Dynamic Pricing in Ridesharing Platforms

A Queueing Approach

Sid Banerjee | *Ramesh Johari* | Carlos Riquelme

Cornell | Stanford | Stanford

rjohari@stanford.edu

With thanks to Chris Pouliot, Chris Sholley, and Lyft Data Science

20 November 2015

Ridesharing and Pricing

Ridesharing platforms

Examples of major platforms: Lyft, Uber, Sidecar

This talk: Pricing and ridesharing

Ridesharing is somewhat unique among online platforms:

The platform sets the transaction price.

Our goal: Understand optimal pricing strategy.

Our contributions

1. A model that combines:
   - Strategic behavior of passengers and drivers
   - Pricing behavior of the platform
   - Queueing behavior of the system
2. What are the advantages of *dynamic* pricing over *static* pricing?
   - *Static*: Constant over several hour periods
   - *Dynamic*: Pricing changes in response to system state; "surge", "prime time"

Related work

Our work sits at a nexus between several different lines of research:

1. *Matching queues* (cf. [Adan and Weiss 2012])
2. *Strategic queueing models* (cf. [Naor 1969])
3. *Two-sided platforms* (cf. [Rochet and Tirole 2003, 2006])
4. *Revenue management* (cf. [Talluri and van Ryzin 2006])
5. *Large-scale matching markets* (cf. [Azevedo and Budish 2013])
6. *Mean field equilibrium* (cf. [Weintraub et al. 2008])

Model

Two types: Strategic and queueing

We need a *strategic model* that captures:

1. Platform pricing
2. Passenger incentives
3. Driver incentives

We need a *queueing model* that captures:

1. Driver time spent idling vs. driving
2. Ride requests blocked vs. served

Preliminaries

1. Focus on a *block* of time (e.g., several hours) over which arrival rates are roughly stable

2. Focus on a single region (e.g., a single city neighborhood)
   - For technical simplicity
   - Insights generalize to networks of regions

3. Focus on throughput: rate of completed rides
   - For technical simplicity
   - Same results for profit, when system is supply-limited
   - Similar numerical results for welfare; theory ongoing

Strategic modeling: Platform pricing

Platforms:

- Earn a (fixed) fraction γ of every dollar spent (e.g., 20%)
- Need *both* drivers (supply) and passengers (demand)
- Use pricing to align the two sides

Load-dependent pricing:
If # of available drivers = A, then price offered to ride = $P(A)$

Strategic model: Platform pricing

In practice:

- Platforms charge a time- and distance-dependent *base price*
- Platforms manipulate price through a *multiplier*
- Base price typically is not varied

In our model:
price ≡ *multiplier.*

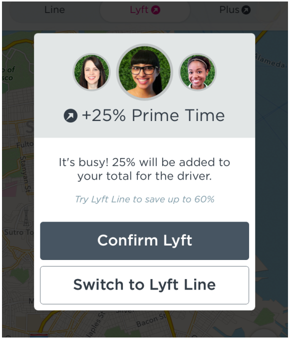

Strategic model: Passengers

How do passengers enter?

- Passenger $\equiv$ one ride request
- Sees *instantaneous* ride price
- Enters if price $<$ reservation value V
- $V \sim \mathsf{F}_V$, i.i.d. across ride requests

$\mu_0 =$ exogenous rate of "app opens".
$\mu =$ actual rate of rides requested.

Then when A available drivers present:

$$\mu = \mu_0 \overline{\mathsf{F}}_V(P(A)).$$

Strategic model: Drivers

How do drivers enter?

