Spatio-Temporal Pricing for Ridesharing Platforms

HONGYAO MA
Columbia University
and
FEI FANG
Carnegie Mellon University
and
DAVID C. PARKES
Harvard University

Ridesharing platforms match drivers and riders to trips, using dynamic prices to balance supply and demand. A challenge is to set prices that are appropriately smooth in space and time, so that drivers will choose to accept their dispatched trips, rather than drive to another area or wait for higher prices or a better trip. We work in a complete information, discrete time, multi-period, multi-location model, and introduce the *Spatio-Temporal Pricing (STP) mechanism*. The mechanism is incentive-aligned, in that it is a subgame-perfect equilibrium for drivers to always accept their trip dispatches. The mechanism is also welfare-optimal, envy-free, individually rational, budget balanced and core-selecting in equilibrium from any history onward. The proof of incentive alignment makes use of the $M^\natural$ concavity of minimum cost flow objectives. We also give an impossibility result, that there can be no dominant-strategy mechanism with the same economic properties. Simulation results suggest that the STP mechanism can achieve significantly higher social welfare than a myopic pricing mechanism.

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

General Terms: Economics, Theory

Additional Key Words and Phrases: Market design, Ridesharing, Incentive compatibility

1. INTRODUCTION

Uber connected its first rider to a driver in San Francisco in the summer of 2010. Within a decade's time, ridesharing platforms such as Uber and Lyft have radically changed the way people get around in urban areas. Comparing with traditional taxi systems, a distinct feature of ridesharing platforms is the emphasis on reliable transportation. For example, Uber's mission is stated as "transportation as ubiquitous and reliable as running water, everywhere, for everyone" [Foroohar 2015], and Lyft's mission is "to provide the best, most reliable service possible by making sure drivers are on the road when and where you need them most" [Lyft 2017]. When demand exceeds supply, these platforms use dynamic "surge" pricing to guarantee rider wait times do not exceed a few minutes [Rayle et al. 2014].

In addition to reliability for riders, the platforms also provide the flexibility for drivers to drive on their own schedule. Uber, for example, advertises itself as "work that puts you first— drive when you want, earn what you need" [Uber 2017], and Lyft promises "To drive or not to drive? It's really up to you" [Lyft 2016]. The

Authors' addresses: `hongyao.ma@columbia.edu`, `feifang@cmu.edu`, `parkes@eecs.harvard.edu`

"real-time flexibility" to decide when and where to drive is an important reason that drivers drive for Uber [Hall and Krueger 2016], and increases both driver supply and driver surplus in comparison to alternative, less flexible arrangements [Chen et al. 2019]. More recently, Uber also started to provide drivers in some markets the option to accept only trips they want, based on the trip destinations and expected earnings [Uber 2019].

Despite their success, there remain a number of problems with the pricing and dispatching rules governing these ridesharing platforms, leading in turn to various kinds of market failure and undercutting the endeavor to provide reliable yet flexible transportation. A particular concern, is that trips may be mis-priced relative to each other, incentivizing drivers to cherry-pick [Cook et al. 2018; Chaudhari et al. 2018; Marshall 2020].[1] Drivers can also increase their earnings through strategic behavior in the following scenarios, where there is spatial imbalance and temporal variation of rider demand:

- (Spatial mis-pricing) When the price is substantially higher for trips that start in location A than an adjacent location B, drivers in location B that are close to the boundary can usefully decline trips. This spatial mis-pricing leads to drivers' "chasing the surge"— turning off a ridesharing app while relocating to another location [Campbell 2016; Chen 2017].
- (Temporal mis-pricing) When large events such as a sports game will soon end, drivers can anticipate that prices will increase substantially in order to balance supply and demand. In this case, many drivers will decline trips and even go off-line in order to wait in place [Gridwise 2017].
- (Network externalities) The origin-based "surge pricing" used by many platforms does not correctly factor market conditions at the destination of a trip. This incentivizes drivers to decline trips to destinations where the continuation payoffs are low, e.g. quiet suburbs with low prices and long wait times [Paul 2018].

