Fair Metric Distortion for Matching with Preferences

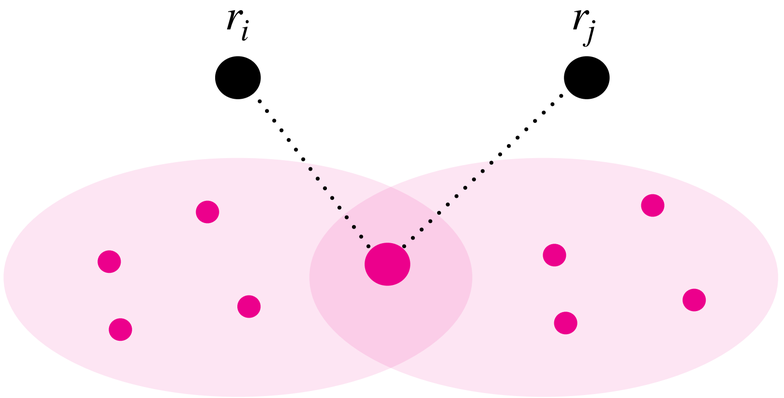

Jabari Hastings and Prasanna Ramakrishnan

Stanford University

Matching is Fundamental

- Matching problems are plentiful. We can use matchings to assign…



Students to Schools



Instructors to Courses

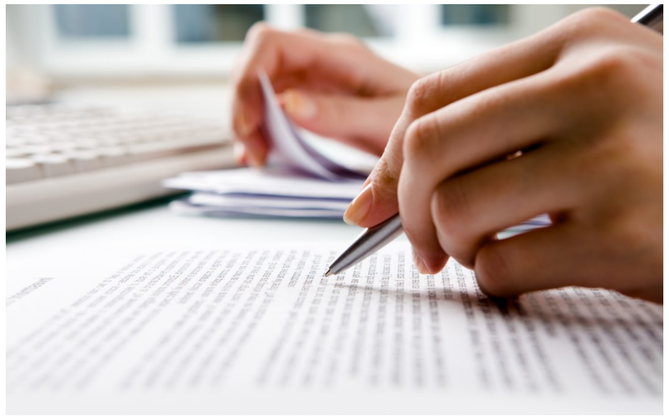

Reviewers to Papers

…

- Settings often involve **partial information** about participants' preferences (e.g., via ordinal rankings).
- Yet, we might want a good fit in terms of the *unknown* **underlying values**.
- There is often a **spatial** way to reason about preferences.

Can we leverage this spatial reasoning?

Our Problem

Metric Matching

- There is some metric space d defined on **agents** $A = \{a_1, \ldots, a_n\}$ and **items** $B = \{b_1, \ldots, b_n\}$.
- The **cost** of the item b for agent a_i is given by the distance $d(a_i, b)$.
- **Goal:** We want a perfect matching $M : A \to B$ with low cost (e.g. minimizing the **sum** of assigned costs)
- **Caveat:** We do not know the metric! We are only given partial information based on agent rankings of items.

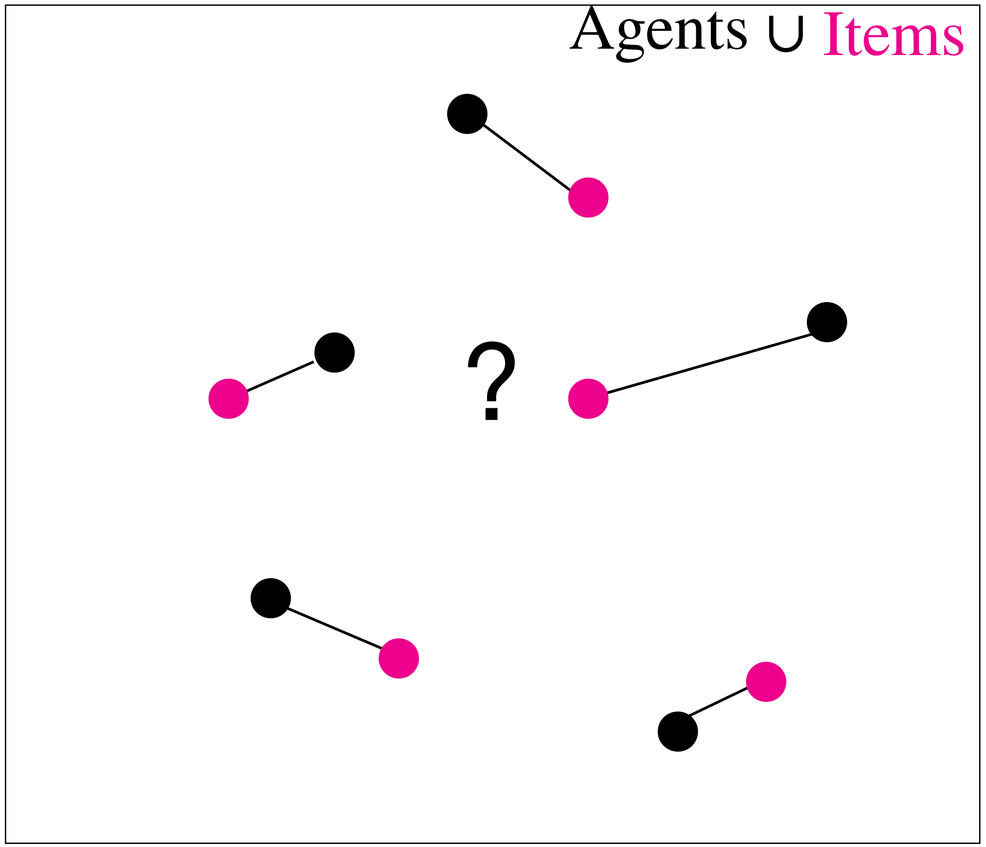

Our Problem

Metric Matching with Preferences

- **Agents** $A = \{a_1, \ldots, a_n\}$ and **Items** $B = \{b_1, \ldots, b_n\}$.
- We are given a **preference table** σ that consists of each agent's ranking over all the items.
- **Goal:** We want a matching M such that for *any consistent metric* d, the matching M is a good approximation of the optimal matching M^*:

$$\text{cost}(M, d) \leq \alpha \cdot \text{cost}\left(M^*, d\right) \quad \forall d \in \rho(\sigma)$$