- Sensitive to *expected earnings over the block*
- Choose to enter if:
  reservation earnings rate $C \times$
  expected total time in system
  $<$ expected earnings while in system
- $C \sim \mathsf{F}_C$, i.i.d. across drivers

$\Lambda_0 =$ exogenous rate of driver arrival.
$\lambda =$ actual rate at which drivers enter.
Then:

$$\lambda = \Lambda_0 \mathsf{F}_C \left(\frac{\text{expected earnings in system}}{\text{expected time in system}} \right)$$

Queueing model

1. Drivers enter at rate λ.
2. When A drivers available,
   ride requests arrive at rate $\mu(A)$.
3. If a driver is available, ride is *served*; else *blocked*.
4. Rides last exponential time, mean τ.
5. After ride completion:
   - With probability q_{exit}: Driver signs out
   - With probability $1 - q_{\text{exit}}$: Driver becomes available

Queueing model: Steady state

Jackson network of two queues: $M/M(n)/1$ and $M/M/\infty$

$\implies$ product-form steady state distribution π.

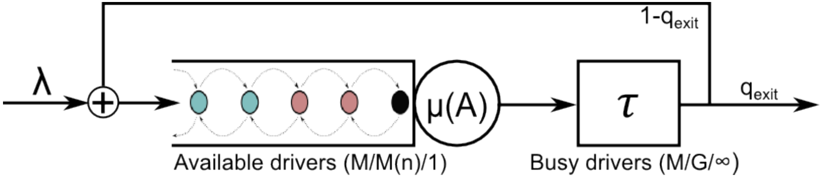

Putting it together: Equilibrium

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

1. π is the steady state distribution, given λ and μ
2. η is the expected earnings per ride, given $P(\cdot)$ and π
3. ι is the expected idle time per ride, given π and λ
4. λ is the entry rate of drivers, given ι and η:

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

5. $\mu(A)$ is the arrival rate of ride requests when A drivers are available, given $P(\cdot)$:

$$\mu = \mu_0 \overline{\mathsf{F}}_V(P(A)).$$

Putting it together: Equilibrium

If price increases when number of available drivers decreases:

- Equilibria always exist under appropriate continuity of $\mathsf{F}_C, \mathsf{F}_V$.
- Equilibria are unique under reasonable conditions

Large Market Limit

The challenge

- To understand optimal pricing, we need to characterize system equilibria.
- In particular, need *sensitivity* of equilibria to changes in pricing policy.
- Our approach: *asymptotics* to simplify analysis.

Large market asymptotics

Consider a sequence of systems indexed by n.

- In n'th system, exogenous arrival rates are $n\Lambda_0, n\mu_0$.
- In n'th system, pricing policy is $P_n(\cdot)$.
- In each system, this gives rise to a system equilibrium.

We analyze pricing by looking at asymptotics of equilibria.

Static Pricing

What is static pricing?

Static pricing means: *price policy is constant.*
Let $P(A) = p$ for all A.

Theorem
Let $r_n(p)$ denote the equilibrium rate of completed rides in the n'th system. Then:

$$r_n(p) \to \hat{r}(p) \triangleq \min\{\Lambda_0 \mathsf{F}_C(\gamma p/\tau)/q_{\mathit{exit}}, \mu_0 \overline{\mathsf{F}}_V(p)\}.$$

Throughput* = min** $\{$ ***available supply, available demand $\}$

Static pricing: Illustration

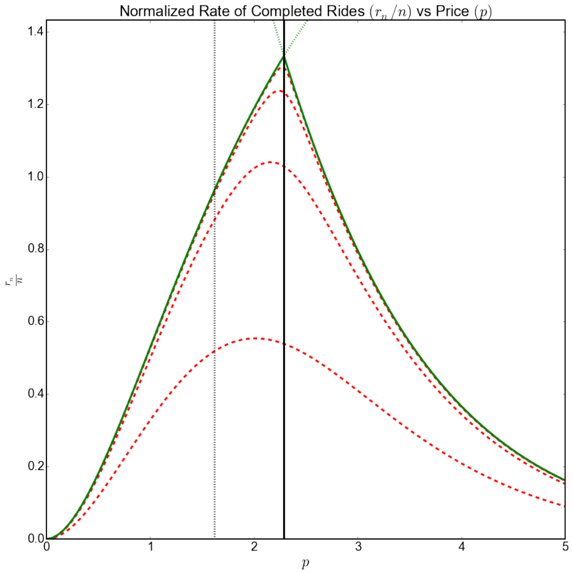

Static pricing: Interpretation

Note that at *any price*, queueing system is always stable:

- When supply < demand:
  Drivers become fully saturated
- When supply > demand:
  Drivers forecast high idle times and don't enter

Balance price p_{bal}: Price where supply $=$ demand

Corollary

The optimal static price is p_{bal}*.*

Dynamic pricing

What is dynamic pricing?

