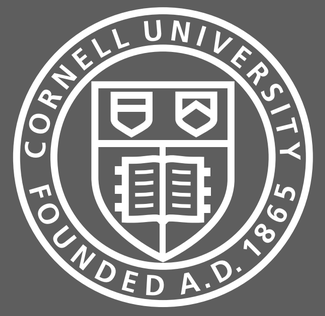
Implicit Utilitarian Voting II

Paul Gölz

Lecture 18, ORIE 4360/5360, Fall ’26



Implicit Utilitarian Voting

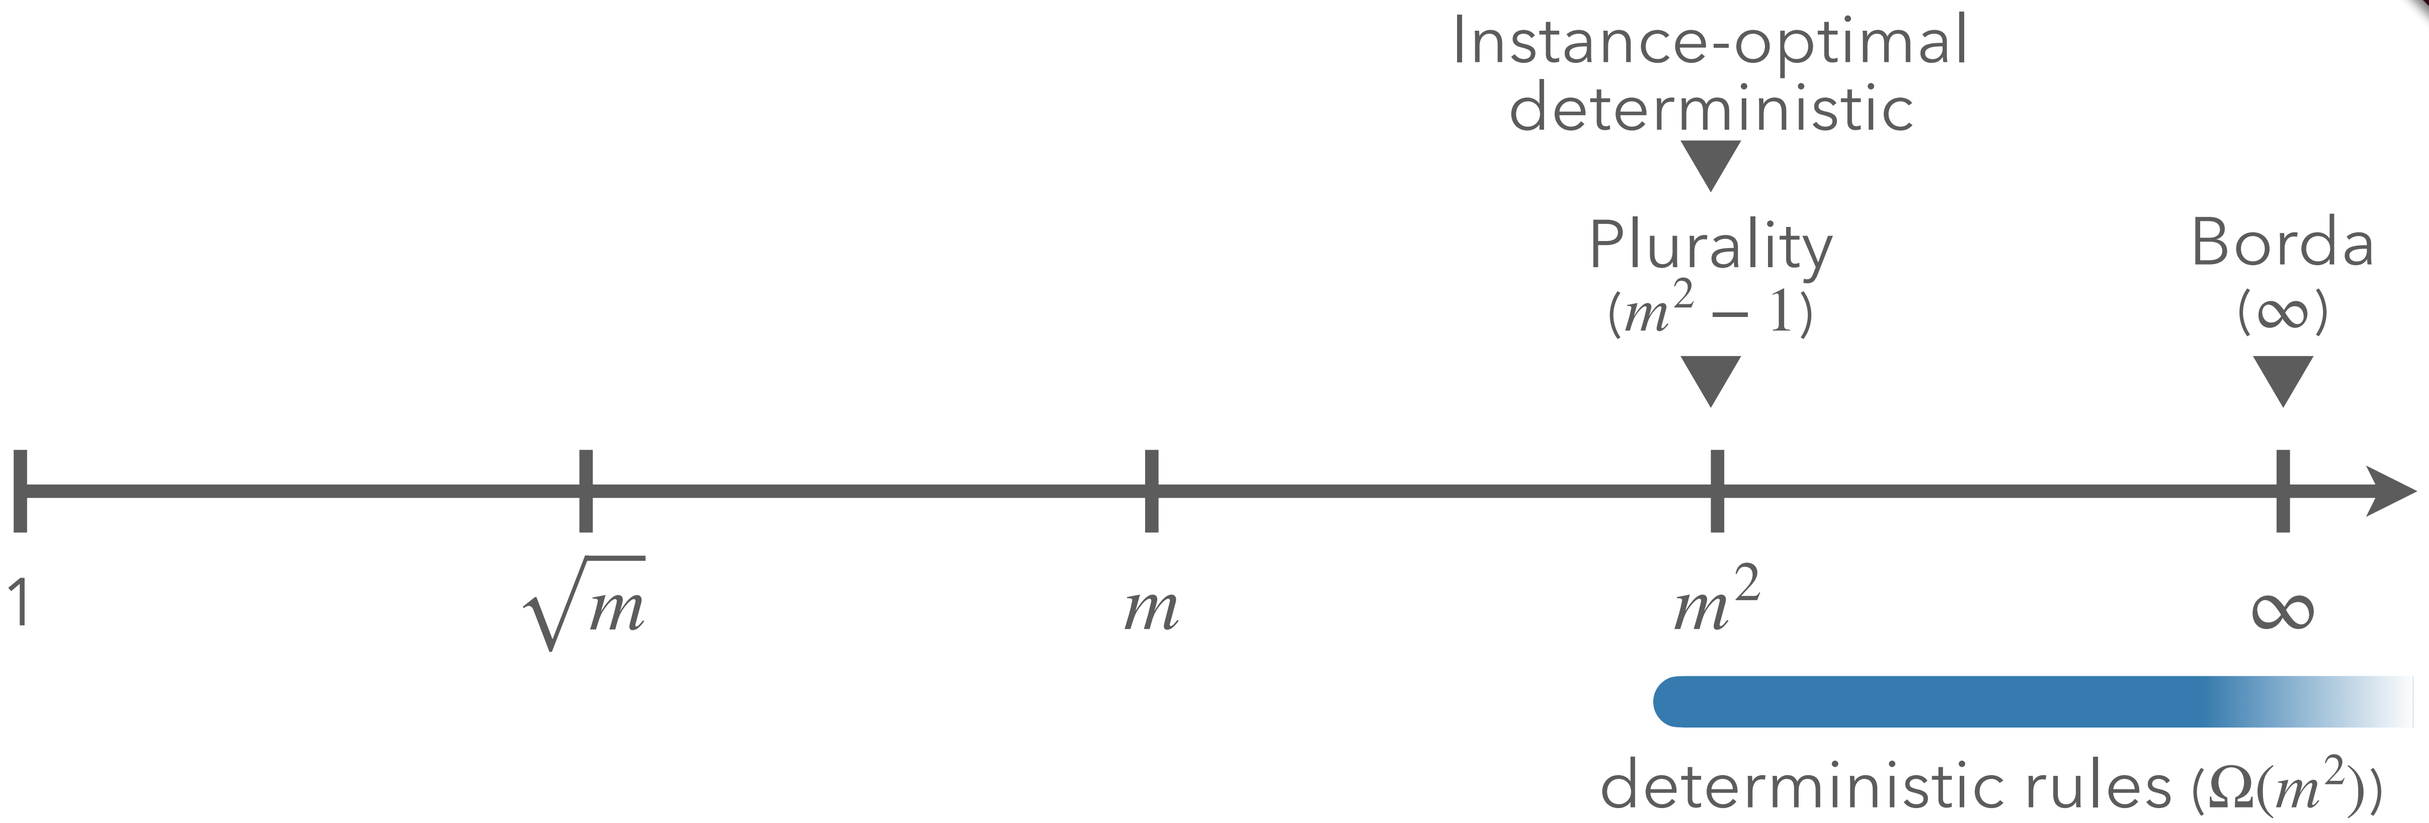
Conceptually: want to maximize utilitarian welfare, but only know the rankings not the utilities ("implicit utilitarian voting").

