ridesharing: the road ahead

Sid Banerjee

RTDM reunion, June 2019

School of ORIE, Cornell University

where the ride started

Dynamic Pricing in Rideshare Platforms

Sid Banerjee

School of Operations Research and Information Engineering
Cornell University

Simons Workshop on Real-Time Decision Making, June 2016

Joint work with Carlos Riquelme and Ramesh Johari (Stanford)
and the Data Science team at Lyft

Sid Banerjee (Cornell ORIE) | Pricing in Rideshare Platforms | Simons Workshop | 1 / 19

The Bigger Picture

> *Over the next 10 years, the major breakthrough of economics will be in applications of market design, which improves the efficiency of markets using a combination of game theory, economics and algorithm design. We've already seen fruitful application in search auctions, spectrum auctions, kidney exchange and school assignment.*
> *(2016 will be the year that) Silicon Valley recognizes that the value of Uber is its marketplace, not the data...*

Preston McAfee

shout-out to all my co-passengers

Daniel Freund

Raga G

Chamsi Hssaine

Ramesh Johari

Yash Kanoria

Thodoris Lykouris

Pengyu Qian

Carlos Riquelme

Samitha Samaranayake

Thibault Séjourné

and a few important acknowledgments

special shout-out to

– the amazing folks in the lyft research science team

– ARO (W911NF-17-1-0094) & NSF (ECCS1847393, DMS1839346) support

(stochastic) model for ridesharing

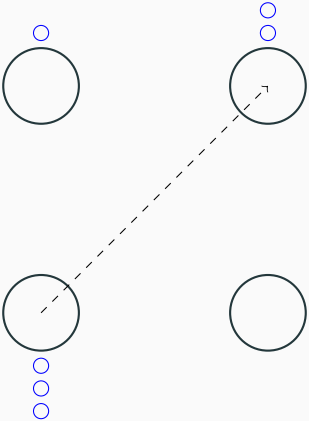

- K units (cars) across n stations (closed network)
- system state $\in S_{n,K} = \{(x_i)_{i\in[n]} | \sum_{i=1}^{n} x_i = K\}$
- $i \to j$ passengers arrive via Poisson process with rate ϕ_{ij}

(stochastic) model for ridesharing

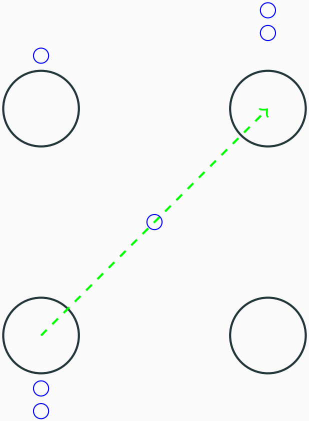

- passenger requests ride if offered price is acceptable
- matched to idle unit, which then travels to destination
- trips have independent travel-times

(stochastic) model for ridesharing

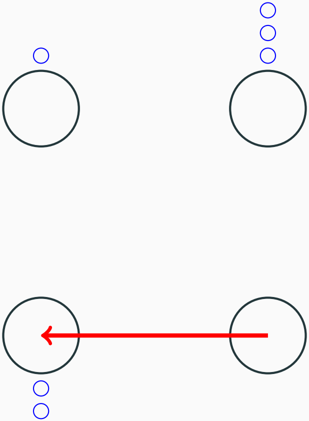

- myopic customers: abandon system if unit unavailable

platform objective

maximize long-term average welfare/revenue

control levers for ridesharing

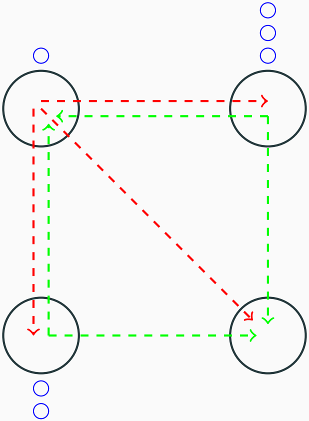

- pricing
  - modulates demand between locations
  - dynamic, state-dependent

control levers for ridesharing

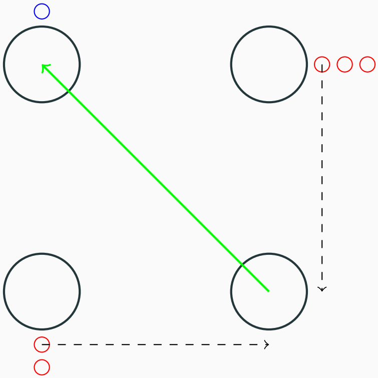

- dispatch: choose ‘nearby’ car to serve demand
  - can use any car within ‘ETA target’

