Between Proportionality and Diversity: Balancing District Sizes under the Chamberlin–Courant Rule

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

Nimrod Talmon
Ben-Gurion University
Be'er Sheva, Israel
talmonn@bgu.ac.il

ABSTRACT

The Monroe and Chamberlin–Courant (CC) multiwinner rules proceed by partitioning the voters into virtual districts and assigning a unique committee member to each district, so that the voters are as satisfied with the assignment as possible. The difference between Monroe and CC is that the former creates equal-sized districts, while the latter has no constraints. We generalize these rules by requiring that the largest district can be at most X times larger than the smallest one (where X is a parameter). We show that our new rules inherit worst-case computational properties from their ancestors; evaluate the rules experimentally (in particular, we provide their visualizations, analyze actual district sizes and voter satisfaction); and analyze their approximability.

KEYWORDS

multiwinner elections; proportionality; diversity; Chamberlin–Courant; Monroe; algorithms; simulations

ACM Reference Format:

Piotr Faliszewski and Nimrod Talmon. 2018. Between Proportionality and Diversity: Balancing District Sizes under the Chamberlin–Courant Rule. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

1 INTRODUCTION

In the setting of multiwinner voting, we are given a group of candidates, a group of voters (or agents), and the task is to select a committee of k candidates that both reflects the voters' preferences and matches the application at hand [1, 14, 28, 34]. For example, if we are to choose some sort of a governing body—such as a parliament—then we should make sure that the committee represents the views of the voters proportionally, but if we want to select a set of movies for passengers on a long-distance flight, then we should select a "committee of movies" that is as diverse as possible, so that every passenger (i.e., every voter) can find something appealing. Indeed, each application may require committees with different properties, chosen according to different principles (see, e.g., the survey of Faliszewski et al. [19] for an overview). In this paper we consider applications requiring committees that achieve a given level of compromise between proportional representation and diversity; specifically, we study the complexity of computing such compromise committees, and we analyze the nature of these committees experimentally. We model the notions of "a diverse committee" and "a proportional committee" through the Chamberlin–Courant [10] (CC) and Monroe [30] voting rules, respectively. Given a set of m candidates, a set of n voters, and the target committee size k, these rules proceed as follows: They partition the voters into k *virtual districts* and associate each district with a representative candidate, so that the voters are as satisfied with their representatives as possible (see Section 2 for formal definitions). The difference between these rules is that Monroe requires the districts to be of equal sizes (give or take a voter), whereas CC puts no such constraints. In consequence, under the Monroe rule each winning candidate represents (almost) the same number of voters, and under CC different candidates may be associated with varying numbers of the voters.

Example 1.1. Consider the following society that wishes to select a parliament of size $k = 5$. There are five parties, A, B, C, D, and E, supported by, respectively, 72%, 11%, 10%, 6%, and 1% of the voters (i.e., 72% of the voters most prefer members of Party A, 11% most prefer members of Party B, and so on. (To be more formal, we would say that each party is a set of candidates and the voters who support a given party rank all its candidates on top, in a given, fixed order.) Under the Monroe rule, we would partition the voters into five virtual districts, each with 20% of the voters, so that, e.g., the first three would contain supporters of Party A only, the next one would contain the remaining 12% of the supporters of Party A and 8% of the voters chosen among the supporters of parties C, D, and E, and the final district would contain the 11% of the supporters of Party B and all the remaining voters. In consequence, members of Party A would win in the first four districts, and a member of Party B would win in the last one.

Example 1.2. Consider an airline that wishes to choose $k = 5$ movies to provide on its long-distance flights. The preferences of the voters are the same as above, except that we replace members of Party A with Hollywood blockbuster movies, members of Party B with family movies, members of Party C with documentaries, members of Party D with artistic movies, and members of Party E with movies on NP-completeness (i.e., we consider the same election as above, but with renamed candidates to illustrate a diversity-oriented setting). The CC rule would create a virtual district for each movie type, consisting exactly of the people who most prefer movies of this type, and would select one movie in each district. Thus the largest district would be 72 times larger than the smallest one.

The results in the above examples are somewhat extreme. In the first one, the supporters of Party C would be disappointed to not be represented, even though there are almost as many of them as supporters of Party B,[1] and 17% of the voters are assigned to

[1] Of course, to some extent this is because our committee is small. This explanation, however, would not prevent the voters from being disappointed.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

representatives that they do not support. In the second example, people who most enjoy Hollywood movies might complain that even though they form an overwhelming majority, they have only one movie to choose from, just like the 1% of the society that enjoys movies on NP-completeness. We address such complaints by introducing a parameterized family of rules where each rule in the family achieves a given level of compromise between the Monroe and Chamberlin–Courant rule: For each $X \geq 1$, we define the X-BalancedCC rule similarly to the CC rule, except that we require that the largest virtual district can be at most X times larger than the smallest one (we refer to X as the *balancedness ratio*); 1-BalancedCC is the Monroe rule and ∞-BalancedCC is the CC rule.

Example 1.3. If we used 2-BalancedCC rule for the election from the first example, then we would, e.g., obtain three virtual districts, each containing 24% of the supporters of Party A, one district with the 11% voters supporting Party B and the 1% of the supporters of Party E, and one district with the 10% voters supporting Party C and the 6% of the voters supporting Party D. Thus, members of Party A would win in three districts, and members of Parties B and C would win in one district each. (In terms of the movie example, we would get three Hollywood blockbusters, one family movie, and one documentary.) The ratio between the size of the largest and smallest district would be $^{24\%}/_{12\%} = 2$, and only 7% of the voters would be represented by candidates they do not like.

Notice that, in the context of parliamentary elections the X-BalancedCC rules implement the idea of *degressive proportionality*, where parties with smaller support obtain more seats than would follow from the proportion of their support (see, e.g., the work of Koriyama et al. [25]), and in the context of movie selection, the rules implement the idea of *support-sensitive diversity*, where we want our committee to be diverse, but we are unwilling to select candidates supported by very few voters.

We can also use X-BalancedCC rules in the context of *resource allocation with soft constraints*. Consider a university department that is required to send all its professors to courses on creative teaching. There are many courses to choose from, but—due to limited budget—the university will pay only for k of them, and each professor needs to attend one course. The department can gather the professors' preferences and then run a multiwinner voting rule to choose which courses to run. If there is a strict requirement that each course should be attended by the same number of people (e.g., because the people running the courses require such balance), then the university should use the Monroe rule. However, if it is acceptable that some courses would have a bit more attendees than others, then some X-BalancedCC rule (for some X not much larger than 1) would be a more natural choice.[2] (It would also be natural to put upper and lower bounds on the numbers of people in each selected course, e.g., to model that classrooms have bounded capacity and that we do not want too few students in a classroom. On the other hand, our X-balancedness requirement is, in spirit, closer to modeling fairness to the course teachers, ensuring that all of them have quite similar numbers of professors to teach.)

