Stable matching as transport

Federico Echenique
Berkeley

Joseph Root
Chicago

Fedor Sandomirskiy
Princeton

Fedor and Joe

Matching with aligned preferences

	y_1	y_2	y_3
x_1	7	6	8
x_2	1	2	0
x_3	4	5	7

Matching with aligned preferences

	y_1	y_2	y_3
x_1	7	6	8
x_2	1	2	0
x_3	4	5	7

Matching with aligned preferences

	y_1	y_2	y_3
x_1	7	6	8
x_2	1	2	0
x_3	4	5	7

Matching with aligned preferences

	y_1	y_2	y_3
x_1	7	6	8
x_2	1	2	0
x_3	4	5	7

Markets with aligned preferences

- Agents have **aligned preferences**:

  if agents with types x and y are matched, both enjoy utility $u(x, y)$

- u is an objective fit, or match-quality.
  - e.g., partners interested in maximizing a common production function

Our contribution

- A general matching model, encompassing finite and infinite markets
- **Connection to optimal transportation theory:**
- Structural properties of optimal matchings
- Stability-fairness-welfare tension
- Extension to many-sided matching, e.g., team formation

Aligned preferences are interesting and realistic!

Main result: optimal transport

Optimization problem
w/parameter α.

$$\begin{array}{ll} \max & f(\mu, \alpha) \\ \text{s.t.} & \mu \text{ is a matching} \end{array}$$

stability ($\alpha = +\infty$)
welfare ($\alpha = 0$)
fairness ($\alpha = -\infty$)

Aligned preferences: motivation

School choice

- Distance is a key component of student preferences (Walters, 2018)
- Distance is a key component of school priorities
- Aligned distance-based preferences is an approximation

Aligned preferences: motivation

Markets with transfers but lack of commitment power

- Becker's (1973) marriage market model:
  - A couple (x, y) generates surplus $s(x, y)$ and can share it as
  $$s(x, y) = \hat{u}(x, y) + \hat{v}(x, y)$$
  - Shares $\hat{u}(x, y)$ and $\hat{v}(x, y)$ are determined at the time of the match
- Transfers are negotiated and committed to, as part of the bargaining over the match
- **Question:** What if no commitment power?
- Partners use Nash bargaining with weights $(1/2, 1/2)$ to split surplus after the match is formed
- Aligned preferences with $u(x, y) = s(x, y)/2$

Matching ⌄

Uber Marketplace

In the seconds after a rider requests a ride, we evaluate nearby drivers and riders in one batch. We then pair riders and drivers in the distribution, aiming to reduce the average wait time for everyone, not just the closest pair. This helps keep things moving and rides reliable across the network.

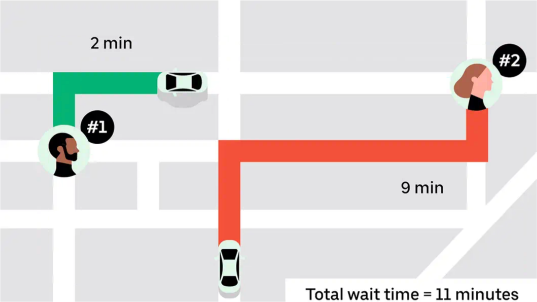

First to request

In the early days, a rider was immediately matched with the closest available driver. It worked well for most riders but sometimes led to long wait times for others. Across a whole city, those longer wait times really added up.

Related literature

- **Aligned Preferences:** decentralized dynamics (**Ferdowsian, Niederle, Yariv** 2020), random preferences (**Lee, Yariv** 2018), greedy algorithms and uniqueness for finite markets (**Eeckhout** 2000, **Clark** 2006, **Galichon, Ghelfi, Henry** 2023)
- **Large Markets:** **Azevedo, Leshno** (2016), **Ashlagi, Shi** (2016); **Leshno, Lo** (2021), **Arnosti** (2022), **Greinecker, Kah** (2021)
- **Optimal Transport in Econ:** markets with transfers (**Galichon, Salani** 2022, **Boerma, Tsyvinski, Wang, Zhang** 2023), mechanism design (**Daskalakis, Deckelbaum, Tzamos** 2015, **Kolesnikov, Sandomirskiy, Tsyvinski, Zimin** 2022, **Perez-Richet, Skreta** 2023), information design (**Malamud, Cieslak, Schrimpf** 2021, **Arieli, Babichenko, Sandomirskiy** 2023)

Outline

- Matching on the line with distance-based preferences
    - Stability and no-crossing property from optimal transport
- General markets with aligned preferences & optimal transport
    - stability, fairness, and welfare as objectives in a transport problem
    - trade-offs and worst-case bounds
- Distance-based matching in $\mathbb{R}^d$

