The Group Perspective on Fairness in Multi-Winner Voting

SUBMISSION 601

In computational social choice theory, many studies use the setting of spatial or metric voting, where both candidates and voters are embedded in a common metric space X representing positions on issues. The "cost" to a voter v of candidate c is their distance $d(v, c)$ in the space X, and it is presumed that voters asked to rank their preferences would do so in order of proximity. Many authors have searched for voting rules that tend to elect "optimal" (lowest-cost) winners, or where rules cost ratios have some guarantee in the form of an upper bound. Most work has focused on optimizing cost in scenarios with worst-case embeddings.

In this paper we study fairness for multi-winner voting rules in a metric setting, focusing our attention on a novel definition we call **group inefficiency** to evaluate representation for select groups of voters. We show that depending upon *who* fairness is measured for, the findings can change significantly and, that overall or worst-off definitions generally do not succeed in identifying the perspectives of the most important voter groups. Using simple metric settings, we evaluate this notion of fairness for many common multi-winner election mechanisms, some of which are currently being used in real elections and have yet to be considered under a similar lens by related work.

Manuscript submitted for review to the 26th ACM Conference on Economics & Computation (EC'25).

1 Motivation

Computational social choice theory is the study of group decisions, using tools from computer science and economics. Because it focuses on voting rules and their properties, it is often used to give insights into democratic mechanisms, such as those that elect political representatives. The motivation for the current project is to take up current problems of interest in computational social choice and show that re-framing fairness measurements to focus on *salient* groups of voters can cause major qualitative changes to the findings. In doing so we aim to bring understanding to representation properties for many commonly used voting rules, giving insight into which of them might be appropriate for use in real world settings.

Throughout the paper, we'll refer to voter groups or blocs as being "salient" if they have some fundamental importance for the analysis of fairness. In real election settings a group might be found to be salient because of the law—for instance, the U.S. Voting Rights Act of 1965 identifies groups with common racial, ethnic, or language features as being covered by a distinctive set of protections. A group might also be salient because of common voting patterns, such as voting in a cohesively similar way in a polarized electoral environment.

Although previous work has also considered the perspective of the bloc or group which is worst-off in terms of fairness, we note that this may leave room for arbitrary or non-salient choices of group which are less meaningful for measurement and comparison. Voting rules are not always designed or used to represent voters in the same way, and therefore the perspectives that we evaluate them from matter greatly. Furthermore, restricting our attention to specified groups opens the door for computational experiments which consider many popular voting rules that have yet to be analyzed under a similar framework.

In particular we focus upon a notion of fairness which we call *inefficiency* that is designed to evaluate the ability of multi-winner voting rules to proportionally represent their voters in a metric space setting. In our experiments, we use simple metric embeddings to compute inefficiency scores for both commonly used voting rules, as well more as recently developed rules from the metric voting literature. We show that from a perspective of overall fairness, or inefficiency for the entire voter population, the Borda and Plurality Veto voting rules do best compromise and elect candidates that are central to the entire population. In sharp contrast, our group perspective, which measures inefficiency for individual, cohesive voter blocs, sees Borda as one of the worst options and instead favors rules such as STV, Monroe, and Chamberlain-Courant to diversely represent the population.

2 Main definitions

2.1 Elections, profiles, and voting rules

Suppose we have a set $\mathcal{V}$ of N voters and a set C of m candidates. An *embedded election* is a metric space X with distance function d, equipped with an embedding map $\mathcal{M} : \mathcal{V} \cup C \to X$ that identifies the voters and candidates with points in the space. This gives us a way to measure how different two candidates are, and how close a voter is to a candidate. The associated *preference profile* $P = P(X, \mathcal{M})$ is a set of rankings $\succcurlyeq_v$ for each $v \in \mathcal{V}$, where for any voter v and pair of candidates c, c', we have $c \succcurlyeq_v c' \iff d(v, c) \leq d(v, c')$, meaning that voters prefer the candidates located closer to themselves. Importantly, the profile is strictly less information than the embedded election, in the sense that the profile can be uniquely derived from the embedding but not vice versa.

If $\mathcal{P}$ is the set of profiles of a given type, then $f : \mathcal{P} \to 2^C$ denotes a multi-winner election mechanism (otherwise known as a *voting rule* or *social choice function*). We will focus on rules designed to elect $|f(P)| = k$ winners.[1] One simple example is *single non-transferable vote* (SNTV), which focuses only upon voters' first-place selections in the profile, and elects the k candidates with the most first-place votes. Another simple example is *plurality block voting*, where we consider voters' top k preferences, awarding one point to a candidate for receiving a mention regardless of its relative position within the top k, and elect the candidates with the most points. Many more examples, both deterministic and randomized, are given below. The study of *metric distortion*, which has attracted great attention since its introduction circa 2017, evaluates the performance of a mechanism on an embedded election by considering the ratio of distances (or "costs," in economic terms) between its chosen winner(s) and those that are optimal in X. The performance of a mechanism on a profile can then be measured by the worst case over all embeddings consistent with that profile. The foundational work is surveyed in the next section.

2.2 Blocs, costs, and candidate assignments

Elections are called *polarized* when there are sufficiently large, disjoint groups of voters whose voting behavior and preferences are sharply different. In the metric setting, we can easily model this by dropping groups of voter points centered at different locations. We call subsets of voters $\mathcal{B} \subseteq \mathcal{V}$ voting *blocs*. For any given bloc $\mathcal{B}$ and set of candidates S, we then define the cost of S to $\mathcal{B}$ as the sum of distances between voters and candidates.

$$\text{cost}(\mathcal{B}, S) = \sum_{v \in \mathcal{B}} \sum_{c \in S} d(v, c) \tag{1}$$

This is in keeping with the dominant choice of objective function in the single-winner metric-voting literature, where a voter's cost for a candidate is given by distance and the societal cost is the sum of the distances from all voters to the winner.

Next, we also consider assigning subset of voters to a subset of preferred, representative candidates. Given a bloc $\mathcal{B} \subseteq \mathcal{V}$ and a candidate set $S \subseteq C$, their *candidate assignment* is a function $\text{Rep} : 2^{\mathcal{V}} \times 2^C \to 2^C$ satisfying $\text{Rep}(\mathcal{B}, S) \subseteq S$. This is thought of as identifying a designated subset of candidates from S that is best representative of the voters in bloc $\mathcal{B}$.

The choice of assignment function Rep may be varied, but in this paper we focus on a greedy candidate assignment keyed to the intuition of proportional representation that we denote by Rep_{prop}. We say that for mechanisms electing k winners, the number of winners proportionally "deserved" by bloc $\mathcal{B}$ has size $\lfloor k|\mathcal{B}|/n \rfloor$. Using this, we can assign to a bloc of voters the cost-minimizing set of candidates of that size. For $T \subseteq C$ we define

$$\text{Rep}_{\text{prop}}(\mathcal{B}, T) := \underset{\substack{S \subseteq T \\ |S| = \lfloor |T| \cdot |\mathcal{B}|/n \rfloor}}{\arg\min} \quad \text{cost}(\mathcal{B}, S).$$

For example, if a voter bloc makes up 30% of the electorate in a 4-winner election, then by assigning those voters to the winning set of candidates using Rep_{prop}, we'll find that the voters are assigned to their single most preferred winner, i.e., the winning candidate who minimized the summed cost to that bloc. Blocs smaller than 25% of voters are assigned to the empty set in a 4-winner scenario.

