The group perspective on fairness in multi-winner voting rules

Kevin Quinn and Moon Duchin

Abstract

In the metric voting literature, both candidates and voters are assumed to be embedded in a common metric space X representing positions on issues. The "cost" to a voter v of a candidate c is their distance $d(v, c)$ in the space X (which is sometimes said to measure the difference in issue positions), and it is presumed that voters with an ordinal ballot style would rank their preferences in order of proximity, even if the metric space is latent (or not explicitly known to the voters). Many authors have searched for voting rules that tend to elect "optimal" (lowest-cost) winners, or where cost ratios have some guarantee in the form of an upper bound. Most work has focused on optimizing cost in worst-case scenarios. In this short paper, we conduct experiments from a different angle: if more general fairness axioms can be cast as metric measurements, then we can survey voting rules in terms of their fairness tendencies. We particularly focus on a novel definition we call group inefficiency, and highlight the finding that the group perspective radically shifts which voting rules seem to deliver on proportional fairness.

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 reframing fairness axioms to be about salient groups of people can cause extreme changes to the findings.

This paper fits squarely into an active research area in computational social choice, but offers novel insights through a change of perspective.

Main definitions

Elections, profiles, and voting rules

Suppose we have a set $\mathcal{V}$ of n voters and a set $\mathcal{C}$ of m candidates. An *embedded election* is a metric space X with distance function d, equipped with a map $f : \mathcal{V} \cup \mathcal{C} \to X$ that identify 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* $P = P(X, f)$ is a set of rankings $\succcurlyeq_v$ for the $v \in \mathcal{V}$, where for each 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 derived from the embedding but not vice versa.

If $\mathcal{P}$ is the set of profiles of a given type, then $\mathcal{M} : \mathcal{P} \to 2^{\mathcal{C}}$ denotes a multi-winner election mechanism (otherwise known as a *voting rule*). We will focus on rules designed to elect $|\mathcal{M}(P)| = k$ winners.[1] One simple example is single non-transferable vote (SNTV), where we only pay attention to voters' first-place selections, and we elect the k candidates with the most first-place votes. Another simple example is plurality bloc voting, where we consider voters' top k preferences, awarding one point to a candidate for receiving one of these regardless of its 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, to use the economics framing) between the 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.

Blocs, costs, and representative assignments

Informally, elections are called *polarized* when there are significant, 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

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

define the cost of S to $\mathcal{B}$ as

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

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

Next, suppose we have a way of assigning to a subset of voters their preferred candidates from among a set of options. So, given a bloc $\mathcal{B} \subseteq \mathcal{V}$ and a candidate set $S \subseteq \mathcal{C}$, a *representative assignment* is a function $\Phi : 2^{\mathcal{V}} \times 2^{\mathcal{C}} \to 2^{\mathcal{C}}$ satisfying $\Phi(\mathcal{B}, S) \subseteq S$. This is thought of as identifying a designated subset of candidates from S that correspond to the voters in bloc $\mathcal{B}$.

The choice of assignment function Φ can be shifted, but in this paper we focus on a greedy proportional representation assignment that we denote by Φ_{prop}. For mechanisms electing k winners, the proportional share of the winner set for 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 proportional size. For $T \subseteq \mathcal{C}$ we define

$$\Phi_{\text{prop}}(\mathcal{B}, T) = \underset{\substack{S \subseteq T \\ |S| = \lfloor |T| \cdot |\mathcal{B}|/n \rfloor}}{\arg\min} \text{cost}(\mathcal{B}, S). \tag{2}$$

For example, if a bloc makes up 30% of the electorate in a 4-winner election, then among the winners they would be assigned to their one favorite winner, i.e., the winning candidate who minimizes the summed cost to that bloc. But for blocs smaller than 25% of voters, they would be assigned the empty set in a 4-winner scenario.

Group-centered fairness

Definition 1 *For a mechanism $\mathcal{M}$, embedded election X, winner set $\mathcal{W} = \mathcal{M}(X)$, and voter bloc $\mathcal{B}$, define the* group inefficiency *by*

$$\mathcal{I}(\mathcal{B}) = \mathcal{I}_{\mathcal{M},X}(\mathcal{B}) = \frac{\text{cost}(\mathcal{B}, \Phi(\mathcal{B}, \mathcal{W}))}{\text{cost}(\mathcal{B}, \Phi(\mathcal{B}, \mathcal{C}))}.$$

We say $\mathcal{M}$ has overall inefficiency α *if it has that level of group inefficiency for the undivided electorate $\mathcal{V}$, which means that $\mathcal{I}(\mathcal{V}) \leq \alpha$, or in other words the cost of the winners is no more than α times the cost of the metrically optimal winner set. (This is the standard notion of distortion.)*

We note that it is far more common in computational social choice papers to study the *worst inefficiency* of a mechanism by bounding the group inefficiency for any possible voter bloc $\mathcal{B} \subseteq \mathcal{V}$. We will see below that this would hide fundamental differences in fairness.

Related work

Our study branches off from a line of recent work that studies voting rules in metric settings. For single-winner voting rules, Anshelevich and Postl (2017); Anshelevich et al. (2018) introduce the metric framework and a notion of seeking low-distortion voting rules, proving bounds for both known and novel election mechanisms. This led to 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).

The case of multi-winner elections sometimes called "committee voting," though we will avoid that terminology here because we are motivated by the case of political representation. Despite being interesting and applicable, multi-winner elections have received much less attention in the metric voting literature, whether axiomatic or descriptive, than the single-winner case. Exceptions include (Faliszewski et al. 2017; Elkind et al. 2017b), as well as those described further below. Results from our work are strongly inspired by the latter paper or Elkind et al., where the authors use simple metric embeddings in order to produce intuitive visualizations, giving qualitative insight into how election mechanisms represent voters' preferences. We focus on measurements that build on their work by giving quantitative degrees of fairness from various points of view.

