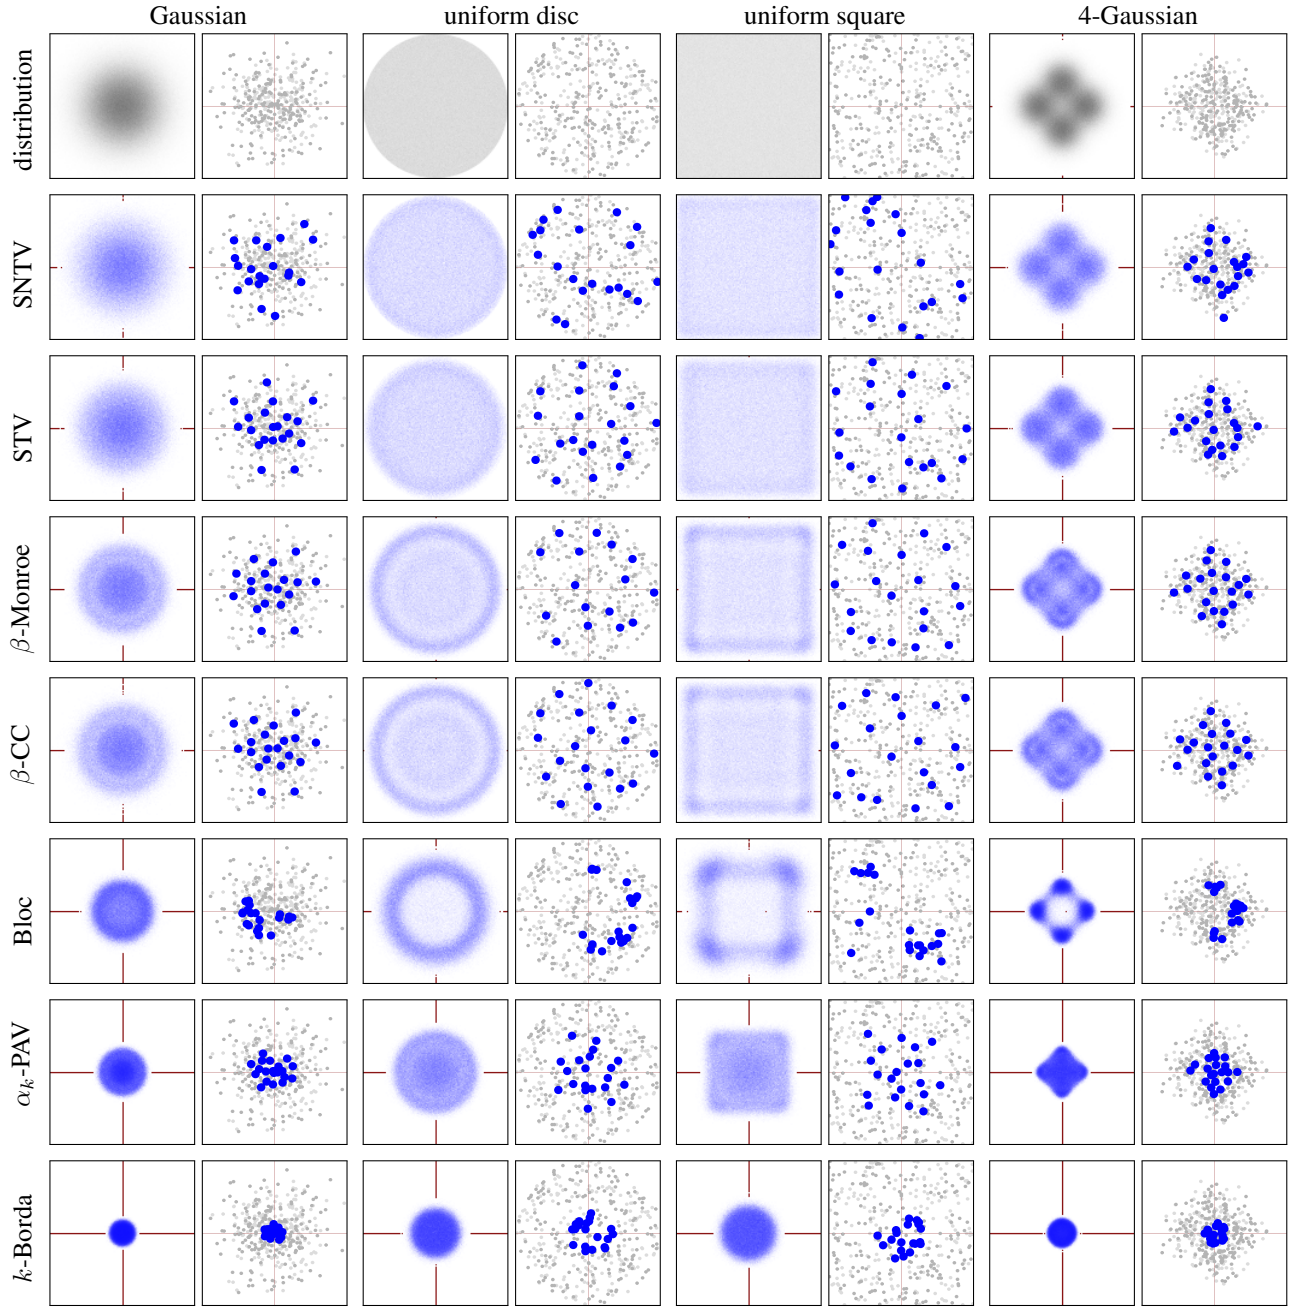


Figure 3: Histograms and sample elections for our rules and distributions. The first row shows the distributions only. For sample election, voters are depicted as dark gray dots, candidates as light gray dots, and the winners as larger blue dots.

tions such areas are distributed evenly, similarly to the distribution of the candidates and voters.

This means that, in addition to considering the histograms, we also need to check if results of individual elections are close to what the histograms show. To this end we have used an indirect approach that, nonetheless, turned out to be very effective. Let us fix some rule \mathcal{R} and one of our distributions. For each generated election, we (1) count how many members of the winning committee are in each of the

four quadrants $[0, \pm\infty) \times [0, \pm\infty)$, (2) collect these numbers in a sequence, and (3) compute the variance of this sequence; Table 1 shows the result of this computation, averaged over all instances. Since all our distributions are symmetric with respect to the x and y axes, for rules that represent voters proportionally *in individual instances* we expect this number to be small. Of course, the converse claim need not be true: Low variance does not guarantee proportional representation. That is, the variance-based approach can be used

rule	square	disc	Gauss.	4 Gauss.
SNTV	3.292	3.219	3.275	2.787
STV	0.994	1.070	1.150	1.043
β -Monroe	0.738	0.797	0.864	0.765
β -CC	0.765	0.820	0.866	0.826
Bloc	17.789	17.146	18.709	9.663
α_k -PAV	1.323	1.391	1.463	1.289
k-Borda	4.605	4.653	4.736	3.653

Table 1: Variance of the number of winners in each quadrant. Bold font indicates rules where this value suggests asymmetric placement of winners on the plane (for k -Borda, this turns out to be a false alarm).

to eliminate ‘bad’ rules rather than to identify ‘good’ rules.

Table 1 clearly identifies a group of rules for which the variance of the number of winners per quadrant is close to or below 1.0, whereas for other rules the variance is significantly higher (in our experiments, typically close to or above 3.0). Thus, the performance of SNTV (close to 3.0) is a strong argument against it. On the other hand, the results for STV (both the shape of histograms and the variance) indicate that it is an exceedingly good rule for selecting parliaments. Indeed, this