[1] More precisely, we consider rules that are guaranteed to elect k winners when candidates and voters are in general position. This allows for rules for which ties are possible, but resolve when the embedding is perturbed.

2.3 Group-centered fairness

Recall that $C, \mathcal{V}$ are the sets of candidates and voters respectively, and let $\mathcal{W} = f(X)$ be the winner set given an embedded election and a voting rule.

Definition 2.1. For a voting rule f, candidate assignment Rep, and voter bloc $\mathcal{B} \subseteq C$ in an embedded election X, define the *group inefficiency* by

$$\mathcal{I}(\mathcal{B}) = \mathcal{I}_{f,\mathrm{Rep},X}(\mathcal{B}) := \frac{\mathrm{cost}(\mathcal{B}, \mathrm{Rep}(\mathcal{B}, \mathcal{W}))}{\mathrm{cost}(\mathcal{B}, \mathrm{Rep}(\mathcal{B}, C))}.$$

We say f has *overall inefficiency* α if it has that level of group inefficiency for the undivided electorate $\mathcal{V}$, which means that $\mathcal{I}(\mathcal{V}) = \alpha$, or in other words the cost of the winners is α times the cost of the metrically optimal winner set. This is also the standard notion of distortion in computational social choice theory.

Additionally, another approach seen in computational social choice papers to study the *worst-bloc inefficiency* of a voting rule by bounding the group inefficiency for any possible voter bloc $\mathcal{B} \subseteq \mathcal{V}$. We will see below that this obscures fundamental differences in fairness that are apparent when considering salient groups.

2.4 Examples

In Figure 1 we display visual examples for the inefficiency definitions from the previous section, computed with a few commonly used multi-winner election mechanisms: Bloc, Borda, and STV. Importantly, our worst-bloc inefficiency is only an estimate since, finding the worst bloc would require searching though all 2^n voter subsets. Therefore, we display an estimate bloc which is found by greedily searching among blocs which are closer to some subset of preferred, non-elected candidates than they are to their designated representatives.

While worst-bloc inefficiency is an appealing notion, we often find that blocs with largest inefficiency scores can be arbitrarily positioned, meaning they may be unmeaningful and/or hard to describe. Depending on the election setting, this may then produce less meaningful fairness measurements. For example, our estimate worst-bloc for the STV rule is a combination of one group along with roughly half of the other. Although both groups appear well represented by the winning candidates, this odd combination of voters taken together is relatively dissatisfied since, in aggregate, they are closer in distance to a more central, less representative subset of candidates. Such a winning set, however, makes little sense if our goal is proportional representation for the two blocs individually. Therefore, to include this bloc in an evaluation of STV—a voting rule designed and used for purposes of diverse and proportional representation—could misrepresent or hide its most valuable qualities. In other words, a worst bloc analysis does not take into account the fact that voting rules may be used for specialized purposes, or that certain blocs of voters have a perspective that is worth singling out for legitimate reasons.

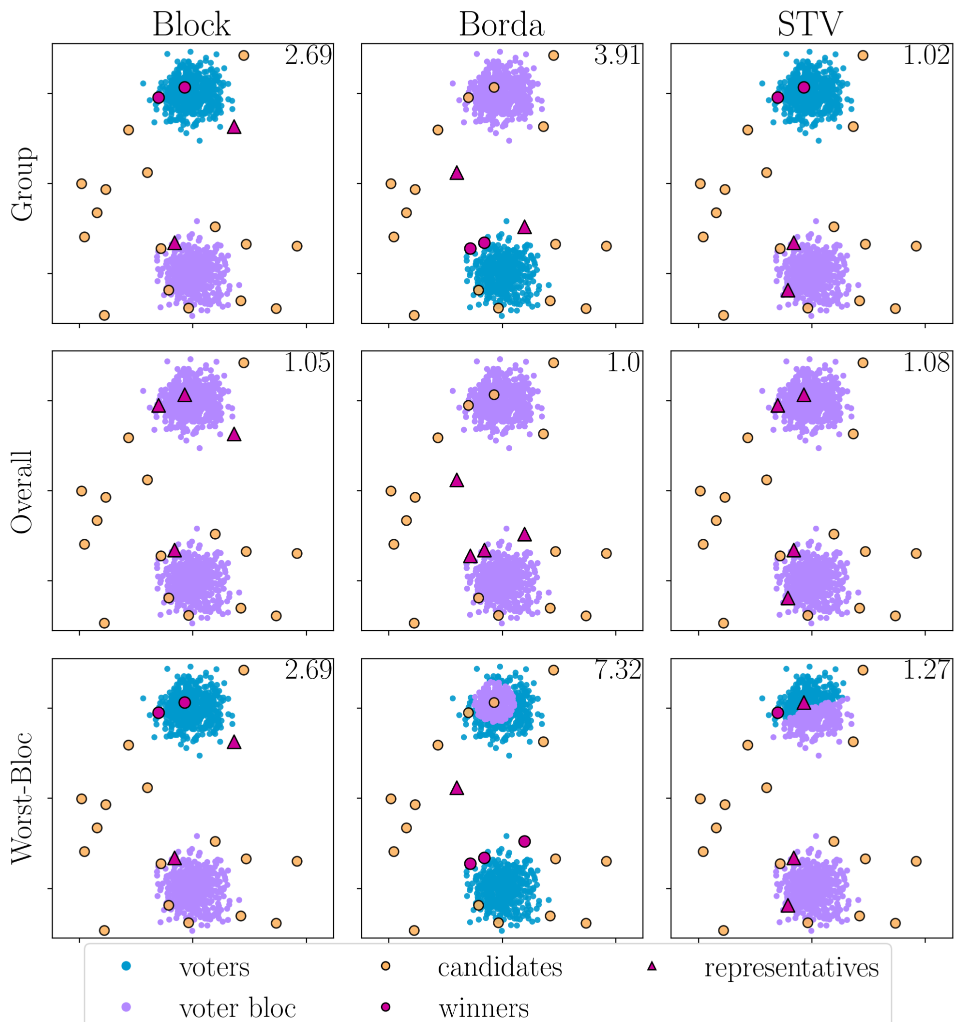

Fig. 1. Visualization for inefficiency measurements in an embedding with two cohesive voter groups (the top and bottom voter blocs) of five hundred voters each, twenty candidates uniformly distributed around them, and four winners. Results for the Bloc, Borda, and STV voting rules are displayed in turn. These blocs and representatives are used to compute group, overall, and (estimated) worst-bloc inefficiencies for each election mechanism. The computed inefficiency values are displayed in the top right corner for each example.

3 Related work

Our study branches off from a line of recent work that evaluates voting rules in metric settings. For single-winner voting rules, Anshelevich et al. [2018], Anshelevich and Postl [2017] introduced the metric framework and the objective of seeking low-distortion voting rules, proving bounds for both known and novel election mechanisms. This opened a vibrant line of work studying metric distortion of voting rules, which in several cases led to the introduction of new rules that are of significant independent interest [Charikar and Ramakrishnan, 2022, Charikar et al., 2024, Kempe, 2020, Kizilkaya and Kempe, 2022]. Given that distortion is computed numerically using distances between voters and candidates, this framework allowed authors to make more fine-grained numerical comparisons between different election mechanisms than had previously been considered with traditional axiomatic approaches.