Warmup: Matching on the line

Matching on the line

Model

- Agents are described by their "types" in $\mathbb{R}$
- Two sets of types $X \subset \mathbb{R}$ and $Y \subset \mathbb{R}$
- Two populations $\mu \in \Delta(X)$ and $\nu \in \Delta(Y)$
- Distance-based preferences: if $x \in X$ and $y \in Y$ match, each get utility

$$u(x, y) = -|x - y|$$

Matchings

Definition

$\pi \in \Delta(X \times Y)$ is a **matching** if it has marginal μ on X and ν on Y
Denote by $\Pi(\mu, \nu)$ the set of all matchings

Example: μ uniform on $[-1, 0]$, and ν uniform on $[0, 1]$

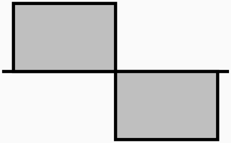

Matchings

Definition

$\pi \in \Delta(X \times Y)$ is a **matching** if it has marginal μ on X and ν on Y
Denote by $\Pi(\mu, \nu)$ the set of all matchings

Example: μ uniform on $[-1, 0]$, and ν uniform on $[0, 1]$

- assortative: $x \to y = x + 1$

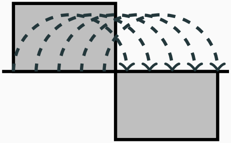

Matchings

Definition

$\pi \in \Delta(X \times Y)$ is a **matching** if it has marginal μ on X and ν on Y

Denote by $\Pi(\mu, \nu)$ the set of all matchings

Example: μ uniform on $[-1, 0]$, and ν uniform on $[0, 1]$

- assortative: $x \to y = x + 1$
- anti-assortative: $x \to y = -x$

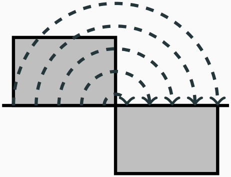

Matchings

Definition

$\pi \in \Delta(X \times Y)$ is a **matching** if it has marginal μ on X and ν on Y
Denote by $\Pi(\mu, \nu)$ the set of all matchings

Example: μ uniform on $[-1, 0]$, and ν uniform on $[0, 1]$

- assortative: $x \to y = x + 1$
- anti-assortative: $x \to y = -x$
- random: $\pi = \mu \times \nu$

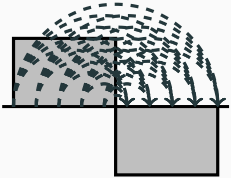

Stability

Definition

A matching π is **stable** if for any for any $(x, y),\ (x', y') \in \text{supp}(\pi)$,

$$u(x, y') \leq \min\left\{u(x, y),\ u(x', y')\right\}$$

At least one member in the mismatched pair (x, y') prefers their current partner, i.e., (x, y') is not a blocking pair

Stability

Example: μ uniform on $[-1, 0]$, and ν uniform on $[0, 1]$

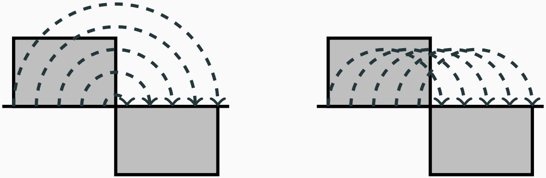

stable unstable

Stability

Example: μ uniform on $[-1, 0]$, and ν uniform on $[0, 1]$

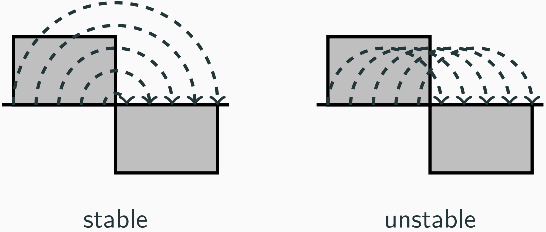

For $u(x, y) = -|x - y|$, stability is related to **no-crossing**

No-crossing

For interval $(z_1, z_2) \subset \mathbb{R}$, denote the circle in $\mathbb{R}^2$ having the interval as the diameter by $O(z_1, z_2)$

Definition

A matching π satisfies **no-crossing** if, for any $(x, y), (x', y') \in \mathrm{supp}(\pi)$, the circles $O(x, y)$ and $O(x', y')$ do not cross

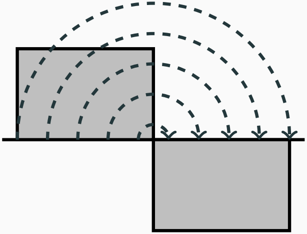

satisfies no-crossing

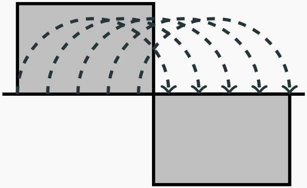

violates no-crossing

Stability and no-crossing

Lemma

Any stable matching satisfies no-crossing

Stability and no-crossing

Lemma

Any stable matching satisfies no-crossing

Proof. We need to rule out the following two patterns in stable matching

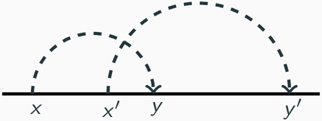

blocked by (x', y)

Stability and no-crossing

Lemma

Any stable matching satisfies no-crossing

Proof. We need to rule out the following two patterns in stable matching

blocked by (x', y) blocked by (x, y')

□

No-crossing II

Structure of no-crossing matching (McCann 1999)

Consider μ, ν w/densities f and g.
Any no-crossing matching is a cvx. comb. of 2 deterministic matchings:

- Match $x = y$ as much as possible.
  - All common mass $h = \min\{f, g\}$ is eliminated
- No-crossing matchings of residual populations $(f - h)$ and $(g - h)$ form a finite number of parametric families

- The no-crossing condition makes the problem parametric!