$$\mathrm{dist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x \in M} \mathrm{SW}(x, \boldsymbol{u})}{\mathrm{SW}(f(\boldsymbol{\sigma}), \boldsymbol{u})}.$$

Supremum means we take worst-case utilities, which make ratio as high as possible. But definition restricted to nonnegative utilities, that add up to 1 for each voter.

Measures how much better we could have done than $f(\boldsymbol{\sigma})$. Ratio at least 1, could be ∞ if denominator is zero.

Distortion Landscape



Can Randomization Help?

Let's allow social choice functions f to be randomized: $f : \mathscr{L}^n \to \Delta(M)$.

I.e., for each preference profile $\boldsymbol{\sigma}$, $f(\boldsymbol{\sigma})$ assigns probability masses to alternatives.

Randomness can "hedge" against lopsided utility profiles: we still have some chance of choosing a high-welfare alternative.

Distortion for Random Voting Rules

Can naturally define distortion of randomized social choice function f:

$$\mathrm{dist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x \in M} \mathrm{SW}(x, \boldsymbol{u})}{\mathbb{E}_{x \sim f(\boldsymbol{\sigma})}[\mathrm{SW}(x, \boldsymbol{u})]},$$

and, as before,

$$\mathrm{dist}(f) := \sup_{\boldsymbol{\sigma} \in \mathscr{L}^n} \mathrm{dist}(f, \boldsymbol{\sigma}).$$

Distortion of Random Alternative

$f_{\text{rand-alt}}(\boldsymbol{\sigma})$ assigns mass $1/m$ to each alternative.