Meant to capture "surge" (Uber)
and "prime time" (Lyft) pricing strategies.

We focus on *threshold pricing*:

- Threshold θ
- High price p_h charged when available drivers $< \theta$
- Low price $p_\ell < p_h$ charged when available drivers $> \theta$

Dynamic pricing: Numerical investigation

- Fix one price, and vary the other price.
- Compare to static pricing.

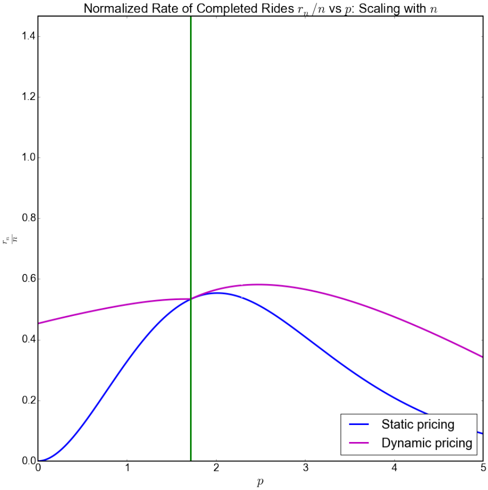

$n = 1$

Dynamic pricing: Numerical investigation

- Fix one price, and vary the other price.
- Compare to static pricing.

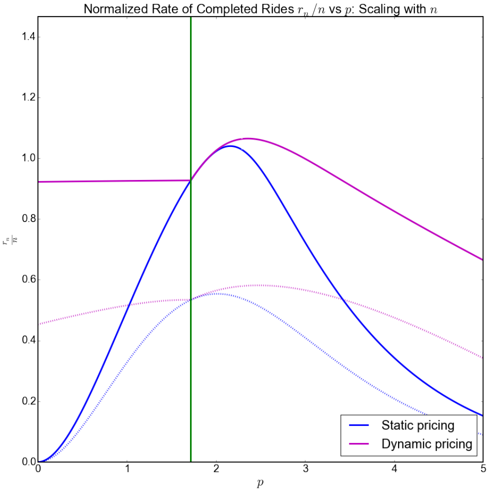

$n = 10$

Dynamic pricing: Numerical investigation

- Fix one price, and vary the other price.
- Compare to static pricing.

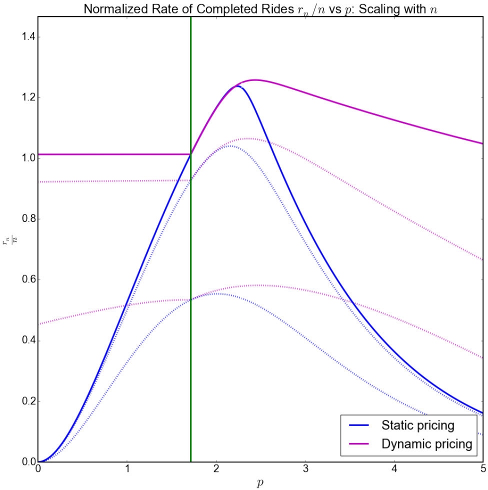

$n = 100$

Dynamic pricing: Numerical investigation

- Fix one price, and vary the other price.
- Compare to static pricing.