An example

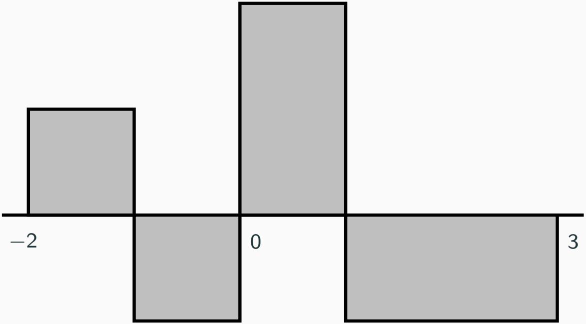

- No crossing matchings form a one-parametric family

An example

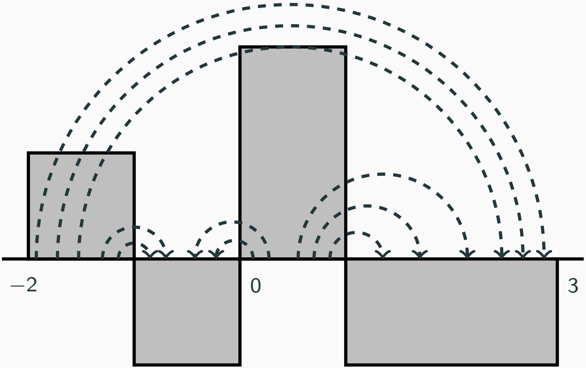

- No crossing matchings form a one-parametric family

An example

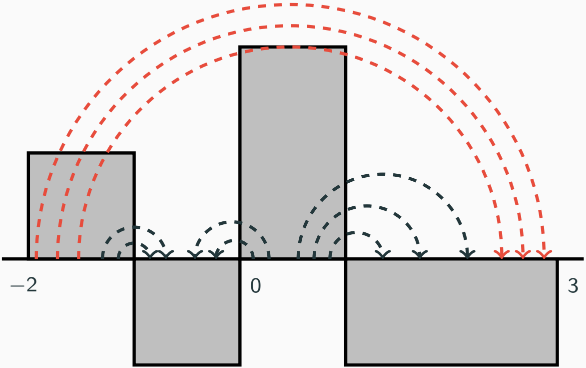

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

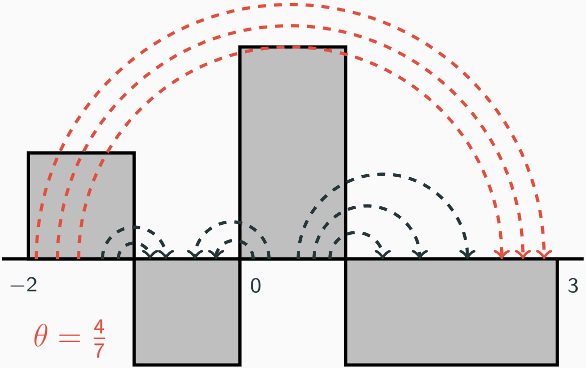

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

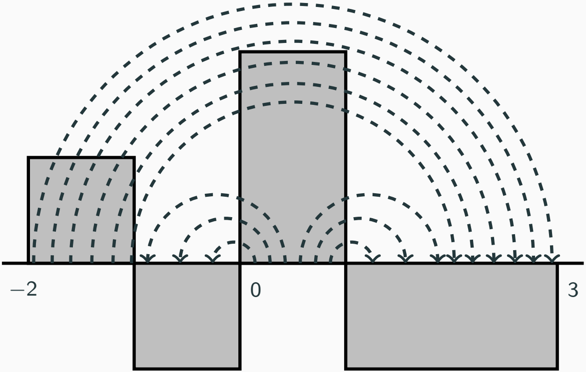

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

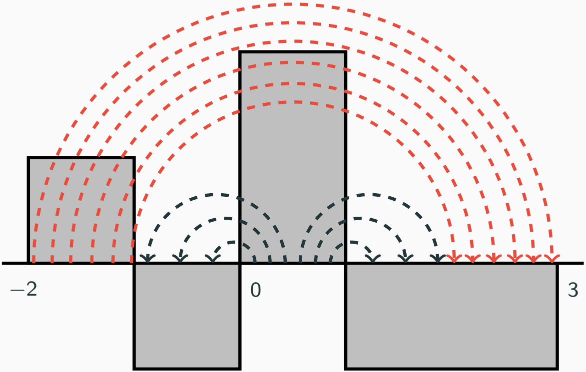

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

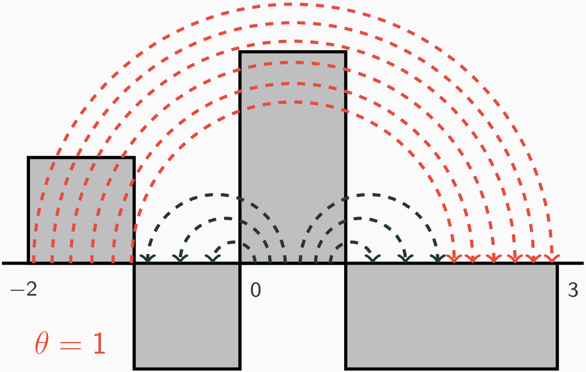

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

