

Pareto-improving Pricing: Why 3 is better than 2

Piotr Dworczak^{*}

Northwestern University; GRAPE

Zi Yang Kang

University of Toronto

Seminar, London School of Economics
November 20, 2025

^{*} Co-funded by the European Union (ERC, IMD-101040122). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council.  

Motivation

Motivation

"personally, i really like my utility, relative to tan's"

Paula Onuchic

Motivation

"personally, i really like my utility, relative to tan's"

Paula Onuchic

Reminder that Pareto improvements are a nice thing!

Motivation

Waiting (due to congestion) is common in resource allocation:

- Internet traffic;
- vehicular traffic;
- security lines at airports;
- wait lists (public housing, healthcare, DMV...)
- ...

Motivation

Waiting (due to congestion) is common in resource allocation:

- Internet traffic;
- vehicular traffic;
- security lines at airports;
- wait lists (public housing, healthcare, DMV...)
- ...

Equity-efficiency trade-off associated with **pricing priority**:

- Gains from allocating priority to the most time-pressed agents...
- ... but drawback of making things worse for poorer agents.

Motivation

The Lexus Lanes – and why they won't work

Instead of rewarding carpools and getting people out of private cars, we are rewarding wealth and encouraging more people to drive. How does this make sense?

By

Zelda Bronstein

-

October 19, 2017



Share



WEALTH CONCENTRATION

HOT lanes make my blood boil: Fast tracks for the rich on the rise

ARTICLES

OCTOBER 31, 2012

by Salvatore Babones

Motivation

MONEY

Why must we pay to have a slightly less miserable time at the airport?

TSA PreCheck, Clear, and how the airport splits travelers into haves and have-nots.

by **Emily Stewart**

May 10, 2022, 7:00 AM CDT



Travelers wait in a long queue at the security checkpoint of Orlando International Airport the day before the Thanksgiving holiday on November 24, 2021. Paul Hennessy/SOPA Images/LightRocket via Getty Images



Emily Stewart covered business and economics for Vox and wrote the newsletter *The Big Squeeze*, examining the ways ordinary people are being squeezed under capitalism. Before joining Vox, she worked for *TheStreet*.

This Paper

We propose a **simple model** of assigning wait times to agents differing in the rate of substitution between time and money.

This Paper

We propose a **simple model** of assigning wait times to agents differing in the rate of substitution between time and money.

We study the properties of priority systems that **Pareto improve** on an unpriced status quo.

This Paper

We propose a **simple model** of assigning wait times to agents differing in the rate of substitution between time and money.

We study the properties of priority systems that **Pareto improve** on an unpriced status quo.

Main Results:

#1. If there is a trade-off between the dispersion and the mean of wait times, **no Pareto improvement with 2** priority tiers exists.

~> 1 fast lane + 1 slow lane are not enough; some agents will be worse off

This Paper

We propose a **simple model** of assigning wait times to agents differing in the rate of substitution between time and money.

We study the properties of priority systems that **Pareto improve** on an unpriced status quo.

Main Results:

#1. If there is a trade-off between the dispersion and the mean of wait times, **no Pareto improvement with 2** priority tiers exists.

~> 1 fast lane + 1 slow lane are not enough; some agents will be worse off

#2. When the trade-off between dispersion and mean is not too severe, **Pareto improvement with 3** priority tiers exists.

(Seed of a) literature review

- Possibility of Pareto improvements in priority service:
Chao and Wilson (1987), Gershkov and Schweinzer (2010), Hall (2018), Gershkov and Winter (2022),...
- Market-design approach to congestion pricing:
Ostrovsky and Schwarz (2018), Ostrovsky and Yang (2024),...
- Inequality-aware Market Design:
Weitzman (1977), Condorelli (2013), Dworczak ^① Kominers ^① Akbarpour (2021), Kang (2023, 2024), Kang and Watt (2025),...
- Recent empirical papers on congestion:
Hanna, Kreindler, Olken (2017), Kreindler (2024), Cook and Li (2024),...

Model

Model

Model

- Unit mass of agents indexed by a privately observed **marginal rate of substitution** r between wait time and money.