[2]This application is interesting also because the assignment of the voters (the professors) to the committee members (the courses) is public. This is not the case for parliamentary elections (due to voter anonymity) nor for the movie-selection example (where the election might also be held among a small fraction of passengers only).

Our Contribution. We study computational properties of our rules, analyze the trade-off between the balancedness ratio and the satisfaction of the voters, and perform experimental evaluation:

- **Computational Properties.** We explore the complexity landscape of our rules (including parameterized complexity) and the ability to solve them using integer linear programming.
- **Experimental Evaluation.** We provide a visual comparison of the X-BalancedCC rules (in the style of Elkind et al. [13]), we analyze the sizes of the virtual districts computed under the X-BalancedCC rules (and, in particular, under CC), and we show the relation between the voter (dis)satisfaction with their representatives and the balancedness ratio. We consider two families of election models, the Polya-Eggenberger urn model and the 2D Euclidean model.
- **Approximability and Voter Satisfaction.** We show initial results regarding approximability of our rules (under favorable conditions, our algorithms may achieve roughly $1 - {}^{1}/_{X}$ approximation ratios for balancedness ratio X).

There is an interesting interplay between our theoretical and experimental results. For example, the theoretical analysis suggests that the sizes of the virtual districts (when considered from the largest to the smallest) may decrease exponentially, and our experiments confirm that this happens in practice, but interestingly, show other types of behavior as well (this is interesting as it may lead to better approximation algorithms). Further, by contrasting our theoretical and experimental results, we obtain a better understanding of Skowron et al.'s [35] approximation algorithm for the CC rule.

Related Work. There is quite an extensive literature on both the Monroe and the Chamberlin–Courant rule. From a computational point of view, it is known that winner determination for these rules is NP-hard [28, 32], but that this problem is in FPT when parameterized either by the number of candidates or by the number of voters [6]. Further, there are efficient approximation algorithms [28, 35] and heuristics [20]. Indeed, our analysis of the voters' satisfaction in Section 5 is largely inspired by the analysis of Algorithm P of Skowron et al. [35]. Finally, there are papers on the complexity of the Monroe and CC rules under various domains of restricted
Monroe and CC rules were first studied by their inventors [10, 30] and then, e.g., by Elkind et al. [14] and Aziz et al. [1]. In particular, Elkind et al. [14] explicitly argued that the Monroe and CC rules have different applications, with Monroe better suited for proportional representation and CC better suited for selecting diverse committees (see also the survey of Faliszewski et al. [19]).

We generalize the Monroe and CC rules by requiring various degrees of balancedness from their virtual districts. Other modifications of these rules include, e.g., their egalitarian variants, studied by Betzler et al. [6], and the generalizations of the CC rule due to Skowron et al. [34] and Elkind and Ismaili [15]. The rules of Skowron et al. [34] take into account that each voter may feel represented by more than a single candidate, and the rules of Elkind and Ismaili [15] form a spectrum of rules between the classic CC rule and its egalitarian variant. In a similar vein, Faliszewski et al. [18] show several spectra of rules between CC and the k-Borda rule [12] (while we view our work as providing a spectrum of rules between proportional representation and diversity, their work can be seen as providing spectra between diverse committees and those consisting of individually excellent ones). Lackner and Skowron [26] suggest similar spectra of rules in the approval-based setting.

Talmon [37] considered the Monroe and CC rules in the context of an underlying social network of the voters (since the group activity selection problem of Darmann et al. [11] can also be seen as a very general extension of the Monroe rule, the works of Igarashi et al. [22, 23] might also be understood in this way).

Finally, as we are interested in the sizes of the virtual districts used by the Monroe and CC rules, our work relates to the extensive studies of districting (thus, to some extent, also to Gerrymandering); we specifically mention the papers of Bachrach et al. [3] and Lewenberg et al. [27], but emphasize that these papers consider geographic districts while we consider virtual districts.

2 PRELIMINARIES

For a positive integer t, we write $[t]$ to mean the set $\{1, \dots, t\}$. An election $E = (C, V)$ consists of a set of candidates $C = \{c_1, \dots, c_m\}$ and a collection of voters $V = (v_1, \dots, v_n)$, where each voter v_i has *preference order* $\succ_{v_i}$, ranking all the candidates from the most to the least desired one (i.e., we assume the ordinal model of elections; committee elections are also often studied in the approval model, where each voter specifies a set of acceptable candidates [1, 2, 8, 24]). We often refer to the preference orders as the votes.

For a voter v and candidate c, by $\mathrm{pos}_v(c)$ we mean the position of c in v's ranking (the most preferred candidate has position 1, the next one has position 2, and so on). A (single winner) scoring function (for the case of m candidates) $\gamma_m : [m] \to \mathbb{R}$ associates each possible position in a vote with a score. For example, the Borda scoring function is defined as $\beta_m(i) = m - i$ and the Plurality scoring function associates score 1 with the top position and score 0 with all the other positions. The γ-score of a candidate c in election (C, V) is defined as $\sum_{v \in V} \gamma(\mathrm{pos}_v(c))$.

Given an election $E = (C, V)$ and a positive integer k, $k \le |C|$, a multiwinner voting rule $\mathcal{R}$ outputs a set $\mathcal{R}(E, k)$ of size-k committees (i.e., size-k subsets of C) that tie as winners of this election. For example, the single non-transferable vote rule (the SNTV rule) outputs committees of candidates with the highest Plurality scores and the k-Borda rule outputs committees of candidates with the highest Borda scores. The Monroe and Chamberlin–Courant rules rely on the notions of assignment functions and voter satisfaction.

Assignment Functions and Voter Satisfaction. Let us fix an election $E = (C, V)$ and committee size k. A k-CC-assignment function Φ is a function $\Phi\colon V \to C$ that associates each voter with one of at most k candidates. We say that an assignment function Φ induces committee W (of size k) if for each voter v it holds that $\Phi(v)$ belongs to W. Given a voter v, we refer to the candidate $\Phi(v)$ as the representative of v. For a candidate c, we refer to the set of voters represented by c (i.e., to the set $\Phi^{-1}(c)$) as his or her *virtual district.* A k-Monroe-assignment function is a k-CC-assignment function where each candidate c either does not represent any voters or has a virtual district of size $|\Phi^{-1}(c)| \in \{\lfloor n/k \rfloor, \lceil n/k \rceil\}$. When speaking about assignment functions, we drop the k-CC- and k-Monroe-prefixes when they are clear from the context.