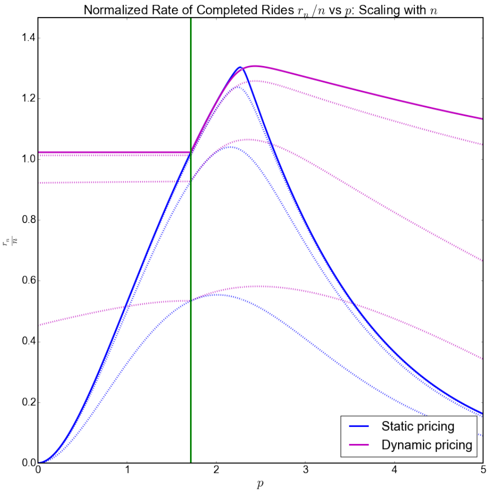

$n = 1000$

Dynamic pricing: Numerical investigation

- Fix one price, and vary the other price.
- Compare to static pricing.

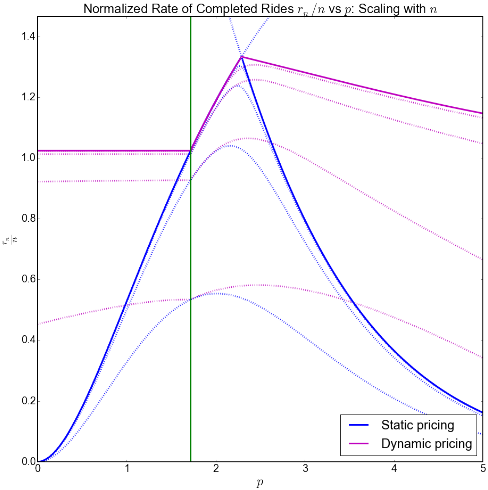

$$n \to \infty$$

Dynamic pricing: Numerical investigation

- Fix one price, and vary the other price.
- Compare to static pricing.

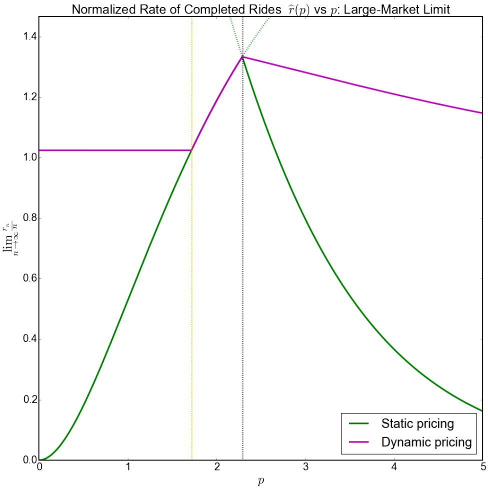

$n \to \infty$

Optimal dynamic pricing

Theorem

Let r_n^* *be the rate of completed rides in the* n*'th system, using the optimal static price.*

Let r_n^{**} *be the rate of completed rides in the* n*'th system, using the optimal threshold pricing strategy.*

Then if F_V has monotone hazard rate,

$$\frac{r_n^* - r_n^{**}}{n} \to 0 \text{ as } n \to \infty.$$

Optimal dynamic pricing

In other words:

In the fluid limit, no dynamic pricing policy yields higher throughput than optimal static pricing.

Optimal dynamic pricing

In other words:

In the fluid limit, no dynamic pricing policy yields higher throughput than optimal static pricing.

This result is reminiscent of similar results in the classical revenue management literature (e.g., [Gallego and van Ryzin, 1994]).

The main differences arise due to the presence of a two sided market.

Proof sketch

Under threshold pricing:

- Drivers are sensitive to *two* quantities: idle time, and price.
- Show that optimal $\theta_n^* \to \infty$, but chosen so that idle time $\to 0$ as $n \to \infty$.
- In this limit, drivers are sensitive to the *average* price per ride:

$$p_{\text{avg}} = \pi_h p_h + \pi_\ell p_\ell,$$

  where π_h, π_ℓ are $\approx$ probabilities of being below or above θ, respectively.
- If p_{avg} decreases, fewer drivers will enter.

Proof sketch (cont'd)

We note that:

1. If $p_\ell < p_h \leq p_{\text{bal}}$, then $p_{\text{avg}} = p_h$.
2. If $p_{\text{bal}} \leq p_\ell < p_h$, then $p_{\text{avg}} = p_\ell$.
3. If $p_\ell < p_{\text{bal}} < p_h$, then $\pi_\ell > 0, \pi_h > 0$.

