Balancing the Tradeoff between Profit and Fairness in Rideshare Platforms During High-Demand Hours

Vedant Nanda, Pan Xu, Karthik Abinav Sankararaman, John P. Dickerson, Aravind Srinivasan

Rideshare

“...participate in an arrangement in which a passenger travels in a private vehicle driven by its owner, for free or for a fee, especially as arranged by means of a website or app.”

- **Two sided** market: passengers and drivers
- Platforms have control over assigning passengers to drivers

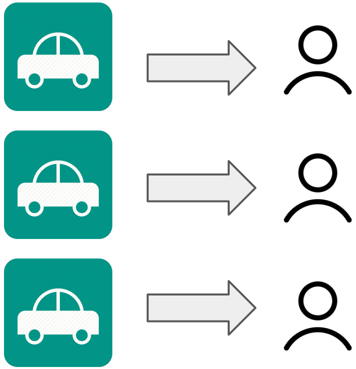

Issues in Rideshare

- Rideshare platforms can allow drivers to reject requests up to a predefined number - exacerbating biases!
- In certain scenarios (e.g., peak hour), drivers have leverage over passengers
- Drivers can choose to reject rides based, e.g., on trip length and starting/ending location
- Can disadvantage those going to “unpopular” destinations

Did Uber Just Enable Discrimination by Destination?

In California, the ride-hailing company is changing a policy used as a safeguard against driver discrimination against low-income and minority riders.

Issues in Rideshare

- UCLA Study (Brown, 2018)
  - Black riders had to wait **1 minute and 43 seconds longer** than their white counterparts for rides and were **4 percent more likely** to have drivers cancel on them
  - Even worse for taxi service
- Other reported incidents of unfair treatment of passengers

Woman Says Uber Driver Denied Her Ride Because of Her Wheelchair

She said she explained to the driver that the chair could be broken down easily and it wasn't much bigger than a piece of luggage but the driver kept insisting there was no space for it

Ont. woman says Uber driver rejected her guide dog

Mitigation by Rideshare Platforms

- Riders' photo and destination are hidden from the driver until they accept/reject the request
  - Recently, Uber has started showing destination to the drivers **before** they accept[1]
- Penalty is imposed if drivers cancel a certain number of trips after initially accepting them
  - Eg: temporarily suspending a driver from the platform

Not Enough!

[1] https://www.uber.com/blog/california/keeping-you-in-the-drivers-seat-1/

Mitigation by Rideshare Platforms - Not Enough!

- Riders’ photo and destination are hidden from the driver until they accept/reject the request

    Some drivers start the trip moments before picking up a passenger to see the destination

- Penalty is imposed if drivers cancel a certain number of trips after initially accepting them

    In some cases, drivers intentionally delay the pickup, thus forcing the rider to cancel the trip

Profit and Fairness Tradeoffs

- In this work we consider the **peak hour**
  - Demand (passengers) >> Supply (drivers)
- Drivers can afford to be selective, thus putting certain groups (eg: those living in “unpopular” parts of a city) in an **unfair** position
- To ensure **fairness**, the platform should aggressively match such trips
- However, this could potentially lead to less profit
  - Eg: Drivers leave the platform because it assigns them unfavorable requests

Can we balance these conflicting goals?

Our Contributions

- Formalize a metric of fairness relevant in this setting
- Present a **provably efficient** online matching algorithm
  - Performs better than a reasonable lower bound on **both** profit and fairness objectives
  - Includes driver's preference for rides
- Evaluate the proposed algorithm on real-world rideshare data

Setup

Set of driver **types** (Offline)
$U = \{u_1, u_2, \ldots\}$

Set of request **types** (Online)
$V = \{v_1, v_2, \ldots\}$

Upon arrival of a request v, an **immediate and irrevocable** assignment must be made