control levers for ridesharing

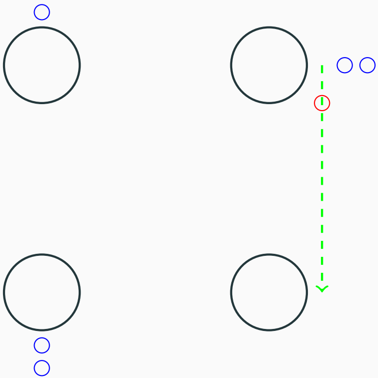

- rebalancing: re-direct free car to empty location
  - incur a cost for moving the car
  - driver ‘nudges’ (heat-maps), autonomous vehicles

key assumptions

assumption 1: timescales of platform operations

# of cars, arrival rates, demand elasticities remain constant over time

– time-varying rates (re-solve policies at change-points...)

– drawbacks: driver entry/exit behavior, bursty arrivals

assumption 2: timescales of strategic interactions

– passengers abandon if price too high/no vehicle

– drivers react at longer timescales

assumption 3: availability of data

platform has perfect knowledge of arrival rates, demand elasticities

– can be relaxed via online learning techniques

flow model for ridesharing

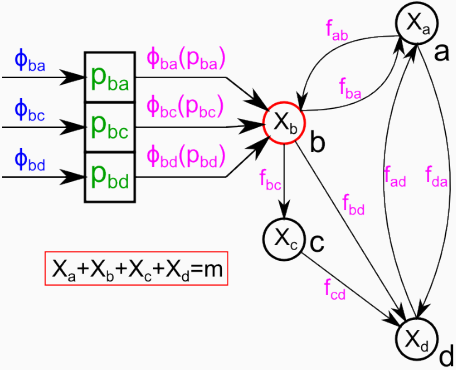

flow model for ridesharing Markov chain

flows of cars in network

– occupied car-flows + empty-car flows (rebalancing)

– satisfies admitted passenger flows on edges

– flow-balance constraints at nodes

from online policy to fluid model

how well can we approximate the flow model via online controls?

theorem [B, Freund & Lykouris 2017]

flow relaxation gives state-independent dispatch policy which is

- $1 + \frac{n-1}{K}$ approximate (with instantaneous trips)
- $1 + O\left(\frac{1}{\sqrt{K}}\right)$ approximate (with travel-times)

from online policy to fluid model

how well can we approximate the flow model via online controls?

theorem [B, Freund & Lykouris 2017]

flow relaxation gives state-independent dispatch policy which is

- $1 + \frac{n-1}{K}$ approximate (with instantaneous trips)
- $1 + O\left(\frac{1}{\sqrt{K}}\right)$ approximate (with travel-times)

- extends to pricing, rebalancing controls, most objectives
- large-market optimality: factor goes to 1 as system scales

theorem [B, Kanoria & Qian 2018]

family of state-dependent dispatch policies which are

- $1 + e^{-\Theta(K)}$ approximate (for large K, instantaneous trips)
- convex program gives optimal exponent

from online policy to fluid model

how well can we approximate the flow model via online controls?

theorem [B, Freund & Lykouris 2017]

flow relaxation gives state-independent dispatch policy which is

- $1 + \frac{n-1}{K}$ approximate (with instantaneous trips)
- $1 + O\left(\frac{1}{\sqrt{K}}\right)$ approximate (with travel-times)

theorem [B, Kanoria & Qian 2018]

family of state-dependent dispatch policies which are

- $1 + e^{-\Theta(K)}$ approximate (for large K, instantaneous trips)
- convex program gives optimal exponent

survey chapter: Ride Sharing, B & Johari
in Sharing Economy, Springer Series in Supply Chain Management

our initial ETA prediction

Summary

Main takeaways

- Throughput of Dynamic Pricing (in the large-market limit) is:
  - Bounded by throughput of optimal static pricing
  - Robust to perturbations in system parameters
- Results extend to revenue, networks of ride-sharing queues