We define the satisfaction of a voter with his or her representative to be the Borda score that the representative receives from the voter. The total voter satisfaction with Φ is defined as the sum of the satisfactions of the particular voters:

$$\mathrm{sat}(\Phi) = \textstyle\sum_{v \in V} \beta_m\big(\mathrm{pos}_v(\Phi(v))\big).$$

The choice of the Borda function is arbitrary and one could use any other scoring function instead. We take the Borda-based approach because it was used in the original definitions of Monroe and CC [10, 30] and it has a very natural interpretation (the Borda-based satisfaction of the voters is proportional to the average number of candidates that a voter ranks below his or her representative).

Monroe and Chamberlin–Courant. Given the above setup, we are ready to define the Monroe and CC rules.

Definition 2.1. Let $E = (C, V)$ be an election and let $k \in [|C|]$ be the desired committee size. For each size-k committee S, the CC-score of S is the highest voter satisfaction achievable by a k-CC-assignment function that induces S (i.e., the CC score of S is $\max\{\mathrm{sat}(\Phi) \mid \Phi \text{ is a } k\text{-CC-assignment that induces } S\}$). The CC rule outputs all the size-k committees with the highest CC-score.

The Monroe rule is defined analogously, except that we consider k-Monroe assignments instead of k-CC-assignments.

Example 2.2. Consider an election $E = (C, V)$ with six candidates, $C = \{a, b, c, d, e, f\}$, and the following six votes:

$$\begin{aligned} &v_1 : a \succ e \succ d \succ c \succ b \succ f, && v_2 : a \succ f \succ d \succ e \succ b \succ c, \\ &v_3 : a \succ d \succ e \succ c \succ b \succ f, && v_4 : a \succ f \succ d \succ e \succ b \succ c, \\ &v_5 : b \succ e \succ d \succ a \succ c \succ f, && v_6 : c \succ d \succ e \succ a \succ b \succ f. \end{aligned}$$

We seek a committee of size $k = 3$. Under CC, the unique winning committee is $\{a, b, c\}$, with $\{v_1, v_2, v_3, v_4\}$ as a's virtual district, and $\{v_5\}$ and $\{v_6\}$ as the districts of b and c, respectively. Under Monroe, the unique winning committee is $\{a, d, e\}$, with a's district set to $\{v_2, v_4\}$, d's district $\{v_3, v_6\}$, and e's district $\{v_1, v_5\}$. Under CC, the voter satisfaction is 30 (highest possible), and under Monroe it is 26.

***X*-BalancedCC.** To define our X-BalancedCC rules, we need the notion of an X-balanced k-CC-assignment.

Definition 2.3. Let $E = (C, V)$ be an election, $k \in [|C|]$ be the committee size, Φ be a k-CC-assignment function, and $X \ge 1$ be a rational number. We say that Φ is X-balanced if the following holds: (a) there are exactly k candidates, $a_1, \dots, a_k$, assigned to non-empty virtual districts (i.e., for each $i \in [k]$, $\Phi^{-1}(a_i) \neq \emptyset$); and (b) for each two candidates a_i and a_j associated with non-empty virtual districts, it holds that a_i's district is at most X times larger than a_j's district (i.e., $|\Phi^{-1}(a_i)| \le X \cdot |\Phi^{-1}(a_j)|$).

In other words, a k-CC-assignment is X balanced if it provides k nonempty virtual districts, of which the largest one is at most X times larger than the smallest one. For each $X \ge 1$, we define the X-BalancedCC rule in the same was as the CC rule, except that instead of considering all k-CC-assignment functions, we consider only the X-balanced ones. By a slight abuse of notation, we say that every k-CC-assignment function is ∞-balanced (and for ∞-balancedness we drop the requirement that there are k-nonempty districts). In consequence, 1-BalancedCC is equivalent to the Monroe rule (for the cases where the committee size divides the number of voters evenly), and ∞-BalancedCC is equivalent to the CC rule.

Example 2.4. Consider the election from Example 2.2. Under 3-BalancedCC, there are two winning committees, $\{a, b, d\}$ and $\{a, c, e\}$ (which, in this case, can be seen as intermediate between the CC and Monroe committees). The virtual districts for the former committee are $\{v_1, v_2, v_4\}$ for a, $\{v_5\}$ for b, and $\{v_3, v_6\}$ for d. For the latter committee they are $\{v_2, v_3, v_4\}$ for a, $\{v_6\}$ for c, and $\{v_1, v_5\}$ for e. Indeed, the largest district is three times larger than the smallest one. The voter satisfaction is 28 for both committees.

3 COMPLEXITY AND EXACT ALGORITHMS

We begin our discussion of the X-BalancedCC rules with an overview of their computational complexity. We inherit most of the results from Monroe and CC, but sometimes care is needed.

Observation 3.1. *For each rational number X, $X \geq 1$, the problem of deciding if there exists a size-k committee and a X-balanced k-CC-assignment function for it with at least a given voter satisfaction (for a given election and committee size k) is* NP*-complete. This problem is in* FPT *for both the parameterization by the number of voters and the parameterization by the number of candidates.*

Proof sketch. NP-hardness follows from the proofs already provided for the case of Monroe and CC [28, 32, 35]. FPT algorithms for Monroe and CC were provided by Betzler et al. [6] and proceed by appropriate brute-force search, which can be adapted to the case of X-BalancedCC rules (see also Observation 3.2 below). □

The FPT algorithms above require the ability to compute optimal assignments for given committees. We use Betzler et al.'s [6] solution of this problem, adapted to the case of X-BalancedCC rules.

Observation 3.2. *There is a polynomial-time algorithm that, given an election $E = (C, V)$, a committee W of size k, and a number X, $X \geq 1$, computes an optimal X-balanced k-CC-assignment for W.*

Proof. Given an election with n voters and a committee of size k, the algorithm of Betzler et al. [6] finds an optimal assignment such that each committee member represents at least $s = \lfloor n/k \rfloor$ and at most $\ell = \lceil n/k \rceil$ voters. The algorithm works for all values of s and ℓ; we try each pair $s \in [|V|/k]$ and $\ell = \lfloor X \cdot s \rfloor$ and output an assignment with the highest voter satisfaction. □

Remark 1. *The reader may wonder if, in the proof of Observation 3.2, we really need to try more than one value of s. Indeed, it seems that we do. For example, consider an election with 18 voters, committee size 5, and $X = 2$. Then, for $s = 2$ we may find an assignment with district sizes $(2, 4, 4, 4, 4)$, but for $s = 3$ we may find one with sizes $(3, 3, 3, 3, 6)$.*

In practice it is useful to be able to compute a winning committee by solving an appropriate integer linear program (ILP). We provide such a program, inspired by those for the cases of Monroe and CC, by Potthoff and Brams [31] and Lu and Boutlier [28].

