ridesharing

Sid Banerjee

School of ORIE, Cornell University



based on work with – D. Freund, T. Lykouris (Cornell),
– C. Riquelme & R. Johari (Stanford),

special thanks to the data science team at Lyft

ridesharing platforms

- critical components of modern urban transit
- crucible for Real-Time Decision Making/Ops Management/EconCS

ridesharing: overview

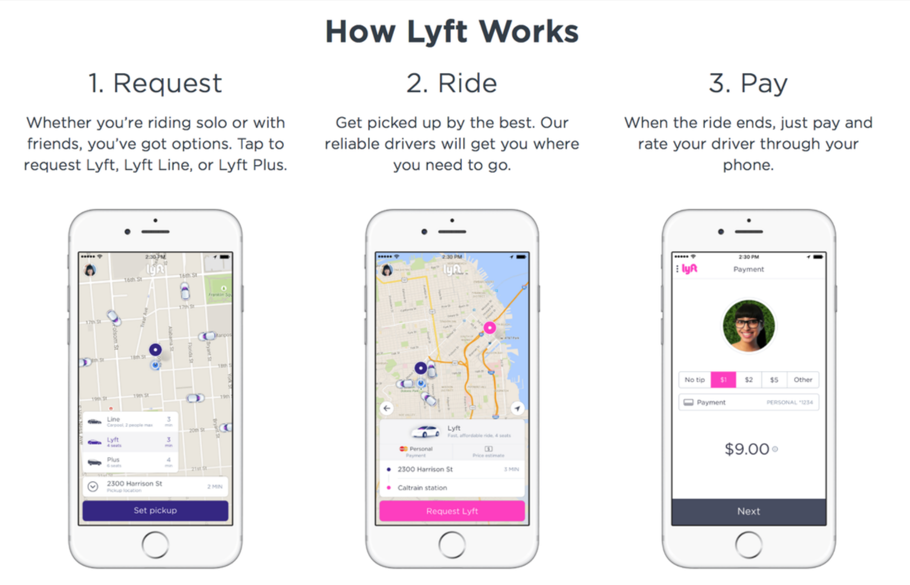

credit: lyft.com

ridesharing: pricing

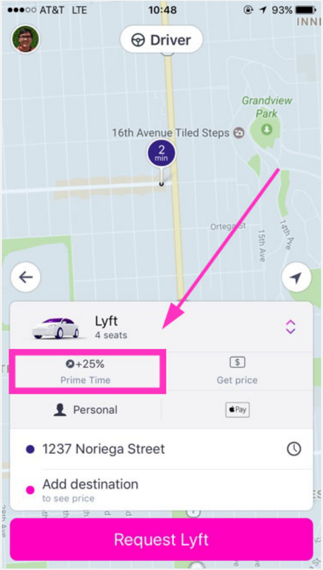

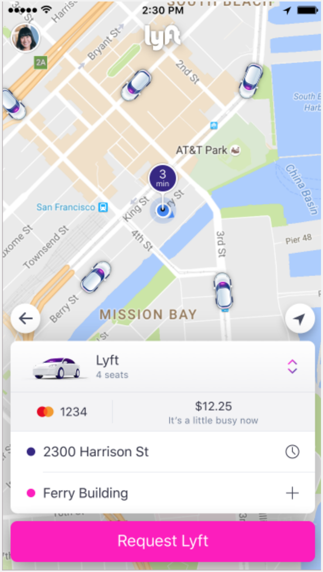

credit: lyft.com

rideshare platforms: pricing

Thanks for riding with [illegible]!

Ride ending January 31 at 12:35 AM

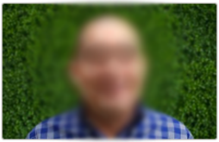

Pickup:	[illegible]
Dropoff:	[illegible]

Ride *2.5 mi & 10 min*:	$8.28
Prime Time*:	$2.07
Trust & Safety Fee:	$1.50
Total charged to [illegible]	**$11.85**

*25% Prime Time was included in your total. Prime Time encourages more people to drive when Lyft gets really busy.

Learn More

credit: lyft.com

rtdm in ridesharing: mapping

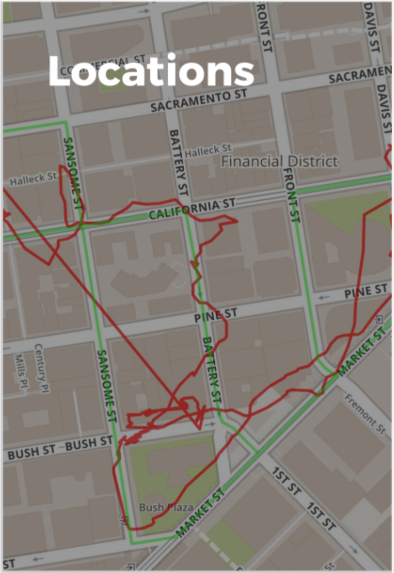

credit: lyft data science team

rtdm in ridesharing: logistics

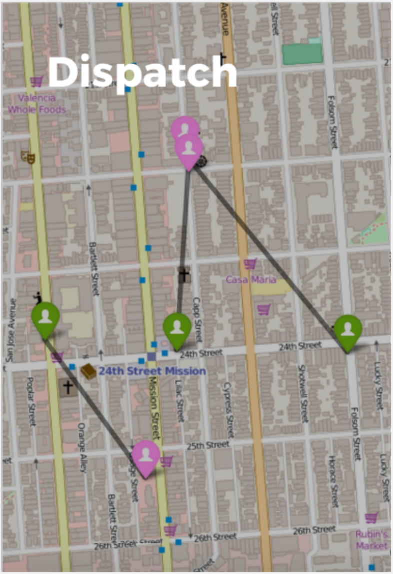

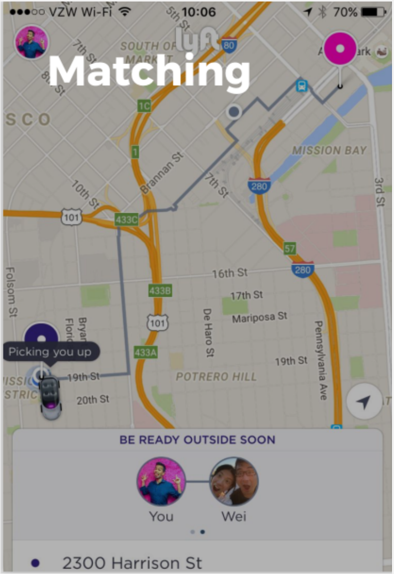

credit: lyft data science team

rtdm in ridesharing: market design

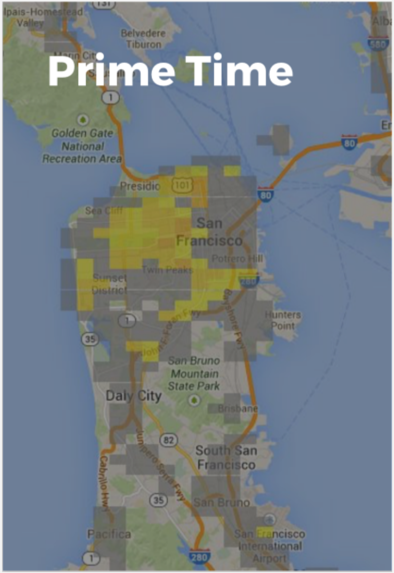

credit: lyft data science team

the bigger picture: on-demand transportation

– fast operational timescales; complex network externalities
– new control-levers: dynamic pricing/dispatch, incentives, pooling
– new(er) challenges: competition, effect on public transit, urban planning