The bigger picture

In many online marketplaces:

platform optimization $\Leftrightarrow$ controlling the equilibrium

- New sources of data, real-time monitoring and control tools
- Many questions: information displays, long-term contracts, competition between platforms, interaction with other systems, etc.

are we there yet?

the price of demand fragmentation

price of fragmentation in ridesharing ecosystems

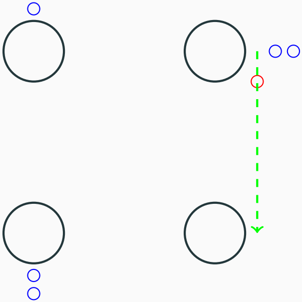

- what is the 'societal cost' of decentralized optimization?
  - multiple platforms with (random) exogenously partitioned demands
  - individual platforms do optimal empty-vehicle rebalancing

price of fragmentation in ridesharing ecosystems

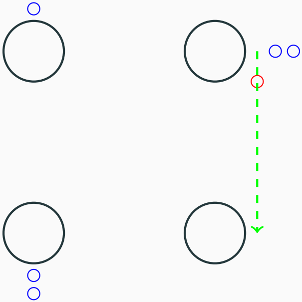

- what is the 'societal cost' of decentralized optimization?
  - multiple platforms with (random) exogenously partitioned demands
  - individual platforms do optimal empty-vehicle rebalancing

price of fragmentation

under exogenous demand split, increase in rebalancing costs of multiple platforms vs. single platform (under large-market scaling)

formal model

- n nodes, distance matrix $D = (d_{ij})_{(i,j)}$
- steady-state customer flows: for each edge (i,j), demand Λ_{ij}
- total demand at node $\Lambda_i = \sum_j (\Lambda_{ij} - \Lambda_{ji}) = (A.\Lambda)_i$

rebalancing cost

$$RC(\Lambda): \quad \min \quad \sum_{(i,j)} d_{ij} x_{ij}$$
$$\text{s.t. } \forall i, \ \sum_j (x_{ij} - x_{ji}) = \Lambda_i$$

formal model

- n nodes, distance matrix $D = (d_{ij})_{(i,j)}$
- steady-state customer flows: for each edge (i,j), demand Λ_{ij}
- total demand at node $\Lambda_i = \sum_j (\Lambda_{ij} - \Lambda_{ji}) = (A.\Lambda)_i$

rebalancing cost

$$RC(\Lambda): \quad \min \quad \sum_{(i,j)} d_{ij} x_{ij}$$
$$\text{s.t. } \forall i, \ \sum_j (x_{ij} - x_{ji}) = \Lambda_i$$

price of fragmentation

operational cost scaling under a duopoly vis-a-vis monopoly:

$$\gamma^\theta = \mathbb{E}_{\lambda^\theta \sim \mathbb{P}^\theta}[RC(\lambda^\theta) + RC(\Lambda - \lambda^\theta)] - RC(\Lambda)$$

$\theta =$ demand scaling parameter $(\lambda^\theta = \theta\Lambda)$
$\mathbb{P}^\theta =$ demand splitting process

counterfactual simulation: NYC taxi data

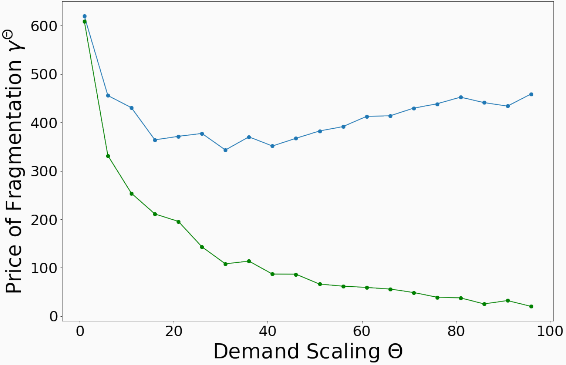

$\log(\gamma^\theta)$ depending on $\log(\theta)$. TLC Data clustered into 40 stations

price of fragmentation in ridesharing markets

what we show

as demand scales, the price of fragmentation undergoes a phase transition based on structure of underlying demand flows

– both regimes observed in NYC data (≈ 10% fragmentation-affected)