Let $\boldsymbol{\sigma}$ be a preference profile with n voters, m alternatives, in which each voter gives the same ranking $a \succ b \succ c \succ \cdots$.

$$\text{dist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{u} \rhd \boldsymbol{\sigma}} \frac{\max_{x \in M} \text{SW}(x, \boldsymbol{u})}{\mathbb{E}_{x \sim f(\boldsymbol{\sigma})}[\text{SW}(x, \boldsymbol{u})]}.$$

What is the distortion of $f_{\text{rand-alt}}$ on $\boldsymbol{\sigma}$?

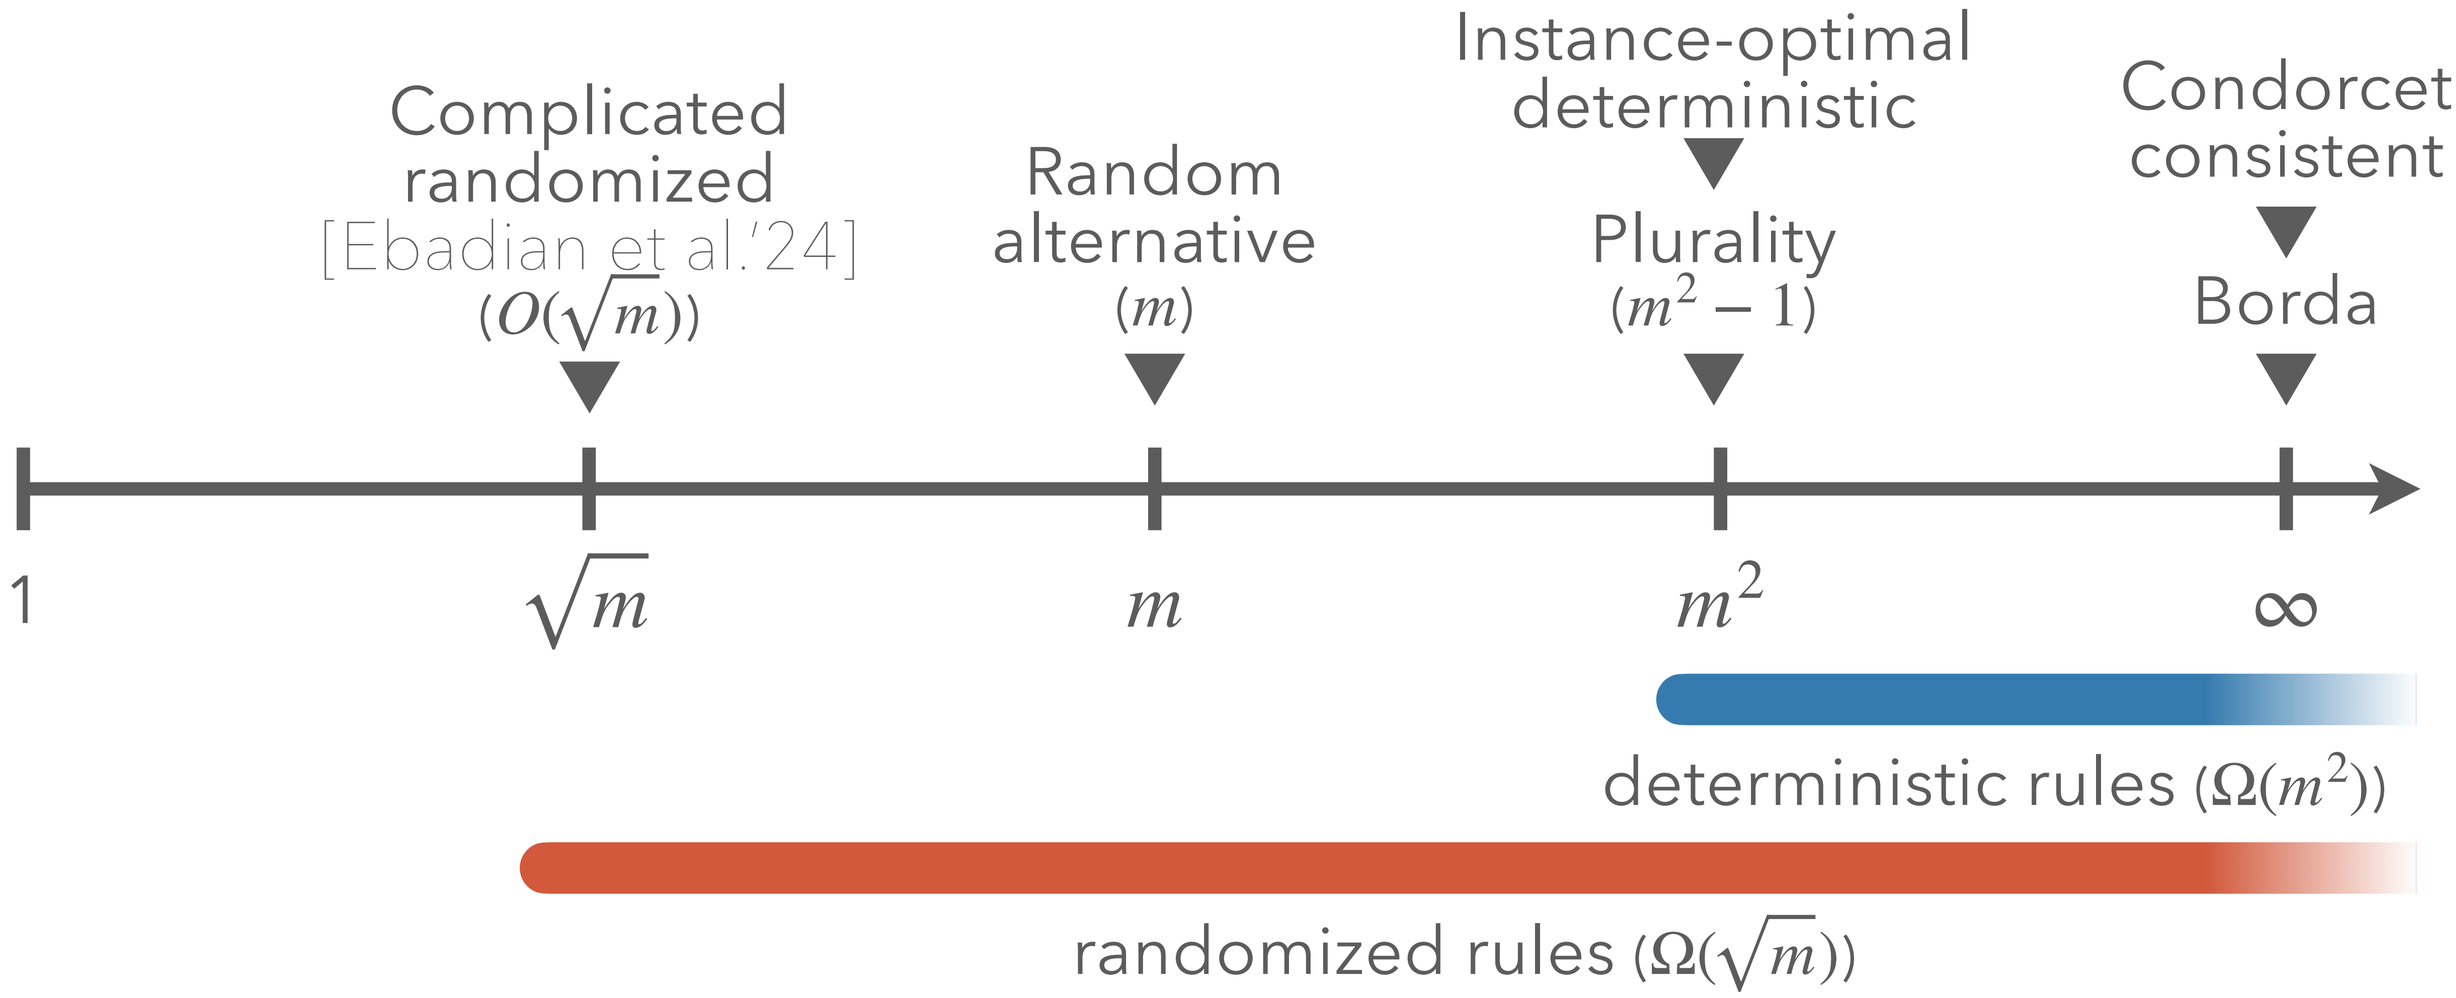
pollev.com/goelz

Overall Distortion of $f_{\text{rand-alt}}$

In fact, the overall distortion of $f_{\text{rand-alt}}$ is m. [Prove on board.]