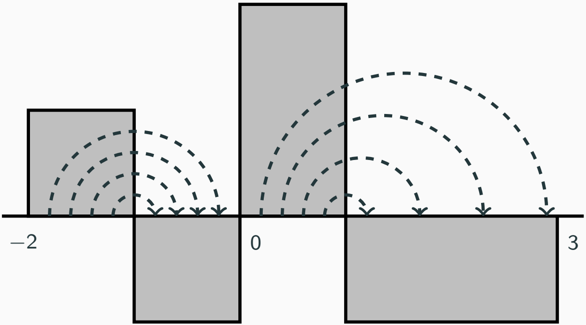

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

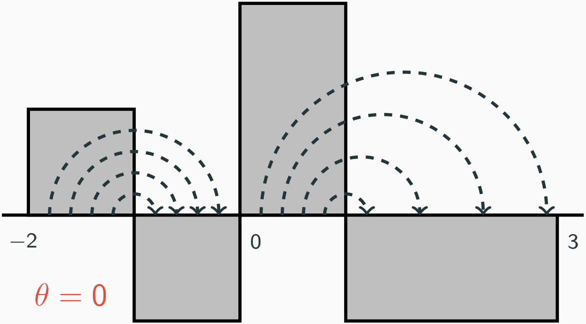

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

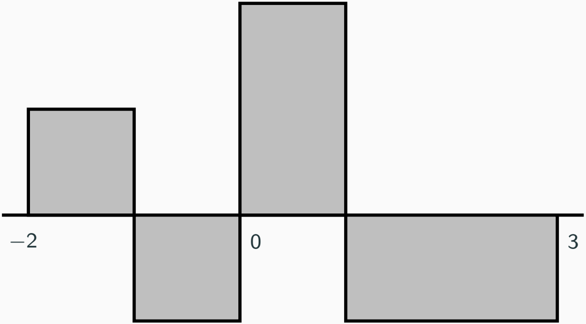

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally

An example

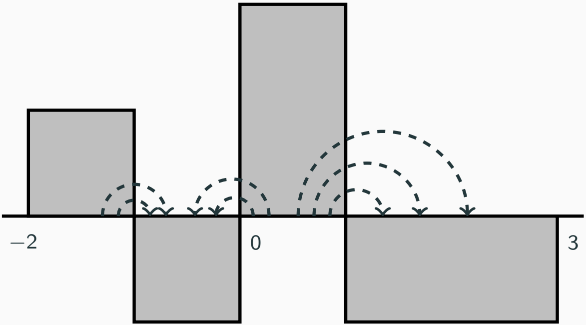

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally
- Stable matching corresponds to $\theta = 4/7$

An example

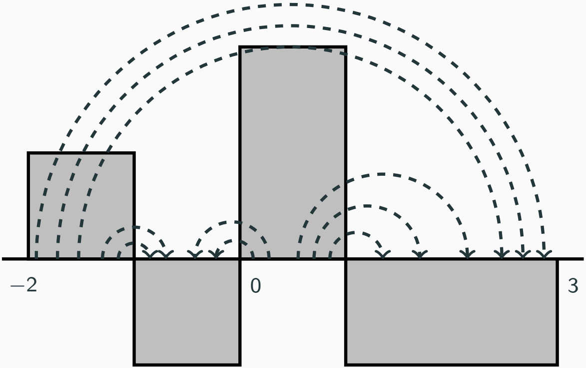

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally
- Stable matching corresponds to $\theta = 4/7$
- Non-local matches $\Rightarrow$ inequality & welfare loss. Quantify later