Model

- Unit mass of agents indexed by a privately observed **marginal rate of substitution** r between wait time and money.
- An agent with type r who obtains some service/good after waiting t and paying p has **utility**

$$V - r \cdot t - p,$$

where V is some known (high enough) constant.

Model

- Unit mass of agents indexed by a privately observed **marginal rate of substitution** r between wait time and money.
- An agent with type r who obtains some service/good after waiting t and paying p has **utility**

$$V - r \cdot t - p,$$

where V is some known (high enough) constant.

- Types r are drawn from a continuous distribution with cdf G and strictly positive pdf g on $[\underline{r}, \bar{r}]$.

Model

- Unit mass of agents indexed by a privately observed **marginal rate of substitution** r between wait time and money.
- An agent with type r who obtains some service/good after waiting t and paying p has **utility**

$$V - r \cdot t - p,$$

where V is some known (high enough) constant.

- Types r are drawn from a continuous distribution with cdf G and strictly positive pdf g on $[\underline{r}, \bar{r}]$.
- Think of t as measuring the **cost** of waiting; we can accommodate the case when $t = c(\tau)$, where c is an increasing function and τ is physical wait time.

Model

- Designer chooses a distribution of average wait times F from a feasible set of wait time distributions $\mathcal{F}$.

Model

- Designer chooses a distribution of average wait times F from a feasible set of wait time distributions $\mathcal{F}$.
- Any choice of F pins down the outcomes of a budget-balanced, incentive-compatible mechanism with transfers:

Model

- Designer chooses a distribution of average wait times F from a feasible set of wait time distributions $\mathcal{F}$.
- Any choice of F pins down the outcomes of a budget-balanced, incentive-compatible mechanism with transfers:
 - Allocation rule is given by $t(r) = F^{-1}(1 - G(r))$.

Model

- Designer chooses a distribution of average wait times F from a feasible set of wait time distributions $\mathcal{F}$.
- Any choice of F pins down the outcomes of a budget-balanced, incentive-compatible mechanism with transfers:
 - Allocation rule is given by $t(r) = F^{-1}(1 - G(r))$.
 - Payments are pinned down by the envelope formula:

$$U(r) \equiv V - rt(r) - p(r) = \bar{U} - \int_{\underline{r}}^r t(s)ds.$$

Model

- Designer chooses a distribution of average wait times F from a feasible set of wait time distributions $\mathcal{F}$.
- Any choice of F pins down the outcomes of a budget-balanced, incentive-compatible mechanism with transfers:
 - Allocation rule is given by $t(r) = F^{-1}(1 - G(r))$.
 - Payments are pinned down by the envelope formula:

$$U(r) \equiv V - rt(r) - p(r) = \bar{U} - \int_{\underline{r}}^r t(s)ds.$$

- Constant $\bar{U}$ pinned down by budget balance.

Running example

Example (lanes on a highway):

- There are N lanes on a highway.

Running example

Example (lanes on a highway):

- There are N lanes on a highway.
- If mass m_k of agents drive in lane k , wait time in lane k is $w(m_k)$, where $w(\cdot)$ is increasing and convex.

Running example

Example (lanes on a highway):

- There are N lanes on a highway.
- If mass m_k of agents drive in lane k , wait time in lane k is $w(m_k)$, where $w(\cdot)$ is increasing and convex.
- Let F_m be the distribution of wait times induced by m (ordered so that $m_i \leq m_{i+1}$), i.e., F assigns mass m_1 of $w(m_1)$, mass m_2 to $w(m_2)$, etc.

Running example

Example (lanes on a highway):

- There are N lanes on a highway.
- If mass m_k of agents drive in lane k , wait time in lane k is $w(m_k)$, where $w(\cdot)$ is increasing and convex.
- Let F_m be the distribution of wait times induced by m (ordered so that $m_i \leq m_{i+1}$), i.e., F assigns mass m_1 of $w(m_1)$, mass m_2 to $w(m_2)$, etc.
- Then, the set of feasible distributions is

$$\mathcal{F} = \left\{ F \in \Delta(\mathbb{R}_+) : F \in \text{MPC}(F_m) \text{ for some } m \in \Delta^{N-1} \right\}.$$

Running example

Unpriced benchmark (laissez-faire allocation):

- Suppose that we set:

$$m = (1/N, \dots, 1/N).$$

Running example

Unpriced benchmark (laissez-faire allocation):

- Suppose that we set:

$$m = (1/N, \dots, 1/N).$$

- The corresponding F_m is a Dirac delta at $w(1/N)$.