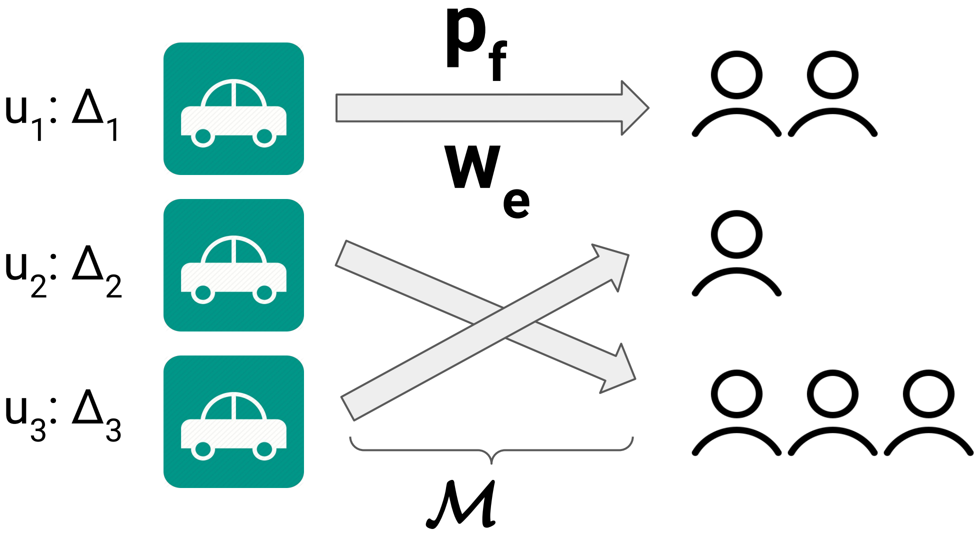

v_1 : $start_1$, end_1, $r_v = 2$

v_2 : $start_2$, end_2, $r_v = 1$

v_3 : $start_3$, end_3, $r_v = 3$

Total Requests (T) = 6

Fairness Measure: $\min_{v \in V} \frac{\mathbb{E}[|\mathcal{M}_v|]}{r_v}$

Profit Measure: $\mathbb{E}[\sum_{e \in \mathcal{M}} w_e]$

Optimal Solution

Consider both sides to be offline

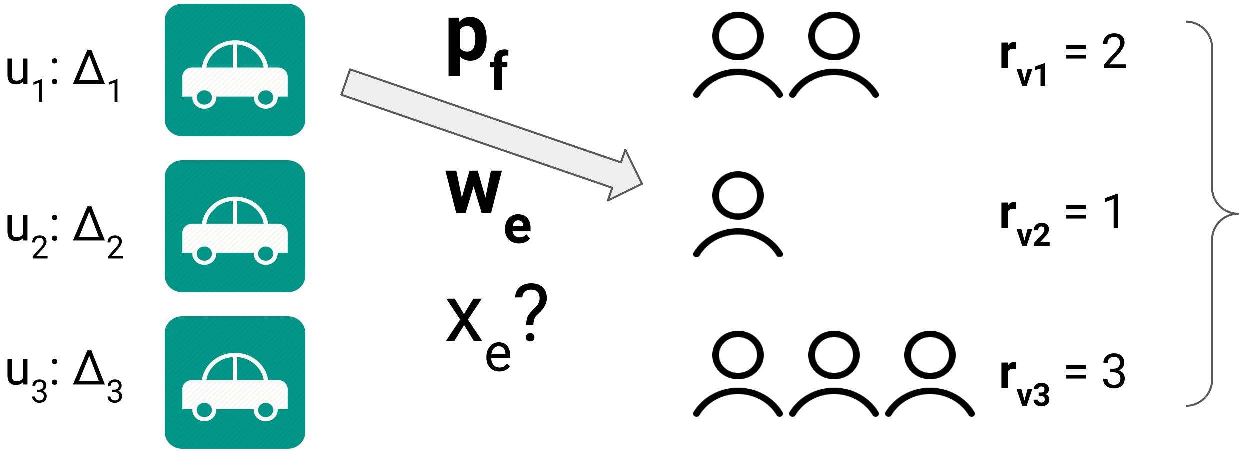

What are the number of rider to driver assignments for the optimal fair solution (x^*_{fair}) and optimal profit solution (x^*_{profit}) respectively?