the bigger picture: on-demand transportation

citi bike operated by motivate

CAR 2GO

lyft

– fast operational timescales; complex network externalities
– new control-levers: dynamic pricing/dispatch, incentives, pooling
– new(er) challenges: competition, effect on public transit, urban planning

this talk

- 'where do we come from?'
  - simple framework for ridesharing: data, state, controls
- 'where are we?'
  - approximate optimal control for ridesharing logistics
  - market mechanisms as a tool for algorithmic self-calibration
- 'where are we going?'

main challenge: rebalancing

demand heterogeneity ⇒ non-uniform supply across space and time

logistical 'solution': rebalance the vehicle fleet
economic 'solution': incentives for passengers and drivers

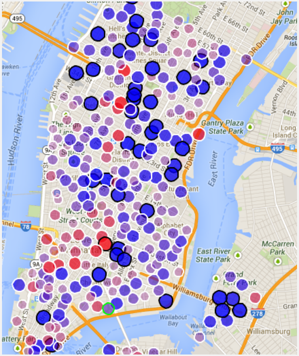

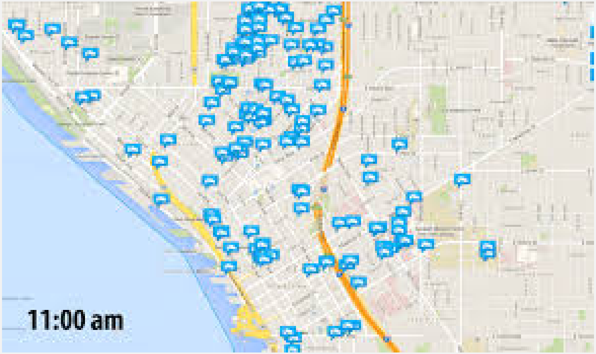

main challenge: rebalancing

demand heterogeneity ⇒ non-uniform supply across space and time

logistical 'solution': rebalance the vehicle fleet
economic 'solution': incentives for passengers and drivers
control-levers: pricing/incentives, dispatch, empty-car rebalancing

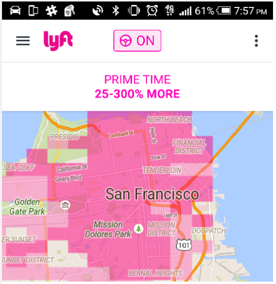

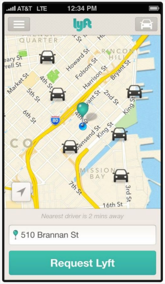

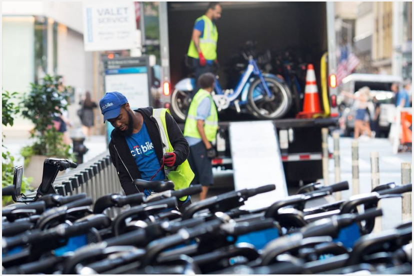

(stochastic-network) model for ridesharing

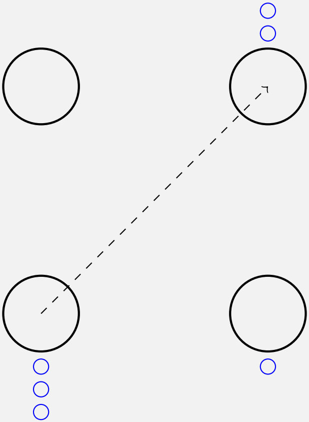

- *m* units (cars) across *n* stations (here, we have $m = 6, n = 4$)
- system state $\in S_{n,m} = \{(x_i)_{i\in[n]} | \sum_{i=1}^{n} x_i = m\}$
- $i \to j$ passengers arrive via Poisson process with rate ϕ_{ij}

(stochastic-network) model for ridesharing

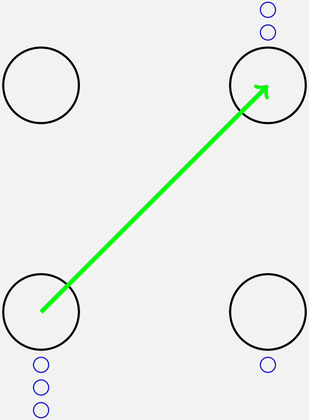

- platform sets state-dependent prices $p_{ij}(\mathbf{X})$
- quantile $q_{ij}(\mathbf{X}) = 1 - F_{ij}(p_{ij}(\mathbf{X}))$: fraction willing to pay $p_{ij}(\mathbf{X})$

(stochastic-network) model for ridesharing

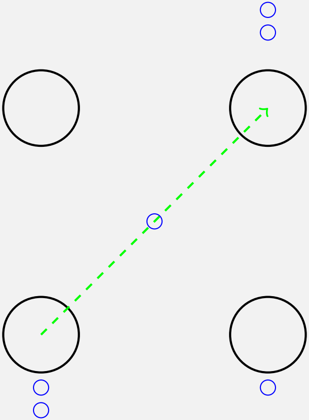

- car travels with passenger to destination
- (this talk: assume travel-times are zero)

