Dynamic Pricing in Ridesharing Platforms

SIDDHARTHA BANERJEE
Cornell University
and
RAMESH JOHARI
and
CARLOS RIQUELME
Stanford University

We study dynamic pricing policies for ridesharing platforms such as Lyft and Uber. On one hand these platforms are two-sided: this requires economic models that capture the incentives of both drivers and passengers. On the other hand, these platforms support high temporal-resolution for data collection and pricing: this requires stochastic models that capture the dynamics of drivers and passengers in the system.

We summarize our main results from [Banerjee et al. 2015], in which we study the role of dynamic pricing in ridesharing platforms using a queueing-theoretic economic model. We build a model of two-sided ridesharing platforms that captures both the stochastic dynamics of the marketplace and the strategic decisions of drivers, passengers and the platform. We show how our model can help explain the success of dynamic pricing in practice: in particular, we argue that the benefit of dynamic pricing over static pricing is not in the optimal performance, but rather, in the robustness of its performance to uncertainty in system parameters.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Science**]: Economics

General Terms: Economics, Market Design, Pricing, Theory

Additional Key Words and Phrases: Ridesharing, Dynamic Pricing, Two-Sided Platforms

1. INTRODUCTION

In this letter, we study pricing in *ridesharing platforms* such as Lyft and Uber. We present a summary of our main results and insights from [Banerjee et al. 2015], and refer interested readers to our full paper for more details.

Since their founding in the last several years, ridesharing platforms have experienced extraordinary growth. At their core, the platforms reduce the friction in matching and dispatch for transportation. A typical transaction on these platforms is as follows: a potential rider opens the app on her phone and requests a ride, and the system matches her to a nearby driver if one is available. There are three key features that make ridesharing platforms unique, and motivate our model and research questions:

(1) Ridesharing platforms typically do not employ drivers, but rather, create a marketplace between passengers and freelance drivers. Drivers can choose when

The authors would like to thank the data science team at Lyft, particularly Adam Greenhall, Ricky Chachra, Chris Pouliot, and Chris Sholley; part of this work was carried out when SB was a technical consultant at Lyft. We gratefully acknowledge support from the National Science Foundation via grants CMMI-1234955 and CNS-1343253, and the DARPA GRAPHS program.
Authors' addresses: `sbanerjee@cornell.edu`, `ramesh.johari@stanford.edu`, `rikel@stanford.edu`

and where to work (or not), and earn a share of the earnings per ride.

(2) An important reason behind the success of ridesharing apps is the minimal friction experienced by passengers in requesting a ride. Critical to this is a design choice made by most platforms wherein matches and prices are determined exclusively by the platform. This is in contrast to many other two-sided marketplaces, where the price of a transaction typically arises via some form of negotiation between agents.

(3) An important role of platform intermediation is thus to control prices to calibrate supply and demand relative to each other, while ensuring relatively high satisfaction to both sides. The key tool available to the platform for this purpose is *dynamic pricing*: the platform can adjust ride prices in real-time, to react to changes in ride requests and available drivers.

Both Lyft and Uber have used dynamic pricing policies (referred to as 'Primetime Pricing' in Lyft and 'Surge Pricing' in Uber) for several years. The main contribution of our work in [Banerjee et al. 2015] is in developing a theoretical framework that captures these three features of ridesharing platforms: their two-sided nature, the strategic reaction of market participants to marketplace policies, and the ability of the platform to price based on real-time state. In this letter, we briefly describe our model, and our findings on how dynamic pricing influences the performance of ridesharing platforms. We conclude by placing our research in the context of related work, and point out some interesting avenues for future research.

2. A MODEL FOR RIDESHARING PLATFORMS

At a high level, our model is based on combining a stochastic model for the dynamics of riders and passengers on the platform, with an equilibrium analysis that captures incentives of both drivers and passengers as well as the objectives of the platform.

To capture the fast-timescale dynamics of ridesharing platforms, we employ a *queueing theoretic* model. In particular, we consider a geographic area divided into *regions*, and track a Markov chain that tracks the number of available and busy drivers in each region. Each ride involves a driver picking up a passenger in one region, and dropping her off in another. For simplicity, we analyze this model first for a single region; classical tools from the theory of reversible queueing networks [Kelly 1979] let us generalize some of our results to networks of regions.

Key to our model is the assumption that there is an *intrinsic timescale separation in the strategic interaction of drivers and passengers*. This is based on the observation that passengers are typically sensitive to the *immediate* availability and price at the moment when they need a ride, while drivers are typically sensitive to the *average* wages earned over a longer period of time (days or weeks), and adjust their activity levels based on their assessment of earnings during the last week [1]. We note that our model does not consider longer-term passenger interactions, for example, demand screening due to persistent low availability or high prices. Understanding the impact of this is an interesting follow-up question.