Running example

Unpriced benchmark (laissez-faire allocation):

- Suppose that we set:

$$m = (1/N, \dots, 1/N).$$

- The corresponding F_m is a Dirac delta at $w(1/N)$.
- The unique mechanism corresponding to F_m charges everyone a price of zero and assigns wait time $w(1/N)$ to everyone.

Running example

Example of non-trivial mechanism:

- Suppose that $N = 3$, $r \sim U[0, 1]$, and we set

$$m = (1/5, 2/5, 2/5).$$

Running example

Example of non-trivial mechanism:

- Suppose that $N = 3$, $r \sim U[0, 1]$, and we set

$$m = (1/5, 2/5, 2/5).$$

- The corresponding F has a mass $1/5$ of wait time $w(1/5)$ and a mass $4/5$ of wait time $w(2/5)$.

Running example

Example of non-trivial mechanism:

- Suppose that $N = 3$, $r \sim U[0, 1]$, and we set

$$m = (1/5, 2/5, 2/5).$$

- The corresponding F has a mass $1/5$ of wait time $w(1/5)$ and a mass $4/5$ of wait time $w(2/5)$.
- Types between $r \in [4/5, 1]$ get wait time $w(1/5)$ and pay p_F , and remaining types $r \in [0, 4/5]$ get wait time $w(2/5)$ and pay p_S .

Running example

Example of non-trivial mechanism:

- Suppose that $N = 3$, $r \sim U[0, 1]$, and we set

$$m = (1/5, 2/5, 2/5).$$

- The corresponding F has a mass $1/5$ of wait time $w(1/5)$ and a mass $4/5$ of wait time $w(2/5)$.
- Types between $r \in [4/5, 1]$ get wait time $w(1/5)$ and pay p_F , and remaining types $r \in [0, 4/5]$ get wait time $w(2/5)$ and pay p_S .
- Relative price pinned down by: $-\frac{4}{5}w\left(\frac{2}{5}\right) - p_S = -\frac{4}{5}w\left(\frac{1}{5}\right) - p_F$.

Running example

Example of non-trivial mechanism:

- Suppose that $N = 3$, $r \sim U[0, 1]$, and we set

$$m = (1/5, 2/5, 2/5).$$

- The corresponding F has a mass $1/5$ of wait time $w(1/5)$ and a mass $4/5$ of wait time $w(2/5)$.
- Types between $r \in [4/5, 1]$ get wait time $w(1/5)$ and pay p_F , and remaining types $r \in [0, 4/5]$ get wait time $w(2/5)$ and pay p_S .
- Relative price pinned down by: $-\frac{4}{5}w\left(\frac{2}{5}\right) - p_S = -\frac{4}{5}w\left(\frac{1}{5}\right) - p_F$.
- Budget balance: $\frac{1}{5}p_F + \frac{4}{5}p_S = 0$.

Assumptions on $\mathcal{F}$

Axiom (Randomization)

The set $\mathcal{F}$ is closed under taking mean-preserving contractions.

Assumptions on $\mathcal{F}$

Axiom (Randomization)

The set $\mathcal{F}$ is closed under taking mean-preserving contractions.

$\rightsquigarrow$ We can always assign wait times randomly.

Assumptions on $\mathcal{F}$

Axiom (Randomization)

The set $\mathcal{F}$ is closed under taking mean-preserving contractions.

$\rightsquigarrow$ We can always assign wait times randomly.

$\rightsquigarrow$ True in the running example by design (as long as we can “randomize” access to the lanes).

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

$\rightsquigarrow$ Giving everyone the same wait time minimizes average wait time.

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

↪ Giving everyone the same wait time minimizes average wait time.

↪ Randomization axiom + $\mathcal{F}$ compact

⇒ some δ_{t_0} minimizes the average wait time.

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

↪ Giving everyone the same wait time minimizes average wait time.

↪ Randomization axiom + $\mathcal{F}$ compact

⇒ some δ_{t_0} minimizes the average wait time.