(stochastic-network) model for ridesharing

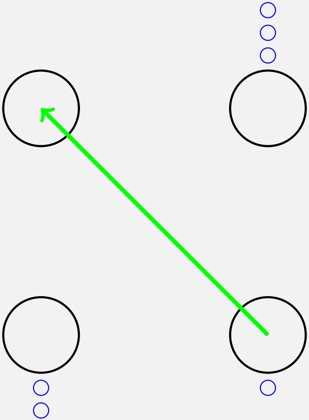

(stochastic-network) model for ridesharing

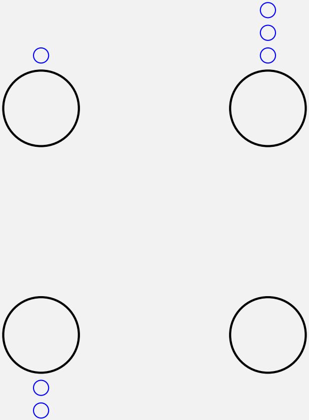

(stochastic-network) model for ridesharing

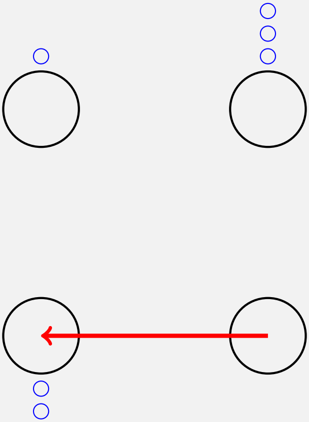

- myopic customers: abandon system if
  – vehicle unavailable

(stochastic-network) model for ridesharing

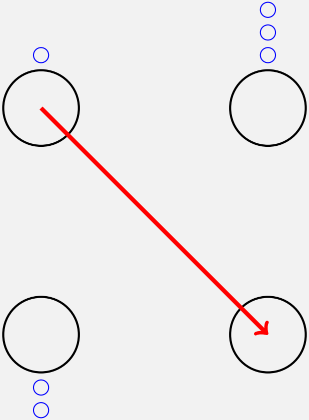

- myopic customers: abandon system if
  – vehicle unavailable or
  – price too high

(stochastic-network) model for ridesharing

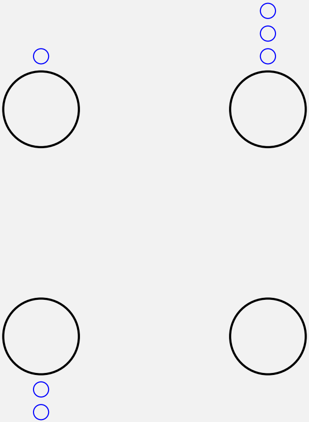

- objective:
  - optimize chosen long-run average system objective
  - objectives: revenue, welfare, customer engagement, etc.

control levers for ridesharing

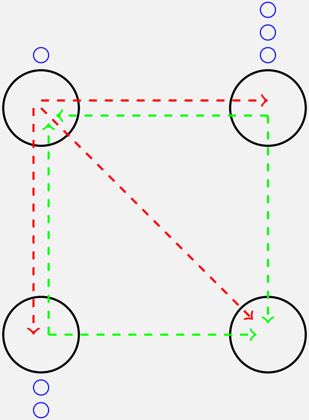

- pricing
  – modulates demand between locations
  – dynamic, state-dependent

control levers for ridesharing

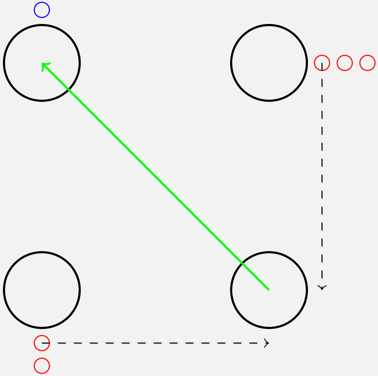

- dispatch: choose 'nearby' car to serve demand
  – can use any car within 'ETA target'

control levers for ridesharing

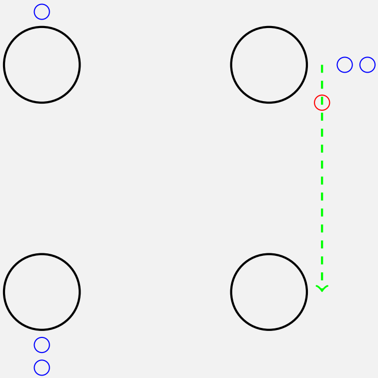

- rebalancing: re-direct free car to empty location
  – incur a cost for moving the car
  – driver 'nudges' (heat-maps), autonomous vehicles

intermezzo: why model?

scales and economics

– need controls that work in real-time, at large-scales

– complex controls need more resources; non-commensurate (?) impact

intermezzo: why model?

scales and economics

– need controls that work in real-time, at large-scales

– complex controls need more resources; non-commensurate (?) impact

known(?) unknowns

– errors in estimation and forecasting

– difficulties in learning demand/supply curves

intermezzo: why model?

scales and economics

– need controls that work in real-time, at large-scales
– complex controls need more resources; non-commensurate (?) impact

known(?) unknowns

– errors in estimation and forecasting
– difficulties in learning demand/supply curves

unknown unknowns

intermezzo: why this model?

assumption 1: timescales of platform operations

number of cars, arrival rates, demand elasticities remain constant over time
– time-varying rates (re-solve policies at change-points…)
– driver entry/exit behavior
– effect of bursty arrivals?

intermezzo: why this model?

assumption 1: timescales of platform operations

number of cars, arrival rates, demand elasticities remain constant over time
– time-varying rates (re-solve policies at change-points. . .)
– driver entry/exit behavior
– effect of bursty arrivals?

assumption 2: timescales of strategic interactions

– passengers abandon if price too high/no vehicle
– drivers react at longer timescales

intermezzo: why this model?

assumption 1: timescales of platform operations

number of cars, arrival rates, demand elasticities remain constant over time
– time-varying rates (re-solve policies at change-points…)
– driver entry/exit behavior
– effect of 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
– is that really true?
– is that really needed?

data-driven optimization for vehicle-sharing

Pricing and Optimization in Shared Vehicle Systems
Banerjee, Freund & Lykouris (2016)
https://arxiv.org/abs/1608.06819

model recap

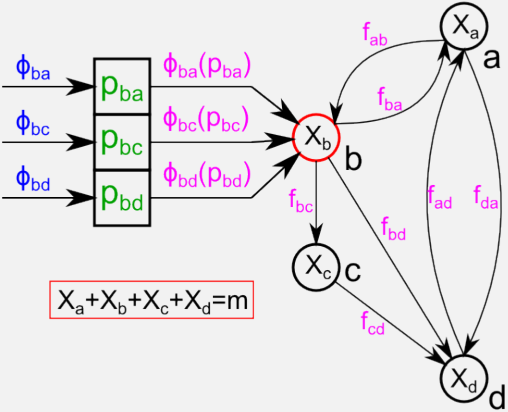

- m units spread across n nodes
- control: state-dependent pricing policy $\vec{p} = \{p_{ij}(\mathbf{x})\}$ (or quantiles $\vec{q}$)
- flows of cars in network: realized via Markov chain dynamics

technical challenges

objective

$$\max_{\mathbf{q}=\{q_e(\mathbf{x})\}} \underbrace{\sum_{\mathbf{x}} \pi_{\vec{q}}(\mathbf{x})}_{\text{long-run avg under control } \mathbf{q}} \left(\sum_{e=(i,j)} \mathbb{E}[\text{reward rate from } i \to j \text{ rides}] \right)$$