- The approximation factor α is called the **distortion**. *Smaller factors are more preferable.*

$$\begin{aligned} a_1 &: b_1 \succ b_2 \succ b_3 \succ b_4 \\ a_2 &: b_2 \succ b_1 \succ b_3 \succ b_4 \\ a_3 &: b_4 \succ b_3 \succ b_2 \succ b_1 \\ a_4 &: b_4 \succ b_3 \succ b_1 \succ b_2 \end{aligned}$$

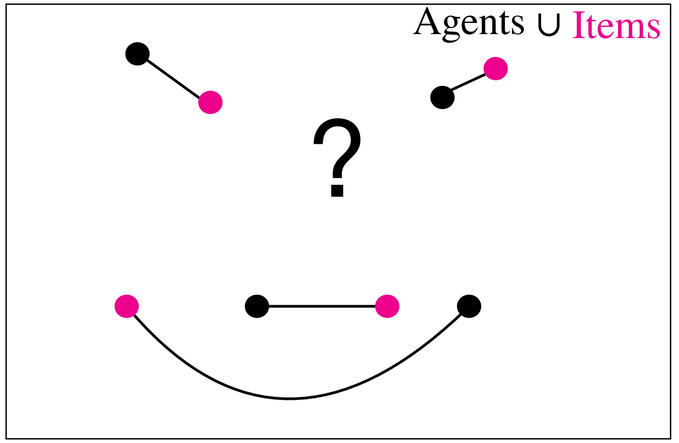

You might be asking…

- **Why the metric assumption?** Preferences are often conveyed along some continuum (e.g., line, plane).

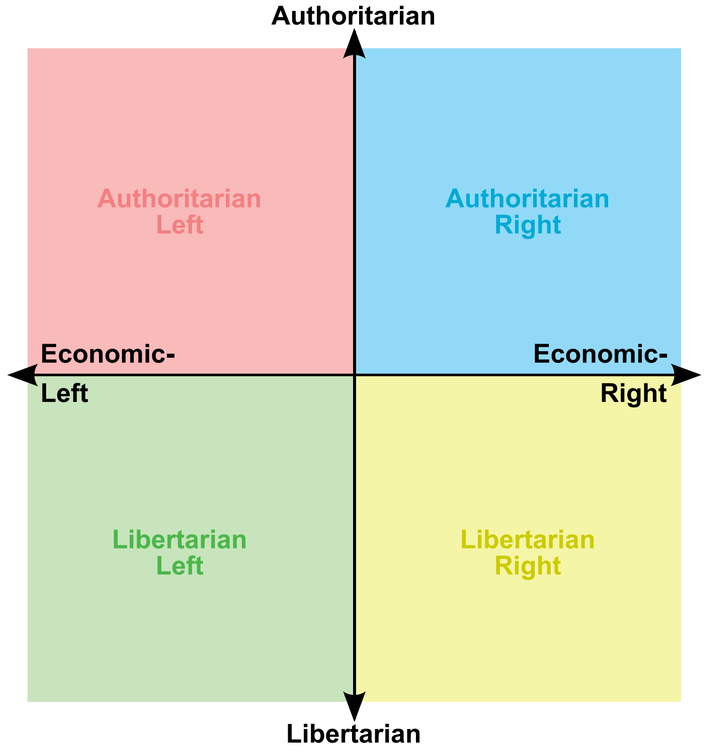

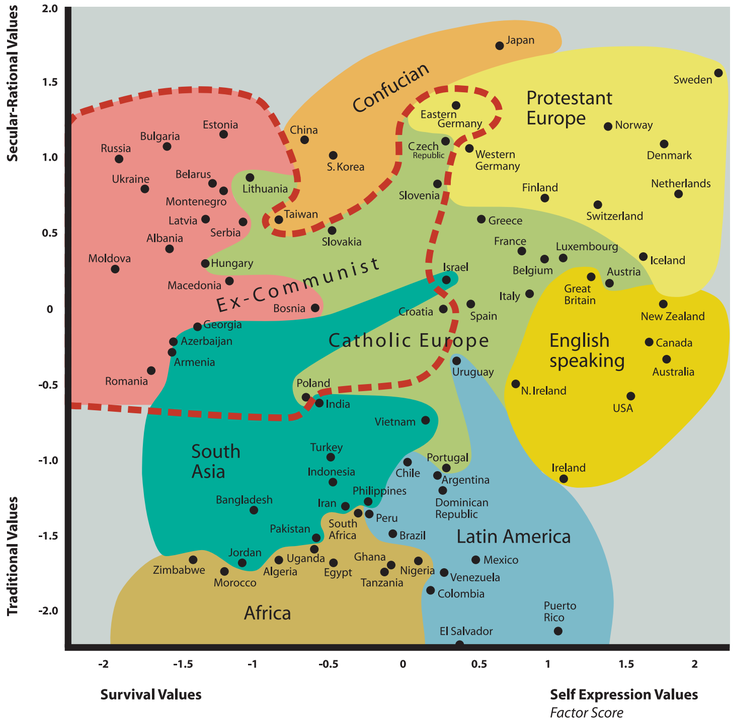

- **Can we reasonably approximate over *all* metrics?** Yes, often quite well!

By Koyos - Own work, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=5459884
By Original: AnonymousTranslation: Zemxer (Türkçe) and OmegaFallon - This file was derived from: Political Compass yLR wo text.svg:, Public Domain, https://commons.wikimedia.org/w/index.php?curid=79649927

Metric Distortion

A Sandbox for Mechanisms

The metric distortion lens has been heavily used in the voting world:

- **Constant** distortion is possible, even for deterministic voting rules.[ABP15, GKM17, MW19, Kem20a, GHS20, KK22, KK23]
- Novel voting rules have been created in the process.[MW19, KK22,KK23, CRWW24]
- Effects of **randomization** and small-group **deliberation** have been studied.[CR22, CRWW24, GS24, GGM25]

The matching literature is far more limited. Can we get similar success stories?