Fairness Measure: $\min_{v \in V} \frac{\mathbb{E}[|\mathcal{M}_v|]}{r_v}$ Profit Measure: $\mathbb{E}[\sum_{e \in \mathcal{M}} w_e]$

Optimal Solution

- We model profit and fairness objectives as maximization problems
  - Subject to feasibility constraints!
- We solve these using linear programming to obtain $\mathbf{x}^*_{\mathbf{fair}}$ and $\mathbf{x}^*_{\mathbf{profit}}$
- Use these optimal solutions to guide the online algorithm

Online Algorithm

- We propose a parameterized matching algorithm **ALG ($\boldsymbol{\alpha}, \boldsymbol{\beta}$)**: $\boldsymbol{\alpha}$ controls profit and $\boldsymbol{\beta}$ controls fairness
- Guided by a linear combination of x^*_{profit} and x^*_{fair} : $\boldsymbol{\alpha} x^*_{\text{profit}} + \boldsymbol{\beta} x^*_{\text{fair}}$

- We prove that **ALG** achieves a competitive ratio of at least $\boldsymbol{\alpha}$**/e** for profit and at least $\boldsymbol{\beta}$**/e** for fairness (for $\alpha + \boldsymbol{\beta} \leq 1$)
- We also prove that no (non-adaptive) algorithm can achieve ($\alpha, \boldsymbol{\beta}$) competitive ratio simultaneously on the profit and fairness with $\alpha + \boldsymbol{\beta} > 1 - 1/e$

Online Algorithm

- We propose a parameterized matching algorithm **ALG ($\boldsymbol{\alpha}, \boldsymbol{\beta}$)**: $\boldsymbol{\alpha}$ controls profit and $\boldsymbol{\beta}$ controls fairness
- Guided by a linear combination of x^*_{profit} and x^*_{fair} : $\boldsymbol{\alpha} x^*_{profit} + \boldsymbol{\beta} x^*_{fair}$

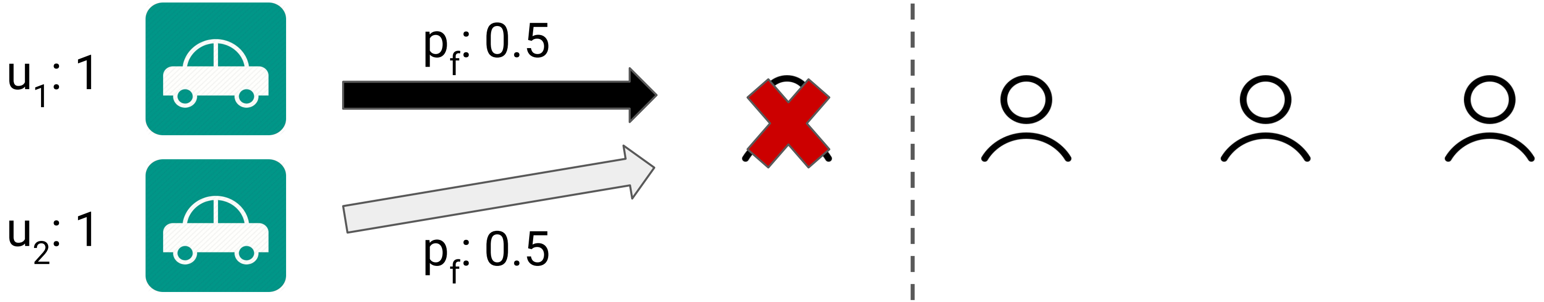

Starting location:	Manhattan	Manhattan	Manhattan	Manhattan
Ending location:	Not-popular	Not-popular	Popular	Popular
Estimated profit:	1	2	10	20