technical challenges

objective

$$\max_{\mathbf{q}=\{q_e(\mathbf{x})\}} \sum_{\mathbf{x}} \pi_{\vec{q}}(\mathbf{x}) \left(\sum_{e=(i,j)} \underbrace{\mathbb{1}_{[x_i>0]}}_{\text{availability at } i} \cdot \ \phi_e q_e(\mathbf{x}) \quad \underbrace{\cdot \quad I_e(q_e(\mathbf{x}))}_{\mathbb{E}[\text{reward for } i \to j \text{ ride}]} \right)$$

assumption: $qI_{ij}(q)$ is concave

true for throughput; welfare; revenue under regular F_{ij}

technical challenges

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_e \phi_e q_e(\mathbf{X}) I_e(q(\mathbf{X}))\right]$$

assumption: $qI_{ij}(q)$ is concave

challenges

- exponential size of policy

technical challenges

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_e \phi_e q_e(\mathbf{X}) I_e(q(\mathbf{X}))\right]$$

assumption: $qI_{ij}(q)$ is concave

challenges

- exponential size of policy
- non-convex problem: even with state-independent q_{ij}

approximately optimal control policies

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q(\mathbf{X}) I_{ij}(q(\mathbf{X}))\right]$$

challenges

- exponential number of states
- non-convex optimization problem

theorem [Banerjee, Freund & Lykouris 2016]

convex relaxation gives state-independent pricing policy with approximation factor of $1 + \frac{\text{number of stations}}{\text{number of cars}}$

approximately optimal control policies

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q(\mathbf{X}) I_{ij}(q(\mathbf{X}))\right]$$

challenges

- exponential number of states
- non-convex optimization problem

theorem [Banerjee, Freund & Lykouris 2016]

convex relaxation gives state-independent pricing policy with approximation factor of $1 + \frac{\text{number of stations}}{\text{number of cars}}$

- extends to dispatch, rebalancing

approximately optimal control policies

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j}\phi_{ij}q(\mathbf{X})I_{ij}(q(\mathbf{X}))\right]$$

challenges

- exponential number of states
- non-convex optimization problem

theorem [Banerjee, Freund & Lykouris 2016]

convex relaxation gives state-independent pricing policy with approximation factor of $1 + \frac{\text{number of stations}}{\text{number of cars}}$

- extends to dispatch, rebalancing
- large-supply/large-market optimality: factor goes to 1 as system scales

proof roadmap

relaxation + resource augmentation

step 1: elevated flow relaxation: convex program that upper bounds performance, encodes essential conservation laws

proof roadmap

relaxation + resource augmentation

step 1: elevated flow relaxation: convex program that upper bounds performance, encodes essential conservation laws

step 2: show EFR is tight for a class of state-independent pricing policies, in the 'infinite-unit system' (i.e., $m \to \infty$)

proof roadmap

relaxation + resource augmentation

step 1: elevated flow relaxation: convex program that upper bounds performance, encodes essential conservation laws
step 2: show EFR is tight for a class of state-independent pricing policies, in the 'infinite-unit system' (i.e., $m \to \infty$)
step 3: bound objective in finite-unit system against infinite-unit system for this simpler class of policies

proof roadmap

relaxation + resource augmentation

step 1: elevated flow relaxation: convex program that upper bounds performance, encodes essential conservation laws

step 2: show EFR is tight for a class of state-independent pricing policies, in the 'infinite-unit system' (i.e., $m \to \infty$)

step 3: bound objective in finite-unit system against infinite-unit system for this simpler class of policies

$$\boxed{OBJ_m(\mathbf{p_m}(\mathbf{X}))} \leq \boxed{EFR(\mathbf{p}^*)} = \boxed{OBJ_\infty(\mathbf{p}_\infty)} \leq \boxed{(1/\alpha_{mn})OBJ_m(\mathbf{p}_\infty)}$$

the elevated flow relaxation

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q(\mathbf{X}) I_{ij}(q(\mathbf{X}))\right]$$

Suppose we knew $\mathbf{q}^\star$: Let $\hat{q}^\star = \mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}[q^\star(\mathbf{X})]$

the elevated flow relaxation

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q(\mathbf{X}) I_{ij}(q(\mathbf{X}))\right]$$

Suppose we knew $\mathbf{q}^\star$: Let $\hat{q}^\star = \mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}[q^\star(\mathbf{X})]$

$$\mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q^\star(\mathbf{X}) I_{ij}(q^\star(\mathbf{X}))\right] \leq \sum_{i,j} \phi_{ij} \hat{q}^\star I_{ij}(\hat{q}^\star) \quad \text{(Jensen's Ineq.)}$$

the elevated flow relaxation

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j}\phi_{ij}q(\mathbf{X})I_{ij}(q(\mathbf{X}))\right]$$

Suppose we knew $\mathbf{q}^\star$: Let $\hat{q}^\star = \mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}[q^\star(\mathbf{X})]$

$$\mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}\left[\sum_{i,j}\phi_{ij}q^\star(\mathbf{X})I_{ij}(q^\star(\mathbf{X}))\right] \leq \sum_{i,j}\phi_{ij}\hat{q}^\star I_{ij}(\hat{q}^\star) \quad \text{(Jensen's Ineq.)}$$

$$\leq \max_{\mathbf{q}\in[0,1]^{|E|}} \sum_{i,j}\phi_{ij}q_{ij}I_{ij}(q_{ij})$$

the elevated flow relaxation

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q(\mathbf{X}) I_{ij}(q(\mathbf{X}))\right]$$

Suppose we knew $\mathbf{q}^\star$: Let $\hat{q}^\star = \mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}[q^\star(\mathbf{X})]$

$$\mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}\left[\sum_{i,j} \phi_{ij} q^\star(\mathbf{X}) I_{ij}(q^\star(\mathbf{X}))\right] \leq \sum_{i,j} \phi_{ij} \hat{q}^\star I_{ij}(\hat{q}^\star) \quad \text{(Jensen's Ineq.)}$$

$$\leq \max_{\mathbf{q}\in[0,1]^{|E|}} \sum_{i,j} \phi_{ij} q_{ij} I_{ij}(q_{ij})$$

this is convex! however, it is too weak

the elevated flow relaxation

objective

$$\max_{\mathbf{q}\in[0,1]^{|E|}} \mathbb{E}_{\pi_{\mathbf{q}}(\mathbf{X})}\left[\sum_{i,j}\phi_{ij}q(\mathbf{X})I_{ij}(q(\mathbf{X}))\right]$$