The case of multi-winner elections is sometimes referred to as "committee voting," though we will avoid that terminology here because we are motivated by the case of political representation. Multi-winner voting rules have recently seen increased attention, with various authors studying their axiomatic properties [Elkind et al., 2017b, Faliszewski et al., 2017]. The study of proportionality axioms for multi-winner rules, in particular, is very active, and we refer readers to Brill and Peters [2023] for both a comprehensive review, as well as recent developments. However, despite being interesting and applicable, multi-winner elections have received much less attention in the metric voting literature. Results from our work are strongly inspired by Elkind et al. [2017a], where the authors use simple metric embeddings in order to produce intuitive visualizations, giving qualitative, geometric insight into how multi-winner election mechanisms represent voters' preferences. We build upon their work by defining fairness measurements and using them to evaluate the performance of several election mechanisms in various metric embeddings.

Other existing work on multi-winner voting in metric settings with similar notions of distortion, cost, and efficiency include the following.

3.1 Worst-bloc fairness

Goel et al. [2018] evaluate the "fairness" of a winner set for a given voter bloc size, b, by computing $\max_{\mathcal{B},|\mathcal{B}|=b} \text{cost}(\mathcal{B}, \mathcal{W})$, the sum of distances from the worst-off bloc to *all* winning candidates. Our work is distinct from theirs in the sense that we take on a notion of proportional representation. That is, we assign voter blocs to smaller subsets of representatives, and the winning set of candidates need not represent all voter blocs simultaneously.

3.2 Distortion with q-cost

Caragiannis et al. [2022] define the q-cost to voter v as their distance to their qth closest candidate among the winner set. That is, if $d_1 \leq d_2 \leq \cdots \leq d_k$ are the distances from v to the members of $\mathcal{W}$ listed in non-decreasing order, then the cost to that voter is d_q. The cost of an election is computed by summing the q-costs for the entire voter set $\mathcal{V}$.

When q is small, this measure rewards proportional outcomes, since each voter is satisfied by their few favorite winning candidates. However, as the authors show, optimizing for q-cost produces intriguing phase transitions (as q varies), and distortion becomes unbounded when $q \leq \frac{k}{3}$ for *any* voting rule. Although there are interesting results for larger q values, we note that increased values lead to outcomes that are less sensitive to proportionality for distinct, cohesive voter blocs. Since their notion of cost sums over the entire voter set, the extreme case of $q = k$, for example, rewards winner sets in which voters are forced to compromise. As noted by Kalayci et al. [2024], this can lead to homogeneous winner sets which lie directly near the center of the entire voter population.

Furthermore, while their results provide an interesting framework, their work does not study how specific voting rules actually perform within it.

3.3 γ-proportional representation

The very recent paper of Kalayci et al. [2024] is closest in spirit to our work. In both their work, as well as ours, sufficiently large blocs of voters are assigned a proportional share of representatives from the winner set, and cost ratios are computed. Their main theorem shows that a relatively new voting rule (Expanding Approvals) guarantees their definition of "proportional fairness" at a constant level.

The fundamental difference between their definition and ours is that (like the worst-bloc definition in Section 2.3 and the worst-bloc fairness notion from Goel et al. [2018]) they look for a theoretical worst-case inefficiency over all possible blocs and metric embeddings, while we focus on specified, salient blocs in experimental settings. Along with the previously mentioned issues for worst-bloc analysis, we argue that using a salient-group perspective along with simple empirical experiments allows us to extend their progress by making practical observations over a wider array of commonly used election mechanisms, including Expanding Approvals.

4 Election mechanisms

For our experiments, we consider a range of existing multi-winner election mechanisms. For all mechanisms we note that if they do arise, ties are broken randomly. A full suite of implemented election mechanisms is publicly available in GitHub.[2]

- **Single non-transferable vote (SNTV):** Each candidate receives a score equal to their number of first-place votes. The k winners are the k candidates with the highest scores. (This is sometimes called multi-winner plurality voting.)
- **Block voting:** Among m candidates, a ranked ballot contributes +1 to the score of each candidate ranked in the top k positions; again we choose the k candidates with the highest scores. (This is sometimes called plurality bloc voting or plurality block voting.)
- **Single transferable vote (STV):** This is a multi-round process that uses a threshold of election, often $\tau = \lfloor \frac{N}{k+1} \rfloor + 1$ as we use here, to choose k winners from N ranked ballots.[3] Any candidate with more than that level of first-place support is elected, and their excess votes are transferred to their supporters with fractional weight (for instance, if a candidate receives 150 votes for a threshold of 100, then the 50 surplus votes are re-distributed proportionally among the voters who contributed to the total; Those voters' ballots are transferred to their next choice with an updated weight $\frac{50}{150} w_v$ where w_v is voter v's original contribution). This continues until k candidates are elected. If in a given round no candidate meets the quota, then the one with the least first-place support is eliminated and the ballots headed by those candidates are transferred (with full weight) to the voter's next choice.
- **Borda:** This chooses winners based on positional scoring. Among m candidates, we default to a scoring where a ranked ballot contributes $m - j$ to the score of the candidate ranked in position $j \in \{1, ..., m\}$; again we choose the k candidates with the highest scores.
- **Chamberlin-Courant:** [Chamberlin and Courant, 1983] This rule elects the group of k candidates which maximizes the sum over voter's Borda scores for their most preferred winner. The voter is thought of as being assigned to and represented by their highest-scoring winner. This rule has significant computational complexity, as finding the CC winner set is

[2] https://github.com/REDACTED

[3] Basing the threshold on $N/(k+1)$ is called the Droop quota, while replacing that with N/k would give the Hare quota. We follow [Elkind et al., 2017a] in our choice of quota.

known to be NP-hard [Procaccia et al., 2008]. Our experiments have few enough candidates and voters that this can be solved exactly with an integer programming formulation described by Skowron et al. [2015]. Ties are broken randomly by searching for up to 1000 optimal solutions and choosing uniformly at random from among that set.