Distortion Bounds

Matching Mechanisms

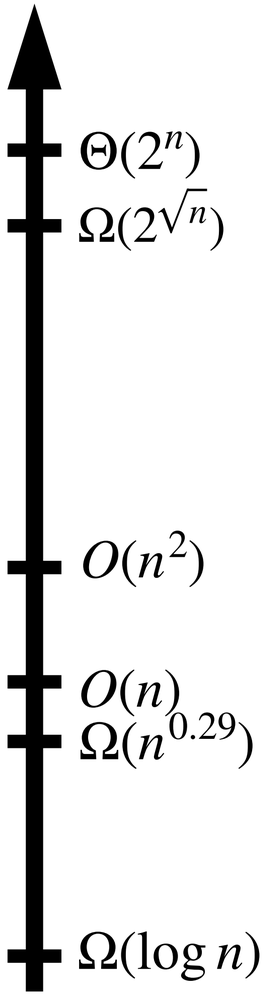

$\Theta(2^n)$ **Serial Dictatorship (SD)** [CFRF+16]

$\Omega(2^{\sqrt{n}})$ **Boston Mechanism** [ACR 23]

$O(n^2)$ **RepMatch Mechanism** [ACR 23] *only deterministic alg with **poly** distortion*

$O(n)$

$\Omega(n^{0.29})$ **Random Serial Dictatorship (RSD)** [CFRF+16]

$\Omega(\log n)$ ***Any* Mechanism** [ACR 23]

Does low distortion imply *fair* outcomes for agents?

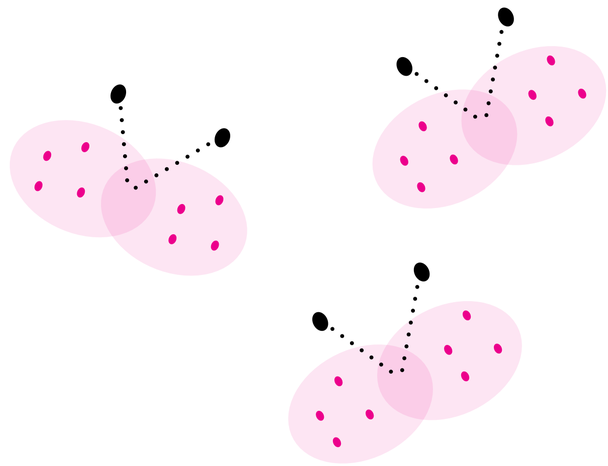

Fair Distortion

Evaluating Outcomes

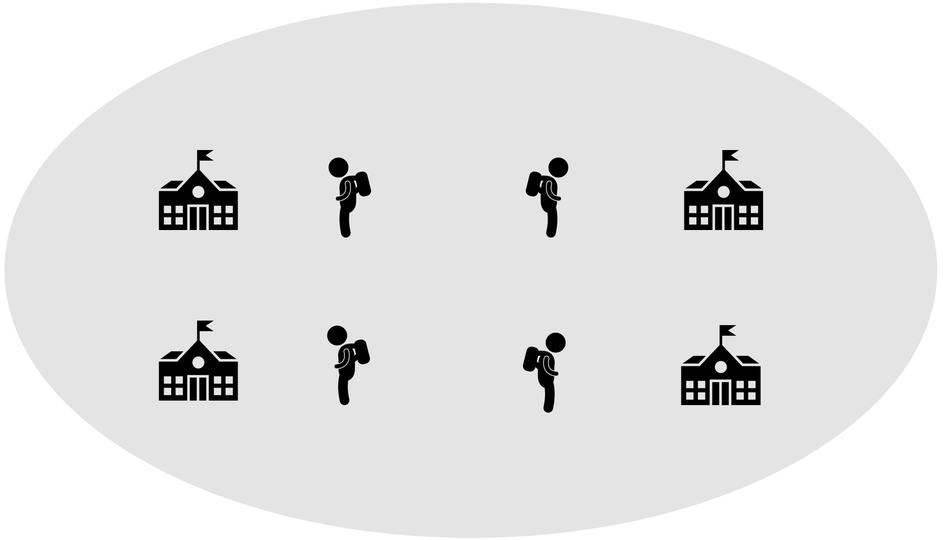

All n children have 1-minute commutes

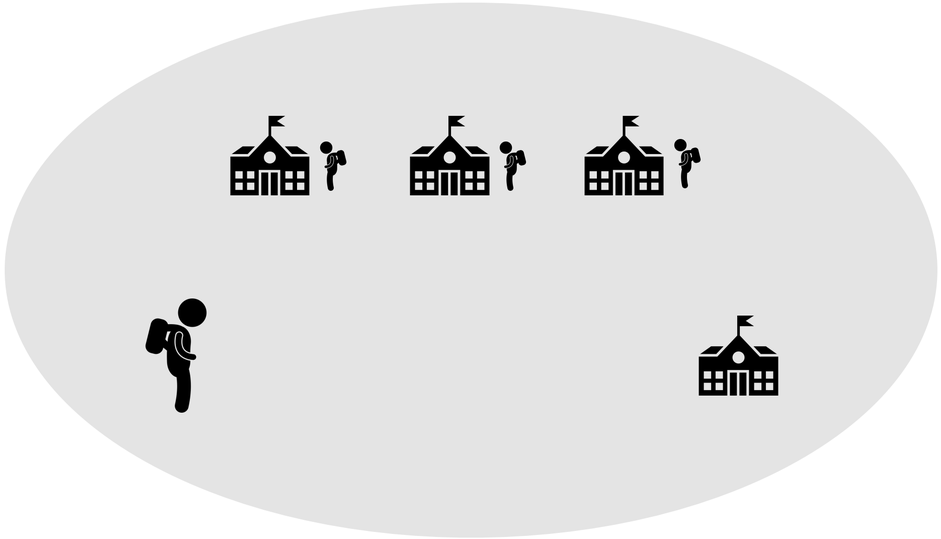

1 child has an n-minute commute

Others have an ϵ-minute commute

- Two outcomes could have same **utilitarian cost** (sum), but vastly different **egalitarian cost** (max).
- Ideally, our matching should have low distortion under both **sum** and **max** objectives.

Fair Distortion

Evaluating Outcomes

The **cost** of a matching M is some function $f : \mathbb{R}^n \to \mathbb{R}$ of induced agent costs (e.g., **sum** or **max**)

$$\text{cost}_f(M) := f\left(d(a_1, M(a_1)), d(a_2, M(a_2)), \ldots, d(a_n, M(a_n))\right)$$

Definition (Fairness Ratio) [GKM17]**.**

A matching M has **fairness ratio** α if $\text{distortion}_f(M) \leq \alpha$ for every **monotone symmetric norm** f.