Suppose we knew $\mathbf{q}^\star$: Let $\hat{q}^\star = \mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}[q^\star(\mathbf{X})]$

$$\mathbb{E}_{\pi_{\mathbf{q}^\star}(\mathbf{X})}\left[\sum_{i,j}\phi_{ij}q^\star(\mathbf{X})I_{ij}(q^\star(\mathbf{X}))\right] \leq \sum_{i,j}\phi_{ij}\hat{q}^\star I_{ij}(\hat{q}^\star) \quad \text{(Jensen's Ineq.)}$$

$$\leq \max_{\mathbf{q}\in[0,1]^{|E|}} \sum_{i,j}\phi_{ij}q_{ij}I_{ij}(q_{ij})$$

this is convex! however, it is too weak

idea: strengthen relaxation by adding additional constraints on $\mathbf{q}$

- circulation: $\sum_j \phi_{ij}q_{ij} = \sum_k \phi_{ki}q_{ki} \quad \forall\, i \in V$
- Little's law: $\mathbb{E}[\text{units in transit}] \leq m$

in summary

$$\boxed{\mathrm{OBJ}_m(\mathbf{p_m}(\mathbf{X}))} \leq \boxed{\mathrm{EFR}(\mathbf{p}^*)} = \boxed{\mathrm{OBJ}_\infty(\mathbf{p}_\infty)} \leq \boxed{(1/\alpha_{mn})\mathrm{OBJ}_m(\mathbf{p}_\infty)}$$

theorem [Banerjee, Freund & Lykouris 2016]

state-independent prices $\vec{p}_\infty$ (from EFR) in m-unit system gives

$$OBJ_m(\vec{p}_\infty) \geq \alpha_{mn} OPT_m \quad , \quad \text{where } \alpha_{mn} = \frac{m}{m+n-1}$$

in summary

$$\boxed{\mathrm{OBJ}_m(\mathbf{p}_m(\mathbf{X}))} \leq \boxed{\mathrm{EFR}(\mathbf{p}^*)} = \boxed{\mathrm{OBJ}_\infty(\mathbf{p}_\infty)} \leq \boxed{(1/\alpha_{mn})\mathrm{OBJ}_m(\mathbf{p}_\infty)}$$

theorem [Banerjee, Freund & Lykouris 2016]

state-independent prices $\vec{p}_\infty$ (from EFR) in m-unit system gives

$$OBJ_m(\vec{p}_\infty) \geq \alpha_{mn} OPT_m \quad , \quad \text{where } \alpha_{mn} = \frac{m}{m+n-1}$$

main takeaway

new technique for optimizing stochastic dynamical system in steady-state

- can extend to more complex settings (?)
  (travel-times, multi-objective, pooling, reservations)

in summary

$$\boxed{OBJ_m(\mathbf{p_m}(\mathbf{X}))} \leq \boxed{EFR(\mathbf{p}^*)} = \boxed{OBJ_\infty(\mathbf{p}_\infty)} \leq \boxed{(1/\alpha_{mn})OBJ_m(\mathbf{p}_\infty)}$$

theorem [Banerjee, Freund & Lykouris 2016]

state-independent prices $\vec{p}_\infty$ (from EFR) in m-unit system gives

$$OBJ_m(\vec{p}_\infty) \geq \alpha_{mn} OPT_m \quad , \quad \text{where } \alpha_{mn} = \frac{m}{m+n-1}$$

main takeaway

new technique for optimizing stochastic dynamical system in steady-state

- can extend to more complex settings (?)
  (travel-times, multi-objective, pooling, reservations)
- but where do we get the demand-rate and price-elasticity estimates?

market design in ride-share platforms

Pricing in Ride-Share Platforms
Banerjee, Johari & Riquelme (2015)
(EC'15: https://ssrn.com/abstract=2568258)

why market design? and why ridesharing?

> *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 and 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...*
>
> R. Preston McAfee

why market design? and why ridesharing?

> *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 and 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...*

R. Preston McAfee

data-driven optimization vs. market design

- default approach for complex operational problems:
  model – calibrate from data – optimize specific problem instance
- market mechanisms self-calibrate to solve the optimization problem
- ridesharing unique among online marketplaces: *platform sets prices*

quasi-static vs. dynamic

for a large block of time (e.g., few hours), region (e.g., city-neighborhood), mean system parameters are constant, predictable.

why not have hourly location-based prices?

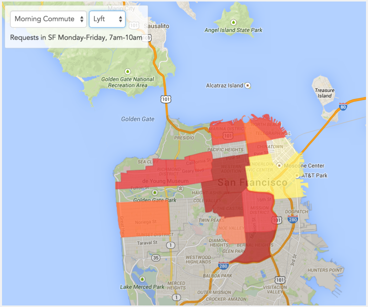

.

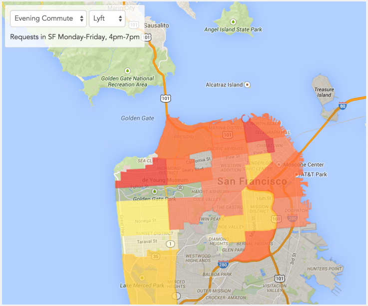

Source: whatsthefare.com

dynamic pricing vs. static pricing

- dynamic: price changes instantaneously, in response to system state
- (quasi) static: constant over several hours (predictably changing)

model for studying rideshare pricing

focus on a single block of time, and a single region.

system state = number of available drivers

assumption 1: mean system parameters stay constant

- state-dependent (dynamic) pricing policy:
  if # of available drivers $= A$, then price for ride $= P(A)$
- platform earns a (fixed) fraction γ of every dollar spent

assumption 2: the two sides react at different time-scales

- myopic passengers: sensitive to instantaneous prices, availability
- drivers are sensitive to long-term (average) earnings and ride-volume

rideshare pricing model: the details

stochastic dynamics + passenger/driver strategic behavior

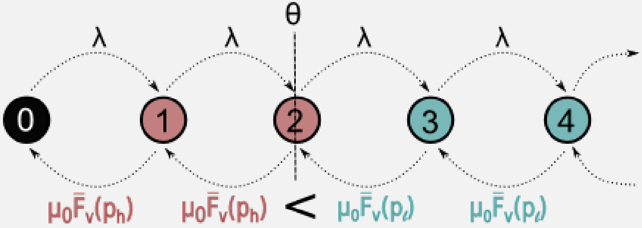

strategic model for passengers

- a (potential) passenger requests a ride iff:
  reservation value $V >$ current price, and driver available
  - $V \sim F_V$, i.i.d. across ride requests
- $\mu_0 =$ exogenous rate of "app opens", $\mu =$ actual rate of requests
  when A drivers present: $\mu = \mu_0 \overline{F}_V(P(A))$