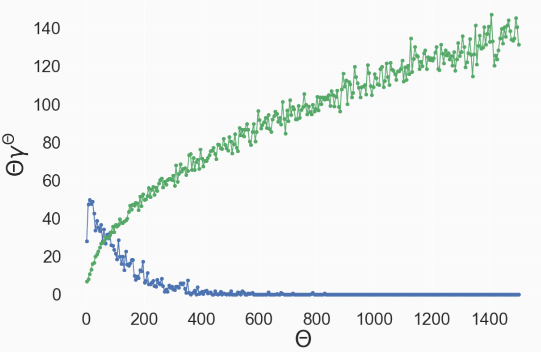

the price of fragmentation: formal result

theorem [Séjourné, Samaranayake & B 2018]

Depending on the spatial homogeneity of market shares and the dual degeneracy of the demand, the PoF undergoes a phase transition from an exponential decay to a square root/ linear divergence

- this results remains true for a wide family of splitting processes
- It is true independently of the number of competing firms

- sufficient condition for dual degeneracy: presence of locally balanced clusters, i.e. disconnected components in the rebalancing flow graph

geometric intuition

dual formulation for RC(Λ)

$$\max \quad \sum_i \alpha_i.\Lambda_i \qquad \text{s.t. } \forall(i,j),\ \alpha_i - \alpha_j \leq d_{ij}$$

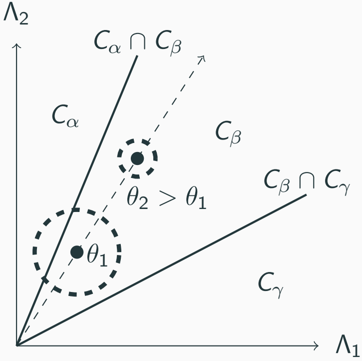

Fragmentation-resilient regime

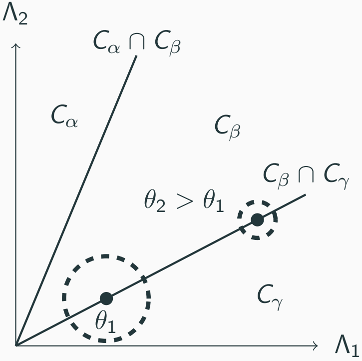

Fragmentation-affected regime

warning: affects numerical simulations in unpredictable ways

fraction of fragmentation-affected regimes depends on data-aggregation granularity(number of stations/time interval)

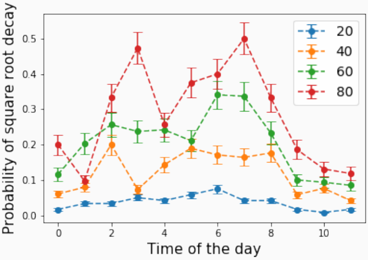

effect of spatial granularity

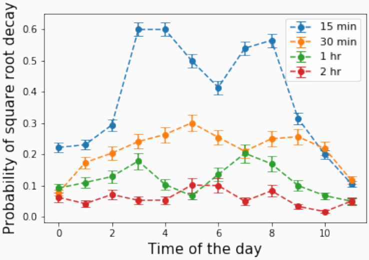

effect of temporal granularity

summary and thoughts

- demand fragmentation has non-trivial effects on overall efficiency
  - phase transition in rebalancing costs
  - depends on structure of demand
- fragmentation reduces when demand has
  - increased spatial homogeneity
  - coarser granularity
- multi-homing is a potential solution
- affected regime caused by populations more suitable for mass transit
- may affect data-driven measurements, numerical studies

designing a transit marketplace

so did ridesharing 'solve' transit?

not yet, but...

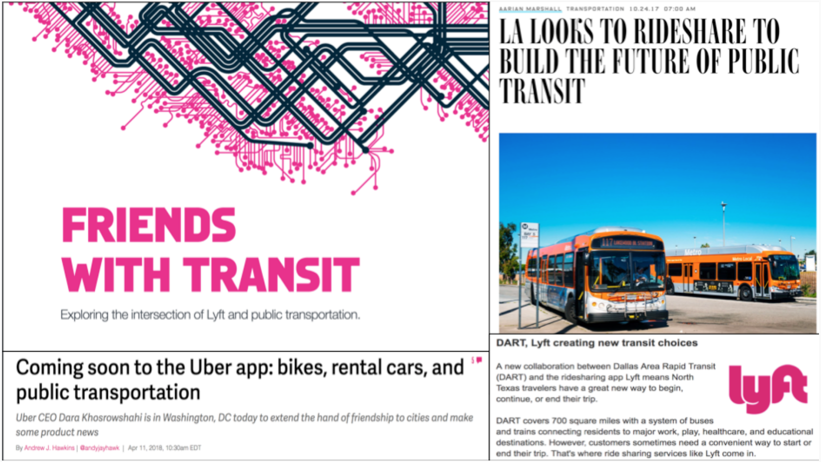

FRIENDS WITH TRANSIT

Exploring the intersection of Lyft and public transportation.