Other notions of cost and efficiency for multi-winner systems include the following.

Worst-bloc fairness

Goel, Hulett, and Krishnaswamy (2018) evaluate the cost of a winner set by the group of voters which is furthest away from it. Specifically, for a blocs of size ℓ the cost is defined as $\text{cost}_\ell(\mathcal{W}) = \max_{\mathcal{B} \subset \mathcal{V}, |\mathcal{B}| = \ell} \sum_{v \in \mathcal{B}} \sum_{c \in \mathcal{W}} d(v, c)$.

Here, the authors define an ℓ-fairness ratio as ${}^{\text{cost}_\ell(\mathcal{W})}/_{\text{cost}_\ell \mathcal{W}^*}$ and search for values of ℓ and worst-case preference profiles P to maximize it.

q-th closest

Caragiannis, Shah, and Voudouris (2022) define the cost to voter v as their distance to their qth closest of the winners. 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 $\Phi_q(v, \mathcal{W}) = d_q$. The q-cost of an election is then computed as $\text{cost}_q(\mathcal{W}) = \sum_{v \in \mathcal{V}} d(v, \Phi_q(v, \mathcal{W}))$.

Optimizing for q-cost produces intriguing phase transitions and some intuitively proportional outcomes, but it is not sensitive to the concept of voter blocs.

γ-proportional representation

The very recent paper of Kalayci, Kempe, and Kher (2024) is closest in spirit to our current setup. In their work, as in 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 is that a relatively new voting rule (expanding approvals) gives a good guarantee, satisfying their definition of proportional fairness at a constant level. The fundamental difference between their definition and ours is that (like the worst-bloc definition above) they look for the greatest inefficiency over all blocs, while we focus on individual groups. As we show below, this can greatly change the findings.

Election mechanisms

For our experiments, we consider a range of existing multi-winner election mechanisms. When needed, we break ties by randomly choosing among candidates. 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. Winners are the k candidates with the highest scores. This is also known as multi-winner plurality voting.
- **Bloc:** 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.
- **Single transferable vote (STV):** This is a multi-round process that uses a threshold of election, typically $\tau = \lfloor \frac{n}{k+1} + 1 \rfloor$, to choose winners. Any candidate with more than that level of first-place support is elected, and their excess votes are transferred to the next selection of their supporters with fractional weight. (For instance, if a candidate receives 150% of the threshold number of votes, then their voters' ballots are transferred to their next choice with weight $50/150 = 1/3$.) 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 their next choice.
- **Borda:** Among m candidates, a ranked ballot contributes $m - j$ to the score of the candidate ranked in position j; again we choose the k candidates with the highest scores.
- **Chamberlin-Courant:** This rule elects the group of k candidates that maximizes the sum over voters of the highest Borda score they award to any winner.

  While appealing for its voter representation properties, finding a winner set with this rule is an NP-hard problem (Procaccia, Rosenschein, and Zohar 2008). For our purposes, however, we use an integer programming formulation described by Skowron, Faliszewski, and Slinko (2015), which is suitable for experiments with small numbers of voters and candidates.
- **Greedy Chamberlin-Courant (GreedyCC):** This is designed as an approximation algorithm for the Chamberlin-Courant problem, which turns out to be submodular (Lu and Boutilier 2011). In an iterative process starting from $\mathcal{W}_0 = \emptyset$ we build a winner set by successively adding the candidate with the highest total Borda score until k winners are selected.
- **Monroe:** This is nearly identical to Chamberlin-Courant, but we add an additional constraint on the number of voters that a single candidate can represent to be between $\lfloor \frac{n}{k} \rfloor$ and $\lceil \frac{n}{k} \rceil$.

  This is again NP-hard and we use a integer programming formulation. A greedy approximation is possible for this problem as well (Skowron, Faliszewski, and Slinko 2015), but has a much more involved implementation than its unconstrained counterpart.
- **Plurality veto** is an interesting election mechanism designed by Kizilkaya and Kempe (2022) as a single-winner rule achieving best-possible metric distortion. Candidates start with initial scores equal to their number of first-place votes. Then voters are arbitrarily ordered (whether randomly and 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 stopping when exactly k candidates remain.
- **Expanding approvals** Another recently designed multi-round election (Kalayci, Kempe, and Kher 2024; Aziz and Lee 2020), where at each 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 ballots have been exhausted and we have elected k candidates.

  Kalayci, Kempe, and Kher (2024) show that this rule has strong guarantees in their framework of all-groups proportionality.

Finally, we close with three multi-winner variants of the fundamental Random Dictator mechanism for single-winner election.

- **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 voted for the winning candidate in that round have their voting power discounted by a fraction $\rho \leq 1$, and the voter distribution is re-normalized for the next round. In the experiments below, we execute DMRD with $\rho = 1/2$.

Empirical results

We conduct experiments in which each voter's position is randomly sampled from a normal distribution and candidate positions are sampled from a uniform distribution on a larger region containing the voter blocs. Following Elkind et al. (2017a), we focus on simple settings that highlight the differences between voting rules; all experiments are drawn in the Euclidean plane $\mathbb{R}^2$.