rideshare pricing model: the details

stochastic dynamics + passenger/driver strategic behavior

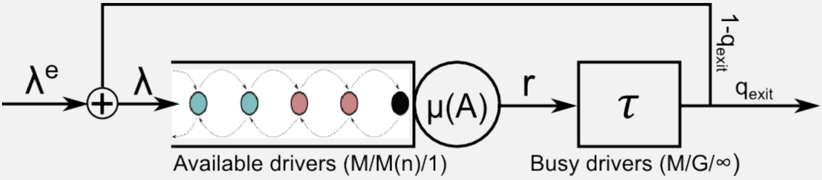

strategic behavior of drivers

- a driver works on the platform iff:
  reservation rate $C \times \mathbb{E}[\text{per-ride time spent}] < \mathbb{E}[\text{per-ride earning}]$
  - $C \sim F_C$, i.i.d. across drivers
- $\Lambda_0 =$ "potential" driver-arrival rate, $\lambda =$ actual driver-arrival rate

$$\lambda = \frac{\Lambda_0}{q_{exit}} F_C \left(\frac{\mathbb{E}[\text{Per-ride earning}]}{\mathbb{E}\left[\text{Idle (waiting) time} + \text{Ride time}\right]} \right)$$

driver decision aids

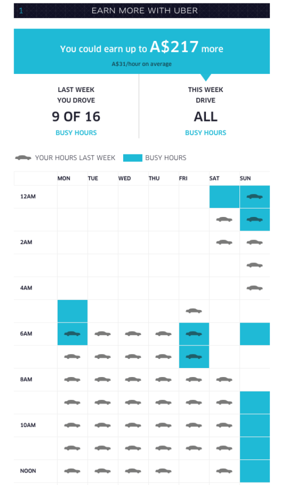

.

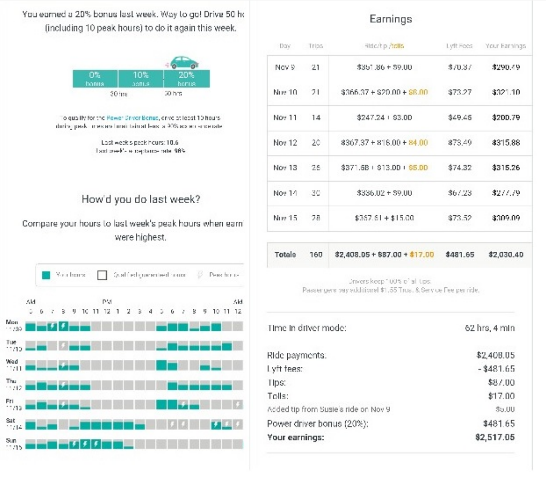

source: therideshareguy.com

rideshare pricing model: overview

putting it together: equilibrium

given pricing policy $P(\cdot)$, equilibrium $(\lambda, \mu, \pi, \eta, \iota)$ such that:

1. μ: passenger-arrival rate, given state A, satisfies:

$$\mu = \mu_0 \overline{F}_V\left(P(A)\right)$$

2. λ: driver-arrival rate λ, given ι, η, satisfies:

$$\lambda = \Lambda_0 F_C\left(\frac{\eta}{\iota + \tau}\right)$$

3. π: steady-state distribution of A given λ, μ
4. η: $\mathbb{E}$[Earning per ride], given $P(.)$ and π
5. ι: $\mathbb{E}$[Idle time per ride], given $P(.)$ and π