- **Greedy Chamberlin-Courant (GreedyCC):** This is designed as an approximation algorithm for the Chamberlin-Courant problem, which exhibits submodularity in its objective [Lu and Boutilier, 2011]. In an iterative process starting from $\mathcal{W}_0 = \emptyset$, we build the winner set by successively adding the candidate with the largest contribution to the Chamberlin-Courant objective. This is repeated until k winners are selected.
- **Monroe:** This rule is nearly identical to Chamberlin-Courant, but with the additional constraint that each winner must represent disjoint sets of $\frac{N}{k}$ voters, rounded up or down to account for fractions. This is again NP-hard, and we use the integer programming formulation described by Skowron et al. [2015]. Ties are broken randomly by searching for up to 1000 optimal solutions and choosing uniformly at random from among that set.
- **Greedy Monroe:** Following Elkind et al. [2017a], we also use a greedy approximation for Monroe under the simplifying assumption that k divides N. This is again an iterative process starting from $\mathcal{W}_0 = \emptyset$, in which winners are successively selected by taking the candidate, along with a size $\frac{N}{k}$ selection of unassigned voters which have the largest contribution to the current objective value. A more thorough definition without simplifying assumptions is described by [Skowron et al., 2015].
- **Plurality veto** was recently introduced by Kizilkaya and Kempe [2022] as a single-winner rule achieving best-possible metric distortion among deterministic rules. Candidates start with initial scores equal to their number of first-place votes. Then voters are arbitrarily ordered (whether randomly or deterministically) and each voter in turn is queried for their least favorite among those candidates with positive scores; that least-favorite candidate loses a point. The winner is the last candidate with a positive score. We extend this rule to the multi-winner setting by initializing with k-approval scores (as is done for the Block rule), and then sequentially iterating through the voter ordering until exactly k candidates remain.
- **Committee veto** is a multi-winner variant of Plurality Veto proposed by [Kizilkaya and Kempe, 2022], in which voters rank size k sets of candidate *committees*. For each voter, the top k candidates in their ranking are taken to be a single candidate committee. The size N set of these committees are then taken to be the 'candidates' in a single winner Plurality Veto election. To create a preference profile, each voter ranks candidate committees according to their distance to the qth closest committee member. For our experiments, we default to using $q = k$.
- **Expanding approvals** is another recently designed multi-round election [Aziz and Lee, 2020, Kalayci et al., 2024]. In round i, all voters are queried in a randomized order for their ith candidate preference. As soon as a candidate c is named $\lceil \frac{N}{k} \rceil$ times, they are immediately elected and all voters who previously voted for c are removed. This repeats until k candidates are elected.

Finally for comparison, we consider three variants of the Random Dictator mechanism for single-winner elections, in which the winner is the first choice of a uniform random voter.

- **Sequential Multi Random Dictator (SMRD):** Randomly select k voters, in order, to act as dictators. Sequentially elect their favorite candidate who has not yet been elected.
- **One-Shot Multi Random Dictator (OMRD):** A single voter is randomly chosen and their top k preferences are immediately elected.

- **Discounted Multi Random Dictator (DMRD):** In every round, a voter is randomly chosen and elects their top (not-yet-elected) preference. Then all voters who would have selected the same candidate in that round have their voting power discounted by a fraction $r \leq 1$, and the voter distribution is re-normalized for the next round. Note that a voter nay be discounted multiple times. In the experiments below, we execute DMRD with $r = 1/2$.

5 Empirical results

Following Elkind et al. [2017a], we focus on simple settings that are intended to have clear perspectives to measure from and that sharply highlight differences between voting rules. Therefore, all experiments are drawn within the Euclidean plane $\mathbb{R}^2$, and are specifically designed to contain distinct, cohesive groups of voters. For each experiment, we randomly sample voter and candidate positions as described in the list below. Once we have sampled a full embedding of voters and candidates, we construct a preference profile P by ranking the distances from each voter to all candidates. Each election is then performed on the preference profile with a given number of winning candidates, and inefficiency measurements are computed for the resulting winner sets. Statistics for the inefficiency values are found by repeating this procedure and aggregating the set of results. Unless otherwise stated, each experiment performs 10,000 random trials, which we found to be sufficient and reasonable given the high computational cost for some of the voting rules.

5.1 Experiment list

(1) **Basic two-bloc experiment.** Two blocs, $n = 1000$ voters, $m = 20$ candidates, $k = 4$ winners. Voter blocs are equally sized with 500 voters each, are normally distributed with means $\mu_1 = (0, -2)$, $\mu_2 = (0, 2)$ respectively, and are given standard deviation $\sigma = 1/3$. Candidates are uniformly distributed on the 6×6 square centered at $(0, 0)$.

(2) **Varying bloc size.** Two blocs of *varying* size, $n = 1000$ voters, $m = 20$ candidates, $k = 4$ winners. Voters and candidates are distributed identically to Experiment 1, but voter bloc sizes are varied, starting with 0 voters in the first bloc and all 1000 in the second and incrementally shifting voters between blocs until all 1000 voters belong to the first bloc. Results for each incremental setting are computed for 1000 random trials.

(3) **Four blocs, two winners.** Four blocs, $n = 1000$ voters, $m = 20$ candidates, $k = 4$ winners. Voter blocs are equally sized with 250 voters each, are normally distributed with means $\mu_1 = (0, -2)$, $\mu_2 = (0, 2)$, $\mu_3 = (-2, 0)$, $\mu_4 = (0, 2)$ respectively, and are given standard deviation $\sigma = 1/3$. Candidates are uniformly distributed on the 6×6 square centered at $(0, 0)$.

(4) **Many candidates, many voters.** Two blocs, $n = 10,000$ voters, $m = 200$ candidates, $k = 4$ winners. Voters and candidates distributed identically to Experiment 1. The much larger number of candidates is unrealistic, but should reduce the small-sample effects of candidate placement and exaggerate the differences between methods.

(5) **Overlapping blocs.** Two blocs, $n = 1000$ voters, $m = 200$ candidates, $k = 4$ winners. Candidates are distributed identically to Experiment 1. The voter distribution, however, is given a larger standard deviation $\sigma = 1$ in order to create overlapping groups of voters.

(6) **Four blocs, five winners.** Four blocs, $n = 1000$ voters, $m = 20$ candidates, $k = 5$ winners. Voters and candidates distributed identically to Experiment 3.

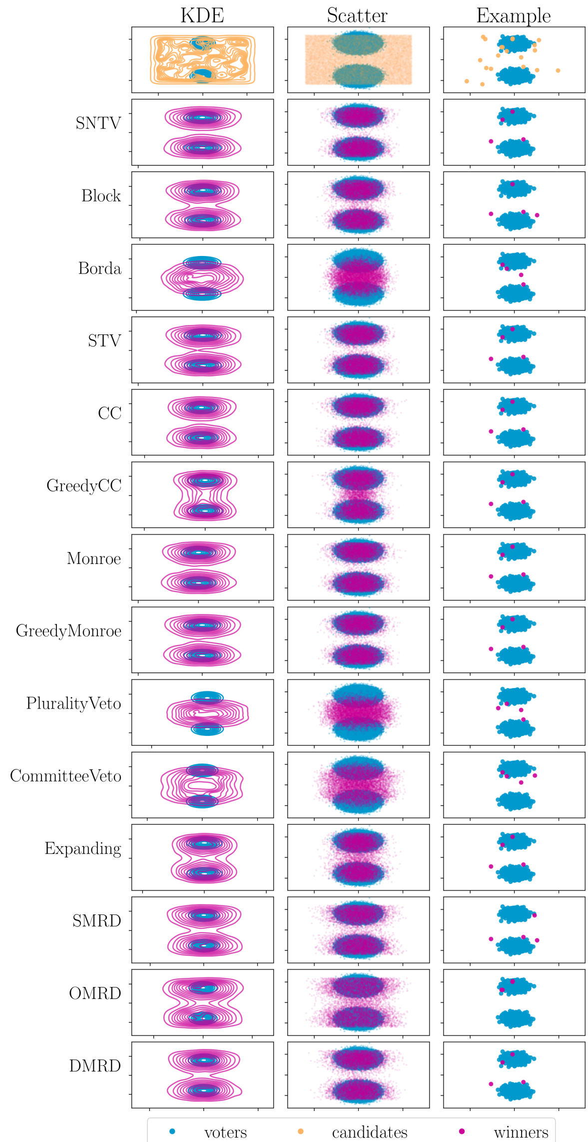

Fig. 2. Experiment 1

Fig. 3. For both experiments, we visualize results with KDE plots (left) to showcase the distribution of voter, candidate, and winner locations; scatter plots (middle) for an aggregate subset of of voter, candidate, and winner locations to highlight density; and examples from a single random trial (right). The top row shows the embedding of voters and candidates, and the remaining rows show only the voter and winner locations for each of our voting rules. This plot replicates the main demonstration from Elkind et al. [2017a] and extends to new voting rules.