↪ Substantial restriction: δ_{t_0} **uniquely** minimizes average wait time.

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

- ↪ Giving everyone the same wait time minimizes average wait time.
- ↪ Randomization axiom + $\mathcal{F}$ compact
 - $\implies$ some δ_{t_0} minimizes the average wait time.
- ↪ Substantial restriction: δ_{t_0} **uniquely** minimizes average wait time.
- ↪ Any mean-preserving spread of δ_{t_0} must take us outside of $\mathcal{F}$.

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

- ↪ Giving everyone the same wait time minimizes average wait time.
- ↪ Randomization axiom + $\mathcal{F}$ compact
 - $\implies$ some δ_{t_0} minimizes the average wait time.
- ↪ Substantial restriction: δ_{t_0} **uniquely** minimizes average wait time.
- ↪ Any mean-preserving spread of δ_{t_0} must take us outside of $\mathcal{F}$.
- ↪ There is a **trade-off between dispersion and average wait time**.

Assumptions on $\mathcal{F}$

Axiom (Friction)

There exists t_0 such that $\delta_{t_0} \in \mathcal{F}$ (the **unpriced benchmark**) **uniquely minimizes** the average wait time $\mathbb{E}_F[t]$ across all feasible $F \in \mathcal{F}$.

- ↪ Giving everyone the same wait time minimizes average wait time.
- ↪ Randomization axiom + $\mathcal{F}$ compact
 - $\implies$ some δ_{t_0} minimizes the average wait time.
- ↪ Substantial restriction: δ_{t_0} **uniquely** minimizes average wait time.
- ↪ Any mean-preserving spread of δ_{t_0} must take us outside of $\mathcal{F}$.
- ↪ There is a **trade-off between dispersion and average wait time**.
- ↪ True in the running example if $w(\cdot)$ is strictly convex (around $1/N$).

Assumptions on $\mathcal{F}$

Axiom (Smoothness)

For every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ and $q \in (0, 1)$ such that

$$F^{-1}(q + \sqrt{\epsilon}) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{q(t_0 - \mathbb{E}_F[t \mid F(t) \leq q])} < \epsilon.$$

Assumptions on $\mathcal{F}$

Axiom (Smoothness)

For every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ and $q \in (0, 1)$ such that

$$F^{-1}(q + \sqrt{\epsilon}) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{q(t_0 - \mathbb{E}_F[t | F(t) \leq q])} < \epsilon.$$

$\rightsquigarrow$ Trade-off between dispersion and average wait time is not too severe.

Assumptions on $\mathcal{F}$

Axiom (Smoothness)

For every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ and $q \in (0, 1)$ such that

$$F^{-1}(q + \sqrt{\epsilon}) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{q(t_0 - \mathbb{E}_F[t | F(t) \leq q])} < \epsilon.$$

$\rightsquigarrow$ Trade-off between dispersion and average wait time is not too severe.

$\rightsquigarrow$ Holds in the running example if $w(\cdot)$ is continuously differentiable in the neighborhood of $1/N$.

Assumptions on $\mathcal{F}$

Proof that Smoothness axiom holds in the running example:

Assumptions on $\mathcal{F}$

Proof that Smoothness axiom holds in the running example:

Stronger property: There exists $q > 0$ such that, for every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ such that

$$F^{-1}(q) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{t_0 - \mathbb{E}_F[t \mid F(t) \leq q]} < \epsilon.$$

Assumptions on $\mathcal{F}$

Proof that Smoothness axiom holds in the running example:

Stronger property: There exists $q > 0$ such that, for every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ such that

$$F^{-1}(q) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{t_0 - \mathbb{E}_F[t \mid F(t) \leq q]} < \epsilon.$$

Let $(1/N) - \delta$ agents use lane 1, and $(1/N) + \delta$ agents use lane N .
Use $q = 1/(2N)$.