Observation 3.3. *For each rational X, $X \geq 1$, the problem of computing a winning X-BalancedCC committee can be expressed as an integer linear program.*

Proof. Let $E = (C, V)$ be the input election, where $C = \{c_1, \ldots, c_m\}$ and $V = (v_1, \ldots, v_n)$, let k be the desired committee size, and let X be the required balancedness ratio. We form an integer linear program, whose goal is to identify a winning X-BalancedCC committee S. For each $j \in [m]$ we have a binary variable x_j and for each $(i, j) \in [n] \times [m]$ we have a binary variable $y_{i,j}$. The intention is that if $x_j = 1$ then c_j belongs to S and if $y_{i,j} = 1$ then c_j represents v_i. Our objective is to maximize $\sum_{i \in [n]} \sum_{j \in [m]} y_{i,j} \cdot \beta_m(\mathrm{pos}_{v_i}(c_j))$. We introduce the following constraints to ensure consistency between our variables.

(1) $\sum_{j \in [m]} c_j = k$, to ensure that the committee has k members.
(2) For each $(i, j) \in [n] \times [n]$ we have constraint $y_{i,j} \leq x_j$, to ensure that only committee members can be representatives.
(3) For each $i \in [n]$ we have a constraint $\sum_{j \in [m]} y_{i,j} = 1$, to ensure that each voter has exactly one representative.

We introduce two integer variables, *max-size* and *min-size*, and add the following constraints:

(4) For each $j \in [m]$ we add two constraints: *max-size* $\geq \sum_{i \in [n]} y_{i,j}$ and *min-size* $\leq \sum_{i \in [n]} y_{i,j} + (1 - x_j) \cdot n$, which ensure that *max-size* is at least the size of the largest district and *min-size* is at most the size of the smallest one.
(5) Finally, we add a constraint *max-size* $\leq X \cdot$ *min-size*, to ensure that the variables $y_{i,j}$ describe an X-balanced assignment.

We obtain the winning committee from variables x_j, and the assignment (if needed) from variables $y_{i,j}$. □

4 EXPERIMENTAL EVALUATION

In this section we present our main results, i.e., experimental analysis of the X-BalancedCC rules. We are interested in three issues. First, we visualize our spectrum of rules between Monroe and CC, using the 2D-histogram approach of Elkind et al. [13]. Second, we are interested in the sizes of the virtual districts created by our rules (and under CC). Third, we evaluate on which position the voters rank their representatives, depending on the balancedness ratio (and their preference distributions).

Experimental Setup. In all our experiments, we consider elections with $m = 100$ candidates, $n = 100$ voters, and committee size $k = 10$. While it certainly would be interesting to see results for elections with other parameters, we believe that doing so would not lead to qualitatively different result, but—with the given space restrictions—would force us to omit some results. These or similar parameters were already used in a number of papers [13, 18, 20].

We consider two types of distributions of voter preferences, the Polya-Eggenberger urn model [5] and the 2D Euclidean model. In the urn model, we assume that there is an urn that contains all $m!$ distinct preference orders over m candidates, and the process of generating votes is as follows. For each voter, we draw a preference order from the urn (this becomes the preference order of the voter) and we return the order to the urn, together with $\alpha m!$ additional copies, where α is a parameter of the distribution. The larger is the value of α, the more correlation there is among the generated preference orders (and, in particular, for $\alpha = 0$ there is no correlation and we obtain the impartial culture model, where each preference order is equally likely). Impartial culture and the urn model are among the most popular models for generating elections. They were used, e.g., by Skowron et al. [35] for the case of Monroe and CC, and in many other papers for other scenarios [16, 29, 38].

In the 2D Euclidean model, each candidate and each voter is associated with a point on a two-dimensional plane (referred to as the candidate's or the voter's ideal point). Each voter forms his or her preference order by sorting the candidates in the order of increasing Euclidean distances of their ideal points from the voter's ideal point. We generate the candidates' and the voters' ideal points by drawing them from appropriate distributions (described in Section 4.1). This model was recently used by Faliszewski et al. [17] and Elkind et al. [13], and is now attracting increased attention.

We compute the results of all our elections by invoking the CPLEX ILP solver for the program described in Observation 3.3.

4.1 Visualization

In our first set of experiments we visualize the results of X-BalancedCC elections generated using the 2D Euclidean model. We consider two types of elections. With the first type (Models A, B, C, and D below) we attempt to capture elections that have an appealing real-life interpretation. With the second one (Models U, V, and W below) we aim to provide elections that illustrate the differences between Monroe, CC, and other X-BalancedCC rules as much as possible. Below we describe how the candidates' and voters' ideal points are generated for each of our election models.

Models A, B, C, and D. In these elections the ideal candidates are always generated uniformly at random from a disc centered at point $(0, 0)$, with radius 3. The ideal points of the voters are generated so that most of them fall in this area. In Model A, they are generated according to a two-dimensional Gaussian distribution with center $(0, 0)$ and standard deviation 1. In Model B, half of the voters are generated as in Model A and half of them are generated according to a two-dimensional Gaussian distribution centered at $(0, 0)$, with standard deviation 0.3. In Model C, $1/5$ of the voters are generated from the same uniform disc as the candidates, while the remaining voters are generated from four Gaussian distributions centered at $(-1, 0)$, $(1, 0)$, $(0, -1)$, and $(0, 1)$, each with standard deviation 0.5 (we generate $1/5$ of the voters from each of the four Gaussians). In Model D, the voters' ideal points are drawn uniformly at random from the same disc as the ideal points of the candidates.

The fact that the ideal points of the candidates are drawn uniformly at random from the disc (which covers the area where almost all voters may appear) models the idea that whenever there are voters, a candidate eventually appears. As opposed to Elkind et al. [13], however, we use different distributions of candidates' and voters' ideal points. We do so for two reasons. First, it seems that whenever these distributions are the same, the results of Monroe and CC are very similar (this is confirmed by the histograms presented by Elkind et al. [13]). The second reason is that we capture the setting where candidates appear in areas where they can attract voters, but tend to avoid competition (i.e., we avoid areas with large concentration of candidates, competing for the same voters).

Models A and B capture societies with a single predominant opinion (in the center point $(0, 0)$) and differ with respect to number of voters that have ideal points close to the center. Model C captures a society with four main opinions (the centers of the four Gaussians), but where there is also a non-negligible number of voters that do not follow such fashions (the $1/5$th of the voters whose ideal points are distributed uniformly on the same disc as the candidates).