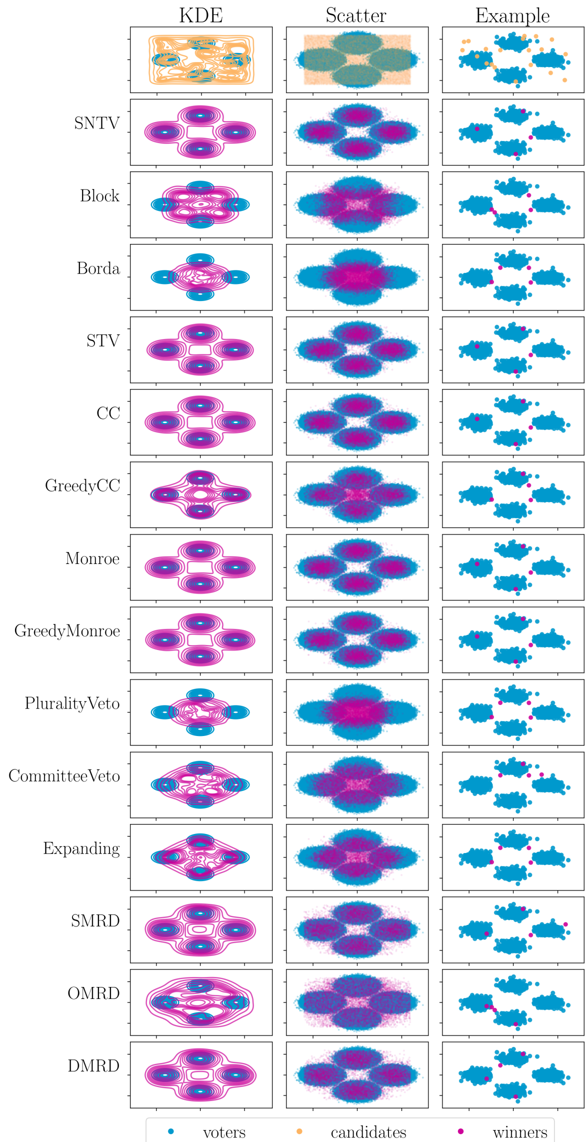

Fig. 4. Experiment 3

Fig. 5. For both experiments, we visualize results with KDE plots (left) to showcase the distribution of voter, candidate, and winner locations; scatter plots (middle) for an aggregate subset of of voter, candidate, and winner locations to highlight density; and examples from a single random trial (right). The top row shows the embedding of voters and candidates, and the remaining rows show only the voter and winner locations for each of our voting rules. This plot replicates the main demonstration from Elkind et al. [2017a] and extends to new voting rules.

5.2 Findings

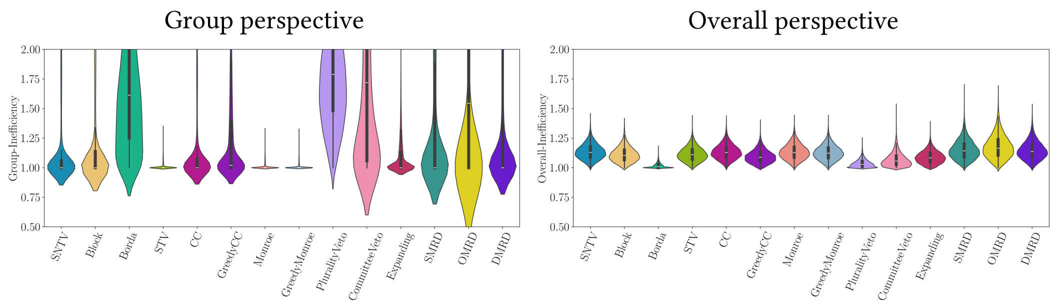

Fig. 6. Experiment 1 (two blocs of equal size). Violin plots show the distribution of values found in our random trials for group inefficiency (left) compared to overall inefficiency (right), giving sharply different views on fairness. Although some violins appear to show values below 1, this is simply an artifact of distribution smoothing and only indicates high variance – inefficiency values are never smaller than 1.

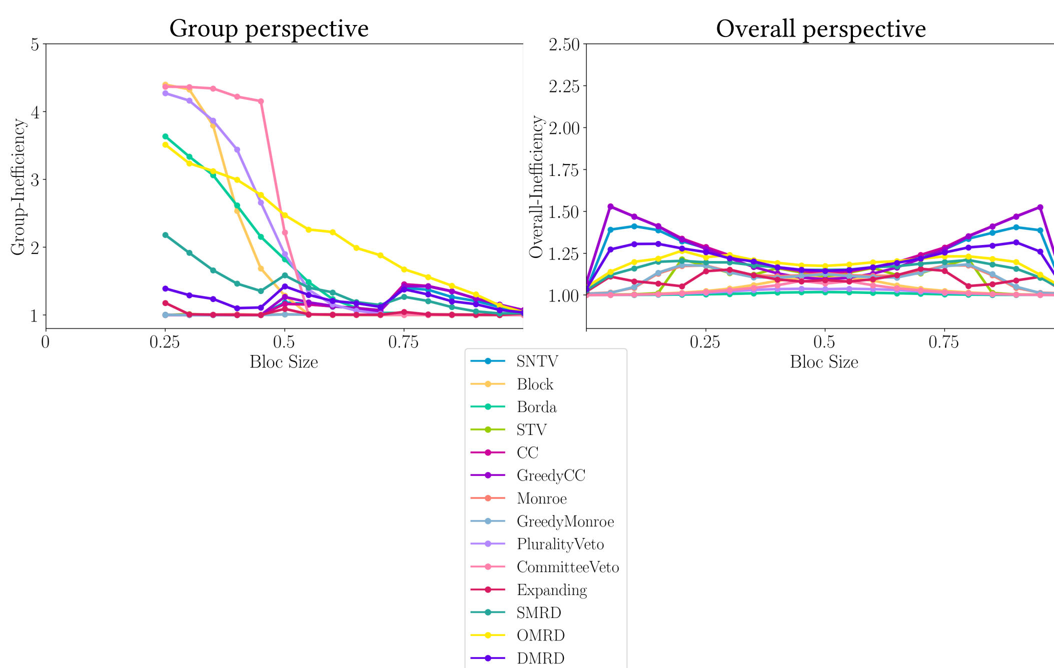

Fig. 7. Experiment 2 (two blocs of varying sizes), we plot inefficiency as the fractional size of one bloc varies with respect to the other (e.g. when one bloc has 25% of the voters, the other has 75%). Note that group inefficiency is undefined when the voter group is below 25% in this 4-winner example. Overall inefficiency once again tells a strikingly different story from the one that emerges when we center fairness to the distinctive voter blocs.