is the only rule with low variance that is computationally tractable. This is quite important, as STV is among just a few nontrivial voting rules used in practice, yet some researchers—including some of us, until recently—consider it unappealing. The axiomatic results of Elkind et al. (2014; 2017) and our experiments provide different arguments in favor of using STV for proportional representation.

The results for β -Monroe are slightly less appealing than those for STV. While the variance of the number of winners per quadrant is low, the histograms are farther from resembling the distributions of candidates and voters. They are very similar to those for β -CC, which should not be too surprising. In our experiments, the only difference between these rules is that β -Monroe is forced to assign exactly 10 voters to each selected committee member, whereas β -CC can choose an optimal assignment, where the number of voters assigned to each committee member may be arbitrary. Nonetheless, for each of the distributions, around 80% of the committee members selected by β -CC were assigned to between 7 and 13 voters each. In effect, the assignments computed by β -CC and β -Monroe were quite similar. Naturally, if the distributions of candidates and voters were not identical, the results would be different as well (we have run initial experiments to confirm this, available in the full version of the paper). Below we discuss the intriguing patterns in the histograms for β -CC (a similar explanation applies to β -Monroe).

Portfolio/Movie Selection. Let us now consider the portfolio/movie selection scenario (Lu and Boutilier 2011; 2015; Elkind et al. 2014; 2017; Skowron, Faliszewski, and Lang 2016). Here we care mostly about the diversity of the committee and, intuitively, we would like to obtain histograms that cover a large chunk of the support of the distribution, but which—as compared to the parliamentary elections

setting—are less responsive to the densities of the candidates and voters.

We first analyze the results for β -CC, a rule that seems to be designed exactly for this scenario. However, it does not quite fit the description above. As we will see, to some extent this is due to the nature of the rule, and to some extent this is because our initial expectations were not entirely reasonable. There are two main issues regarding β -CC.