Coming soon to the Uber app: bikes, rental cars, and public transportation

Uber CEO Dara Khosrowshahi is in Washington, DC today to extend the hand of friendship to cities and make some product news

By Andrew J. Hawkins | @andyjayhawk | Apr 11, 2018, 10:30am EDT

AARIAN MARSHALL TRANSPORTATION 10.24.17 07:00 AM

LA LOOKS TO RIDESHARE TO BUILD THE FUTURE OF PUBLIC TRANSIT

DART, Lyft creating new transit choices

A new collaboration between Dallas Area Rapid Transit (DART) and the ridesharing app Lyft means North Texas travelers have a great new way to begin, continue, or end their trip.

DART covers 700 square miles with a system of buses and trains connecting residents to major work, play, healthcare, and educational destinations. However, customers sometimes need a convenient way to start or end their trip. That's where ride sharing services like Lyft come in.

the transit marketplace

transit marketplace: the network

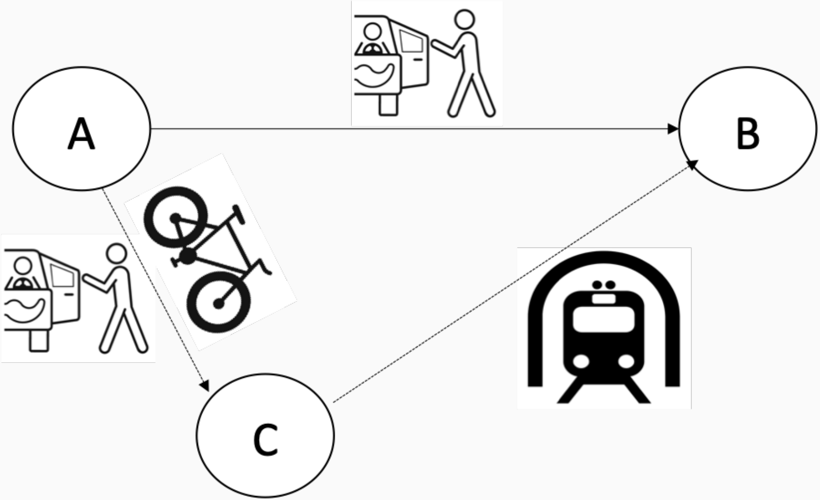

each commuter has a public type
– type = vector of valuations, one for each multi-modal option
– we normalize transit value to 0

market chooses price menu: price for each multi-modal option

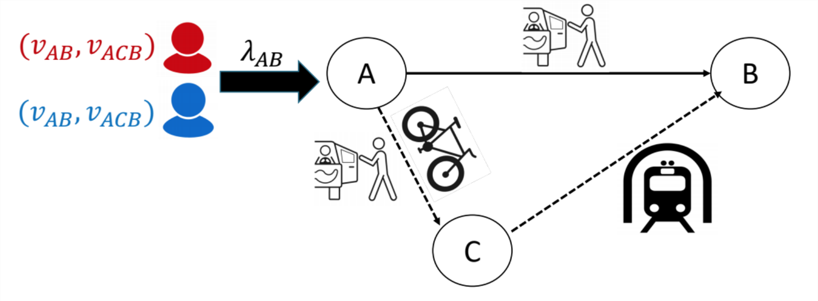

transit marketplace: objectives

operational objective

reduce frictions, improve reliability for multi-modal trips

transit marketplace: objectives

operational objective

reduce frictions, improve reliability for multi-modal trips

economic objective

set prices to maximize overall social welfare

transit marketplace: objectives

operational objective

reduce frictions, improve reliability for multi-modal trips

economic objective

set prices to maximize overall social welfare

is this sufficient?

pareto improvement as a desiderata for markets

Buy-in from *all* parties (providers and commuters) necessary for success

The Boston Globe

COMPUTERS CAN SOLVE YOUR PROBLEM. YOU MAY NOT LIKE THE ANSWER.