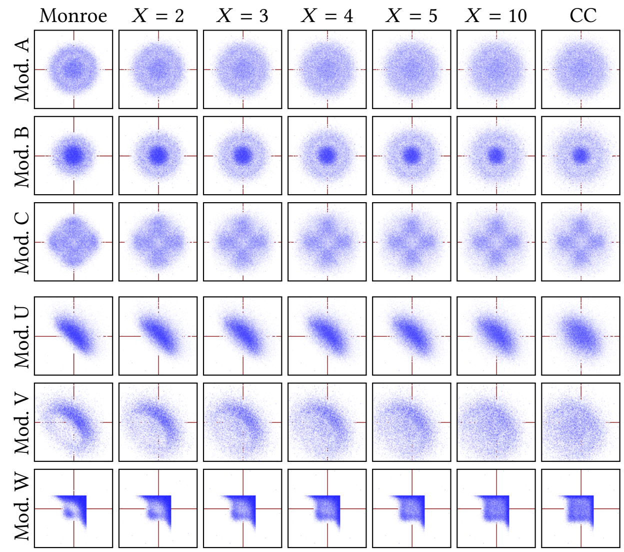

Figure 1: Histograms showing how frequently committee members from given areas are selected in models A-C, U-W.

Models U, V, and W. Here the areas with large concentrations of candidates and voters are only overlapping. In Models U and V, candidates are generated from a Gaussian distribution centered at $(-1, -1)$, with standard deviation 1. The voters are generated from Gaussian distribution centered at $(1, 1)$ with standard deviation 1 in Model U and standard deviation 2 in Model V. In Model W, the candidates are distributed uniformly on a $[-3, 1] \times [-3, 1]$ square and the voters are distributed uniformly on a $[-1, 3] \times [-1, 3]$ square (this distribution was also used by Elkind et al. [13]).

Histograms. For each of our 2D election models, we have generated 5000 elections and computed their X-BalancedCC results for $X \in \{1, 2, 3, 4, 5, 10, 100\}$ ($X = 1$ means using the Monroe rule and, in our case, $X = 100$ means using the CC rule). For each of the election models and each rule, in Figure 1 we present an histogram that indicates how frequently winners from a certain location of the plane are selected. Specifically, the histograms show the $[-3, 3] \times [-3, 3]$ square divided into 120×120 cells, where for each election model and rule we computed the number of times a committee member from a given cell was selected; the darker (the more blue) a given cell is, the more winners were selected from it (i.e., we use the same setup as [13]; see their work for the exact formula for translating numbers of winners in a cell into colors).

Analysis. There are three main types of behavior of our histograms. For Model A, we see only very little difference between the results computed for Monroe, CC, and the rules in between. This is not too surprising, as even under the CC rule, the average balancedness ratio is quite low (see Table 1). This is even more pronounced for Model D (the corresponding histograms for Monroe and CC, presented by Elkind et al. [13], are nearly identical; the same holds for their other distributions, except for Model W).

For Models B and W, there is a noticeable difference between the results for Monroe and CC, but already the results of X-BalancedCC for small X (between 2 and 4, say) are, at least visually, rather close

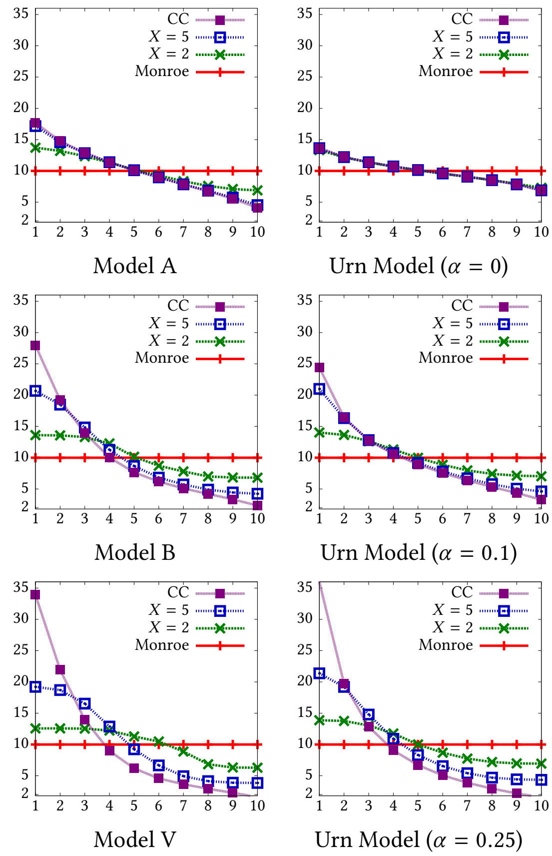

Figure 2: The average sizes of the virtual districts under X-BlancedCC rules for $X \in \{1, 2, 5\}$ and for CC. The values on the x-axis refer to the largest virtual district, the second largest district, and so on.

to those of CC. Interestingly, at least in the case of Model F this cannot be easily explained by low balancedness ratio for the case of CC. It seems that in these elections there seems to be some sort of (perhaps "soft") phase transition between requiring "strict balancedness" (as in Monroe) and even slightly relaxed balancedness.

Finally, for Models U, V, and W, we observe gradual change in the results as we increase the allowed balancedness ratio (for example, for Model C this is visible by gradually more and more visible peaks of the four Gaussians; for Model U it is visible by gradually increasing number of winners in the $[-3, 0] \times [-3, 0]$ quadrant). The results for Model C are particularly intriguing. Intuitively, one might expect that in the most proportional setting (Monroe), the peaks of the Gaussians should be most visible, but, in fact, we have most committee members in the area "surrounded" by the Gaussians and very few of them outside of this area. As we increase the allowed balancedness ratio, candidates close to peaks of the Gaussians still win frequently, but those in this "surrounded area" give more and more way for the candidates outside it (thus, increasing the diversity of the result). One possible explanation is that Monroe chooses candidates within a "smaller disc" than CC, because the peaks of the Gaussians account for nearly 80% of the voters, forcing Monroe to place 8 out of 10 committee members on the "smaller disc" (half of the 20% remaining voters are there as well; the other half is too spread to get representatives).

Table 1: Balancedness ratios for the case of the CC rule.

2D Euclidean models							urn model (different α values)			
A	B	C	D	U	V	W	0	0.10	0.25	0.5
4.76	14.10	4.98	2.64	41.21	24.72	21.57	2.01	9.18	29.88	47.66

4.2 District Sizes

In our second set of experiments we analyze the sizes of the virtual districts used by the X-BalancedCC rules, depending on their required balancedness ratio and the distribution of the voters' preference orders. We consider the six 2D models from the previous section and the urn model, with α values taken from the set $\{0, 0.05, 0.10, 0.25, 0.5, 1\}$.

Balancedness Ratios for CC. First, to get some idea about the nature of our committees, we computed the average balancedness ratios achieved under the CC rule for our elections. The averages are taken from computing the results of 500 elections and are reported in Table 1.