In first two cases, *de facto* static pricing.

Proof sketch (cont'd)

We explore the third case.
Suppose that we start with $p_\ell < p_h = p_{\mathsf{bal}}$ (so $p_{\mathsf{avg}} = p_h$).

Now increase p_h:

- Before $\pi_\ell = 0$, but now $\pi_\ell > 0$, so some customers pay p_ℓ; this lowers p_{avg}.
- p_h higher, so customers arriving when $A < \theta$ pay more; this increases p_{avg}.

When F_V is MHR, we show that the first effect dominates the second, so throughput falls.

Robustness

The value of dynamic pricing

How does dynamic pricing help?

- When system parameters are *known*, performance does not exceed static pricing.
- When system parameters are *unknown*, dynamic pricing naturally "learns" them.

Robustness: Illustration

What happens to *static pricing* in a demand shock?

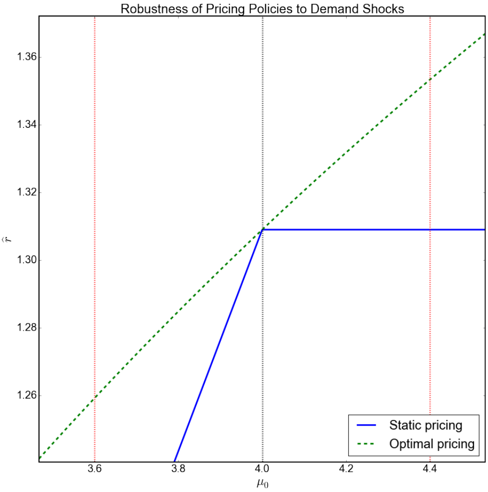

Robustness: Illustration

What happens to *dynamic pricing* in a demand shock?

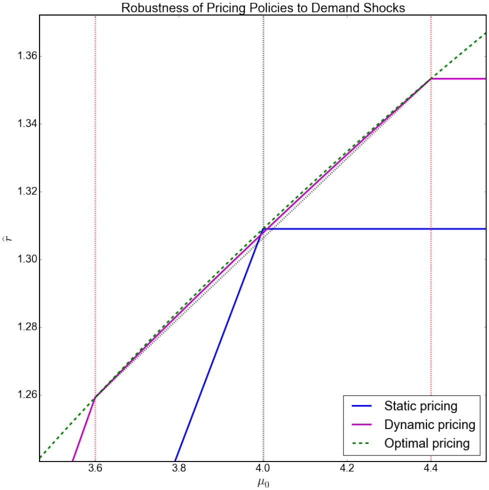

Robustness: Dynamic pricing

We can formally establish the observation in the previous illustration:

- Suppose F_C is logconcave, and $\mu_0^{(1)} < \mu_0^{(2)}$ are fixed.
- Let $p_{\mathsf{bal},n}^{(1)}, p_{\mathsf{bal},n}^{(2)} =$ optimal static prices in the n'th system.
- Let $r_n^{(1)}, r_n^{(2)} =$ optimal throughput in the n'th system.
- Suppose now the true $\mu_0 \in [\mu_0^{(1)}, \mu_0^{(2)}]$.
- Using both prices $p_{\mathsf{bal},n}^{(1)}, p_{\mathsf{bal},n}^{(2)}$ is robust:
  - There exists a sequence of threshold pricing policies with throughput at any such μ_0 (in the fluid scaling) $\geq$ the linear interpolation of $r_n^{(1)}$ and $r_n^{(2)}$.

(Same holds w.r.t. Λ_0.)

Conclusion

Platform optimization

This work is an example of *platform optimization*:
Requires understanding *both* operations and economics.
Other topics under investigation:

1. Network modeling (multiple regions):
   Our main insights generalize
2. Effect of pricing on aggregate welfare
3. Modeling driver heat maps
4. Fee structure: changing the percentage
5. Effect of changing the matching algorithm