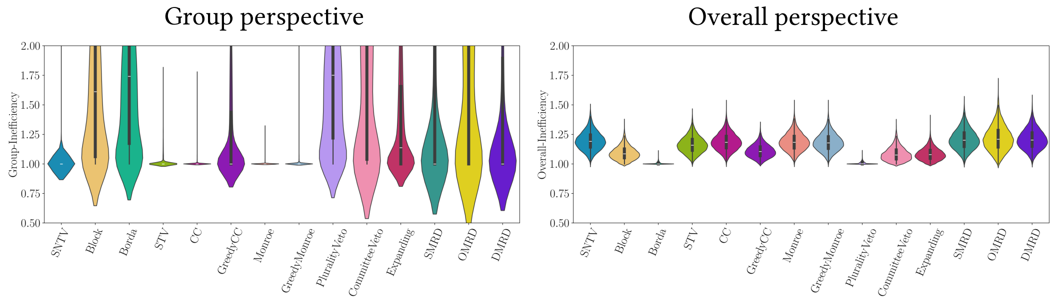

Fig. 8. Experiment 3 (four blocs of equal size). Violin plots show the distributions values found in our random trials for group inefficiency values (left) compared to overall inefficiency values (right). See Figure 6 for a note on values below 1. Notably, the inefficiency values for Block voting are much larger as compared to Experiment 1, but this appears to be a result of random tiebreaking – if two groups have the same number of votes for their top candidate ties are broken randomly, and there is a 1/2 chance that one of them gets left with no representatives at all.

These experiments invite many observations.

- **Borda stands out.** The Borda rule, along with Plurality Veto, and to some degree Committee Veto, have a clearly different pattern of winners from the others when aggregated over many trials, as seen by both experiments in Figure 5. Even the completely unreasonable rule OMRD, where one voter selects the entire winner set, resembles the more proportional rules in the sense that non-central candidates are more likely to be elected.
- **Some rules favor compromise candidates.** There are subtle differences in the tendency of voting rules to select compromise candidates, homogeneously located in the centers between voting blocs. In addition to a blatant tendency from the previously mentioned rules, Greedy-CC, Expanding Approvals, and OMRD often choose candidates in the compromise zone as well. Their tendency to compromise also intuitively explains their small inefficiency values for the overall, undivided electorate perspective visualized in Figures 6 and 8. Candidates in the compromise zone near the mean are also the ones which minimize the sum of distances to the set of all voters. This, however, works against them from the group perspective, where rules like Borda show steep inefficiency values since they do not elect candidates that represent groups individually. In a Borda election, non-central groups of voters are ill-represented by the winning results.
- **STV and Monroe are best from the group perspective.** The distributions in Figure 5 show that these methods do well to directly represent the distinct voter blocs. Likewise, Figure 6 shows that a focus on individual, salient groups gives a radically different view of fairness than we get from the undivided electorate (which is used in the usual metric distortion definitions). STV, Monroe, and Greedy Monroe perform extremely strongly for each bloc of voters, with Expanding Approvals close behind. All of these rules, however, appear ordinary relative to other elections in the whole-electorate, overall perspective – despite still performing well with very small inefficiency values.
- **A worst-bloc analysis loses important information.** A definition which considers the worst case over all subsets of voters, such as the proportional representation definition of Kalayci et al. [2024], would hide the markedly strong group inefficiency performance of STV, Monroe, and Expanding Approvals in a polarized setting, since the whole-electorate inefficiency is worse than for the natural choice of blocs.

- **Block voting and Borda are especially costly for small groups.** Figure 7 extends the theme that the group perspective is crucial to a deeper understanding of fairness. When one bloc has less than half of the voters Block voting, Plurality Veto, Committee Veto, OMRD, and Borda stand out as most unfair with very large inefficiency values. When one group dominates in size Block voting, for example, provides no check on their voting power. If a group were to have even just a slight majority, once they agree sufficiently on a set of k candidates, they will be able to elect them all. We note again, that this is not reflected by measurements from the overall perspective.
  Notably, OMRD is designed as an intentionally unreasonable system and Borda voting is rarely used for political representation, but Block voting is one of the most frequently used systems for local election in the United States, such as for city councils and county commissions.[4]
- **Tie-breaking strategies are important.** In Experiment 3 we find a dramatic increase in group inefficiency values for the Block voting rule. This appears mainly as a result of a more pronounced tendency to elect central, compromise candidates. However, we also note that extreme values in the tail of the distribution for both experiments are likely to be a result of our random tie-breaking strategy. Some voting rules, including Block and SNTV, appear more susceptible to this effect than others. For example, consider a Block election with two winners and two equivalently sized blocs of voters who share the exact same preferences. Under a random tie-breaking mechanism there is a 1/6 chance a single bloc will take all the winning seats, since each bloc's first and second preferences receive the same number of votes. When the setting is split into four blocs deserving a single representative each, it becomes even more likely that a single party will not elect any winners.
- **Results are robust across a variety of settings.** Experiments 4, 5, and 6 are displayed in the appendix. Respectively, these results increase the scale of the election, decrease bloc saliency, and change the parity for the number of winning seats. There are small various to the results. In Experiment 4 we notice much larger group-inefficiency values for elections such as Borda and Block (since groups now have more candidates to choose from), and sharply defined winner distributions in Figure 9. In Experiment 5 we find that some rules such as Expanding Approvals show a stronger tendency to elect central candidates and therefore find larger group inefficiency values. Finally in Experiment 6 GreedyCC shows a better ability to represent voters with smaller group-inefficiency scores. Our main theme, however, remains present in all of these settings. That is to say, the relative patterns among voting rules is consistent with our previous experiments: rules such STV and Monroe perform strongly in group-inefficiency measurements, while Borda comes with a steep cost.
- **This analysis has real-world implications.** STV is currently used, or is being considered for adoption, in many parts of the world. In the United States, reformers often claim that STV will provide stronger proportional representation for minority groups than legacy systems such as SNTV and Block voting.[5] The current demonstrations give supporting evidence for this assertion from the perspective of distinctive voter blocs.

[4]For example, block voting is still used across Colorado, including in Cortez, Superior, Lyons, Lafayette, Boulder, and in the at-large seats in Denver and Longmont.
[5]This was a claim made by reform advocates in Portland, Oregon, which recently moved to STV for the election of its city council.

6 Identifying salient groups

Real-world voting often involves measurable polarization; for instance, in the decades that followed the passage of the Voting Rights Act of 1965 (VRA) in the United States, a lineage of statistical techniques has been created and refined for quantifying the polarization along racial and ethnic lines. Minority groups who have distinct candidate preferences may be entitled to protection under the VRA if their preferences are consistently blocked by a cohesive majority. For instance, this is what the U.S. Supreme Court recently held for Black voters in Alabama in *Allen v. Milligan* (2023), sending the state back to the drawing board to make new electoral districts.

From a perspective closer to home for computer scientists and economists, a relevant notion of salience can be cashed out by looking for groups with simple descriptions. One way of elaborating that, for instance, would be to choose several pairwise comparisons that classify the voters into groups—for instance, preferring both of candidates A, B to both of candidates C, D may be a strong enough description to pick out important subgroups of voters.

More generally, this is a special case of using a *small decision tree* to pick out groups, as for instance in [Hebert-Johnson et al., 2018], where such groups are called "computationally identifiable."