- For any **metric voting** mechanism, the **sum**-distortion and **max**-distortion differ by a constant [GHK18].
- For any **metric matching** mechanism, does low **sum**-distortion imply low **max**-distortion?

Fair Distortion

Sum vs Max

By standard norm arguments, the max-sum discrepancy is at most a factor of n. *Does it have to be this large?*

Proposition (Max-Sum discrepancy can be linear).

(1) A matching mechanism can have **sum**-distortion α and **max**-distortion $\Omega(\alpha n)$ in some metric space.

(2) A matching mechanism can have **max**-distortion α and **sum**-distortion $\Omega(\alpha n)$ in some metric space.

Thus, it is also worthwhile to study the **max** objective.

Sum-Distortion Bounds

Matching Mechanisms

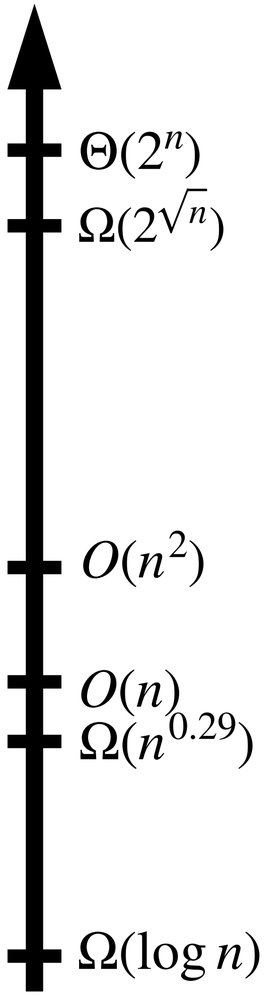

$\Theta(2^n)$ **Serial Dictatorship (SD)** [CFRF+16]

$\Omega(2^{\sqrt{n}})$ **Boston Mechanism** [ACR 23]

$O(n^2)$ **RepMatch Mechanism** [ACR 23]

*Does **max**-distortion become $\Omega(n^3)$?*

$O(n)$

$\Omega(n^{0.29})$ **Random Serial Dictatorship (RSD)** [CFRF+16]

$\Omega(\log n)$ ***Any* Mechanism** [ACR 23]

Our Results

Result 1 (RepMatch *improves* under max-objective)

(1) The RepMatch mechanism has **max**-distortion $O(n^{1.58})$.

(2) Furthermore, its distortion with respect to any **monotone symmetric norm** is at most $O(n^2)$.

Result 2 (New distortion lower bounds)

(1) Any deterministic mechanism must have **max**-distortion at least 7.

(2) The RepMatch mechanism has **max**-distortion $\Omega(n)$ and **sum**-distortion $\Omega(n^2)$.

Distortion of RepMatch

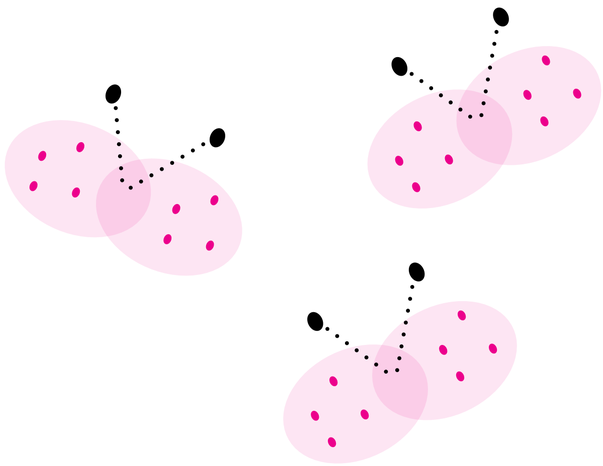

RepMatch

Basic Intuition

If two agents have similar favorite items, then a good item for one is likely good for the other…

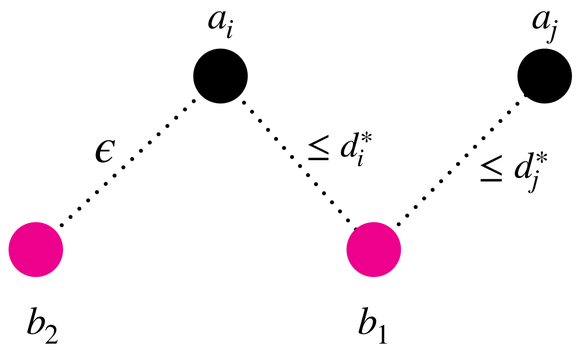

$a_i : b_1 \succ b_2$

$a_j : b_1 \succ ?$

*… so we allow one agent to act as a **representative** for the other agents who have similar favorite items.*

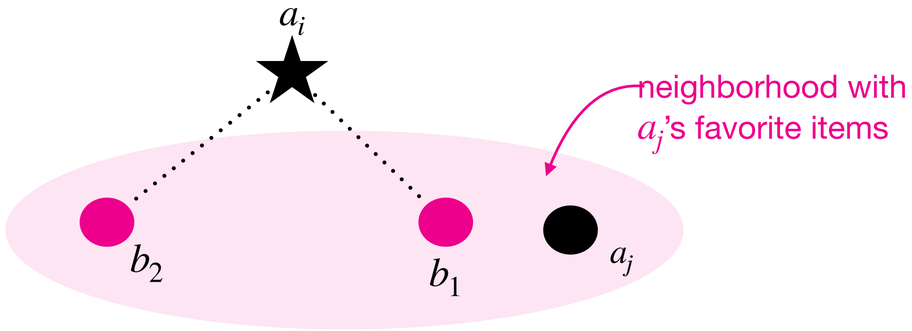

*If r_i represents k agents, then we try to assign her k **favorite items** to the coalition…*

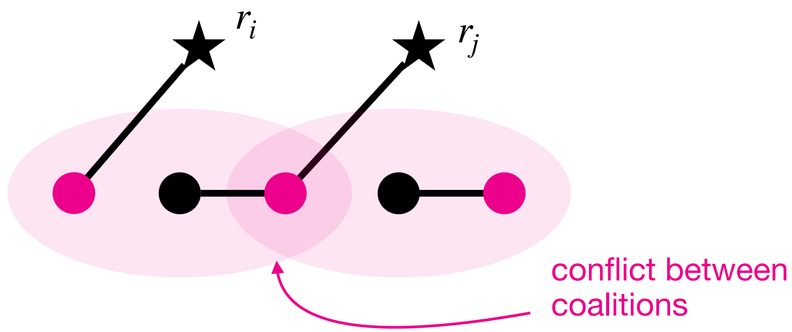

*…if we can't do this, then we **merge** coalitions and choose a new representative.*