So, randomization (even a dumb way of randomizing) helps a lot in reducing distortion. How much better can we get?

Distortion Landscape



Take-Aways?

One thing attractive about distortion: assigns a number to each voting rule, “which is better”?

But the ordering is at odds with which voting rules are good in other settings. (Why?)

Also, distortion is still really high; methods designed for low distortion aren’t that natural.

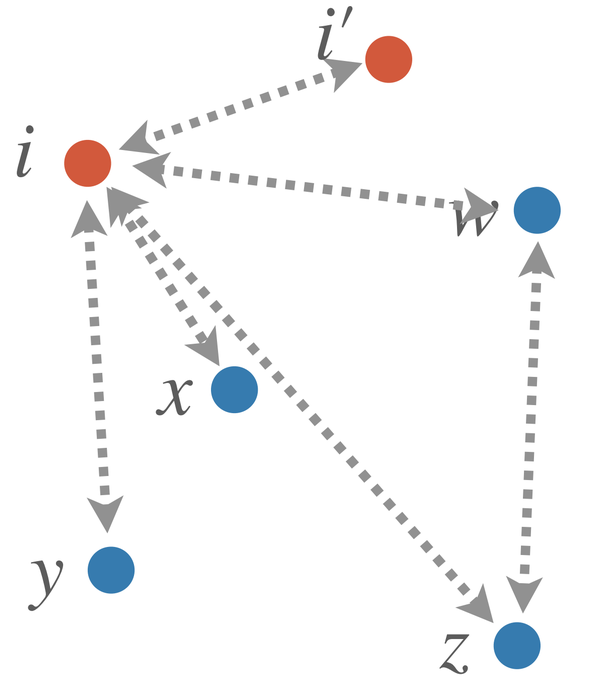
Let’s look at a variant of the distortion setting that’s been more successful along these axes.

Metric Distortion

Let voters i and alternatives x be points in some metric space.

$d(\cdot, \cdot)$ measures distance between points. Assume symmetric and triangle inequality: $d(a, c) \leq d(a, b) + d(b, c)$.

$u_i(x) = -d(i, x)$, so voters prefer closer alternatives.