We illustrate our single-region model in Fig. 1. The main object of study is a birth-death chain that tracks the number of available drivers. From the passengers'

[1]The drivers are aided in this by information provided by the platforms of their weekly performance, as well as the typical earnings of drivers at different times and locations.

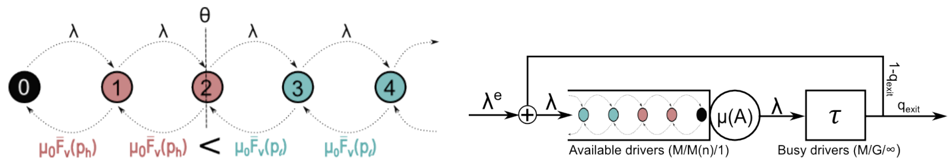

(a) Markov chain for available drivers   (b) Flow of drivers in single-region model

Fig. 1. *Queueing model for ridesharing platform: Figure 1(a) shows the birth-death chain for the number of available drivers in a region. In particular, we have shown a* single-threshold *pricing policy, where the platform uses a 'base' fare-multiplier p_ℓ when the number of drivers is greater than a threshold θ, else charges a 'primetime' price-multiplier $p_h > p_\ell$ (hence the queue drains slower when there are $\leq \theta$ drivers). Figure 1(b) shows the flow of drivers in the network: exogenous drivers arrive to the platform at an (equilibrium) rate λ^e and join the available-drivers queue; when matched with a passenger, they transition to a busy-drivers queue; after completing a ride, they either exit the platform, or return to the available-drivers queue.*

viewpoint, we consider a loss system, where passengers arrive (i.e., open their apps) at a rate μ_0, and each incoming passenger either requests a ride (depending on availability and price), or immediately exits the system. Specifically, each passenger has a private value drawn i.i.d. from a distribution F_V, and requests a ride only if a driver is available, and the price quoted is lower than their value. [2] In contrast, we assume that drivers' entry decisions are made by comparing their expected lifetime earnings with their expected lifetime (i.e., total time spent in waiting for and giving rides); if this wage rate exceeds a driver-specific reservation rate (drawn i.i.d. from distribution F_C), then the driver chooses to join the platform. The price for a ride $P(A)$ can be *dynamically* adjusted by the platform based on the current number of available drivers A. Finally, we assume that a fraction γ of the price goes to the driver, while the remaining is kept by the platform.

The pricing policy, along with the demand and supply functions together determine an equilibrium arrival rate of driver-arrivals λ and passenger arrivals $\mu(A)$, as well as system throughput and revenue. Suppose the maximum potential rate of exogenous driver-arrivals is Λ_0. Based on the above discussion, the equilibrium rates of passenger requests and exogenous driver-arrivals must satisfy:

$$\mu(A) = \mu_0(1 - F_V(P(A))) \qquad , \qquad \lambda^e = \lambda q_{exit} = \Lambda_0 F_C\left(\frac{\eta}{\iota + \tau}\right),$$

where η denotes the expected per-ride earnings, ι the expected waiting time for a driver between rides, and τ the expected ride time. Exact expressions for these can be computed from standard Markov chain theory; note though that η and ι depend on λ and μ, as well as the pricing policy $P(A)$. Details of these computations, and of the existence/uniqueness of the equilibrium, are given in [Banerjee et al. 2015].

[2] In reality, passengers are typically shown the nearest driver, and the request decision is a joint function of the price and pickup time. Our model focusses on the effect of the price, and uses unavailability as a proxy for passenger loss due to long pickup times. Understanding the effect of pickup times is an important question, but outside the scope of our model.

3. A SUMMARY OF OUR RESULTS

Given the above model of a ridesharing platform, we focus on studying two pricing policies: (i) static pricing, where the price is fixed as a function of mean system parameters, but not instantaneous state [3]), and (ii) *threshold* dynamic pricing policies, where the platform raises the price whenever the number of available drivers in a region falls below a threshold.

The Large-Market Scaling: One technical problem is that our model does not admit a closed-form expression for the equilibrium driver/passenger arrival rates, which are necessary for designing optimal pricing policies. To circumvent this, we consider the system under a *large-market scaling*. Formally, we consider a sequence of systems parametrized by n, wherein $\Lambda_0(n) = \Lambda_0 n$ and $\mu_0(n) = \mu_0 n$, and all other parameters $(\tau, q_{exit}, \gamma, F_C, F_V)$, as well as the price p, are held fixed. We then let n approach ∞, and study the *normalized* equilibrium state, i.e. $\lim_{n\to\infty} \lambda(n)/n$, of the limiting system. For dynamic pricing policies, in addition to scaling Λ_0 and μ_0, we keep the price vector fixed, but allow the threshold $\theta(n)$ to scale with n.