These kinds of pricing problems undercut the mission of reliable transport, with even high willingness-to-pay riders unable to get access to reliable service for certain trips, such as trips leaving the stadium before a game ends, and trips going to a quiet suburb. This can also lead to inequity, with demonstrated learning effects leading to differences in drivers' long-run earnings (e.g. a gender gap in driver hourly earnings [Cook et al. 2018]), with potential consequences around driver churn from the platform. Simple fixes by limiting drivers' flexibility are not fully effective. For example, when a platform hides trip destinations from drivers before the pick-up, experienced drivers will call riders to ask about trip details, and cancel those trips that are not worthwhile [Cook et al. 2018]. Nor does the imposition of penalties on drivers solve these problems, since drivers may decide to go offline, or choose not to participate in the platform from certain locations or times.

We conceptualize many of the problems with current platforms as arising from prices failing to be appropriately "smooth" in space and time— if prices for trips are higher in one location then they should be appropriately higher in adjacent

[1] There are also other incentive problems, including inconsistencies across classes of service, competition among platforms, drivers' bonuses and off-platform incentives. In the interest of simplicity, we only model a single class of service and ignore cross-platform competition.

locations; if demand would soon increase in a location then the current prices should already be appropriately higher; and if destinations differ in continuation payoffs then trip prices to these destinations need to reflect this. With appropriately smooth prices that correctly reflect the on-trip and network costs of completing each trip, drivers who retain the flexibility to decide how to work will nevertheless choose to accept any trip they are dispatched. Fixing problems with pricing and dispatching also matters in a future with drivers who are employees of the platform, or when trips are completed by autonomous vehicles controlled by the platform. This is because many of the strategic behaviors on today's platforms are symptoms of inefficiencies in pricing and dispatching, for example when dispatching drivers to low-priced trips that send them away from a sports stadium five minutes before a game ends. Correctly designed, ridesharing platforms can succeed in optimally orchestrating trips, without having the power to tell drivers what to do.

2. THE SPATIO-TEMPORAL PRICING MECHANISM

We propose the *Spatio-Temporal Pricing* (STP) mechanism for dispatching and pricing in the context of a ridesharing platform, addressing the problem of providing efficient and reliable transportation while leaving drivers with the flexibility to decide how to work. The STP mechanism achieves:

- Welfare-optimality: maximizing total rider values minus driver costs.
- Incentive-alignment: the prices are appropriately smooth in space and time, such that drivers will always choose to accept any dispatched trips.
- Envy-freeness: drivers at the same location and time do not envy each other's future payoff; riders requesting the same trips do not envy each other's outcomes.
- Core-selecting: no coalition of riders and drivers can make a strictly better plan among themselves.
- Robustness: the mechanism updates the plans after deviations from the original dispatches.
- Temporal-consistency: plans are computed and updated based on the current state but not past history, without using penalties or time-extended contracts.

Welfare optimality and incentive alignment are standard desiderata. For the others, we consider envy-freeness and core-selecting to be of special importance for sharing economy systems such as ridesharing platforms: an envy-free mechanism is fair, removing the variation in drivers' income that depends on lucky dispatches or from learning by doing; a core outcome guarantees that no competing platform can easily enter and take over the more lucrative parts of the market. Even given incentive alignment, robustness is important in the face of unmodeled effects, erroneous predictions, or mistakes by participants. A robust mechanism ensures the other properties from any history onward, and without robustness any solution would be necessarily brittle and poorly suited to practice. Finally, temporal-consistency is important, since using penalties, or threatening to fire drivers or shut down the system are incompatible with the spirit of the sharing economy, and the real-time flexibility of being able to choose how to work.

We work in a complete information, discrete time, multi-period and multi-location model. Thus, the challenge addressed here is one of promoting desirable behavior in the absence of time-extended contracts, and not one of information asymmetry. At the beginning of each time period, based on the history, current positioning of drivers, and current and future demand, the STP mechanism dispatches each available driver to a rider trip, or to relocate, or to exit the platform for the planning horizon. The mechanism also determines a payment to be made if the driver follows the dispatch. Each driver then decides whether to follow the dispatch, or to decline and stay, or to relocate to any location, or to exit. After observing the driver actions in a period, the mechanism collects payments from the riders and makes payments to the drivers. The main assumptions that we make are:

(i) Complete information about supply and demand over a planning horizon,
(ii) Impatient riders, with a value for being picked-up at a particular time and location (and without preferences over drivers), and
(iii) Drivers who each face the same costs for completing the same trip from some origin to some destination at some particular time (and without idiosyncratic preference over riders or locations), and who are willing to provide trips until the end of the planning horizon.