[2]https://github.com/REDACTED/Metric-Vote-Representation

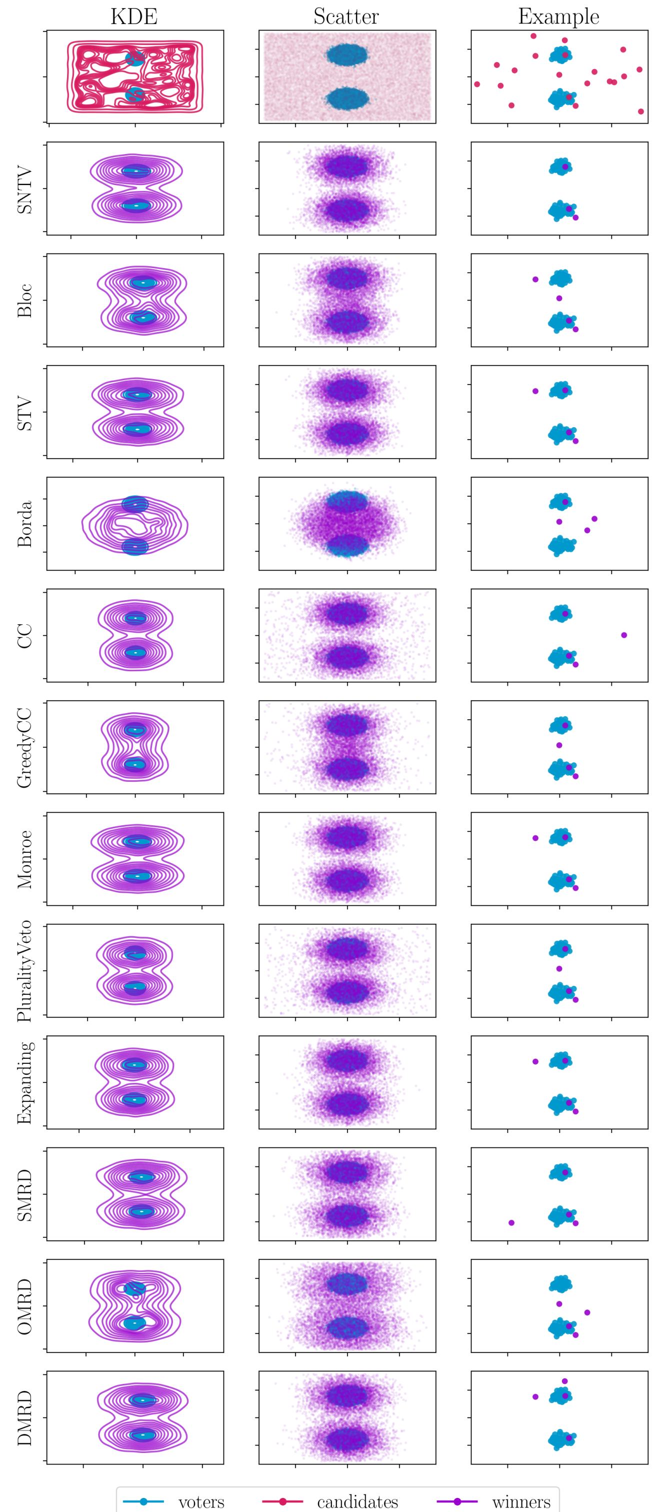

Figure 1: Here, the top row shows the embedded election setup for Experiment 1 (two blocs of equal size), and the remaining rows show the outcomes of various multiwinner voting rules. Election results can be visualized with KDE plots (left column) or scatterplots (middle column) over 10,000 trials. Examples of single trials are shown in the right column. This plot replicates the main demonstration from Elkind et al. (2017a) and extends to more voting rules.

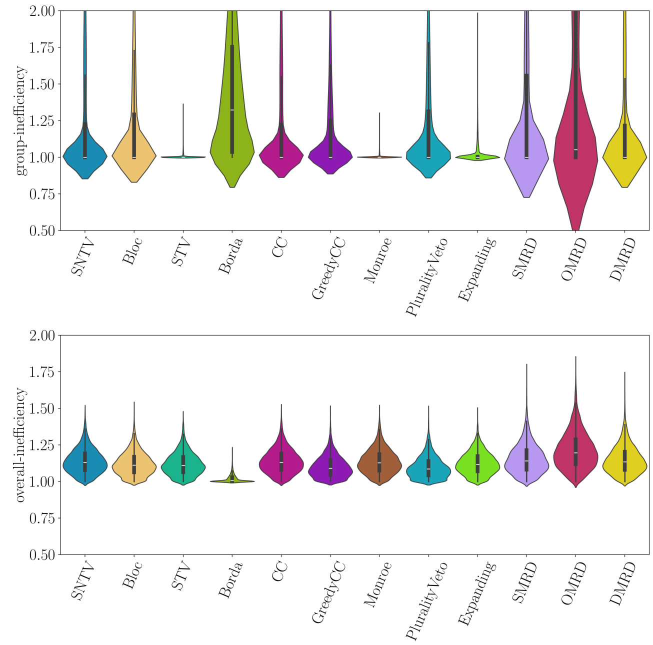

Figure 2: In Experiment 1 (two blocs of equal size), violin plots show group inefficiency compared to overall inefficiency, giving sharply different views on fairness.

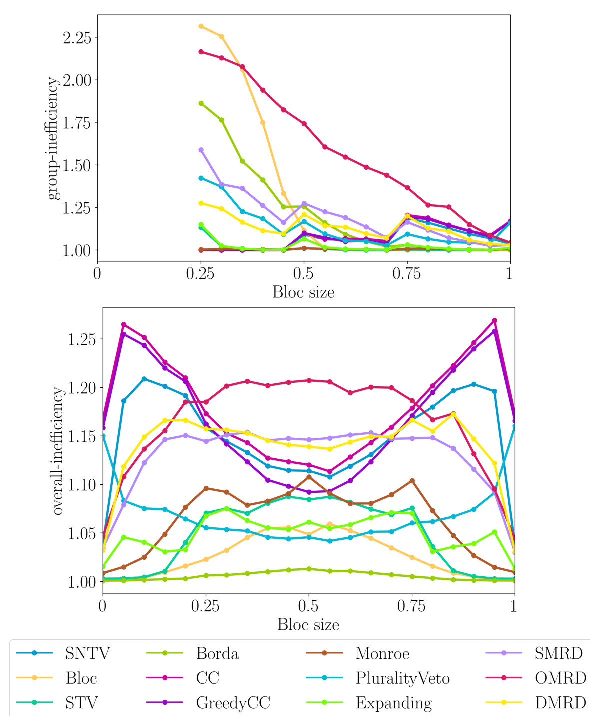

Figure 3: In Experiment 2 (two blocs of varying sizes), average group inefficiency values are shown as bloc sizes vary. Overall inefficiency once again tells a strikingly different story from the fairness to the distinctive voter blocs.