An example

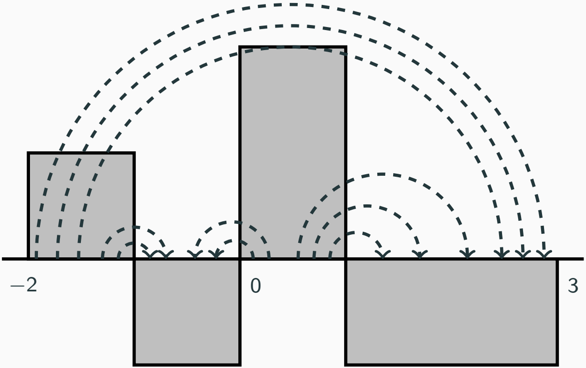

- No crossing matchings form a one-parametric family
- $\theta \in [0, 1]$ is the fraction of the interval $[-2, -1]$ matched non-locally
- Stable matching corresponds to $\theta = 4/7$
- Non-local matches $\Rightarrow$ inequality & welfare loss. Quantify later
- **Angrist, Gray-Lobe, Idoux, Pathak (2022):** Deferred Acceptance in NYC and Boston $\Rightarrow$ 50% increase in travel expenditure

Stable matching on the line

Corollaries of no-crossing:

- A stable matching=a convex combination of two deterministic ones:
  x is matched with the ideal partner $y = x$ or at most one other y'
- It can be searched for within a finite number of parametric families

Bad news: The number of families blows up exponentially with the number of times $\mu - \nu$ changes sign

Stable matching on the line

Proposition

For non-atomic $\mu, \nu \in \Delta(\mathbb{R})$, a stable matching exists and is unique, and can be constructed via a simple algorithm. For piecewise-constant densities with m intervals of constancy, it requires $O(m^2)$ operations

Proof idea. Find a "simple independent submarket"

- is to be matched independently of the rest of the population
- a no-crossing matching is unique and thus is stable
- after eliminating, the number of sign changes decreases by 1

Repeat □

Optimal transport and general markets with aligned preferences

Optimal transport

General optimal transport problem

Given measurable spaces X, Y, distributions $\mu \in \Delta(X),\ \nu \in \Delta(Y)$, payoff $p\colon X \times Y \to \mathbb{R}$, find a matching $\pi \in \Pi(\mu, \nu)$:

$$\pi \in \Pi(\mu, \nu) \quad : \quad \int_{X \times Y} p(x, y)\, \mathrm{d}\pi(x, y) \quad \to \quad \max$$

Optimal transport

- Often formulated for cost minimization ($c = -p$)
- **Standard interpretation:** μ and ν are spatial distributions of production and demand, π is the cheapest way to transport
- An archetypal problem of optimal correlation between two distributions $\Rightarrow$ omnipresent in Math, OR, and (gradually) Econ

—**McCann (1990):** $X, Y \subset \mathbb{R}$, convex $p \Rightarrow$ no-crossing π
—Stability for $\mathbb{R}$ and $u(x, y) = -|x - y| \;\Rightarrow$ no crossing π

Question: Any direct connection between stability and transport?

Yes, and it is not limited to $\mathbb{R}$ and distance-based utility

General markets with aligned preferences

Model

- X and Y are Polish spaces with Borel σ-algebra.
- Two populations $\mu \in \Delta(X)$ and $\nu \in \Delta(Y)$.
- If x and y match, both obtain utility/payoff $u(x, y)$.
- Assume $u\colon X \times Y \to \mathbb{R}$ is cont. and bounded.
    - often, measurability is enough (in the paper)
    - “acyclicity” of ordinal preferences $\Rightarrow$ existence of u (in the paper)

Designer's objectives

Criteria for matchings

- Approximate stability.
- Approximate egalitarianism.
- Utilitarian welfare.

Approximate stability

Definition

A matching π is ε-**stable** with $\varepsilon \geq 0$ if for any $(x, y), (x', y') \in \text{supp}(\pi)$,

$$u(x, y') \leq \max\left\{u(x, y),\ u(x', y')\right\} + \varepsilon$$

- At least one partner in any mismatched pair cannot benefit from leaving their current partner by more than ε
- for $\varepsilon = 0$, get the familiar notion of stability
- ε-stability $\simeq$ stability in the presence of ε-friction

Approximate egalitarianism

- For each matching $\pi \in \Pi(\mu, \nu)$ define

$$U_{\min}(\pi) = \min_{(x,y) \in \mathrm{supp}(\pi)} u(x, y)$$

  - Well-defined for compact X and Y
  - For non-compact, replace minimum with infimum
- Egalitarian lower bound

$$U^*_{\min}(\mu, \nu) = \max_{\pi \in \Pi(\mu,\nu)} U_{\min}(\pi)$$

Approximate egalitarianism

Definition

A matching $\pi \in \Pi(\mu, \nu)$ is ε-**egalitarian** if there is a subset $S \subset X \times Y$ with $\pi(S) \geq 1 - \varepsilon$ such that

$$u(x, y) \geq U^*_{\min}(\mu, \nu) - \varepsilon \qquad \text{for all} \quad (x, y) \in S$$

- All agents except ε-fraction have utilities above the ε-relaxed egalitarian bound

Utilitarianism

- The utilitarian welfare of a matching π by

$$W(\pi) = \int_{X \times Y} u(x, y) \, \mathrm{d}\pi(x, y)$$

- Optimal welfare

$$W^*(\mu, \nu) = \max_{\pi \in \Pi(\mu, \nu)} W(\pi)$$

- Welfare-max. $\simeq$ opt. transport w/payoff $p = u$
- The other objectives correspond to p equal to a **transformation** of u