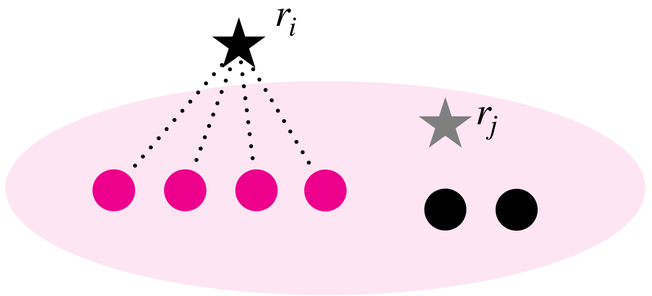

Agents ∪ Items

RepMatch

Formal Description

Algorithm (RepMatch).

Initialize representatives $(r_i, S_i)_{i \in [n]}$, where $S_i = \{r_i\}$

While there exist sets S_i, S_j such that some item is both one of r_i's favorite $|S_i|$ items and one of r_j's favorite $|S_j|$ items:

- Merge S_i and S_j.
- The new representative for $S_i \cup S_j$ is r_i if $|S_i| \geq |S_j|$ and r_j otherwise.

Match the agents in each S_i to the favorite $|S_i|$ choices of r_i.

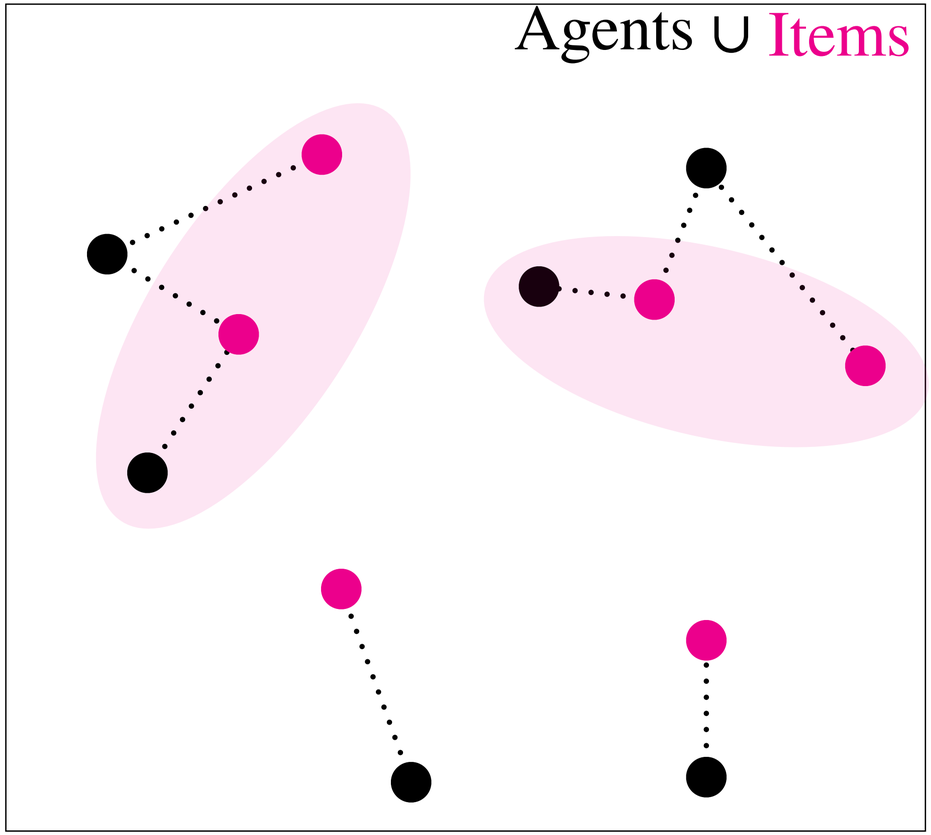

RepMatch

Distortion Guarantees

Theorem [ACR ’23].

The RepMatch mechanism guarantees **sum-**distortion $O(n^2)$.

Our Result (RepMatch *improves* under max-objective)

(1) The RepMatch mechanism has **max**-distortion $O(n^{1.58})$.

(2) Furthermore, its distortion with respect to any **monotone symmetric norm** is $O(n^2)$.

We benefit from a more refined analysis of a key invariant of the algorithm.

RepMatch
Proof Overview

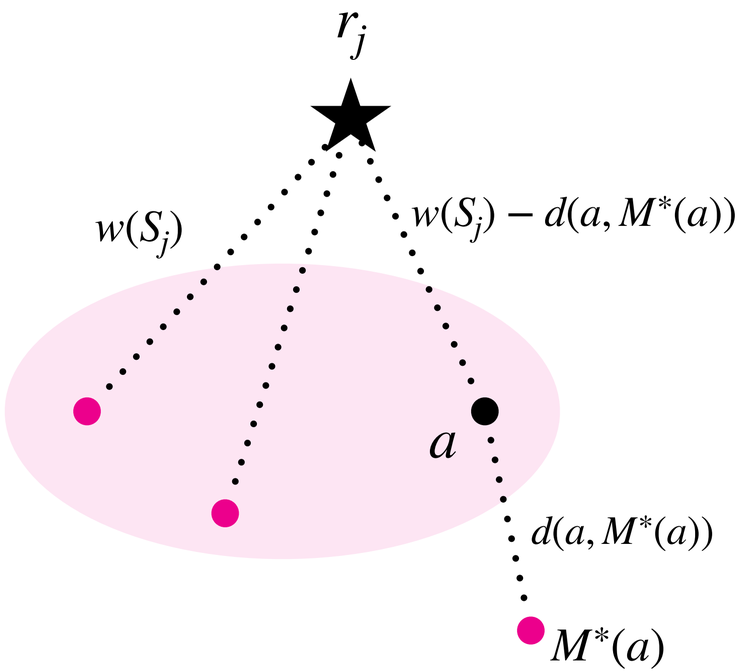

Key Invariant.