Experiment list

1. Two blocs of equal size, each normally distributed, candidates uniform on square. 100 voters, 20 candidates, 4 winners.
2. Two blocs of varying size, each normally distributed, candidates uniform on square. 100 voters, 20 candidates, 4 winners. Bloc sizes $|\mathcal{B}| = 5, 10, \ldots 100$.
3. Two equal blocs, each normally distributed, candidates uniform on square. 10,000 voters, 20 candidates, 4 winners.
4. Four equal blocs, all normally distributed, candidates uniform on square. 100 voters, 20 candidates, 4 winners.
5. Four equal blocs, all normally distributed, candidates uniform on square. 10,000 voters, 20 candidates, 4 winners.
6. Four equal blocs, all normally distributed, candidates uniform on square. 100 voters, 20 candidates, 5 winners.
7. Four equal blocs, all normally distributed, candidates uniform on square. 10,000 voters, 20 candidates, 5 winners.

Due to space constraints, most outputs are included in the supplemental material. Note that the reason to conduct very small experiments (with $n = 100$ voters) is to allow comparisons including computationally intensive voting rules like Chamberlin-Courant; by juxtaposing with the outputs from larger elections, we see that this does not cause a significant change in our main findings for other voting rules.

Findings

Many observations are available from these experiments.

- Only the Borda rule has a clearly different pattern of winners from the others when aggregated over many trials, as seen in Figure 1. Even the completely unreasonable rule OMRD, where one voter selects the entire winner set, resembles the more proportional rules in the aggregate.
- There are subtle differences in the tendency of voting rules to select compromise candidates, located between the two voting blocs, in addition to a blatant tendency from Borda rule to focus on those. Bloc voting, Greedy-CC, and (surprisingly) OMRD are the most likely to choose candidates in the compromise zone.
- The violin plots in Figure 2 show that the focus on individual salient groups gives a radically different view of fairness than we get from the undivided electorate (i.e., the usual metric distortion) or from the worst-case choice of groups. STV and Monroe perform extremely strongly for each bloc of voters, with expanding approvals close behind, but all three look ordinary from the whole-electorate perspective. And both groups pay a steep cost in a Borda election, while that voting rule performs peerlessly for the electorate at large.
- Figure 3 extends the theme that the group perspective is crucial to a deeper understanding of fairness. The one-shot random dictator rule OMRD, together with Bloc voting, stand out as most unfair for small voter blocs, with Borda voting not far behind. Notably, OMRD is designed as an intentionally unreasonable system and Borda voting is rarely used for political representation, but Bloc voting is one of the most frequently used systems for local election in the United States, such as for city councils and county commissions.
- A definition which considers the worst case over all subsets of voters, such as the proportional representation definition of Kalayci, Kempe, and Kher (2024), would hide the markedly strong group inefficiency performance of STV, Monroe, and expanding approvals in a polarized setting, because the whole-electorate inefficiency is worse than for the natural choice of blocs.
- 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 Bloc voting.[3] The current demonstrations give supporting evidence for this assertion from the perspective of distinctive voter blocs.

Conclusions

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 consideration under the VRA if their preferences are consistently blocked by a cohesive majority, as the U.S. Supreme Court recently held for Black voters in Alabama, sending the state back to the drawing board to make new electoral districts.

In the present paper, simple toy models of polarization show that fairness for distinctive voting blocs is not well captured by standard definitions of metric distortion, or even by existing notions of proportional representation that take a worst-case measurement over all groups.

[3] For example, this was a claim made by reform advocates in Portland, Oregon, which recently moved to STV for the election of its city council.

References

Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264: 27–51.

Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58: 797–827.

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

Caragiannis, I.; Shah, N.; and Voudouris, A. A. 2022. The metric distortion of multiwinner voting. *Artificial Intelligence*, 313: 103802.

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

Charikar, M.; Ramakrishnan, P.; Wang, K.; and Wu, H. 2024. Breaking the metric voting distortion barrier. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1621–1640. SIAM.

Elkind, E.; Faliszewski, P.; Laslier, J.-F.; Skowron, P.; Slinko, A.; and Talmon, N. 2017a. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017b. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48: 599–632.

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

Goel, A.; Hulett, R.; and Krishnaswamy, A. K. 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.

Kalayci, Y.; Kempe, D.; and Kher, V. 2024. Proportional representation in metric spaces and low-distortion committee selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9815–9823.

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

Kizilkaya, F. E.; and Kempe, D. 2022. Plurality veto: A simple voting rule achieving optimal metric distortion. *arXiv preprint arXiv:2206.07098*.

Lu, T.; and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *IJCAI*, volume 11, 280–286.

Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2008. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3): 353–362.

Skowron, P.; Faliszewski, P.; and Slinko, A. 2015. Achieving fully proportional representation: Approximability results. *Artificial Intelligence*, 222: 67–103.