Online Algorithm

- We propose a parameterized matching algorithm **ALG ($\boldsymbol{\alpha}, \boldsymbol{\beta}$)**: $\boldsymbol{\alpha}$ controls profit and $\boldsymbol{\beta}$ controls fairness
- Guided by a linear combination of x^*_{profit} and x^*_{fair} : $\boldsymbol{\alpha} x^*_{profit} + \boldsymbol{\beta} x^*_{fair}$

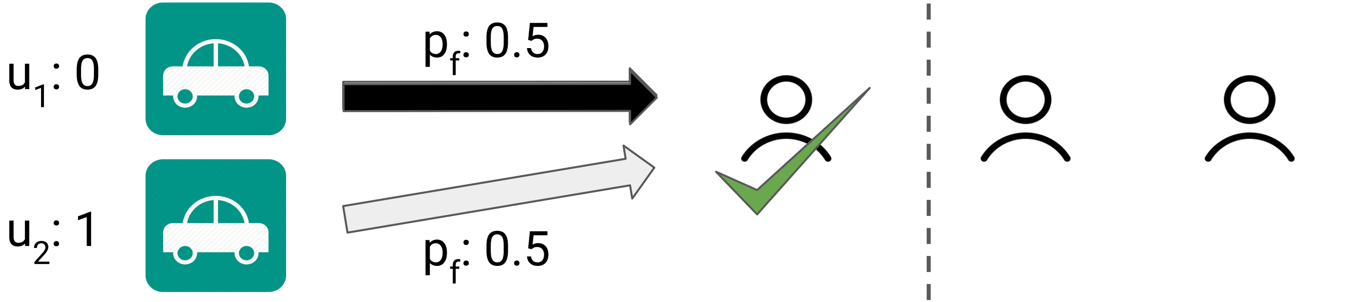

Starting location:	Manhattan	Manhattan	Manhattan
Ending location:	Not-popular	Popular	Popular
Estimated profit:	2	10	20

Online Algorithm

- We propose a parameterized matching algorithm **ALG ($\boldsymbol{\alpha}, \boldsymbol{\beta}$)**: $\boldsymbol{\alpha}$ controls profit and $\boldsymbol{\beta}$ controls fairness
- Guided by a linear combination of x^*_{profit} and x^*_{fair} : $\boldsymbol{\alpha} x^*_{profit} + \boldsymbol{\beta} x^*_{fair}$

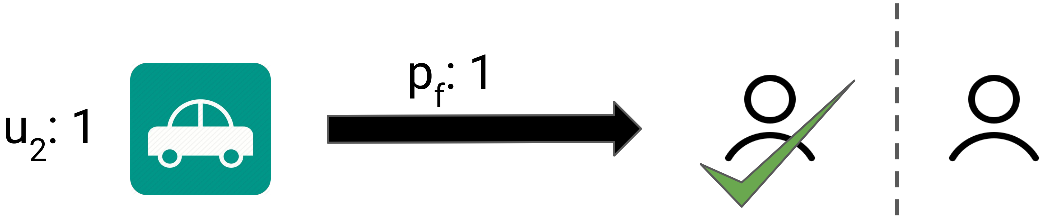

Starting location:	Manhattan	Manhattan
Ending location:	Popular	Popular
Estimated profit:	10	20

Online Algorithm

- We propose a parameterized matching algorithm **ALG ($\boldsymbol{\alpha}, \boldsymbol{\beta}$)**: $\boldsymbol{\alpha}$ controls profit and $\boldsymbol{\beta}$ controls fairness
- Guided by a linear combination of x^*_{profit} and x^*_{fair} : $\boldsymbol{\alpha} x^*_{\text{profit}} + \boldsymbol{\beta} x^*_{\text{fair}}$