The first one concerns what we call the *edge effect* and the *corner effect*. Let us consider the uniform square distribution. If a candidate is located far from the edges, then he or she is also surrounded by a relatively large number of other candidates with whom he or she needs to compete for a high position in voters’ preference orders. On the other hand, if a candidate is located near an edge (or, better yet, near a corner) then the competition is less stiff. However, if a candidate is close to the edge/corner, the number of voters for whom he or she would be a representative also decreases. In effect, for the uniform square and uniform disc distributions, we see increased frequencies of winners near (but not exactly on) the edges and corners. The edge and corner effects are visible also for SNTV and STV (though to a lesser extent), and they are very prominent for Bloc (especially in conjunction with cases where an area near edge/corner has an above-average density of voters).

The second issue regarding β -CC is that when some candidate is included in the committee, other candidates that are very close to him or her are unlikely to be selected; indeed, this behavior is quite desirable when one wants to maintain diversity of the committee. This explains why for the uniform square and uniform disc distributions the near-edge area with increased frequencies is surrounded by an area with lower frequencies. This effect also explains the interesting pattern for the 4-Gaussian distribution. Since there are many voters in the centers of the four Gaussians, candidates from these locations are likely to be included in the committee. But this very fact strongly decreases the chances of the candidates that are located just a bit further away from the centers of the Gaussians.

Our visual inspection of the election results for β -CC shows that every single committee appears to be diverse and appealing for the portfolio/movie selection problem (this is also supported by the low value of the variance of the number of winners per quadrant). However, the histograms show that the rule also has an implicit, systematic bias against certain candidates (the nature of this bias depends on the distribution) that users of the rule should take into account.

α_k -PAV also appears to be a very interesting rule for the portfolio/movie selection task (and, perhaps, even for parliamentary elections). In our experiments, α_k -PAV chose committees distributed fairly uniformly in the central areas, ignoring candidates with extreme opinions.

Shortlisting. Here our guiding principle is that the committee should consist of similar candidates (i.e., located close to each other). For this criterion, k -Borda is our rule of choice. In all of the experiments it consistently chose candidates located in the center, close to each other. Table 1 indicates that k -Borda has high variance of the number of win-

ners per quadrant. We believe that this is caused not by any faults of the rule itself, but by a fairly natural statistical property of our distributions. Since k -Borda selects 20 candidates from the center, due to random perturbations, sometimes the central candidates are not distributed over the quadrants in a perfectly balanced way, and our variance-based measure does not take into account the candidates' centrality.

The Strange Case of Bloc. In the situation where k candidates are to be selected (e.g., to a city council), it is quite common to ask the voters to come up with k names (ranked or non-ranked). Bloc, in particular, is quite a popular rule. Our histograms show that Bloc is very sensitive to the edge and corner effects (the pattern is similar to that for β -CC, but the effects are much stronger). Worse yet, Table 1 shows very high variance of the number of winners in each quarter and, indeed, the example elections for Bloc in Figure 3 show very asymmetric placements of the winners. These two arguments by themselves make Bloc a questionable voting rule.

Bloc is also the only rule in our collection that shows the following *inversion* effect: For the Gaussian distribution, the frequencies of the cells near the center (i.e., near the mean of the Gaussian distribution) are lower than the frequencies of the cells in the ring surrounding it. This is a very counter-intuitive and unexpected phenomenon: The most popular views in the society are represented less frequently than the not-so-popular ones. We believe that the mechanism behind this effect is similar to that behind the edge/corner effect: Even though the center has the highest density of the voters, it also has the highest density of the candidates, who therefore “steal points away” from each other. As a consequence, the slightly less popular candidates in the ring get enough support (both from some of the voters in the center and from those on the ring and beyond) to be elected.³

Robustness of the Results

So far we have considered elections with $m = 200$ candidates, $m = 200$ voters, and committee size $k = 20$ only. Thus it is natural to wonder if our conclusions remain valid as we vary these parameters.

Except for STV and β -Monroe, all our rules belong to the class of committee scoring rules (Elkind et al. 2014; 2017; Faliszewski et al. 2016), i.e., they define a per-voter score of each possible committee and select committees for which the sums of these scores are the highest. In consequence, the results for these rules should not change significantly with the number of voters (unless this number becomes very small). Since STV and β -Monroe are similar in spirit to committee scoring rules (indeed, STV is similar to SNTV and β -Monroe is very closely related to β -CC), the results for them should be similarly robust.

We also do not expect strong qualitative differences in our results for different numbers of candidates or different committee sizes (again, except for very small values). Nonetheless, we do observe quantitative differences.

³Indeed, this can be seen as a type of approximate cloning (see the discussion in the papers of Tideman (1987), Laffond et al. (1996), and Elkind et al. (2011)).