Supplemental figures

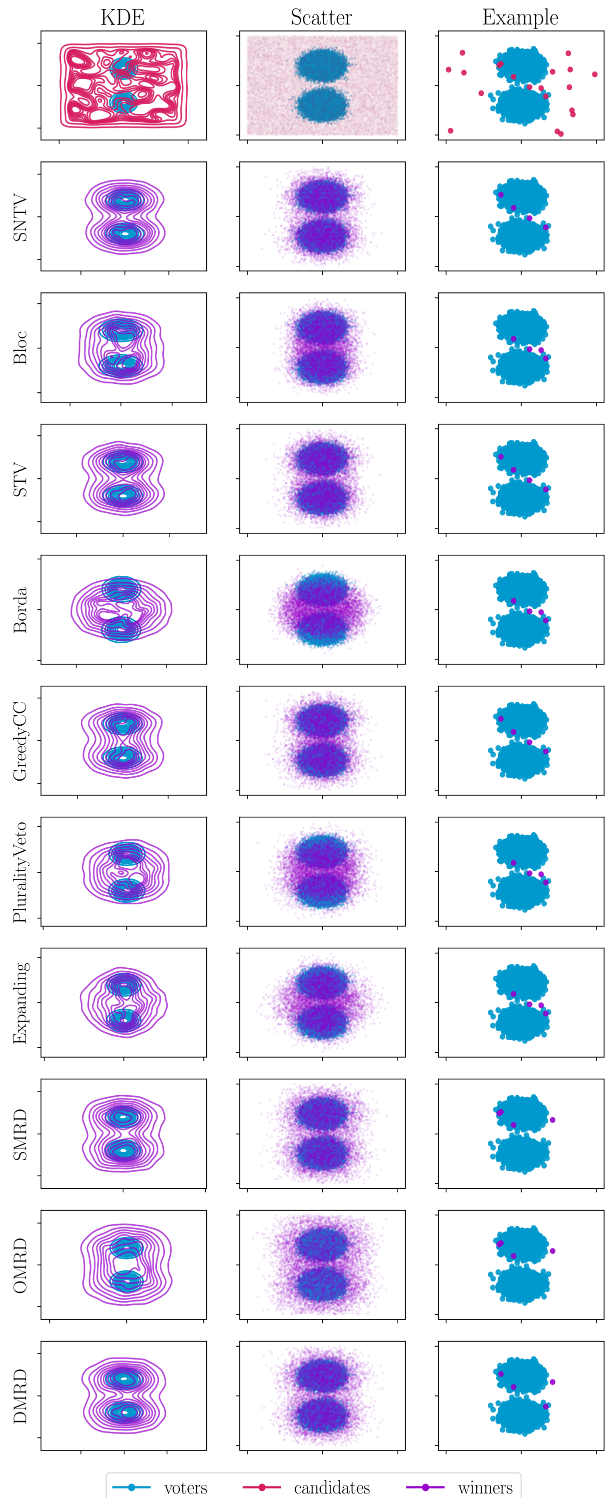

Figure 4: Experiment 3 (2 blocs of equal size, 10,000 voters, 4 winners).

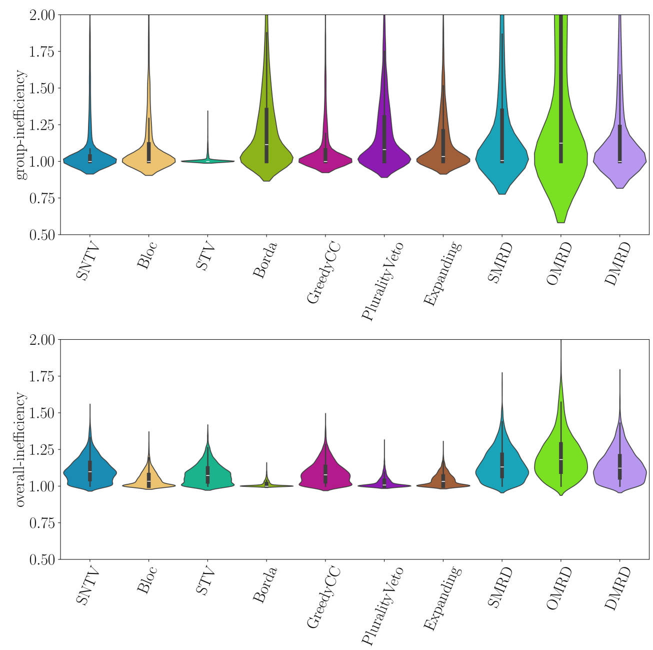

Figure 5: Experiment 3 (two blocs, 10,000 voters, 4 winners).

