1 Introduction

We study social choice in the **metric distortion** framework, where voters and candidates are embedded in an unknown metric space, and voters prefer candidates that are closer to them. A social choice rule sees only the induced rankings and must pick a single winner; its distortion is the worst-case ratio between the social cost, given by the sum of distances to voters, of the chosen winner and that of the optimal candidate. It is well known that any deterministic rule that only uses *ordinal rankings* (e.g., voter v ranks candidate A above candidate B) has distortion at least 3, and deterministic tournament rules face a stronger lower bound of around 3.11.

A natural way to go beyond these limits is to enrich the information available to the rule. Ordinal rankings exhaust all *qualitative* information (for instance, whether voter v prefers A over B) but do not touch upon *quantitative* information (for instance, how strongly v prefers A over B). On the other hand, it is practically impossible to ask voters themselves to reveal precise cardinal utilities: it is unrealistic to expect a voter to answer a question such as "do you prefer candidate A at least $\sqrt{2}$ times as much as candidate B?" A natural workaround, therefore, is to ask questions that act as valid proxies to extract additional cardinal information. In this work, we focus on **two-person deliberation**: for a pair of voters (u, v) and a pair of candidates (A, B), comparing $d(u, A) + d(v, A)$ with $d(u, B) + d(v, B)$ reveals coarse geometric information about the latent metric, essentially indicating on which side of the bisector of (A, B) the pair (u, v)'s barycenter lies. We describe the exact procedure to form such pairwise deliberations later.

This project focuses on the **deliberation-via-matching** protocol and its geometric analysis. On a high level, the write-up is structured as follows.

- We first review the metric and tournament preliminaries in Section 2, and describe the deliberation-via-matching protocol in Section 3. The protocol involves *parsimonious* pairwise deliberations, where each voter participates in at most one deliberation for each candidate. In most real-world voting instances, where the number of candidates m is small compared to the number of voters n, this per-voter overhead is modest.

- In Section 4, we apply the deliberation-via-matching protocol to the two-candidate setting. We show that the deterministic deliberation-via-matching rule achieves distortion 2, which is optimal among *any* deterministic rule that has access to ordinal rankings and outcomes of *any* pairwise deliberations. We then introduce a randomized variant for two candidates and show that it obtains distortion at most 1.53, close to the $3/2$ lower bound, again, for *any* randomized rule that reads ordinal rankings and outcomes of pairwise deliberations.

- In Section 5, we extend this protocol to general instances with m candidates via a weighted *uncovered-set* tournament rule, and show that deliberation-via-matching achieves an overall distortion of 3. This is significant in three ways: (i) it breaks a previously known lower bound of 3.11 for tournament rules that only use ordinal rankings, and (ii) it conceptually shows that tournament rules are just as powerful as general deterministic rules (which are lower bounded by 3) given *minimal* additional cardinal information, and finally, (iii) to our best knowledge, unlike previous literature in this line of work, our proof is the first to be analytically tractable and is *governed by clean geometric intuitions* throughout.