One of the most interesting results regards Model D (where both candidates' and voters' ideal points come from the same distribution, uniform on a disc). We computed the average balancedness ratio under CC for the other three settings of Elkind et al. [13] where candidates' and voters' ideal points come from the same distribution, and we always get average balancedness ratio close to 2.6. It is interesting on its own, but it also explains why they did not observe much difference between the results for Monroe and CC.

The results for the other 2D models are quite intuitive. For Models B, U, V, and W, the Monroe histograms in Figure 1 show areas with high concentrations of winners, suggesting that in these areas there are so many voters that Monroe has to put several committee members there, to satisfy all the voters. On the other hand, CC can place fewer committee members in these areas, as it is not limited by the district sizes. These intuitions are confirmed by the fact that under CC, Models B, U, V, and W have each fairly high balancedness ratios, whereas the other models have much lower ratios.

The results for the urn model are intuitive as well. The larger the α value (i.e., the more correlated the votes are), the larger the balancedness ratio. Indeed, elections generated with larger α values are more clustered, containing very large and very small groups of voters with identical preference orders, who are most satisfied when represented by the same candidate.

District Sizes. For each of our election models, and for each X value in $\{1, 2, 3, 4, 5, 10, 100\}$, we computed the average size of the largest virtual district, the second largest district, and so on, until the smallest district (each average is from 500 election instances). We present these results (for the most representative cases) in Figure 2.

For Models A, C, D, and for the urn model with $\alpha = 0$ (i.e., for the impartial culture model), our results are nearly the same Models C and D are omitted due to space restrictions but are available upon request): The average sizes of the districts change linearly as we

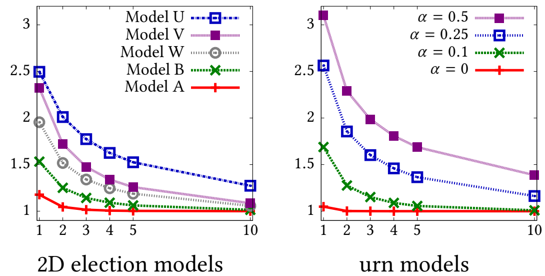

Figure 3: Rank-loss ratios under the 2D election models (on the left) and under the urn models (right). The results for Model C are essentially the same as for Model A.

move from the largest district toward the smallest one (see the top two plots in Figure 2 for examples).

For the remaining election models (i.e., Models D, U, V, and W, and the urn model with $\alpha \in \{0.1, 0.25, 0.5\}$) the results are qualitatively very similar to each other, but differed in terms of the particular values (see the middle and bottom rows of Figure 2). Under CC, in all these cases the sizes of the districts drop exponentially as we move from the largest one to the smallest one. Yet, for intermediate values of X (in particular for $X = 2$) we see that the rules try to cover sufficiently large districts for as many voters as possible, and then make the remaining districts as small as possible (this is most visible for Model V and $X = 2$; in this model, on the average, the 2-BalancedCC rule creates seven districts with size above 10, holding about 72 voters, whereas the CC rule creates three districts with size above 10, holding about 70 voters).

It is quite intriguing that for each of our 2D models there is a corresponding value of α for which the urn model for this α achieves a very similar distribution of district sizes (for example, for Model U, urn model with $\alpha = 0.5$ achieves nearly the same results; not shown in Figure 2 due to space restrictions). One possible explanation is that, perhaps, in real-life elections (or, at least, in typical models of elections) there are only three types of distributions of district sizes: the linear one (as on top of Figure 2), the exponential one (as in the middle and on the bottom of Figure 2 for CC), and the sigmoidal one (as in the middle and on the bottom of Figure 2 for $X \in \{2, 5\}$).

4.3 Ranks of the Representatives

In our final set of experiments, we compare the average positions of the representatives of the voters, depending on the election model and the balanced ratio X. For each election model, we generated 500 elections and computed committees under the CC rule and X-BalancedCC rules for $X \in \{1, 2, 3, 4, 5, 10\}$. Then, for each election model we calculated the average ranks of the voters' representatives (e.g., if under some election model half of the voters ranked their representatives on position 2 and half of the voters ranked their representatives on position 3, then the average representative rank would be 2.5). Finally, for each election model and value of X, we computed the ratio of the average representative rank under X-BalancedCC and under CC. We refer to this value as the *rank-loss* ratio. We present the values of the rank-loss ratios for our election models in Figure 3. Naturally, the rank-loss ratio is highest for $X = 1$, where the rules are most constrained.

For election models where the balancedness ratio is low by nature (i.e., Models A and C and urn model for $\alpha = 0$), the rank-loss ratio is already very low even for $X = 1$ and drops to nearly 1 very quickly with X. For the other models, the rank-loss ratio also seems to drop very quickly with X (and, as expected, the largest drop is between 1-balancedness and 2-balancedness).

Remark 2. *The rank-loss ratio is most interesting for the settings where the assignment of the committee members to the voters affects the voters directly. For example, this is the case in the course assignment example from the introduction. For the cases where voters do not learn the assignment (as in the movies example or in the parliament example), it may be more natural to measure the ratio between how highly, on the average, a voter ranks his or her most preferred committee members computed under X-BalancedCC and the rank of his or her representative computed under CC. Our preliminary simulation results suggest that these two types of ratios behave quite similarly, with the second type having lesser magnitude.*

5 APPROXIMATING VOTER SATISFACTION

So far, our computations of X-BalancedCC committees were exact and optimal. In our final set of results, we consider computing X-balanced assignments that may not be optimal, but which nonetheless achieve high voter satisfaction.

We start by recalling Algorithm P, an approximation algorithm for the CC rule due to Skowron et al. [35]. Let $E = (C, V)$ be an election, where $C = \{c_1, \ldots, c_m\}$ and $V = (v_1, \ldots, v_n)$, and let k be the desired committee size. The algorithm proceeds as follows:

(1) Set $\lambda = {}^{\mathrm{W}(k)}/_{k}$, where the numerator is the value of the Lambert's W function[3] (we will use other values $\lambda \in [0, 1]$).
(2) Repeat k iterations of: Find a candidate c ranked among the top λm positions most frequently and let V' contain the voters that rank c among top λm positions. Assign c to the voters in V' and remove these voters from consideration.
(3) Output the computed assignment (which is guaranteed to give voter satisfaction at least $nm(1 - {}^{2\mathrm{W}(k)}/_{k})$).[4]

In their analysis, Skowron et al. [35] show that, if after the i-th iteration there are $n_i = n(1 - \lambda)^i$ unassigned voters, then after the $(i+1)$-st one there will be at most $n_{i+1} = n(1-\lambda)^{i+1}$ of them. Since $n_0 = n$, this lets us derive an algorithm for X-BalancedCC rules.

