

Relating Metric Distortion and Fairness of Social Choice Rules

Extended Abstract*

Ashish Goel
Stanford University
ashishg@stanford.edu

Reyna Hulett
Stanford University
rmhulett@stanford.edu

Anilesh K. Krishnaswamy
Stanford University
anilesh@stanford.edu

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory and mechanism design**; *Facility location and clustering*; Random projections and metric embeddings;

KEYWORDS

social choice, metric distortion, fairness

ACM Reference Format:

Ashish Goel, Reyna Hulett, and Anilesh K. Krishnaswamy. 2018. Relating Metric Distortion and Fairness of Social Choice Rules: Extended Abstract. In *NetEcon '18: Economics of Networks, Systems and Computation, June 18, 2018, Irvine, CA, USA*. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3230654.3230658>

Voting is at the core of many important decisions, not least the election of government officials. Voting, or social choice, theory studies the aggregation of agent preferences into a single collective decision via a social choice rule. Usually, each agent submits a ranking over the candidates, and a social choice rule maps any instance of these preferences to one or more alternatives. But how do we choose the social choice rule? Traditionally, these rules have been evaluated based on simple axiomatic criteria—e.g., the majority winner criterion—but unfortunately even small sets of axioms can be impossible to satisfy simultaneously.

A more recent approach goes beyond axioms to take a *utilitarian* view of social choice. In this model, although agents express their preferences via an ordinal ranking, they are assumed to have latent *cardinal* preferences over the alternatives. In the *metric distortion* model [1], we assume that the agents and alternatives lie in an arbitrary, unknown metric space, and an agent’s cost for an alternative equals the distance between the two. With only this simple assumption, it has been shown that some social choice rules, like Copeland, can achieve a constant factor *distortion*—that is, the worst-case ratio of the social cost of the chosen candidate to that of the optimal candidate chosen omnisciently [1].

Within the metric distortion framework, we can also ask how quantitatively “fair” a social choice rule is. To this end, we consider the *fairness ratio* [2], which is a simultaneous bound, over all k , on the ratio of the k largest agent costs for the chosen candidate

versus the optimal candidate in the unknown metric. This notion of fairness, based on approximate majorization, generalizes both total cost minimization (when k equals the total number of agents), and max-min fairness (when $k = 1$) [2]. The fairness ratio also translates to an approximation guarantee on a wide class of convex objective functions [2].

In our work, we demonstrate that these notions—the distortion and fairness ratio—are in fact closely linked. The fairness ratio is a stronger notion than distortion, so it is always at least as large. However, we show that the fairness ratio can only exceed the distortion by at most an additive factor of 2 (and that this gap is tight). In essence, we use the triangle inequality to argue that if the worst-off k agents are far away from the chosen candidate (compared to the optimal candidate), the remaining agents can’t be too close either. We consider the “diamond” formed by: the distance from the chosen candidate to the worst-off k agents (measuring fairness); from the worst-off agents to the optimal candidate; from the optimal candidate to another, arbitrary agent; and from that agent to the chosen candidate (measuring distortion). When comparing the fairness and distortion, the additive factor of 2 arises from the two sides of the diamond connected to the optimal candidate.

This result immediately gives bounds on the fairness ratio of several social choice rules, such as STV, for which only bounds on distortion were previously known. While the difference between the distortion and fairness ratio can still be relevant when the distortion is a small constant, it can be argued that this close relationship serves to further validate distortion as a useful measure of the quality of social choice rules.

Finally, we consider a relatively less well-studied question in the metric distortion model: that of multi-winner social choice rules. Choosing multiple alternatives can be valuable for applications from electing committees to choosing where to build fire stations. Specifically, we focus on the setting where agents’ preferences are additive—the total cost for an agent is just the sum of her costs over the chosen set of alternatives. In this setting, we establish that any single-winner rule can be recursively applied to obtain a multi-winner rule with the same bound on distortion. The fairness ratio does not generalize from single-winner to multi-winner rules so directly, but by extending our main result (that the fairness ratio is at most the distortion plus 2) to multi-winner rules, we demonstrate that a constant distortion *and* fairness ratio can be achieved, for instance, by the recursive Copeland rule.

REFERENCES

- [1] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. 2015. Approximating Optimal Social Choice under Metric Preferences. *Association for the Advancement of Artificial Intelligence, 15th Conference of the* (2015).
- [2] Ashish Goel, Anilesh K Krishnaswamy, and Kamesh Munagala. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 287–304.

*The full version of this paper is available at http://web.stanford.edu/~anilesh/publications/GoHuKr_Relating.pdf