There is some "weight" function w such that

$$d(a, r_j) \leq w(S_j) - d(a, M^*(a)) \qquad \forall S_j \ni a, r_j$$

Thus, for every agent $a \in S_j$,

$$\begin{aligned} d(r_j, M^*(a)) &\leq d(r_j, a) + d(a, M^*(a)) \\ &\leq w(S_j)\,. \end{aligned}$$

So there are at least $|S_j|$ distinct items b such that

$$d(r_j, b) \leq w(S_j)\,.$$

In particular, this holds when b is one of r_j's favorite $|S_j|$ items.

RepMatch
Proof Overview

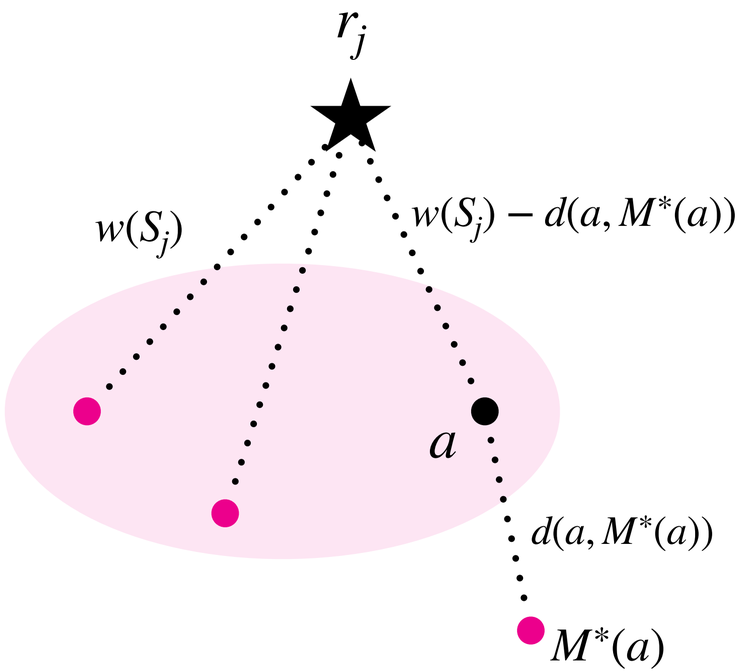

Key Invariant.

There is some "weight" function w such that

$$d(a, r_j) \leq w(S_j) - d(a, M^*(a)) \qquad \forall S_j \ni a, r_j$$

So there are at least $|S_j|$ distinct items b such that

$$d(r_j, b) \leq w(S_j)\,.$$

In particular, this holds when b is one of r_j's favorite $|S_j|$ items.

So if our matching M assigns agent a to any such b, we get

$$d(a, M(a)) \leq 2w(S_j)\,.$$

Bounding $w(S_j)$ w.r.t. M^* and n, we get a **max-**distortion bound.

RepMatch
Proof Overview

- When sets **merge**, we must bound the distance from the **new representative** to the coalition.

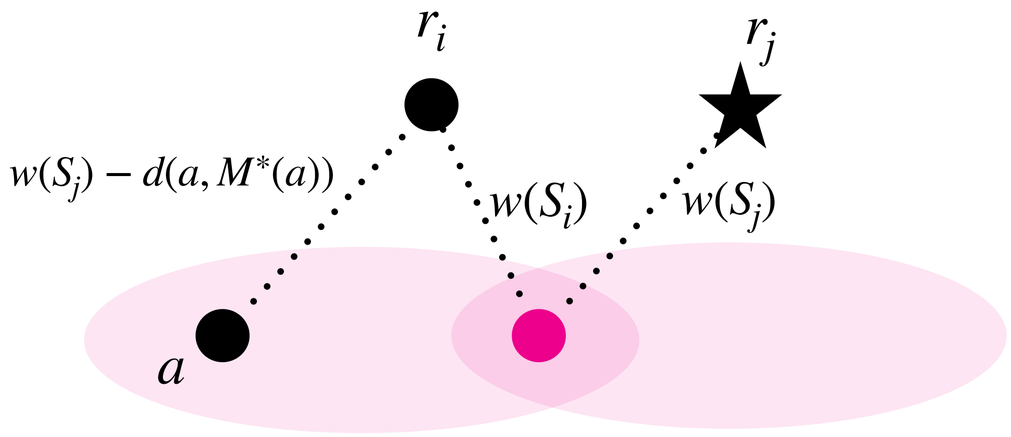

Key Invariant.

There is some "weight" function w such that

$$d(a, r_j) \leq w(S_j) \cdot \text{cost}\left(M^*|_{S_j}\right) - d(a, M^*(a)) \quad \forall S_j \ni a, r_j$$

We need our choice of w to satisfy

$$2w(S_i) \cdot \text{cost}\left(M^*|_{S_i}\right) + w(S_j) \cdot \text{cost}\left(M^*|_{S_j}\right) \leq w(S_i \cup S_j) \cdot \text{cost}\left(M^*|_{S_i \cup S_j}\right)$$

Fact.

Let $f(x) = x^{\log 3}$. If $0 < x \leq y$, then

$$2f(x) + f(y) \leq f(x + y).$$

For **max**, we can set $w(S_i \cup S_j) := 2w(S_i) + w(S_j)$.

We can also show that $w(S) \leq |S|^{\log 3} = |S|^{1.58}$

RepMatch

Summary + Extensions

- For **max objective,** we choose a function w so that for every $a \in S$

$$d(a, M(a) \leq w(S) \cdot \text{cost}(M^*) \leq |S|^{1.58} \cdot \text{cost}(M^*)$$

- For sum **objective,** the same choice of w is not good enough. It would imply

$$\sum_{a \in S} d(a, M(a) \leq |S| \cdot |S|^{1.58} \cdot \text{cost}(M^*)$$

- For **other objectives**, we also need to choose w appropriately to keep scaling bounded above by $|S|^2$.