For each $i \in [k]$, we define $d_i = n_{i-1} - n_i = n\lambda(1 - \lambda)^{i-1}$. We modify Algorithm P so that in the i-th iteration, $i \in [k]$, it assigns a representative to exactly d_i voters, so that these voters rank the representative among top $\lambda m + (i - 1)$ positions (this is possible due to the above-mentioned property shown by Skowron et al.; however, we cannot restrict our attention to top λm positions as in Algorithm P because representatives selected in the preceding iterations may still be present in the tops of the remaining voters). Since $d_1 \geq d_2 \geq \cdots \geq d_k$ holds, it follows that the balancedness ratio of the modified algorithm is at least:

$$B(\lambda) = {}^{d_1}/_{d_k} = {}^{n\lambda}/_{n\lambda(1-\lambda)^{k-1}} = {}^{1}/_{(1-\lambda)^{k-1}}.$$

[3]Its values are defined so that $x = \mathrm{W}(x) \cdot e^{\mathrm{W}(x)}$.

[4]At this point, there still are unassigned voters. Algorithm P assigns each of them to the candidate that this voter ranks highest, but we will handle this step differently.

Table 2: Approximation guarantees of the GreedyMonroe algorithm (GM), the algorithms described in this section (for various values of X), and Algorithm P. In parentheses we report the balancedness ratio obtained by the given algorithm (for Algorithm P we do not balance the unassigned voters, because it assigns them to the most preferred committee member, disregarding the balancedness).

	GM	$X = 1.5$	$X = 2$	$X = 3$	$X = 4$	$X = 10$	Alg. P
$m = n = 100$	0.687	0.658	0.685	0.710	0.710	0.716	0.726
$k = 10$	(1)	(1.44)	(1.87)	(2.57)	(2.57)	(4.20)	(6)
$m = n = 500$	0.808	0.704	0.748	0.787	0.802	0.815	0.813
$k = 20$	(1)	(1.48)	(1.95)	(2.94)	(3.66)	(6.80)	(9.33)
$m = n = 1000$	0.903	0.718	0.763	0.840	0.866	0.909	0.948
$k = 100$	(1)	(1.44)	(1.87)	(2.85)	(3.83)	(9.66)	(34)

(There are also n_k unassigned voters. We assign them to the virtual districts in the following way: As long as some unassigned voter remains, we put him or her in the currently smallest district. We refer to this process as *balancing the unassigned voters*; after this balancing, all the voters are assigned but for simplicity we will still refer to the voters to whom the balancing process was applied as *unassigned* and to the others as *assigned*.)

The voter satisfaction is at least:

$$\sum_{i=1}^{k}(m - m\lambda - (i-1))d_i \geq (m - m\lambda - k + 1)(n - n_k)$$
$$= mn(1 - \lambda - {}^{k-1}/_{m})(1 - (1 - \lambda)^k).$$

Since the highest possible voter satisfaction is $(m-1)n$ (when each voter ranks his or her representative on top), the approximation ratio of our modified algorithm is at least $(1 - \lambda - \frac{k-1}{m})(1 - (1-\lambda)^k)$.

It remains to note that if we set the λ value to be $1 - \sqrt[k-1]{1/X}$, then we obtain balancedness ratio $B(\lambda) \leq X$. With this λ, the achieved approximation ratio of our algorithm is:

$$\left(\sqrt[k-1]{1/X} - \tfrac{k-1}{m}\right)\left(1 - \left(\sqrt[k-1]{1/X}\right)^k\right) = \left(\sqrt[k-1]{1/X} - \tfrac{k-1}{m}\right)\left(1 - 1/X\left(\sqrt[k-1]{1/X}\right)\right).$$

The value $\sqrt[k-1]{1/X}$ approaches 1 quite quickly as k increases, so as long as k is not too small, but the number of candidates is large enough for $\frac{k-1}{m}$ to be close to 0, the approximation ratio is close to $1 - 1/X$. (Our analysis disregarded the fact that λm and several other values may not be integers, which means that our computation of the approximation ratio is slightly off; we ignore this as in the following discussion we use precise approximation guarantees computed by simulating the worst-case behavior of our algorithm).

The above result is both appealing and disappointing. The reason for the disappointment is that the GreedyMonroe algorithm of Skowron et al. [35], which computes 1-balanced assignments, has approximation ratio $\approx 1 - \frac{k-1}{2(m-1)} - \frac{H_k}{k}$ (where m is the number of candidates, k is the committee size, and H_k is the k-th harmonic number); e.g., for elections with 100 candidates, 100 voters, and committee size 10, its approximation guarantee is ≈ 0.687. Our result matches this guarantee only starting with $X \approx 4.1$.

This low performance of our algorithm is due to the unassigned voters; for example, for $X = 2$ we have $n_k = n(\sqrt[k-1]{0.5})^k \approx 0.5n$ of them. Yet, since we balance the unassigned voters, instead of using the λ values computed in the above analysis, we can use a value for which we still obtain an X-balanced assignment and which leads to the best approximation ratio (for a balancedness ration up to X). In Table 2 we show several approximation guarantees that we can achieve, depending on the election parameters, using GreedyMonroe, using the just-described algorithms (taking into account the process of balancing the unassigned voters), and using Algorithm P. To compute the approximation guarantees, we have simulated the worst-case performance of the algorithms (and, indeed, the ratios in Table 2 are true approximation guarantees that hold for all elections with given sizes).

The results in Table 2 show that to beat GreedyMonroe, we need relatively large balancedness ratios. Nonetheless, these ratios are notable smaller than those that Algorithm P may achieve. Nonetheless, we view our results as mostly negative: It may be a better idea to start from GreedyMonroe and not from Algorithm P to obtain good approximate solutions for X-BalancedCC rules.

Remark 3. *Our analysis explains a certain peculiar feature of Algorithm P. Elkind et al. [13] ran this algorithm on 2D Euclidean elections (e.g., on elections generated according to Model D) and noted that, even though its approximation ratio is good, its histograms are substantially different from that of CC.*[5] *Thus, instead of using* $\lambda = \mathrm{W}(k)/k$, *they run the algorithm for all values* $\lambda \in \{1/m, 2/m, \ldots, m/m\}$ *and take the committee with the highest satisfaction; they refer to this algorithm as RangingCC. The histograms for RangingCC look almost identically to those of CC, but Elkind et al. [13] could not explain why. Our analysis shows that RangingCC actually is trying to match the balancedness ratio implicit in the voters' preferences with appropriate* λ *value. E.g., Algorithm P for Model D computes a committee with average balancedness ratio* ≈ 7.65, *but RangingCC achieves ratio* ≈ 2.96, *much closer to the* 2.64 *ratio of the actual CC rule.*

6 CONCLUSIONS

We considered the X-BalancedCC spectrum of rules, which achieve a tradeoff between proportionality and diversity by generalizing CC and Monroe, through tweaking the sizes of the districts they create. We argued that our rules have many applications and that, computationally, they are as difficult as Monroe and CC. We illustrated their results, analyzed their virtual districts' sizes, and studied the satisfaction they give to the voters. We conclude by asking: What other natural spectra of rules of this type exist? Could one use, e.g., the Phragmen rule [7] as a basis for one?