Assumptions on $\mathcal{F}$

Proof that Smoothness axiom holds in the running example:

Stronger property: There exists $q > 0$ such that, for every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ such that

$$F^{-1}(q) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{t_0 - \mathbb{E}_F[t \mid F(t) \leq q]} < \epsilon.$$

Let $(1/N) - \delta$ agents use lane 1, and $(1/N) + \delta$ agents use lane N .
Use $q = 1/(2N)$.

$$\frac{\mathbb{E}_{F_\delta}[t] - t_0}{t_0 - \mathbb{E}_{F_\delta}[t \mid F_\delta(t) \leq q]}$$

Assumptions on $\mathcal{F}$

Proof that Smoothness axiom holds in the running example:

Stronger property: There exists $q > 0$ such that, for every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ such that

$$F^{-1}(q) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{t_0 - \mathbb{E}_F[t \mid F(t) \leq q]} < \epsilon.$$

Let $(1/N) - \delta$ agents use lane 1, and $(1/N) + \delta$ agents use lane N .
Use $q = 1/(2N)$.

$$\frac{(\frac{1}{N} - \delta) w(\frac{1}{N} - \delta) + (\frac{1}{N} + \delta) w(\frac{1}{N} + \delta) - \frac{2}{N} w(\frac{1}{N})}{w(\frac{1}{N}) - w(\frac{1}{N} - \delta)}$$

Assumptions on $\mathcal{F}$

Proof that Smoothness axiom holds in the running example:

Stronger property: There exists $q > 0$ such that, for every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ such that

$$F^{-1}(q) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{t_0 - \mathbb{E}_F[t \mid F(t) \leq q]} < \epsilon.$$

Let $(1/N) - \delta$ agents use lane 1, and $(1/N) + \delta$ agents use lane N .
Use $q = 1/(2N)$.

$$\xrightarrow{\delta \rightarrow 0} \frac{w\left(\frac{1}{N}\right) + \frac{1}{N} w'\left(\frac{1}{N}\right) - w\left(\frac{1}{N}\right) - \frac{1}{N} w'\left(\frac{1}{N}\right)}{w'\left(\frac{1}{N}\right)} = 0.$$

Assumptions on $\mathcal{F}$

We will sometimes refer to the opposite case:

Axiom (No friction)

If $F \in \mathcal{F}$ and F' is a mean-preserving spread of F , then $F' \in \mathcal{F}$.

Assumptions on $\mathcal{F}$

We will sometimes refer to the opposite case:

Axiom (No friction)

If $F \in \mathcal{F}$ and F' is a mean-preserving spread of F , then $F' \in \mathcal{F}$.

$\rightsquigarrow$ Note that “No friction” implies (stronger version of) “Smoothness”:

There exists $q > 0$ such that, for every $\epsilon > 0$ small, there exist $F \in \mathcal{F}$ such that

$$F^{-1}(q) < t_0 \text{ and } \frac{\mathbb{E}_F[t] - t_0}{t_0 - \mathbb{E}_F[t \mid F(t) \leq q]} < \epsilon$$

because we can find F such that $\mathbb{E}_F[t] = t_0$.

Results

Results

Results

- We can **Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) \geq U(r)$, for all r , strictly for some r .

Results

- We can **Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) \geq U(r)$, for all r , strictly for some r .
- We can **strictly Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) > U(r)$, for all r .

Results

- We can **Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) \geq U(r)$, for all r , strictly for some r .
- We can **strictly Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) > U(r)$, for all r .
- Despite quasi-linearity, Pareto optimality is **not** equivalent to maximizing the sum of utilities.

Results

- We can **Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) \geq U(r)$, for all r , strictly for some r .
- We can **strictly Pareto improve** upon a mechanism delivering utilities $U(r)$ if there exists a feasible mechanism delivering utilities $\tilde{U}(r)$ such that $\tilde{U}(r) > U(r)$, for all r .
- Despite quasi-linearity, Pareto optimality is **not** equivalent to maximizing the sum of utilities.
- Inequality-aware Market Design $\implies$ mechanisms on the constrained Pareto frontier may involve **random allocation**.

Results

Definition

We say that $F \in \mathcal{F}$ offers n (priority) tiers if $|\text{supp}(F)| = n$.

Results

Definition

We say that $F \in \mathcal{F}$ offers n (priority) tiers if $|\text{supp}(F)| = n$.

Theorem

Assume Randomization, Friction, and Smoothness.

There exists a (strict) Pareto improvement over the unpriced benchmark that uses 3 priority tiers.

It is not possible to (weakly) Pareto improve over the unpriced benchmark with 2 priority tiers.