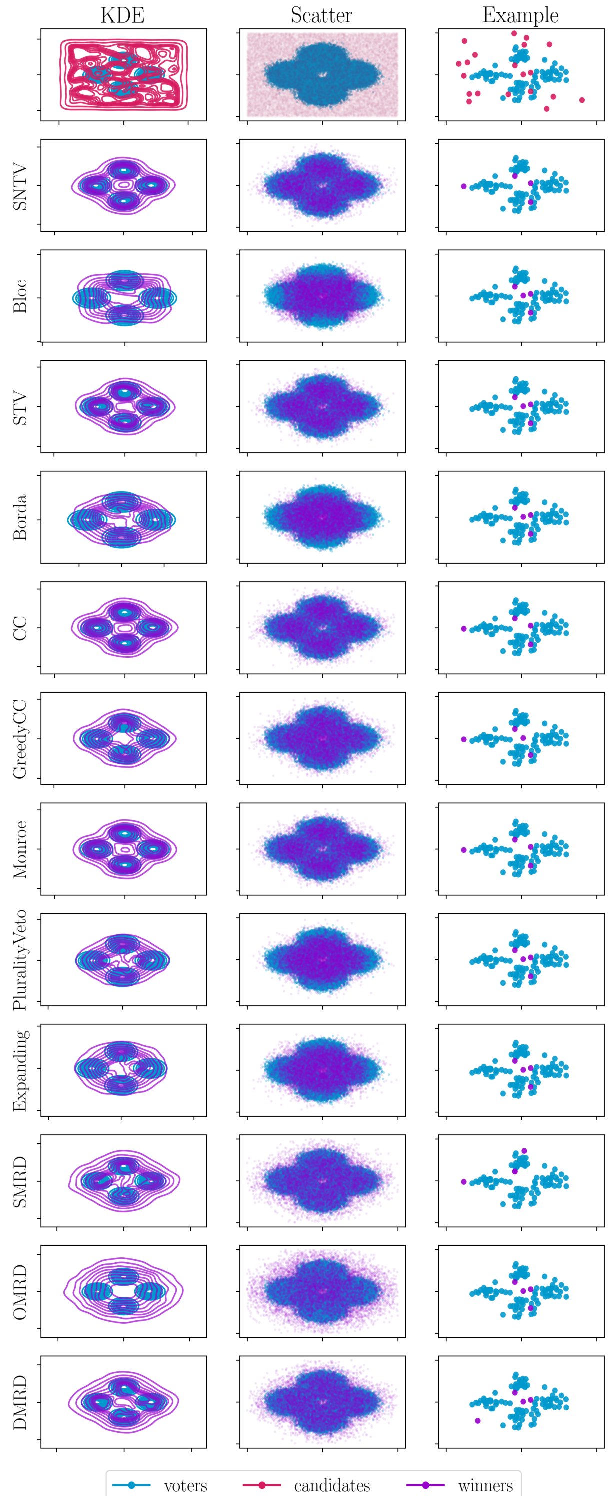

Figure 6: Experiment 4 (four blocs of equal size, 100 voters, 4 winners).

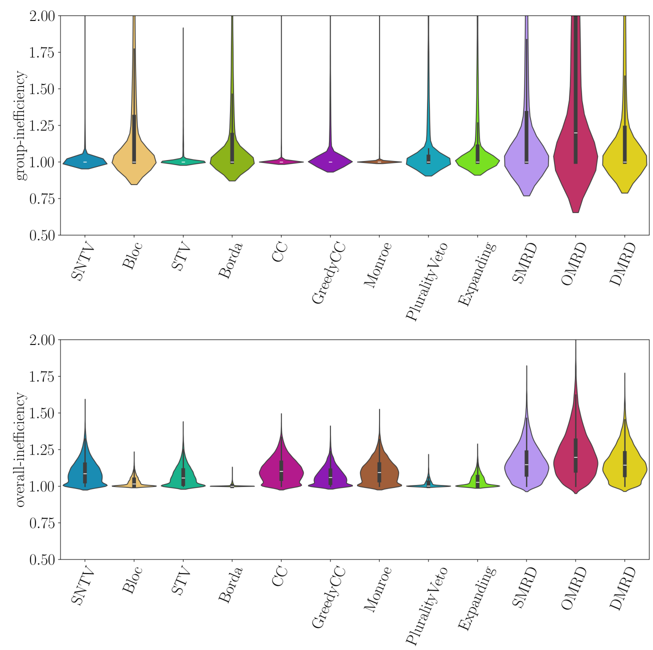

Figure 7: Experiment 4 (four blocs, 100 voters, 4 winners).

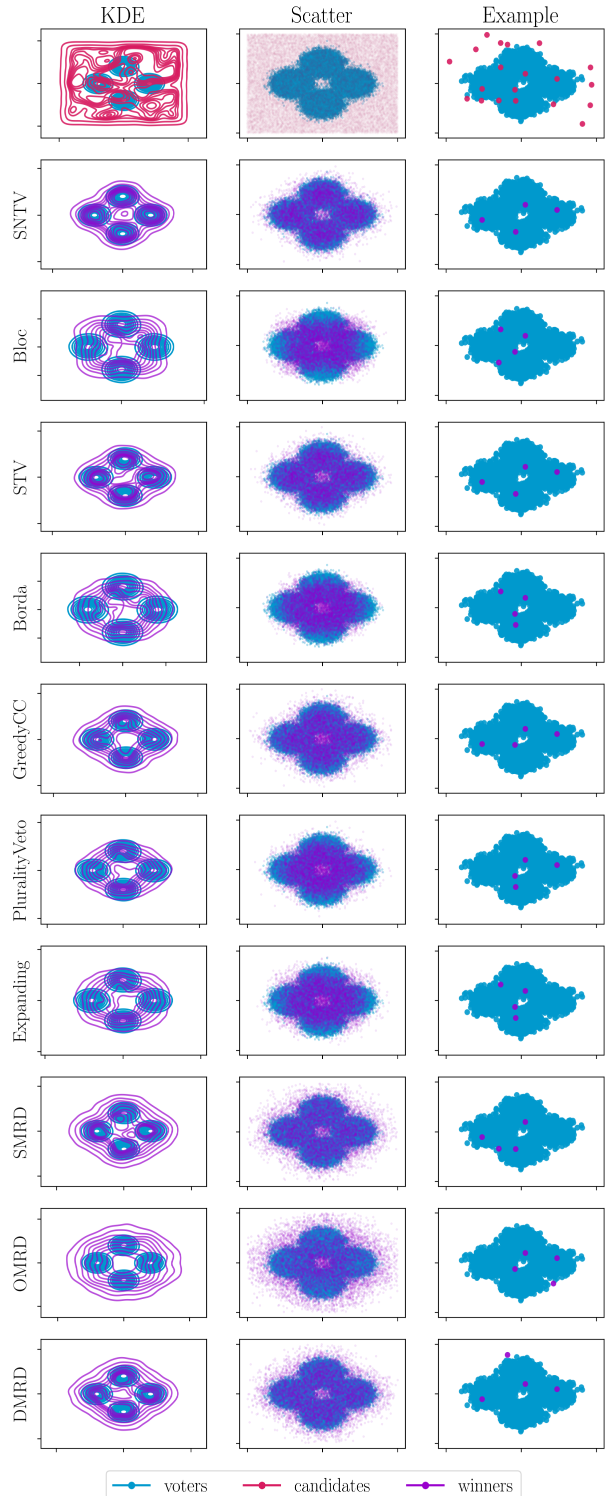

Figure 8: Experiment 5 (four blocs of equal size, 10,000 voters, 4 winners).