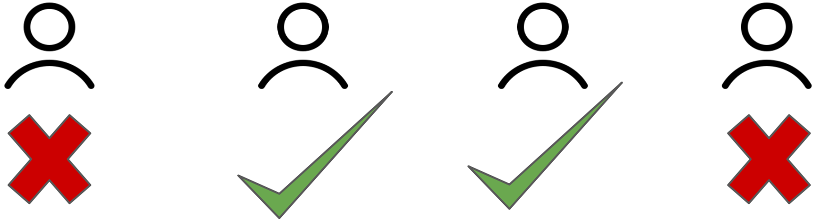

Starting location:	Manhattan	Manhattan	Manhattan	Manhattan
Ending location:	Not-popular	Not-popular	Popular	Popular
Estimated profit:	1	2	10	20

Experiments

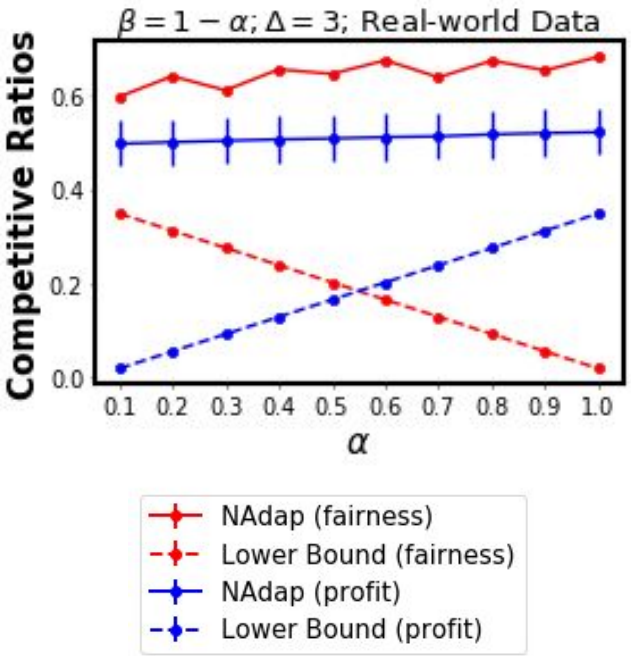

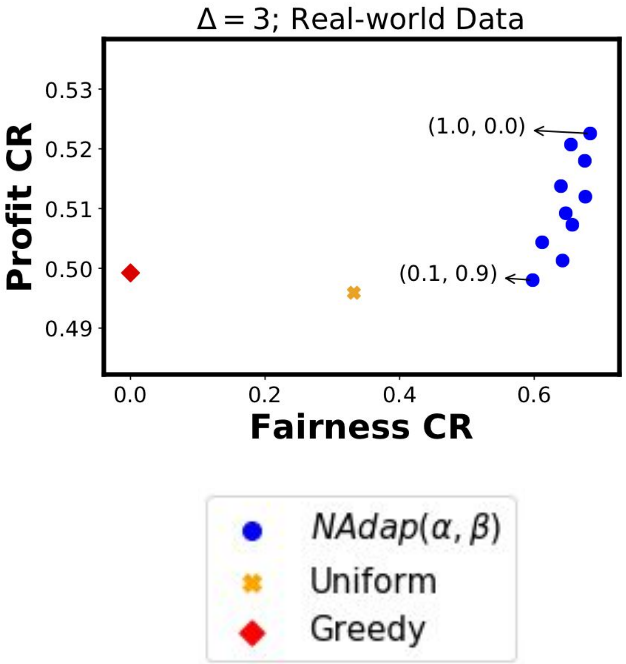

Conclusion

- We propose a fairness metric in the setting of peak-hour
- We propose a provable, efficient and flexible online algorithm that can achieve reasonable performance on both profit and fairness objectives simultaneously
- We validate our theoretical results with experiments on real world data

Questions?

Code: tinyurl.com/rideshare-code Paper: tinyurl.com/rideshare-fairness

vedant@cs.umd.edu @_nvedant_