Results

Definition

We say that $F \in \mathcal{F}$ offers n (priority) tiers if $|\text{supp}(F)| = n$.

Theorem

Assume Randomization, Friction, and Smoothness.

There exists a (strict) Pareto improvement over the unpriced benchmark that uses 3 priority tiers.

It is not possible to (weakly) Pareto improve over the unpriced benchmark with 2 priority tiers.

The first result only uses “Randomization” + “Smoothness”;
the second result only uses “Friction”.

Results

Discussion:

- In applications, we often see priority systems with 2 priority tiers, e.g., fast lane+slow lane, single fast-track option (e.g., Clear).

Results

Discussion:

- In applications, we often see priority systems with 2 priority tiers, e.g., fast lane+slow lane, single fast-track option (e.g., Clear).
- The theorem implies that these priority systems must make some agents worse off.

Results

Discussion:

- In applications, we often see priority systems with 2 priority tiers, e.g., fast lane+slow lane, single fast-track option (e.g., Clear).
- The theorem implies that these priority systems must make some agents worse off.
- “Friction” results in a trade-off between average wait time and dispersion.

Results

Discussion:

- In applications, we often see priority systems with 2 priority tiers, e.g., fast lane+slow lane, single fast-track option (e.g., Clear).
- The theorem implies that these priority systems must make some agents worse off.
- “Friction” results in a trade-off between average wait time and dispersion.
- Intuitively, 2 priority tiers are insufficient to resolve this trade-off while improving for all agents.

Results

Discussion:

- In the running example, there exists a deterministic Pareto-improving mechanism if there are three or more lanes.

Results

Discussion:

- In the running example, there exists a deterministic Pareto-improving mechanism if there are three or more lanes.
- If there are only two lanes, no deterministic mechanism can be a Pareto improvement.

Results


Discussion:

- In the running example, there exists a deterministic Pareto-improving mechanism if there are three or more lanes.
- If there are only two lanes, no deterministic mechanism can be a Pareto improvement.
- However, even with two lanes, Pareto improvement is possible with a stochastic mechanism (that offers three options).


Proof

Proof of main result


Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof




Proof







Proof




Proof



Proof



Proof



Additional applications and extensions

Additional applications

Waiting in line for a service:

- Consider a stationary queuing model where ρ agents can be processed per unit of time.

Additional applications

Waiting in line for a service:

- Consider a stationary queuing model where ρ agents can be processed per unit of time.
- Physical wait times can be “spread out” without changing the average.

Additional applications

Waiting in line for a service:

- Consider a stationary queuing model where ρ agents can be processed per unit of time.
- Physical wait times can be “spread out” without changing the average.
- Formally, utilities are given by $-rc(t) - p$, and

$$\mathcal{F} = \left\{ F : F(\cdot) = F_0(c^{-1}(\cdot)) \text{ for some } F_0 \succ^{MPS} \delta_{1/\rho} \right\}.$$

Additional applications

Waiting in line for a service:

- Consider a stationary queuing model where ρ agents can be processed per unit of time.
- Physical wait times can be “spread out” without changing the average.
- Formally, utilities are given by $-rc(t) - p$, and

$$\mathcal{F} = \left\{ F : F(\cdot) = F_0(c^{-1}(\cdot)) \text{ for some } F_0 \succ^{MPS} \delta_{1/\rho} \right\}.$$

- “Friction” and “Smoothness” axioms hold if $c(t)$ is increasing, strictly convex, and differentiable.

Additional applications

We can add a constant to utility functions:

$$r - rt - p$$

Additional applications

We can add a constant to utility functions:

$$r - rt - p = r(1 - t) - p$$

Additional applications

We can add a constant to utility functions:

$$r - rt - p = r(1 - t) - p =: rq - p,$$

where $q := 1 - t$ can be interpreted as **quality**.

Additional applications

Mussa-Rosen'78 model with a constraint on total resources:

- Designer has total capacity S , which can be divided into qualities q : Distributions of quality F_0 satisfy $\int q dF_0(q) = S$.