Under the large-market scaling, we are able to characterize the equilibrium rates for the limiting system in closed form, for both static and dynamic pricing. This can be seen graphically in Fig. 2(a), where we have plotted the normalized equilibrium throughput (i.e., rate of rides) vs. static price p (the green curves), and also, for a class of dynamic pricing policies (the maroon curves) where we keep one price fixed at 1.75 (the red vertical dotted line). The dotted curves are numerically computed for $n \in \{1, 10, 100, 1000\}$, and can be seen to be monotonically converging up to the solid curves, which plot our theoretical large-market limits.

Optimal Performance of Pricing Policies: One surprising aspect of Fig. 2(a) is that the optimal throughput in the large-market limit over the dynamic pricing policies we consider appears to coincide with that obtained under static pricing. This however turns out to hold for *all* threshold dynamic pricing policies under fairly weak conditions; in [Banerjee et al. 2015], we prove the following:

Theorem 3.1. *Given $(\Lambda_0, \mu_0, \gamma, q_{exit}, \tau)$ and continuous distributions F_C, F_V. Suppose F_V has an increasing hazard rate. Then the optimal normalized throughput in the large-market limit under dynamic pricing collapses to that obtained under the optimal static pricing policy.*

In other words, *the platform cannot increase throughput by employing dynamic pricing.* Similar results hold for revenue, and also for multi-threshold pricing policies.

Note though that dynamic pricing does better than static pricing in the pre-limit, which is unsurprising as it includes static pricing as a special case. The non-trivial aspect is that the difference in performance vanishes in the limit. Note also that the performance of a dynamic pricing policy with prices (p_ℓ, p_h) is not identical to the performance with static price p_ℓ or p_h (in particular, in the regime $p_\ell < p_{bal} < p_h$, where p_{bal} is the balance price under static pricing). This indicates that passengers experience both prices in the large-market limit (in fact, it can be shown that in the limit, passengers sample between the two prices in an i.i.d. manner).

[3]More specifically, we consider quasi-static policies, where the price remains fixed for blocks of time on the order of hours, but can be changed over slower timescales to reflect change in average demand/supply. Such policies were common before the rise of Lyft and Uber.

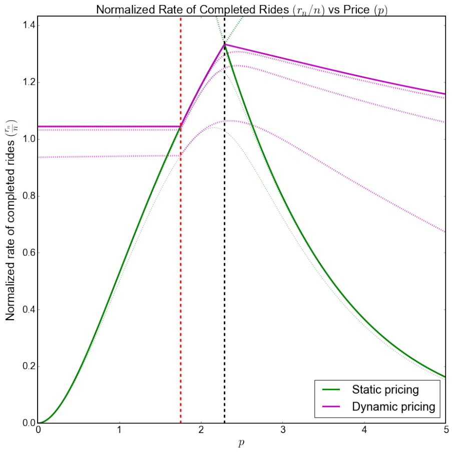

(a) Pricing policies in the large-market limit

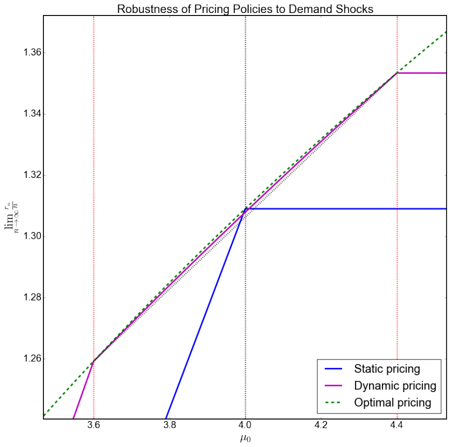

(b) Robustness of dynamic pricing

Fig. 2. *Static vs. dynamic pricing in ridesharing platforms: Figure 2(a) depicts the normalized equilibrium throughput in a ridesharing platform under static pricing (in green), and under dynamic pricing (in maroon) with one price fixed at* 1.75 *(the red vertical dotted line). The dotted lines show the throughput curve for different values of the market-scaling parameter* n*, with higher curves corresponding to higher values of* n*. The solid curves plot our theoretical large-market limits. Note that in the large-market limit, the optimal throughput under both policies is the same (indicated by the black vertical dotted line). Figure 2(b) demonstrates the sensitivity of static and dynamic pricing to demand uncertainty: For a fixed* Λ_0*, we consider* $\mu_0 \in 4 \pm 10\%$*, and compare the normalized throughput under* (*i*) *the optimal static policy with* $\mu_0 = 4$ *(indicated by the black vertical dotted line), and* (*ii*) *the dynamic-pricing policy which sets* p_ℓ *based on* $\mu_0 = 3.6$*, and* p_h *based on* $\mu_0 = 4.4$ *(indicated by the red vertical dotted lines). The dashed green curve shows the performance of the optimal static-pricing corresponding to the actual* μ_0.