7 Conclusion

In the present paper, simple models of polarization show that fairness results can be radically different depending on the point of view. This observation fits well into current work in algorithmic fairness, where several authors have recognized that content-neutral definitions fail to capture cases of particular interest.

Furthermore, our results give insight into the ability of election mechanisms to fairly represent their voters. Rules such as Borda find compromise among the entire voter population, electing central, homogenous sets of winners. STV and Monroe, on the other hand, do well at electing diverse winning committees, likely to secure representation for each relevant voter group. Likewise, some mechanisms such as Expanding Approvals appear moderately strong from either perspective, while others such as Block suffer especially in the presence of a small but cohesive minority.

Our study brings an empirical perspective to the evaluation of proportional fairness in commonly used voting rules. Metric voting has exploded in popularity, but its extension to the multi-winner setting is still new, and most recently proposed rules have yet to be analyzed from a proportionality point of view. An immediate direction for future work, following [Kalayci et al., 2024], could seek worst-embedding bounds to accompany the trends observed here.

The present findings make it clear that the perspective of a cohesive group can be lost in some of the popular fairness formulations. We also find that rules that shine in some kinds of elections are measurably unfair in others, suggesting that the choice of a fair voting rule may be situational and even fluid. Nevertheless, certain voting rules, including popular ones in current practice like Block plurality voting, look unfair across the board.

References

Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.

Elliot Anshelevich and John Postl. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58 (2017), 797–827.

Haris Aziz and Barton E Lee. 2020. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare* 54, 1 (2020), 1–45.

Markus Brill and Jannik Peters. 2023. Robust and verifiable proportionality axioms for multiwinner voting. *arXiv preprint arXiv:2302.01989* (2023).

Ioannis Caragiannis, Nisarg Shah, and Alexandros A Voudouris. 2022. The metric distortion of multiwinner voting. *Artificial Intelligence* 313 (2022), 103802.

John R Chamberlin and Paul N Courant. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77, 3 (1983), 718–733.

Moses Charikar and Prasanna Ramakrishnan. 2022. Metric distortion bounds for randomized social choice. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2986–3004.

Moses Charikar, Prasanna Ramakrishnan, Kangning Wang, and Hongxun Wu. 2024. Breaking the metric voting distortion barrier. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1621–1640.

Edith Elkind, Piotr Faliszewski, Jean-François Laslier, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017a. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.

Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017b. Properties of multiwinner voting rules. *Social Choice and Welfare* 48 (2017), 599–632.

Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice* 74, 2017 (2017), 27–47.

Ashish Goel, Reyna Hulett, and Anilesh K Krishnaswamy. 2018. Relating metric distortion and fairness of social choice rules. In *Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation*. 1–1.

Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. 2018. Multicalibration: Calibration for the (Computationally-Identifiable) Masses. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 1939–1948. https://proceedings.mlr.press/v80/hebert-johnson18a.html

Yusuf Kalayci, David Kempe, and Vikram Kher. 2024. Proportional representation in metric spaces and low-distortion committee selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9815–9823.

David Kempe. 2020. Communication, distortion, and randomness in metric voting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 2087–2094.

Fatih Erdem Kizilkaya and David Kempe. 2022. Plurality veto: A simple voting rule achieving optimal metric distortion. *arXiv preprint arXiv:2206.07098* (2022).

Tyler Lu and Craig Boutilier. 2011. Budgeted social choice: From consensus to personalized decision making. In *IJCAI*, Vol. 11. 280–286.

Ariel D Procaccia, Jeffrey S Rosenschein, and Aviv Zohar. 2008. On the complexity of achieving proportional representation. *Social Choice and Welfare* 30, 3 (2008), 353–362.

Piotr Skowron, Piotr Faliszewski, and Arkadii Slinko. 2015. Achieving fully proportional representation: Approximability results. *Artificial Intelligence* 222 (2015), 67–103.

A Experiment 4: Many candidates, many voters

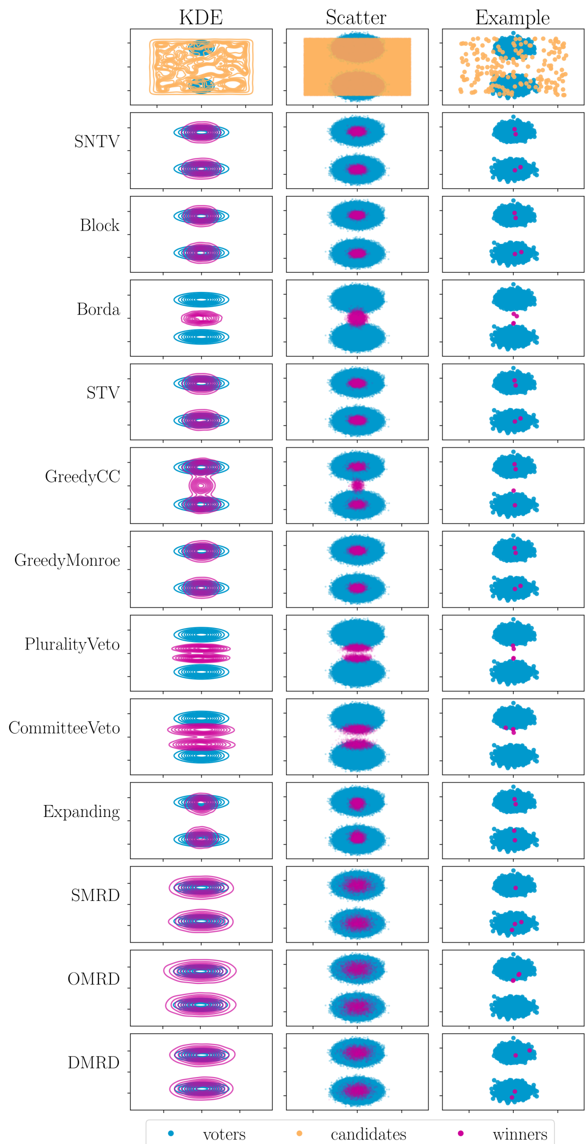

Fig. 9. Experiment 4, two overlapping blocs scaled to 10,000 voters and 200 candidates. Note that Chamberlin Courant and Monroe are left out of this experiment, which was much more computationally demanding.