- We actually analyze the top_k **objective** and use majorization to extend to all **monotone symmetric norms.**

Future Directions

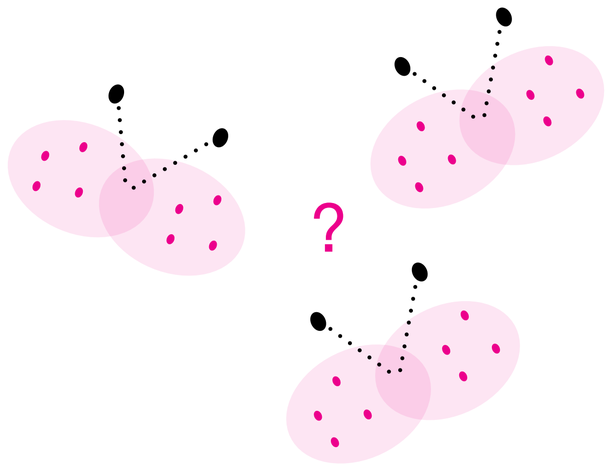

Future Directions

There are many questions to explore:

- Improving **upper/lower bounds** for deterministic mechanisms.
- Is **constant** max-distortion possible?

- What happens if we require **truthfulness**, or allow **randomization**?
- Matching in simpler metric spaces (e.g., in $\mathbb{R}^2$).

Thanks for listening!

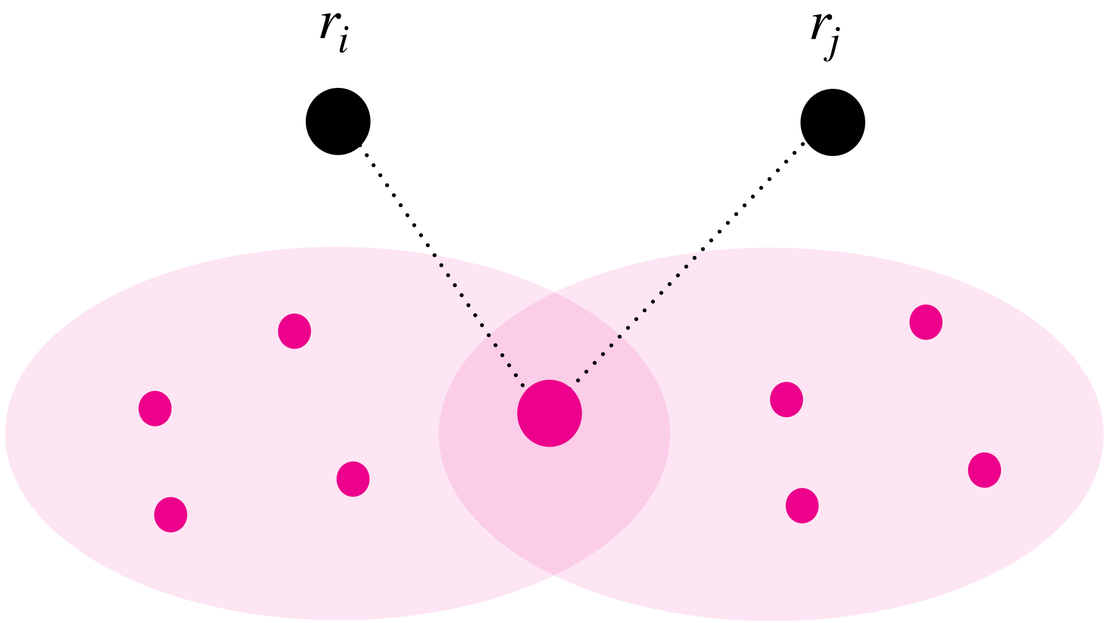

Additional Details

Max vs Sum

Does the choice really matter?

Proposition (Max-Sum discrepancy can be linear).

(1) A matching mechanism can have **sum**-distortion α and **max**-distortion $\Omega(\alpha n)$ in some metric space.

(2) A matching mechanism can have **max**-distortion α and **sum**-distortion $\Omega(\alpha n)$ in some metric space.

For (2), Consider the following $(n+1)$-gon.

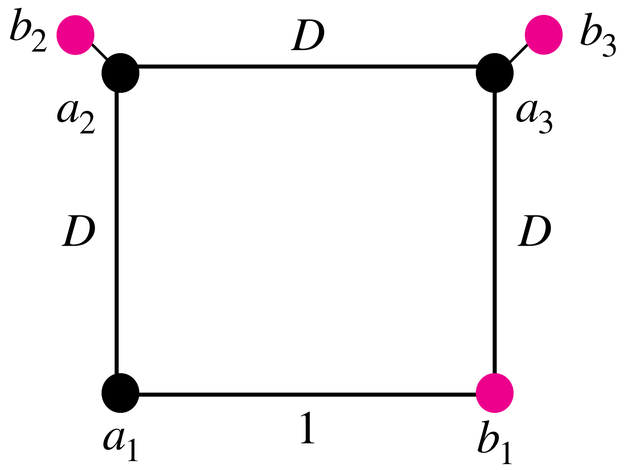

$$M(a_n) = b_1, \quad M(a_t) = b_{t+1} \quad 1 \le t \le n-1$$

$$\text{cost}_{\max}(M) = D, \quad \text{cost}_{\text{sum}}(M) = nD$$

$$\text{cost}_{\max}(M^*_{\max}) = \text{cost}_{\text{sum}}(M^*_{\text{sum}}) = 1$$

RepMatch

Distortion Lower Bounds

The mechanism can suffer from a **domino effect** on the line with exponentially increasing distances.

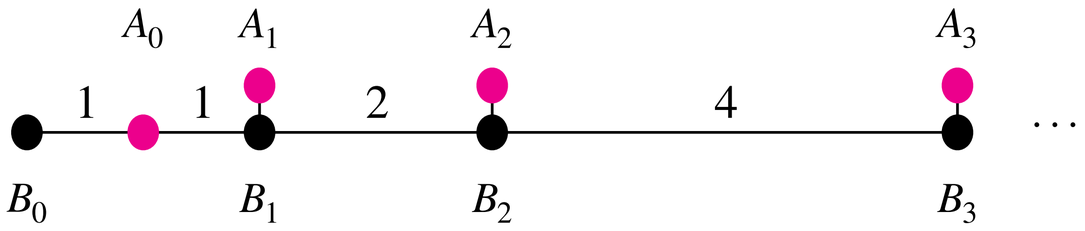

$$|A_0| = 1, \quad |A_t| = 2^{t-1} \quad \forall 1 \leq t \leq k$$

$$|B_0| = 1, \quad |B_t| = 2^{t-1} \quad \forall 1 \leq t \leq k$$

Agents break ties in favor of items on the right.