We believe that (variants of) our rules could be fruitfully employed in the participatory budgeting setting [4, 9, 21, 33], where the candidates (the projects) have costs and we have to select a "committee" of them that does not exceed a given budget. Exploring issues of proportionality, diversity, and district sizes in participatory budgeting seems important and promising.

ACKNOWLEDGMENTS

Piotr Faliszewski was supported by the National Science Centre, Poland, under project 2016/21/B/ST6/01509.

[5] These results are in the full version of their paper, available on the authors' websites.

REFERENCES

[1] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[2] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*. 107–115.

[3] Y. Bachrach, O. Lev, Y. Lewenberg, and Y. Zick. 2016. Misrepresentation in District Voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*. 81–87.

[4] G. Benade, S. Nath, A. Procaccia, and N. Shah. 2017. Preference Elicitation For Participatory Budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 376–382.

[5] S. Berg and D. Lepelley. 1994. On Probability Models in Voting Theory. *Statistica Neerlandica* 48, 2 (1994), 133–146.

[6] N. Betzler, A. Slinko, and J. Uhlmann. 2013. On the Computation of Fully Proportional Representation. *Journal of Artificial Intelligence Research* 47 (2013), 475–519.

[7] M. Brill, R. Freeman, S. Janson, and M. Lackner. 2017. Phragmén's Voting Methods and Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. To appear.

[8] M. Brill, J. Laslier, and P. Skowron. 2017. Multiwinner Approval Rules as Apportionment Methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 414–420.

[9] Y. Cabannes. 2004. Participatory Budgeting: A Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.

[10] B. Chamberlin and P. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review* 77, 3 (1983), 718–733.

[11] A. Darmann, E. Elkind, S. Kurz, J. Lang, J. Schauer, and G. Woeginger. 2012. Group Activity Selection Problem. In *Proceedings of the 8th International Workshop On Internet And Network Economics (WINE-12)*. 156–169.

[12] B. Debord. 1992. An Axiomatic Characterization of Borda's k-Choice Function. *Social Choice and Welfare* 9, 4 (1992), 337–343.

[13] E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. 2017. What Do Multiwinner Voting Rules Do? An Experiment Over the Two-Dimensional Euclidean Domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 494–501.

[14] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.

[15] E. Elkind and A. Ismaili. 2015. OWA-Based Extensions of the Chamberlin–Courant Rule. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT-15)*. 486–502.

[16] G. Erdélyi, M. Fellows, J. Rothe, and L. Schend. 2015. Control Complexity in Bucklin and Fallback Voting: An Experimental Analysis. *J. Comput. System Sci.* 81, 4 (2015), 661–670.

[17] P. Faliszewski, J. Sawicki, R. Schaefer, and M. Smolka. 2016. Multiwinner Voting in Genetic Algorithms. *IEEE Intelligent Systems* 32, 1 (2016), 40–48.

[18] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2017. Multiwinner Rules on Paths From k-Borda to Chamberlin–Courant. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*. 192–198.

[19] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. 2017. Multiwinner Voting: A New Challenge for Social Choice Theory. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access Foundation.

[20] P. Faliszewski, A. Slinko, K. Stahl, and N. Talmon. 2016. Achieving Fully Proportional Representation by Clustering Voters. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-16)*. 296–304.

[21] A. Goel, A. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. 2016. Knapsack Voting for Participatory Budgeting. (2016). Manuscript.

[22] A. Igarashi, R. Bredereck, and E. Elkind. 2017. On Parameterized Complexity of Group Activity Selection Problems on Social Networks. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-17)*. 1575–1577.

[23] A. Igarashi, D. Peters, and E. Elkind. 2017. Group Activity Selection on Social Networks. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-17)*. 565–571.

[24] M. Kilgour. 2010. Approval Balloting for Multi-Winner Elections. In *Handbook on Approval Voting*, J. Laslier and R. Sanver (Eds.). Springer. Chapter 6.

[25] Y. Koriyama, J. Laslier, A. Macé, and R. Treibich. 2013. Optimal Apportionment. *Journal of Political Economy* 121, 3 (2013), 584–608.

[26] M. Lackner and P. Skowron. 2017. *Consistent Approval-Based Multi-Winner Rules*. Technical Report arXiv:1704.02453 [cs.GT].

[27] Y. Lewenberg, O. Lev, and J. Rosenschein. 2017. Divide and conquer: Using geographic manipulation to win district-based elections. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-17)*. 624–632.

[28] T. Lu and C. Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*. 280–286.

[29] J. McCabe-Dansted and A. Slinko. 2006. Exploratory Analysis of Similarities Between Social Choice Rules. *Group Decision and Negotiation* 15, 1 (2006), 77–107.

[30] B. Monroe. 1995. Fully proportional representation. *American Political Science Review* 89, 4 (1995), 925–940.

[31] R. Potthoff and S. Brams. 1998. Proportional representation: Broadening the Options. *Journal of Theoretical Politics* 10, 2 (1998), 147–178.

[32] A. Procaccia, J. Rosenschein, and A. Zohar. 2008. On the Complexity of Achieving Proportional Representation. *Social Choice and Welfare* 30, 3 (2008), 353–362.

[33] E. Shapiro and N. Talmon. 2017. *A Condorcet-Optimal Participatory Budgeting Algorithm*. Technical Report arXiv:1709.05839.

[34] P. Skowron, P. Faliszewski, and J. Lang. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241 (2016), 191–216.

[35] P. Skowron, P. Faliszewski, and A. Slinko. 2015. Achieving Fully Proportional Representation: Approximability Result. *Artificial Intelligence* 222 (2015), 67–103.

[36] P. Skowron, L. Yu, P. Faliszewski, and E. Elkind. 2015. The complexity of fully proportional representation for single-crossing electorates. *Theoretical Computer Science* 569 (2015), 43–57.

[37] N. Talmon. 2018. Structured Proportional Representation. *Theoretical Computer Science* 708 (2018), 58–74.

[38] T. Walsh. 2011. Where Are the Hard Manipulation Problems. *Journal of Artificial Intelligence Research* 42, 1 (2011), 1–29.