platform equilibrium under static pricing

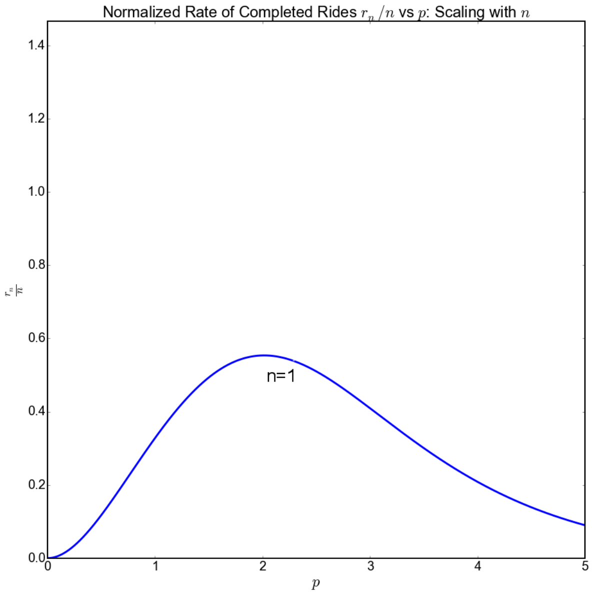

platform equilibrium under static pricing

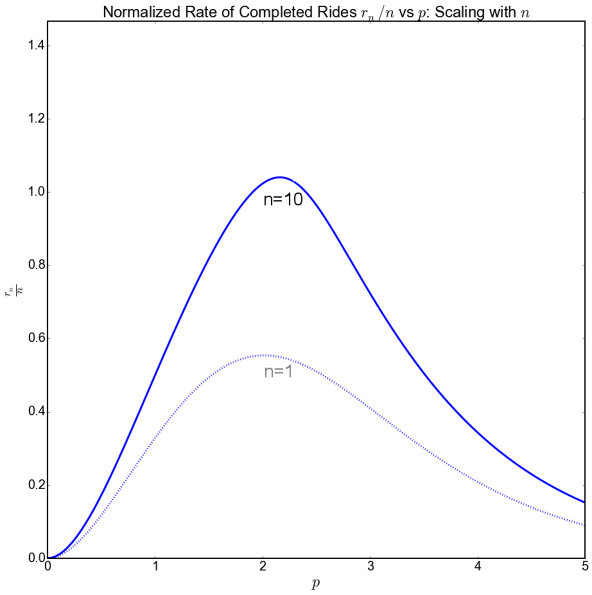

platform equilibrium under static pricing

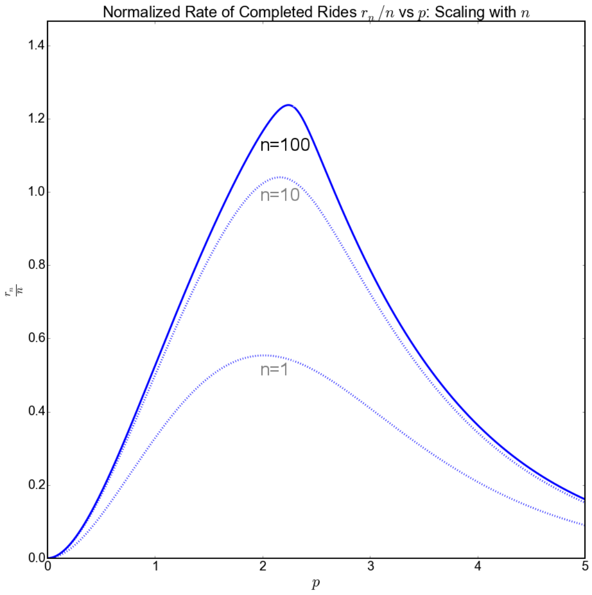

platform equilibrium under static pricing

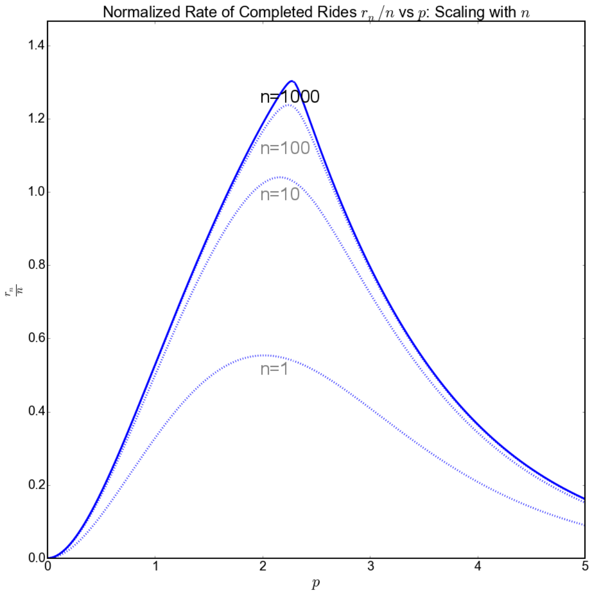

platform equilibrium under static pricing

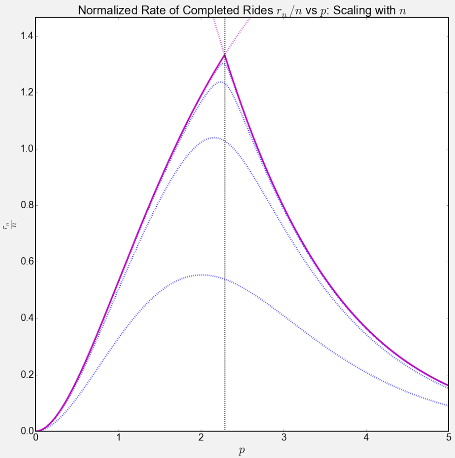

theorem: static pricing in large-market limit $\Rightarrow$ demand-supply curve

rate of rides in large-market limit = min{available supply, available demand}

Platform Equilibrium under Static Pricing

Theorem: Static pricing in large-market limit

Under static pricing (i.e., $P(A) = p \,\forall\, A$), let $r_n(p)$ denote the equilibrium rate of completed rides in the n^{th} system. Then:

$$r_n(p) \to \widehat{r}(p) \triangleq \min\left\{\frac{\Lambda_0}{q_{exit}} F_C\left(\frac{\gamma p}{\tau}\right), \mu_0(1 - F_V(p))\right\}$$

Some intuition:

- At any price, queueing system is always stable (else idle times blow up)
- If supply $<$ demand: Drivers become fully saturated
- If supply $>$ demand: Drivers forecast high idle times and don't enter