B Experiment 5: Two Overlapping Blocs

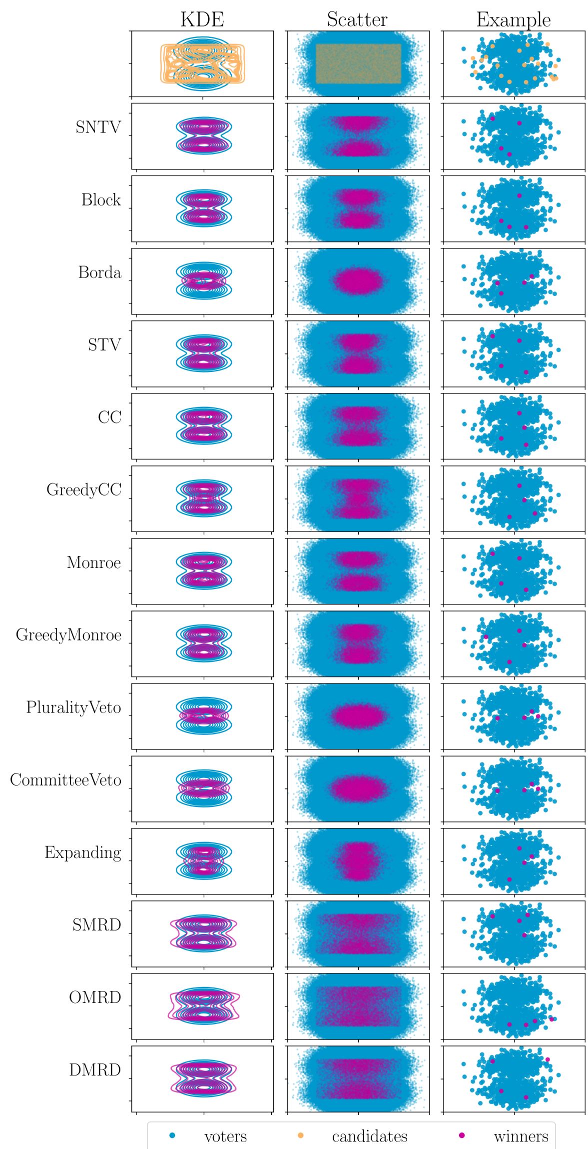

Fig. 10. Experiment 5, two overlapping blocs.

C Experiment 6: Four Blocs, Five Winners

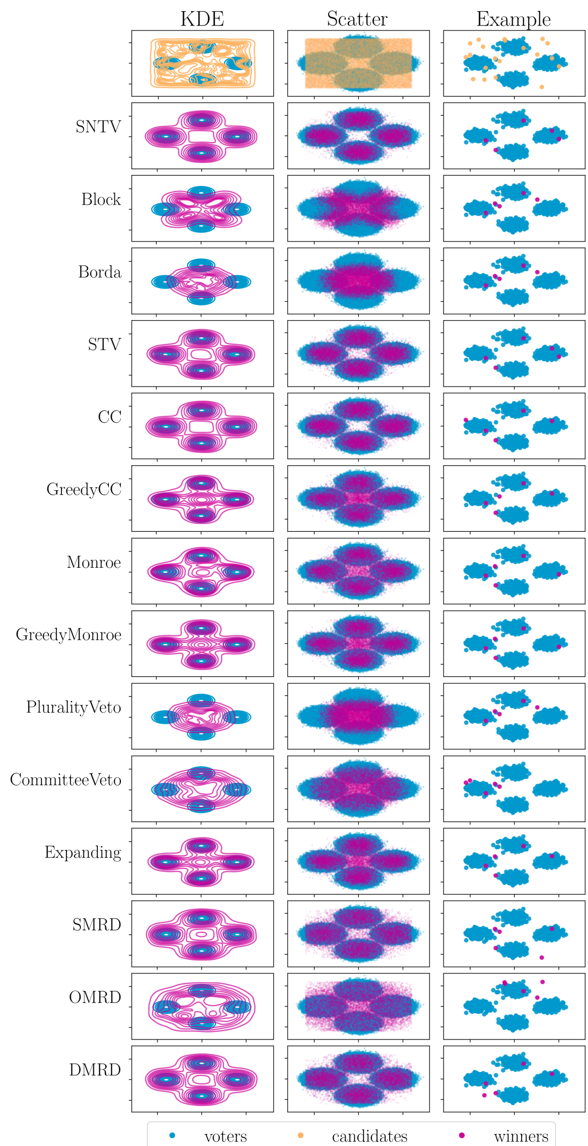

Fig. 11. Experiment 6, four blocs with five winners.

D Summarizing group vs. overall perspective

For Experiments 4-6, we can now juxtapose the summary plots to emphasize the trend: the overall perspective consistently hides the enormously varying information from the point of view of the blocs from which the elections were generated. Violin plots show the distributions values found in our random trials for group inefficiency values (left) compared to overall inefficiency values (right). In all of these, violin plots may create the appearance of values below 1.

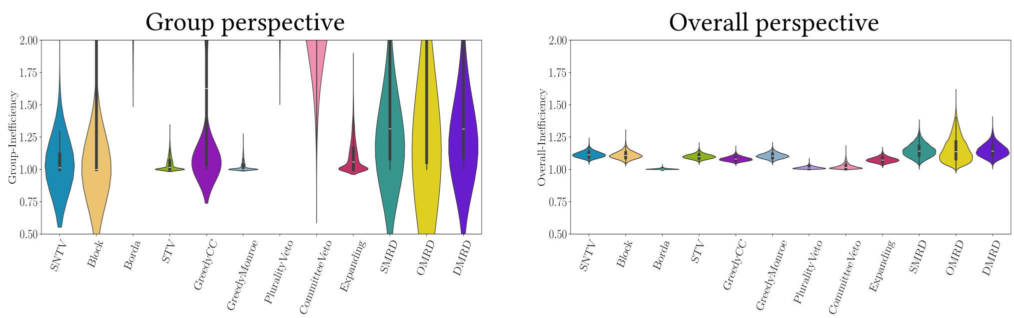

Fig. 12. Experiment 4 (many candidates, many voters).

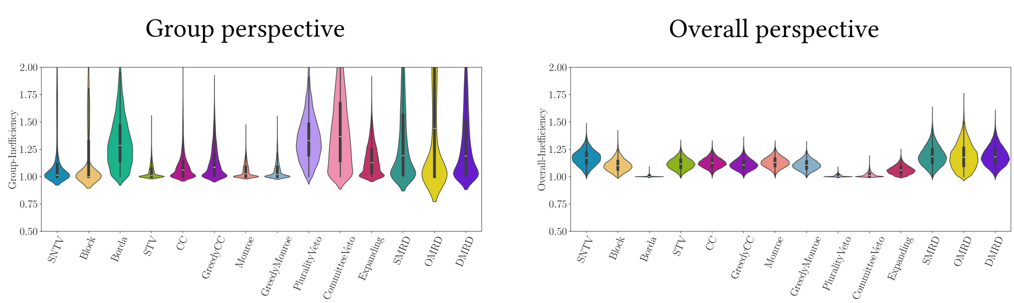

Fig. 13. Experiment 5 (two overlapping groups).

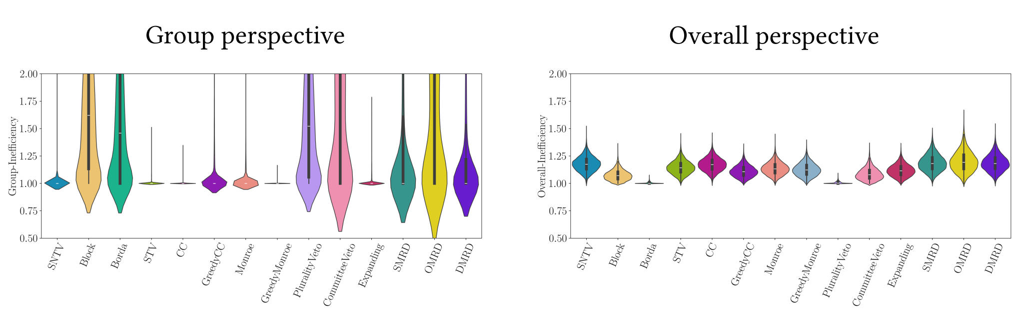

Fig. 14. Experiment 6 (four blocs with five winners).