Additional applications

Mussa-Rosen'78 model with a constraint on total resources:

- Designer has total capacity S , which can be divided into qualities q : Distributions of quality F_0 satisfy $\int q dF_0(q) = S$.
- Agents' utilities are $rv(q) - p$, and

$$\mathcal{F} = \{F : F \text{ is a MPC of } F_0(v^{-1}(\cdot)) \text{ where } \int q dF_0(q) = S\}.$$

Additional applications

Mussa-Rosen'78 model with a constraint on total resources:

- Designer has total capacity S , which can be divided into qualities q : Distributions of quality F_0 satisfy $\int q dF_0(q) = S$.

- Agents' utilities are $rv(q) - p$, and

$$\mathcal{F} = \{F : F \text{ is a MPC of } F_0(v^{-1}(\cdot)) \text{ where } \int q dF_0(q) = S\}.$$

- “Friction” and “Smoothness” axioms hold if $v(q)$ is increasing, strictly concave, and continuously differentiable.

Additional applications

Extension of the running example:

- Consider an extension: “Lane” k has a disutility of congestion $w_k(m_k)$ when m_k agents use it.

Additional applications

Extension of the running example:

- Consider an extension: “Lane” k has a disutility of congestion $w_k(m_k)$ when m_k agents use it.
- Dependence on k can capture lane $\times$ (time of day) pairs.

Additional applications

Extension of the running example:

- Consider an extension: “Lane” k has a disutility of congestion $w_k(m_k)$ when m_k agents use it.
- Dependence on k can capture lane $\times$ (time of day) pairs.
- Another interpretation: “Lanes” 2 through N are lanes on a highway while lane 1 represents public transit (perhaps $w_1(m_1) = \bar{w}_1$).

Additional applications

Extension of the running example:

- Note that **laissez-faire** equalizes wait times (for lanes that are in use):

$$w_k(m_k) = w_j(m_j), \forall k, j \text{ s.t. } m_k, m_j > 0.$$

Additional applications

Extension of the running example:

- Note that **laissez-faire** equalizes wait times (for lanes that are in use):

$$w_k(m_k) = w_j(m_j), \forall k, j \text{ s.t. } m_k, m_j > 0.$$

- This **no longer minimizes average wait time** in general:

$$\min \sum_{k=1}^N m_k w_k(m_k) \text{ subject to } \sum_{k=1}^N m_k = 1$$

$$\text{FOC: } m_k^* w'_k(m_k^*) + w_k(m_k^*) = \lambda$$

Additional applications

Extension of the running example:

- If the “Friction” axiom fails, it means that there must exist a distribution F that has some dispersion and generates **lower average wait time than the laissez-faire**.

Additional applications

Extension of the running example:

- If the “Friction” axiom fails, it means that there must exist a distribution F that has some dispersion and generates **lower average wait time than the laissez-faire**.
- But this means that “**Smoothness**” axiom holds trivially, since $\mathbb{E}_F[t] \leq t_0$.

Additional applications

Extension of the running example:

- If the “Friction” axiom fails, it means that there must exist a distribution F that has some dispersion and generates **lower average wait time than the laissez-faire**.
- But this means that “**Smoothness**” axiom holds trivially, since $\mathbb{E}_F[t] \leq t_0$.
- It is **possible to strictly Pareto improve**.

Additional applications

Extension of the running example:

- If the “Friction” axiom fails, it means that there must exist a distribution F that has some dispersion and generates **lower average wait time than the laissez-faire**.
- But this means that “**Smoothness**” axiom holds trivially, since $\mathbb{E}_F[t] \leq t_0$.
- It is **possible to strictly Pareto improve**.
- It may be possible to Pareto improve with two priority tiers.

Additional applications

Extension of the running example:

- We can ask a more ambitious question: Is it possible to **strictly Pareto improve** upon the Dirac delta distribution that minimizes average wait time? (it exists by the “Randomization” axiom and can be implemented by quotas.)

Additional applications

Extension of the running example:

- We can ask a more ambitious question: Is it possible to **strictly Pareto improve** upon the Dirac delta distribution that minimizes average wait time? (it exists by the “Randomization” axiom and can be implemented by quotas.)
- **Yes**, if w_k ’s are differentiable (“Smoothness” holds).