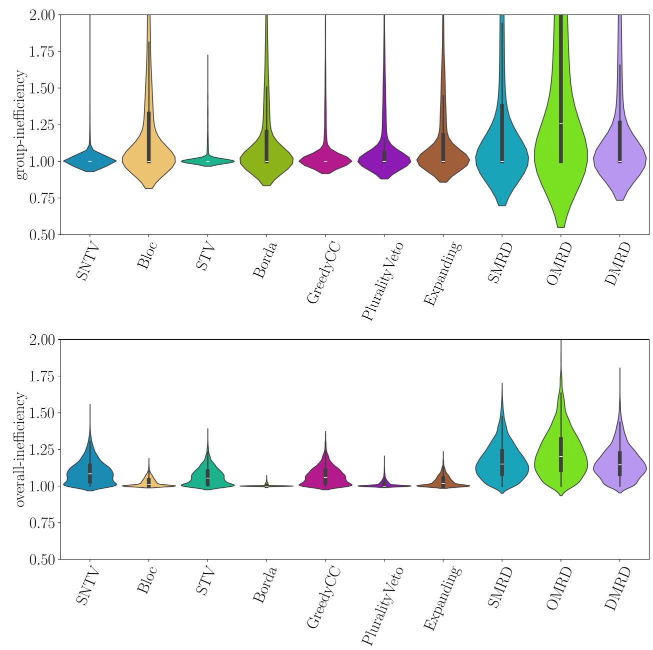

Figure 9: Experiment 5 (four blocs of equal size, 10,000 voters, 4 winners).

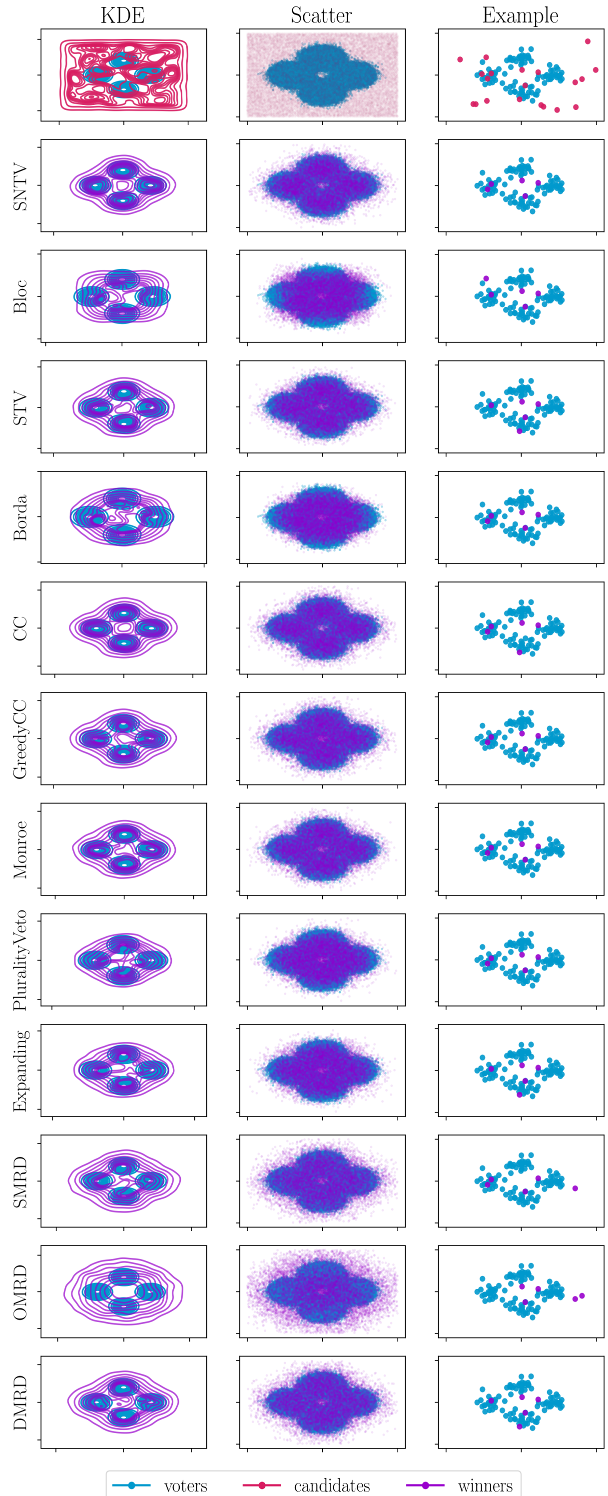

Figure 10: Experiment 6 (four blocs of equal size, 100 voters, 5 winners).