What happened when Boston Public Schools tried for equity with an algorithm

The New York Times

'Airbnb Tax' in N.J. Opens New Front in Battle Over Internet Economy

transit marketplace: objectives

operational objective

reduce frictions, improve reliability for multi-modal trips

economic objective

set prices to maximize overall social welfare AND ensure pareto improvement for all

transit marketplace: incorporating PI constraints

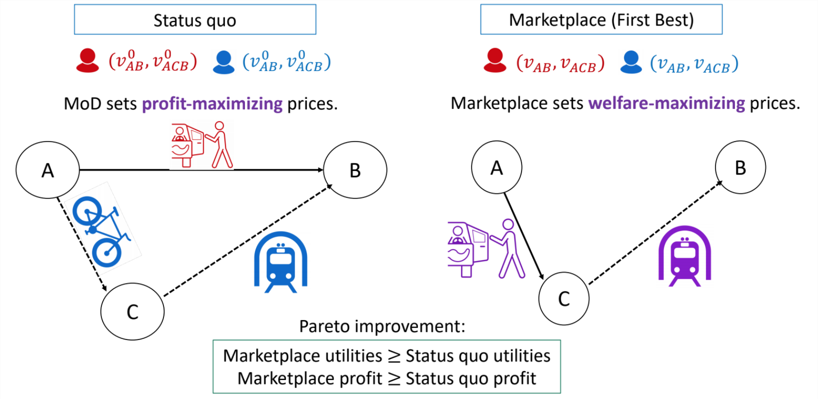

problem: this may be incompatible! (Myerson-Satterthwaite)

transit marketplace: formal model

$$
\begin{aligned}
\max_{\boldsymbol{p}} \quad & \mathcal{W}(\boldsymbol{p}) \\
\text{subject to} \quad & \text{Optimal for MoD to serve full demand} \\
& \sum_{m} (v_{\theta m} - p_m)\, x_{\theta m}(\boldsymbol{p}) \geq U_\theta^0 \quad \forall\, \theta \\
& \sum_{(s,t)} \sum_{m} (p_m - c_m)\phi_m(\boldsymbol{p}) - \sum_{(i,j)} c_{ij} r_{ij}(\boldsymbol{p}) \geq \Pi^0
\end{aligned}
$$

Poly-time (LP)

Pareto improvement constraints introduce *non-convexity*

transit marketplace: example result

1. Find welfare-maximizing prices (via LP).
2. Raises prices until one of three things happens:
   i. Efficient allocation is changed
   ii. Commuter-PI is violated
   iii. Enough money is raised

System of linear inequalities
Linear in size of input

3. Final prices:

$$p_m = p_m^* + \boxed{\Delta_{st}}$$

"Marketplace surcharge"

transit marketplace: example result

1. Find welfare-maximizing prices (via LP).
2. Raises prices until one of three things happens:
   i. Efficient allocation is changed
   ii. Commuter-PI is violated
   iii. Enough money is raised
3. Final prices:

System of linear inequalities
Linear in size of input

$$p_m = p_m^* + \boxed{\Delta_{st}}$$

"Marketplace surcharge"

Informal Theorem. If there exist surcharges such that, for all commuters allocated a mode in the efficient allocation:

Worst-case welfare generated by commuter in marketplace

$$\max_m \left\{ \boxed{v_{\theta m}} - \boxed{\sum_{(i,j)\in E_m} (c_{ij} + c_{ji})} \right\} \geq \text{Status quo utility} + \text{Surcharge}$$

Valuation of commuter for mode m

Cost of m + local rebalancing

Welfare generated by commuter in status quo

and the surcharges make up the status quo profit, then First Best is Pareto-improving.

– can also get bicriteria approximations in general settings

and so we hit the road again

our journey so far

- unified flow-based models for ridesharing
  – capture main operational and market-design details – guide for designing good online controls
  – sandbox for studying more complex problems
- operational losses from demand fragmentation
- the promise and challenges of multi-modal transit marketplaces

open questions

- detailed scaling limits of control policies; elementary proofs of scaling
- value of information: prediction-guided control, maxmin guarantees
- appropriate mix of drivers and autonomous cars
- ridesharing + society

Thanks!