Utility transformation

For a matching market with utility u, define the transformation

$$p_\alpha(x, y) = \frac{\exp(\alpha \cdot u(x, y)) - 1}{\alpha}$$

- p_α is convex in u for $\alpha > 0$ and concave for $\alpha < 0$
- for $\alpha = 0$, the limit $p_0(x, y) = u(x, y)$

Main result

Consider the transportation problem with payoff p_α

$$\pi \in \Pi(\mu, \nu) \quad : \quad \int_{X \times Y} p_\alpha(x, y)\, \mathrm{d}\pi(x, y) \quad \to \quad \max$$

- For $\alpha = 0$, this is welfare-maximization
- What do we get for $\alpha \neq 0$?

Theorem 1 proof

Let π be a solution to the optimal transport problem with payoff p_α

- If $\alpha > 0$ then π is ε-stable, with $\varepsilon = (\ln 2)/\alpha$.
- If $\alpha < 0$ then π is ε-egalitarian, with $\varepsilon = \max\{1, \ln|\alpha|\}/|\alpha|$

Main result

Corollaries:

- Existence of stable and egalitarian matchings (weak limit, $\alpha \to \pm\infty$)
- Changing α, we interpolate between the three objectives:

  fairness ($\alpha = -\infty$), welfare ($\alpha = 0$), stability ($\alpha = +\infty$)
- Fairness and stability are on the opposite sides of the spectrum
- Stability with aligned preferences $\simeq$ an **inequality-loving designer** prioritizing high-utility agents & ignoring externalities on low-utility ones

The result extends to *k*-sided markets: replace "$\ln 2$" with "$\ln k$"

Welfare and fairness of stable matching

How dramatic can be welfare loss and inequality of stable matching?

Theorem 2 proof

If utility $u \geq 0$ and a matching π is ε-stable, then

- π guarantees approximately half of optimal welfare:
$$W(\pi) \geq \frac{1}{2}\left(W^*(\mu, \nu) - \varepsilon\right)$$
- π is ε'-egalitarian with $\varepsilon' = \max\{1/2,\ \varepsilon\}$

- Any stable matching guarantees 1/2 of the optimal welfare and is 1/2-egalitarian
- These conservative bounds are concerned with ε-stable matchings with lowest welfare or that are least egalitarian
- "2" is the number of sides of the market

Distance-based matching in $\mathbb{R}^d$

Distance-based matching in $\mathbb{R}^d$: fairness-welfare tension

- $X = Y = \mathbb{R}^d$, utility $u(x, y) = -\|x - y\|$
- The payoff

$$p_\alpha(x, y) = \frac{\exp(\alpha \cdot \|x - y\|) - 1}{\alpha}$$

  is convex in the distance for $\alpha > 0$ and concave for $\alpha < 0$
- Optimal transport with $p(x, y) = f(\|x - y\|)$ is well-understood for convex/concave f

Let's focus on $d = 1$:

- Concave $f \Rightarrow$ assortative matching
- Thus $\alpha < 0 \Rightarrow$ assortative matching

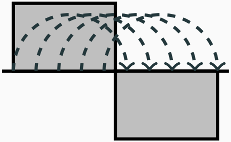

Corollary

For $d = 1$, there is no fairness-welfare tension. Both objectives are attained by the assortative matching.

- For $d > 1$, fairness-welfare tension emerges

Distance-based matching in $\mathbb{R}^d$: stability-fairness tension

Is there stability-fairness tension for $d = 1$? **Yes**

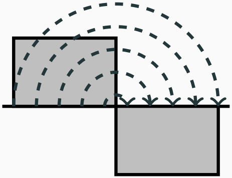

Both have the same welfare. Maybe there is no stability-welfare tension?

Distance-based matching in $\mathbb{R}^d$: stability-welfare tension

McCann (1999):

- For $d = 1$ and $p(x, y) = f(|x - y|)$ with strictly convex f, the optimal matching satisfies no-crossing
- If $\mu - \nu$ changes sign at most twice, a no-crossing matching is unique

For $\alpha > 0$ and ≤ 2 sign changes, the optimum does not depend on α

Corollary

If $\mu - \nu$ changes sign at most twice, there is no stability-welfare tension

Example:

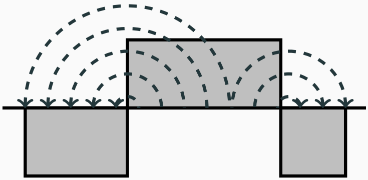

The conclusion extends to a round city in $\mathbb{R}^2$

Distance-based matching in $\mathbb{R}^d$: stability-welfare tension II

If there are ≥ 3 sign changes, stability-welfare tension emerges

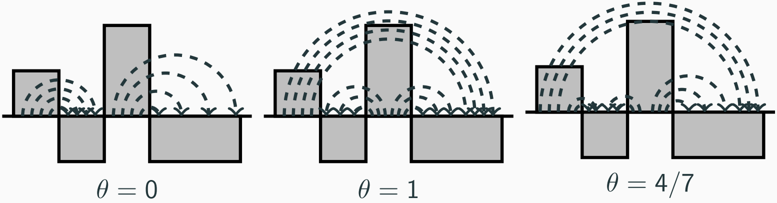

The optimal θ depends on α in the optimal transport problem

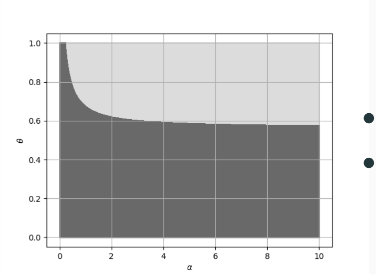

- stability $\Rightarrow \theta = 4/7 \approx 0.57$
- welfare-maximization $\Rightarrow \theta = 1$

Conclusion

- Aligned preferences emerge when
  - match quality is common to both sides (distance in school choice)
  - there are transfers but no commitment power
- Connection to transport: a parametric family of objectives captures

$$\text{stability } (\alpha = +\infty), \text{ welfare } (\alpha = 0), \text{ fairness } (\alpha = -\infty)$$

- Stability $\simeq$ prioritizing high-utility matches over low-utility ones
- Welfare and fairness losses, at most $1/2$ of each
- For particular spatial distributions no loss in welfare
  - Stability is OK if low-utility agents are compensated

Thank you!

Proof of Theorem 1: ε-stability as transport

back to theorem

Definition: Given $p\colon X \times Y \to \mathbb{R}$, a set $\Gamma \subset X \times Y$ is p-**cyclic monotone** if

$$\sum_{i=1}^{n} p(x_i, y_i) \geq \sum_{i=1}^{n} p(x_i, y_{i+1})$$

for all $(x_1, y_1), \dots, (x_n, y_n) \in \Gamma$ with $y_{n+1} = y_1$

Proof of Theorem 1: ε-stability as transport

Theorem (Beiglbock, Goldstern, Maresch, Schachermayer 2009)

If π solves an optimal transport problem

$$\pi \in \Pi(\mu, \nu) \quad : \quad \int_{X \times Y} p(x, y) \, \mathrm{d}\pi(x, y) \quad \to \quad \max,$$

then $\operatorname{supp}(\pi)$ is p-cyclic monotone

Proof of Theorem 1: ε-stability as transport

Use cyclic monotonicity for

$$p_\alpha(x, y) = \exp\left(\alpha \cdot u(x, y)\right)$$

On the support of the optimal matching π,

$$p_\alpha(x_1, y_2) + p_\alpha(x_2, y_1) \leq p_\alpha(x_1, y_1) + p_\alpha(x_2, y_2)$$

Equivalently,

$$\exp\left(\alpha \cdot u(x_1, y_2)\right) + \exp\left(\alpha \cdot u(x_2, y_1)\right) \leq \exp\left(\alpha \cdot u(x_1, y_1)\right) + \exp\left(\alpha \cdot u(x_2, y_2)\right)$$

Proof

Drop the second term on the LHS

$$\exp\big(\alpha \cdot u(x_1, y_2)\big) \leq \exp\big(\alpha \cdot u(x_1, y_1)\big) + \exp\big(\alpha \cdot u(x_2, y_2)\big)$$

Proof

Drop the second term on the LHS

$$\begin{aligned}\exp\left(\alpha \cdot u(x_1, y_2)\right) &\leq \exp\left(\alpha \cdot u(x_1, y_1)\right) + \exp\left(\alpha \cdot u(x_2, y_2)\right) \\ &\leq 2 \cdot \max\left\{ \exp\left(\alpha \cdot u(x_1, y_1)\right), \quad \exp\left(\alpha \cdot u(x_2, y_2)\right)\right\}\end{aligned}$$

Proof

Drop the second term on the LHS

$$\begin{aligned}\exp\left(\alpha \cdot u(x_1, y_2)\right) &\leq \exp\left(\alpha \cdot u(x_1, y_1)\right) + \exp\left(\alpha \cdot u(x_2, y_2)\right) \\ &\leq 2 \cdot \max\left\{ \exp\left(\alpha \cdot u(x_1, y_1)\right), \quad \exp\left(\alpha \cdot u(x_2, y_2)\right)\right\} \\ &= 2 \cdot \exp\Big(\alpha \cdot \max\{u(x_1, y_1), \quad u(x_2, y_2)\}\Big)\end{aligned}$$

Proof

Drop the second term on the LHS

$$\begin{aligned}\exp\left(\alpha \cdot u(x_1, y_2)\right) &\leq \exp\left(\alpha \cdot u(x_1, y_1)\right) + \exp\left(\alpha \cdot u(x_2, y_2)\right) \\ &\leq 2 \cdot \max\left\{ \exp\left(\alpha \cdot u(x_1, y_1)\right), \quad \exp\left(\alpha \cdot u(x_2, y_2)\right)\right\} \\ &= 2 \cdot \exp\Big(\alpha \cdot \max\{u(x_1, y_1), \quad u(x_2, y_2)\}\Big)\end{aligned}$$

Take logarithm and divide by α

$$u(x_1, y_2) \leq \max\left\{u(x_1, y_1), \quad u(x_2, y_2)\right\} + \frac{\ln(2)}{\alpha} \qquad \square$$

Proof of Theorem 2 back to theorem

Let π be an ε-stable matching.