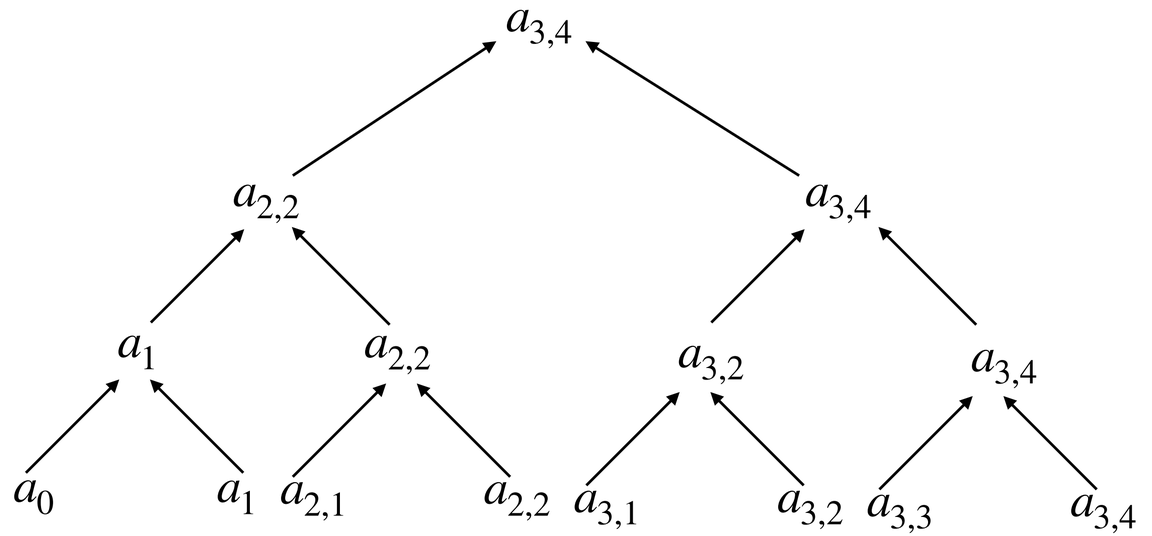

Claim.

Exactly **one** representative could remain after merging.

RepMatch

Distortion Lower Bounds

The mechanism can suffer from a domino effect on the line with exponentially increasing distances.

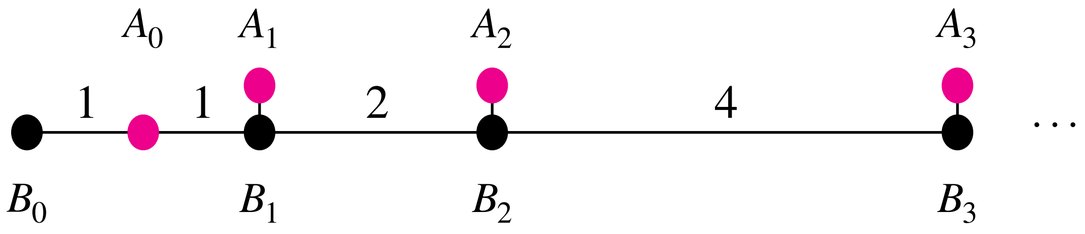

$$|A_0| = 1, \quad |A_t| = 2^{t-1} \quad \forall 1 \le t \le k$$

$$|B_0| = 1, \quad |B_t| = 2^{t-1} \quad \forall 1 \le t \le k$$

Agents break ties in favor of items on the right.

Claim.

RepMatch could return an arbitrary perfect matching.

- If $M(a) \notin B_k$ for some $a \in A_k$, then

$$d(a, M(a)) = 2^{k-1} = n/2.$$

- In the worst case, $M(a) \notin B_k$ for all $a \in A_k$.

General Lower Bounds

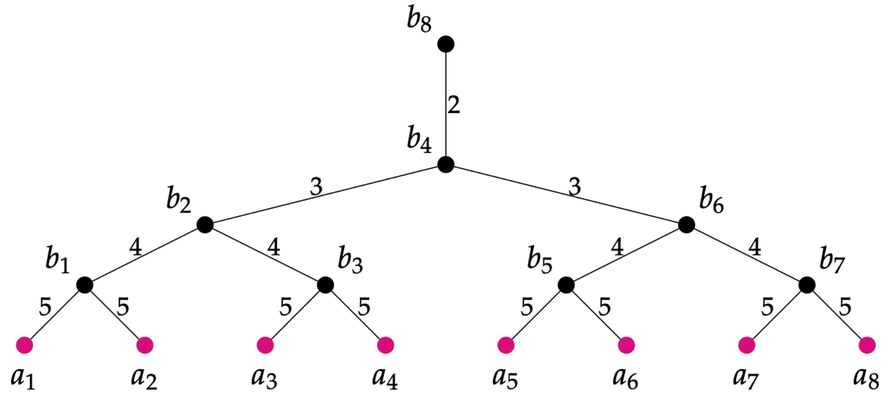

Figure 4: The metric space that defines the preferences in our ordinal matching instance for $n = 8$. b_8 is the root of the tree. The black nodes are the items and the red (leaf) nodes are the agents.

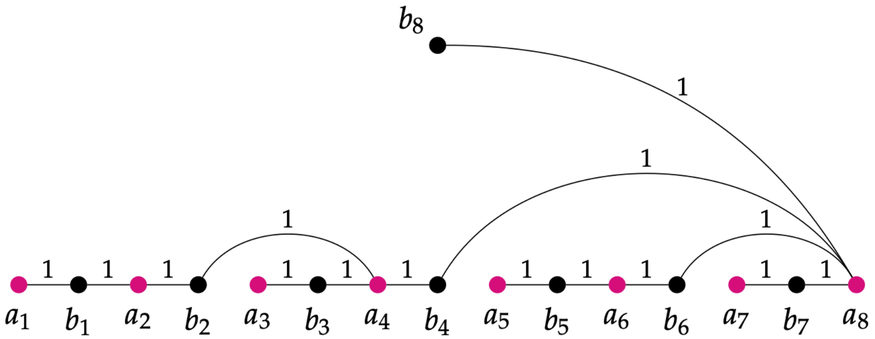

Figure 5: The metric space that defines the preferences in the consistent metric d_1 for $n = 8$. b_8 is the root of the tree. The black nodes are the items and the red nodes are the agents.