platform equilibrium under dynamic pricing

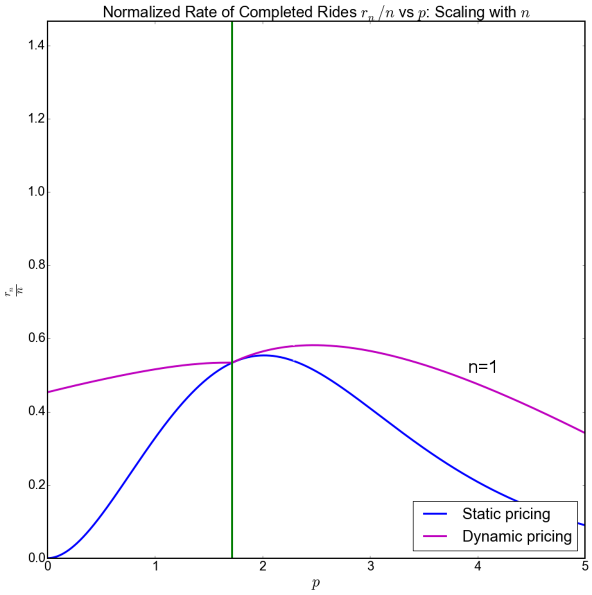

platform equilibrium under dynamic pricing

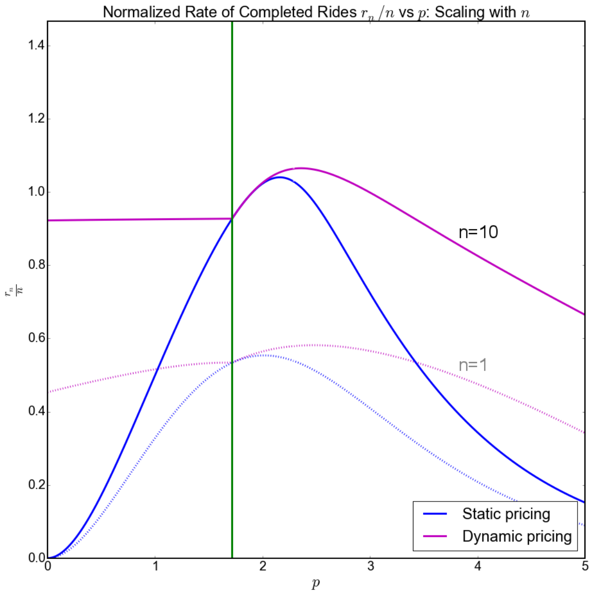

platform equilibrium under dynamic pricing

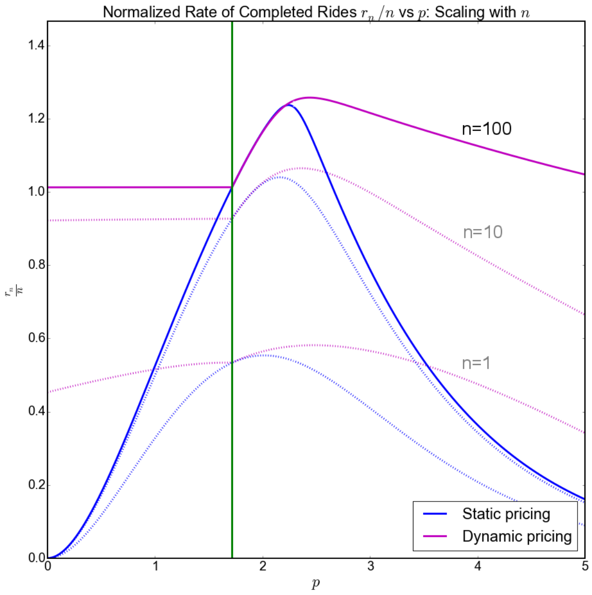

static vs. dynamic pricing: optimality

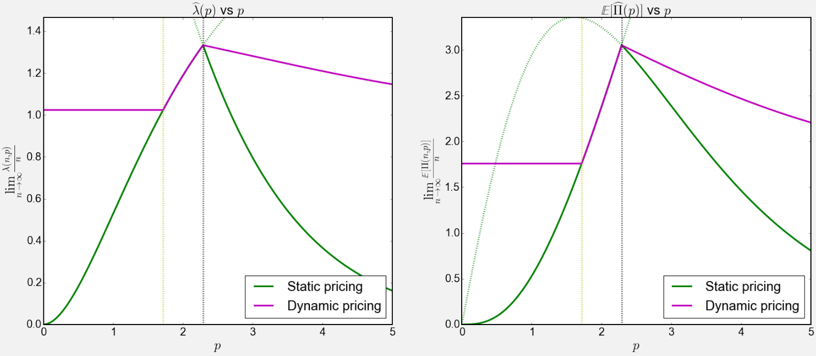

theorem [Banerjee, Johari & Riquelme 2015]

if $\mathbf{F}_V$ has increasing hazard rate: then

rate of rides for any dynamic policy $\leq$ rate of rides under optimal static pricing.

static vs. dynamic pricing: sensitivity to parameters

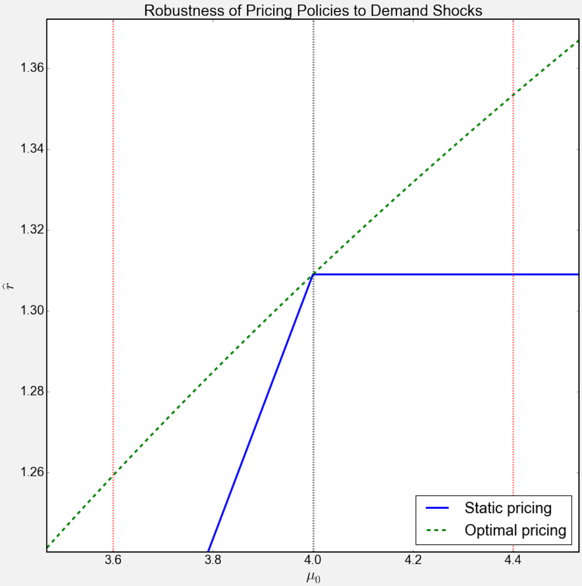

static vs. dynamic pricing: sensitivity to parameters

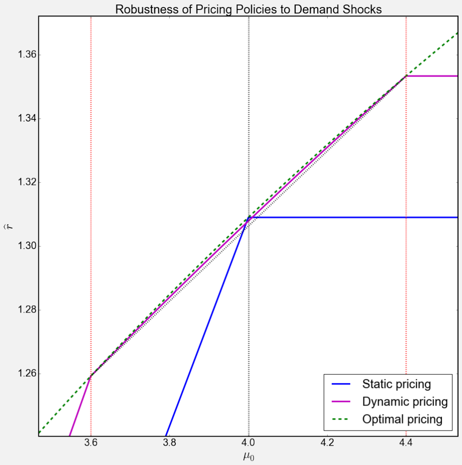

theorem [Banerjee, Johari & Riquelme 2015]

dynamic pricing $\geq$ 'linear approximation' of optimal static-pricing throughput

summary, and the road ahead

main takeaway

ridesharing platforms: crucible for real-time decision making

- well modeled by steady-state stochastic models
- approximate control via new convex relaxation techniques
- algorithm self-calibration via market mechanisms

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

- impact of platform competition

price of fragmentation in ridesharing markets

(with Thibault Séjourné (Ecole Polytechnique), S. Samaranayake (Cornell))

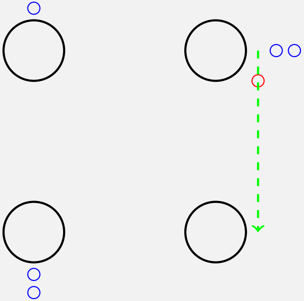

- what is the 'societal cost' of decentralized optimization?
  – multiple platforms with (random) exogenously partitioned demands
  – individual platforms do optimal empty-vehicle rebalancing

price of fragmentation in ridesharing markets

(with Thibault Séjourné (Ecole Polytechnique), S. Samaranayake (Cornell))

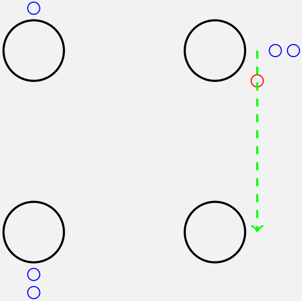

- what is the 'societal cost' of decentralized optimization?
  – multiple platforms with (random) exogenously partitioned demands
  – individual platforms do optimal empty-vehicle rebalancing

price of fragmentation

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

price of fragmentation in vehicle-sharing markets

result (in brief)

as demand scales, the price of fragmentation undergoes a phase transition based on structure of underlying demand flows
– both regimes observed in NYC taxi-data ($\approx$ 10% fragmentation-affected)

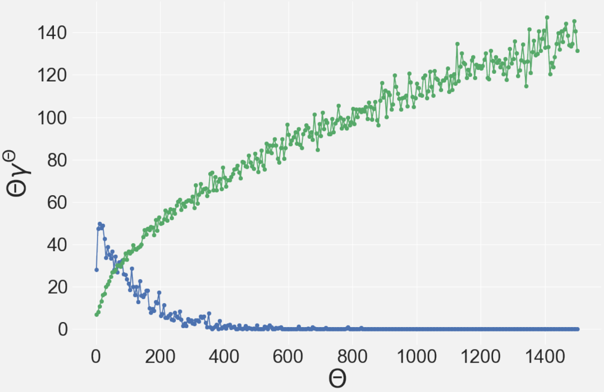

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

- impact of platform competition

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

- impact of platform competition
- the value of information: forecasting vs. self-calibration

the road ahead

some short term targets

- the value of state-dependent controls
  – for general controls, objectives: no improvement possible
  – for dispatch: can achieve exponential decay in m!
  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

- impact of platform competition
- the value of information: forecasting vs. self-calibration
- ridesharing + public transit

the road ahead

some short term targets

- the value of state-dependent controls
  - for general controls, objectives: no improvement possible
  - for dispatch: can achieve exponential decay in m!

  (joint work with Pengyu Qian and Yash Kanoria (Columbia))
- non-stationary and/or bursty arrivals
- algorthms for more complex problems
  (policies for ride-pooling, reservation mechanisms)

going further beyond

- impact of platform competition
- the value of information: forecasting vs. self-calibration
- ridesharing + public transit
- appropriate mix of employees, freelancers and autonomous cars