Figure 4: Histograms for our rules under the disc distribution, for committee sizes 10, 20, and 30. For α_k -PAV ($k \in \{10, 30\}$) and Monroe ($k = 30$) we computed only 5000 elections. Due to technical issues, for β -Monroe with $k = 10$ we computed only about 500 elections.

In Figure 4 we present histograms for our rules with respect to the disc distribution, for committee sizes 10, 20, and 30 (the histogram for committee size 20 is the same as in Figure 3; we repeat it for the sake of comparison). We note that the results for SNTV and STV are nearly the same irrespective of the committee size.⁴ The results for Bloc, α_k -PAV, and k -Borda also look very similar, and the differences are only in the radii of the discs/rings generated by these rules (this is especially natural for k -Borda; as we choose more and more of the centrally located candidates,

⁴For $k = 30$, the quota for STV is $q = \lfloor \frac{200}{31} \rfloor + 1 = 7$. Thus, in the first 28 stages we remove 196 voters, so the 29th candidate is chosen by 4 voters and the 30th candidate is selected randomly.

they form a larger and larger disc). The results for β -CC and β -Monroe for different committee sizes also look similar, but for $k = 10$ (especially for the case of β -CC) the artifacts in the histograms become much more visible (e.g., for $k = 10$ and β -CC, there are two very clearly visible consecutive rings). This indicates that our observations about β -CC and β -Monroe do not necessarily carry over to the case of very small committees.

Conclusions

Our results lead to several interesting observations. Foremost, within the framework of our study STV stands out as an exceptionally good rule for parliamentary elections. On the other hand, the Monroe rule, which is also an appealing rule for this application, did not do quite as well. We also found that the Monroe and Chamberlin–Courant rules may have (somewhat surprising) implicit biases against some candidates. Further, we discovered that in our experiments α_k -PAV tends to ignore extremist candidates and fairly uniformly covers central areas (this seems quite related to the results of Aziz et al. (2015a; 2017) on justified representation). We confirmed that k -Borda has good properties as a shortlisting rule and provided strong arguments against the Bloc rule.

Acknowledgments. Edith Elkind and Piotr Skowron were supported by the ERC grant 639945 (ACCORD), Piotr Faliszewski was supported by the NCN grant 2016/21/B/ST6/01509, Arkadii Slinko was supported in part by the Marsden Fund 3706352 of The Royal Society of New Zealand, and Nimrod Talmon was supported by a postdoctoral fellowship from I-CORE ALGO. Jean-François Laslier thanks the ANR project ANR13-BSH1-0010 DynaMITE.

References

- 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 Multi-agent Systems*, 107–115.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare*. To appear.
- Barberà, S., and Coelho, D. 2008. How to choose a non-controversial list with k names. *Social Choice and Welfare* 31(1):79–96.
- Betzler, N.; Slinko, A.; and Uhlmann, J. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47:475–519.
- 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. 1992. An axiomatic characterization of Borda’s k -choice function. *Social Choice and Welfare* 9(4):337–343.
- Diss, M., and Doghmi, A. 2016. Multi-winner scoring election methods: Condorcet consistency and paradoxes. Technical Report WP 1613, GATE Lyon Saint-Étienne.
- 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.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare*. To appear.
- Elkind, E.; Faliszewski, P.; and Slinko, A. 2011. Cloning in elections: Finding the possible winners. *Journal of Artificial Intelligence Research* 42:529–573.
- Enelow, J. M., and Hinich, M. J. 1984. *The spatial theory of voting: An introduction*. CUP Archive.
- Enelow, J. M., and Hinich, M. J. 1990. *Advances in the spatial theory of voting*. Cambridge University Press.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2016. Committee scoring rules: Axiomatic classification and hierarchy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, 250–256.
- Felsenthal, D., and Maoz, Z. 1992. Normative properties of four single-stage multi-winner electoral procedures. *Behavioral Science* 37:109–127.
- Kilgour, M. 2010. Approval balloting for multi-winner elections. In *Handbook on Approval Voting*. Springer. Chapter 6.
- Laffond, G.; Laine, J.; and Laslier, J. 1996. Composition consistent tournament solutions and social choice functions. *Social Choice and Welfare* 13(1):75–93.
- 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.
- Lu, T., and Boutilier, C. 2015. Value-directed compression of large-scale assignment problems. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 1182–1190.
- Monroe, B. 1995. Fully proportional representation. *American Political Science Review* 89(4):925–940.
- Procaccia, A.; Rosenschein, J.; and Zohar, A. 2008. On the complexity of achieving proportional representation. *Social Choice and Welfare* 30(3):353–362.
- Skowron, P.; Faliszewski, P.; and Lang, J. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241:191–216.
- Skowron, P.; Faliszewski, P.; and Slinko, A. 2015. Achieving fully proportional representation: Approximability result. *Artificial Intelligence* 222:67–103.
- Tideman, T. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare* 4(3):185–206.