Additional applications

Extension of the running example:

- We can ask a more ambitious question: Is it possible to **strictly Pareto improve** upon the Dirac delta distribution that minimizes average wait time? (it exists by the “Randomization” axiom and can be implemented by quotas.)
- **Yes**, if w_k ’s are differentiable (“Smoothness” holds).
- May be possible to **weakly Pareto improve with two priority tiers**.

Extensions

In the baseline model, **can we Pareto improve** on a two-tier priority system?

Extensions

In the baseline model, **can we Pareto improve** on a two-tier priority system?

- Unfortunately, the answer in general is “no”.

Extensions

In the baseline model, **can we Pareto improve** on a two-tier priority system?

- Unfortunately, the answer in general is “no”.
- How do we know? Take the running example with $N = 2$. For any choice of m , there are welfare functions for which **assortative matching is optimal** (Akbarpour ^① Dworczak ^① Kominers, 2024).

Extensions

In the baseline model, **can we Pareto improve** on a two-tier priority system?

- Unfortunately, the answer in general is “no”.
- How do we know? Take the running example with $N = 2$. For any choice of m , there are welfare functions for which **assortative matching is optimal** (Akbarpour ^① Dworczak ^① Kominers, 2024).
- This means that **some two-tier priority systems** are optimal for some welfare weights, hence **on the Pareto frontier!**

Extensions

In the baseline model, **can we Pareto improve** on a two-tier priority system?

- Unfortunately, the answer in general is “no”.
- How do we know? Take the running example with $N = 2$. For any choice of m , there are welfare functions for which **assortative matching is optimal** (Akbarpour ^① Dworczak ^① Kominers, 2024).
- This means that **some two-tier priority systems** are optimal for some welfare weights, hence **on the Pareto frontier**!
- Intuition: A two-tier priority system is potentially very good for agents with low or high willingness to pay.

Extensions

We argued that there exists a Pareto improvement. But **what mechanism should we use** in the end?

Extensions


We argued that there exists a Pareto improvement. But **what mechanism should we use** in the end?

One approach is to maximize a particular welfare function subject to the Pareto improvement constraint:

Extensions

We argued that there exists a Pareto improvement. But **what mechanism should we use** in the end?

One approach is to maximize a particular welfare function subject to the Pareto improvement constraint: In the waiting in line example,



Concluding Remarks

Concluding Remarks

Concluding Remarks

- In a simple model of **priority pricing**, a **Pareto improvement** over an unpriced benchmark is possible with **3 priority tiers**, but not with 2.

Concluding Remarks

- In a simple model of **priority pricing**, a **Pareto improvement** over an unpriced benchmark is possible with **3 priority tiers**, but not with 2.
- A concrete **policy implication**: Create at least three options, with the middle option being free and replicating the wait time in the unpriced benchmark.

Concluding Remarks

- In a simple model of **priority pricing**, a **Pareto improvement** over an unpriced benchmark is possible with **3 priority tiers**, but not with 2.
- A concrete **policy implication**: Create at least three options, with the middle option being free and replicating the wait time in the unpriced benchmark.
- Next steps?

Concluding Remarks

- In a simple model of **priority pricing**, a **Pareto improvement** over an unpriced benchmark is possible with **3 priority tiers**, but not with 2.
- A concrete **policy implication**: Create at least three options, with the middle option being free and replicating the wait time in the unpriced benchmark.
- Next steps?
 - Conditions for Pareto improvement over an arbitrary mechanism.

Concluding Remarks

- In a simple model of **priority pricing**, a **Pareto improvement** over an unpriced benchmark is possible with **3 priority tiers**, but not with 2.
- A concrete **policy implication**: Create at least three options, with the middle option being free and replicating the wait time in the unpriced benchmark.
- Next steps?
 - Conditions for Pareto improvement over an arbitrary mechanism.
 - Robustness to extensive margin and richer preferences.

Concluding Remarks

- In a simple model of **priority pricing**, a **Pareto improvement** over an unpriced benchmark is possible with **3 priority tiers**, but not with 2.
- A concrete **policy implication**: Create at least three options, with the middle option being free and replicating the wait time in the unpriced benchmark.
- Next steps?
 - Conditions for Pareto improvement over an arbitrary mechanism.
 - Robustness to extensive margin and richer preferences.
 - Empirical tests.

Concluding Remarks

$$3 > 2$$