For any $(x_1, y_1), (x_2, y_2) \in \text{supp}(\pi)$,

$$u(x_1, y_2) \leq \max\{u(x_1, y_1),\ u(x_2, y_2)\} + \varepsilon.$$

By non-negativity of u, we get

$$u(x_1, y_2) \leq u(x_1, y_1) + u(x_2, y_2) + \varepsilon.$$

Proof of Theorem 2

Let π' be any other matching with marginals μ and ν.

Consider $\lambda \in \mathcal{M}_+\big((X \times Y) \times (X \times Y)\big)$ s.t. the marginals of λ on (x_1, y_1) and on (x_2, y_2) are equal to π and the marginal on (x_1, y_2) is π'.

We get

$$\begin{aligned} W(\pi') &= \int_{X\times Y} u(x_1, y_2)\, \mathrm{d}\pi'(x_1, y_2) = \int_{(X\times Y)\times(X\times Y)} u(x_1, y_2)\, \mathrm{d}\lambda(x_1, y_1, x_2, y_2) \\ &\leq \int_{(X\times Y)\times(X\times Y)} (u(x_1, y_1) + u(x_2, y_2) + \varepsilon)\, \mathrm{d}\lambda(x_1, y_1, x_2, y_2) = \\ &= \int_{X\times Y} u(x_1, y_1)\, \mathrm{d}\pi(x_1, y_1) + \int_{X\times Y} u(x_2, y_2)\, \mathrm{d}\pi(x_2, y_2) + \varepsilon = \\ &= 2W(\pi) + \varepsilon. \end{aligned}$$

So:

$$W(\pi) \geq \frac{1}{2}\left(W(\pi') - \varepsilon\right)$$

for any matching π'. In particular, this inequality holds for π' maximizing welfare. Thus $W(\pi) \geq \frac{1}{2}\left(W^*(\mu, \nu) - \varepsilon\right)$.

Existence of a potential

Definition

A weak order (a complete and transitive binary relation) is termed a **preference**. If the weak order $\succeq$ is over a topological space Z, then we say that it is **continuous** if the upper contour sets $U_{\succeq}(z) = \{z' \in Z : z' \succ z\}$ and lower contour sets $L_{\succeq}(z) = \{z' \in Z : z' \prec z\}$ are open.

Existence of a potential

Primitives are a tuple $(X, Y, \succeq_X, \succeq_Y)$ in which:

- X and Y are topological spaces.
- $\succeq_X = \{\succeq_x : x \in X\}$, where for each $x \in X$, $\succeq_x$ is a continuous preference on Y;
- $\succeq_Y = \{\succeq_y : y \in Y\}$ for each $y \in Y$, $\succeq_y$ is a continuous preference on X.

A function $u \colon X \times Y \to \mathbb{R}$ is a **potential** for $(X, Y, \succeq_X, \succeq_Y)$ if

- $u(x, y) \geq u(x, y')$ iff $y \succeq_x y'$ for all x, y, y'
- and $u(x, y) \geq u(x', y)$ iff $x \succeq_y x'$ for all x, y, x'

Existence of a potential

The environment $(X, Y, \succeq_X, \succeq_Y)$ is **acyclic** if, for any sequence of couples,

$$(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n),$$

with $n > 2$ and $(x_n, y_n) = (x_1, y_1)$, so that each couple (x_{i+1}, y_{i+1}) has exactly one agent in common with the preceding couple (x_i, y_i), whenever all the common agents prefer their partner in (x_{i+1}, y_{i+1}) to their partner in (x_i, y_i), all common agents are, in fact, indifferent between the two partners

Existence of a potential

1. **Continuity with respect to the agent:** If $b \succ_a b'$ then there is a neighborhood N_a of a for which $b \succ_c b'$ for any $c \in N_a$
2. **Local strictness:** If $b' \succeq_a b$ and $b \succeq_{a'} b''$ with $a \neq a'$ and $b \neq b', b''$, then, in any neighborhood of b, there exists $\hat{b}$ with $b' \succ_a \hat{b}$ and $\hat{b} \succ_{a'} b''$

Theorem

Let $(X, Y, \succeq_X, \succeq_Y)$ be such that X and Y are complete, separable and connected topological spaces. Suppose that acyclicity and properties (1) and (2) are satisfied. Then there is a potential for $(X, Y, \succeq_X, \succeq_Y)$.