Robustness of Pricing Policies: Theorem 3.1 is counterintuitive, and belies the success of dynamic pricing in practice on ridesharing platforms. Our second main result reveals a significant benefit that dynamic pricing holds over static pricing: *robustness*. Specifically, suppose the system operator chooses the optimal threshold dynamic (resp., static) pricing policy assuming system parameters Λ_0, μ_0. Now if the true parameters deviate from the assumed parameters, we show that *dynamic pricing maintains a much higher share of the optimal throughput relative to the optimal static pricing.* This property is graphically depicted in Fig. 2(b); refer to [Banerjee et al. 2015] for a formal geometrical characterization.

4. DISCUSSION AND RELATED WORK

Our work sits at the intersection of research in economics on two-sided marketplaces, and stochastic modeling approaches arising from the revenue management and queueing literatures. Combining both approaches is essential to investigating the role of dynamic pricing (and more generally, other dynamic mechanisms) in two-sided platforms. Moreover, the interaction of the two drives our main conclusion: that dynamic pricing is not fundamentally better than static pricing in terms of performance, but rather, its utility lies in helping discover the 'correct' static price, by stochastically mixing between higher prices (in low supply conditions) and lower prices (when there is low demand).

From an economics standpoint, our paper is in the spirit of the literature on the price theory of two-sided platforms [Rochet and Tirole 2003], [Weyl 2010]. This line of work typically studies the design two-sided markets under exogenously specified utility functions for agents. Our approach instead is to build up the market model from microfoundations; moreover, having a dynamic model allows us to study dynamic pricing, which is not possible in the standard economic models of two-sided platforms.

The question of static vs. dynamic pricing is widely studied in the revenue management literature [Talluri and Van Ryzin 2006]. In particular, our results are very similar in spirit to those in the seminal work of [Gallego and Van Ryzin 1994] on dynamic pricing based on current inventory levels. However, while the primary concern of revenue management is monopolist pricing in a one-sided platform, our focus is on the equilibrium effects of pricing policies in a two-sided marketplace.

From a technical perspective, our work adapts classical *queueing-network models* [Kelly 1979] to settings with matching constraints and strategic agent behavior, and uses a *large-market scaling* approach to derive tractable analysis and insights. There is a long line of work on strategic queueing models for single-queue settings with a fixed number of servers, and strategic customers; see [Hassin and Haviv 2003] for a good overview. More recent work has led to tractable models for more general *matching queues* [Adan and Weiss 2012], as well as simplified fluid models of service systems with strategic servers [Gurvich et al. 2014]. Large-market scaling techniques have also grown in importance in recent years; see [Kojima and Pathak 2009] for an example applying this to matching markets.

REFERENCES

Adan, I. and Weiss, G. 2012. Exact fcfs matching rates for two infinite muti-type sequences. *Operations Research 60*, 475–489.

Banerjee, S., Johari, R., and Riquelme, C. 2015. Pricing in ride-sharing platforms: A queueing-theoretic approach. *Available at SSRN 2568258*.

Gallego, G. and Van Ryzin, G. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science 40,* 8, 999–1020.

Gurvich, I., Lariviere, M., and Moreno, A. 2014. Staffing service systems when capacity has a mind of its own. *Available at SSRN 2336514*.

Hassin, R. and Haviv, M. 2003. *To queue or not to queue: Equilibrium behavior in queueing systems.* Vol. 59. Springer Science & Business Media.

Kelly, F. P. 1979. *Reversibility and stochastic networks.* Cambridge University Press.

Kojima, F. and Pathak, P. A. 2009. Incentives and stability in large two-sided matching markets. *The American Economic Review*, 608–627.

Rochet, J.-C. and Tirole, J. 2003. Platform competition in two-sided markets. *Journal of the European Economic Association 1,* 4, 990–1029.

Talluri, K. T. and Van Ryzin, G. J. 2006. *The theory and practice of revenue management.* Vol. 68. Springer Science & Business Media.

Weyl, E. G. 2010. A price theory of multi-sided platforms. *The American Economic Review*, 1642–1672.