$x \succ_{\sigma_i} y \succ_{\sigma_i} w \succ_{\sigma_i} z$

Distortion for Costs

Instead of maximizing nonnegative utilities, more natural to minimize total cost

$$\text{cost}(x) = \sum_{i \in N} d(i, x).$$

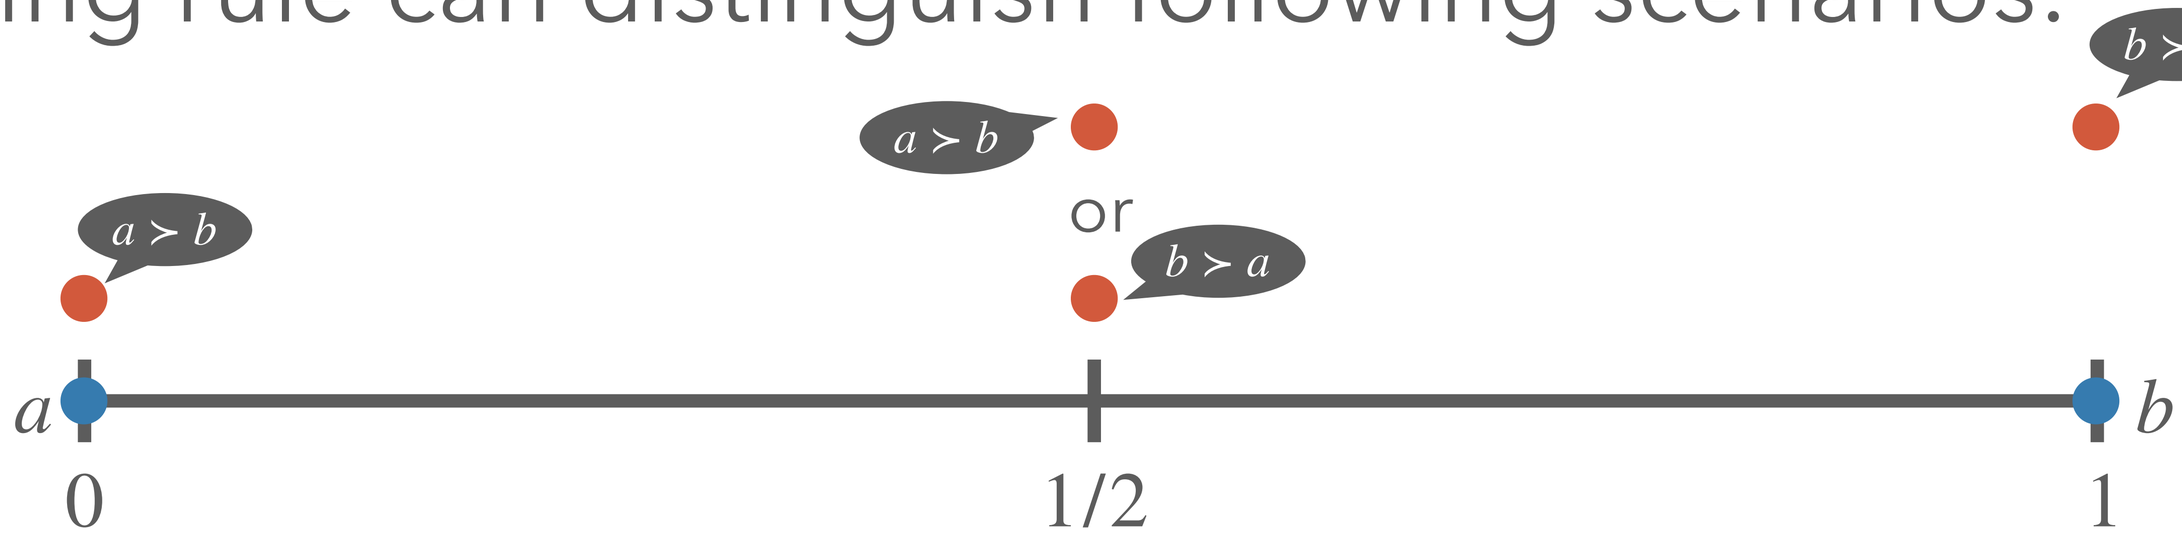
Metric distortion is defined as follows:

$$\text{mdist}(f, \boldsymbol{\sigma}) := \sup_{\boldsymbol{\sigma} \rhd d} \frac{\text{cost}(f(\boldsymbol{\sigma}))}{\min_x \text{cost}(x)}.$$

(As before, $\text{mdist}(f) = \sup_{\boldsymbol{\sigma}} \text{mdist}(f, \boldsymbol{\sigma})$.)

Lower Bound Question

If voters & alternatives on a line, no deterministic voting rule can distinguish following scenarios:



Which lower bound on metric distortion does this imply?

Plurality Veto

New social choice function:

- each alternative starts with as many points as its plurality score
- Go through voters in arbitrary order: voter i selects their lowest-ranked alternative with nonzero points, takes one point away.
- Last alternative to reach zero is the winner.

Plurality Veto

Theorem (Kizilkaya & Kempe ’22)**.** The metric distortion of plurality veto is 3.

[Proof on blackboard.]

Metric Distortion Landscape

Plurality, Borda	$\Theta(m)$
STV	$\Omega(\sqrt{\log m}), O(\log m)$
Copeland	5
Plurality Veto	3

This time, ordering makes some sense, and led us to reasonable new voting rule!

Randomized rules less natural, $> 2.11, < 2.13$.

Bibliography

Boutilier et al.: *Optimal social choice functions*. EC (2012).

Ebadian et al.: *Optimized Distortion and Proportional Fairness in Voting*. TEAC (2024).

Kizilkaya & Kempe: *Plurality Veto*. IJCAI (2022).

Ramakrishnan: Optimal Metric Distortion for Voting – A Proof from the Book. Blog post (2023).