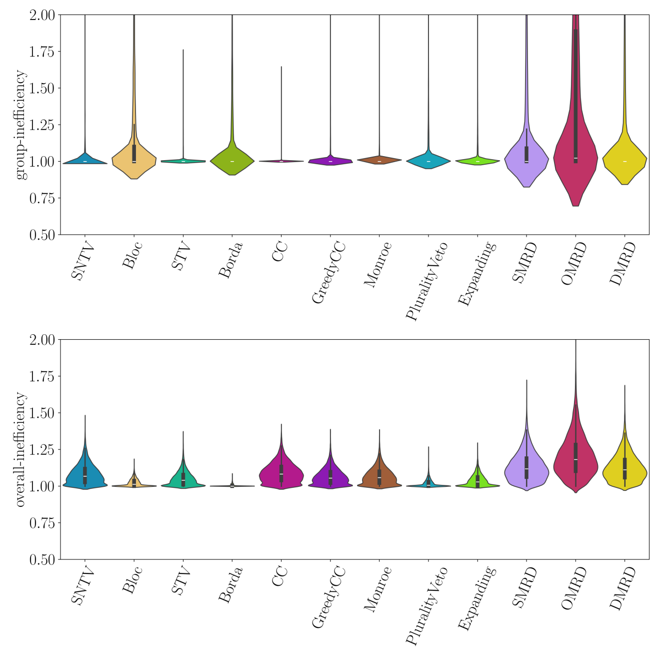

Figure 11: Experiment 6 (four blocs of equal size, 100 voters, 5 winners).

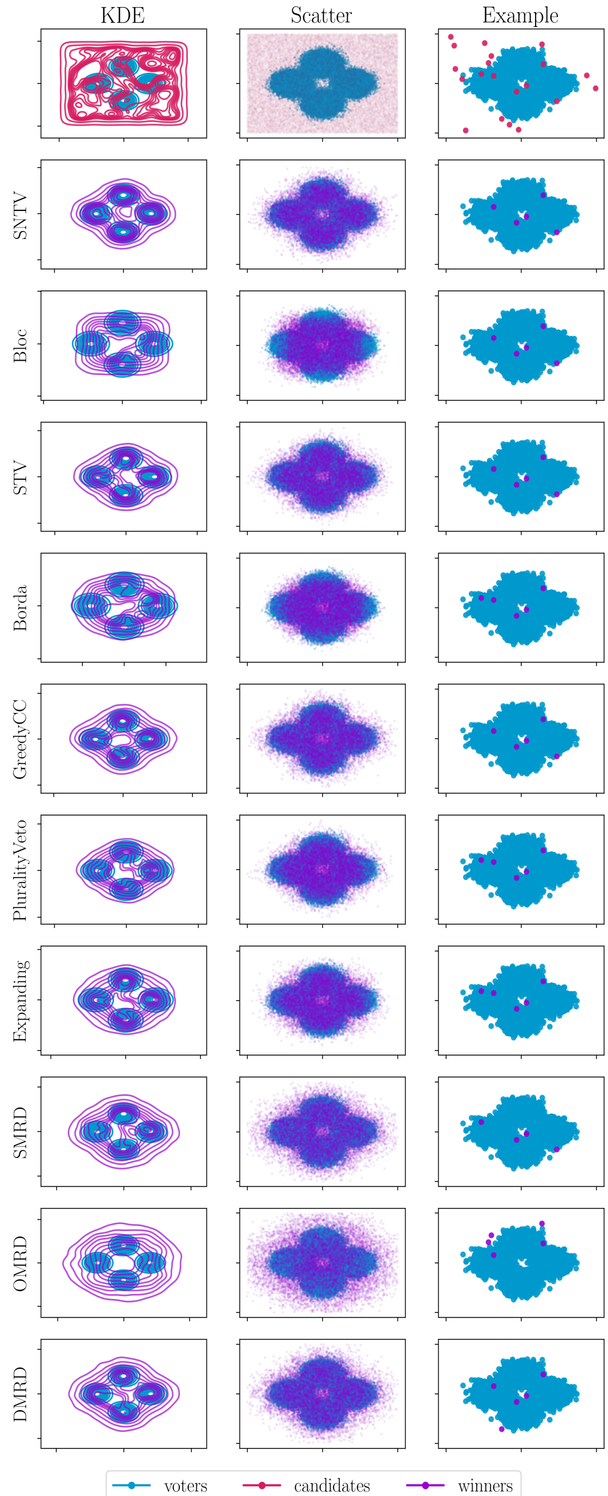

Figure 12: Experiment 7 (four blocs of equal size, 10,000 voters, 5 winners).

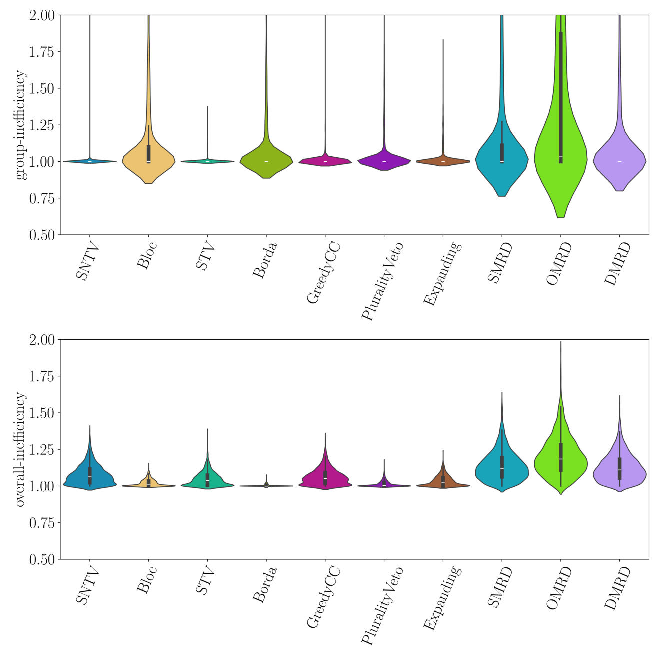

Figure 13: Experiment 7 (four blocs of equal size, 10,000 voters, 5 winners).