We do allow for heterogeneity in rider values and trip details (the origin, destination, and time of a trip). For drivers, we allow them to become available at different times and locations, and we also model the distinction between a driver who is already driving in the platform (for example, finishing a trip), and a driver who has not yet joined and thus needs to make an entry decision (for example, dropping off a child at school at a specific location and time, and willing to drive afterwards). We also allow a driver who is asked to exit the platform earlier than their intended exit time to incur a one-time cost, modeling the forgone opportunity of outside options after the driver has been driving in the platform for some time.

We first show that the computation of welfare-optimal dispatches can be reduced to a minimum cost flow (MCF) problem. The integrality of the linear program (LP) of the MCF guarantees the existence of anonymous, origin-destination, competitive equilibrium (CE) prices, allowing the price of a trip to depend on market conditions at both the origin and destination. The lattice structure of the dual LP also implies that drivers' total utilities among all CE plans form a lattice.

The STP mechanism uses *driver-pessimal CE prices*, where each trip is priced at the welfare contribution of an extra driver at the origin of the trip, minus the welfare contribution of an extra driver at the destination of the trip, plus the costs for a driver to complete this trip. The mechanism computes a driver-pessimal CE plan at the beginning of the planning horizon, as well as after any deviations from the current plan. This induces an extensive-form game among the drivers, where the total payoff to each driver is determined by the mechanism's dispatch and payment rules, as well as the actions taken by the other drivers.

The main result is that the STP mechanism satisfies all the desiderata outlined above. Somewhat surprisingly, the use of driver-pessimal CE prices (vs., for example, driver-optimal CE prices as in Vickrey-Clarke-Groves mechanisms) is essential for achieving our main result, that accepting the mechanism's dispatches at all times forms a subgame-perfect equilibrium among the drivers. The proof of incentive alignment makes use of the $M^\natural$ concavity of the objectives of the MCF problems [Murota 2003], which implies that there is a stronger substitution among drivers at the same location at the same time, in comparison to drivers at different locations or different times.

We also provide an impossibility result, that no dominant-strategy mechanism has the same economic properties. For three stylized scenarios (the end of a sports event, the morning rush hour, and trips to and from the airport with unbalanced flows), we compare the STP mechanism in simulation with a myopic pricing mechanism that simply clears the market for each location at each time without taking future demand and supply into consideration. Extensive simulation results suggest that the STP mechanism can achieve substantially higher social welfare, and highlight the failure of incentive alignment and envy-freeness due to non-smooth prices in myopic mechanisms.

REFERENCES

Campbell, H. 2016. Advice for new uber drivers - don't chase the surge! https://maximumridesharingprofits.com/advice-new-uber-drivers-dont-chase-surge/. [Online; accessed 8-January-2018].

Chaudhari, H. A., Byers, J. W., and Terzi, E. 2018. Putting data in the driver's seat: Optimizing earnings for on-demand ride-hailing. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. ACM, 90–98.

Chen, K. 2017. Dynamic pricing in a labor market: Surge pricing & flexible work on uber. https://www.youtube.com/watch?v=qI094xqMiPA. [Online; accessed 8-January-2018].

Chen, M. K., Rossi, P. E., Chevalier, J. A., and Oehlsen, E. 2019. The value of flexible work: Evidence from uber drivers. *Journal of Political Economy 127,* 6, 2735–2794.

Cook, C., Diamond, R., Hall, J., List, J., Oyer, P., et al. 2018. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. Tech. rep., Stanford University.

Foroohar, R. 2015. Time person of the year 2015 runner-up: Travis kalanick. https://time.com/time-person-of-the-year-2015-runner-up-travis-kalanick/. [Online; accessed 1-September-2017].

Gridwise. 2017. This is how rideshare drivers are taking advantage of football season. https://gridwise.io/this-is-how-rideshare-drivers-are-taking-advantage-of-football-season. [Online; accessed 1-September-2018].

Hall, J. V. and Krueger, A. B. 2016. An analysis of the labor market for uber's driver-partners in the united states. *NBER Working Paper No. 22843*.

Lyft. 2016. 6 reasons to drive with lyft. https://blog.lyft.com/posts/reasons-to-drive-with-lyft. [Online; accessed 12-December-2017].

Lyft. 2017. Pricing when it's busy. https://help.lyft.com/hc/en-us/articles/115012926227. [Online; accessed 1-September-2017].

Marshall, A. 2020. Uber changes its rules, and drivers adjust their strategies. https://www.wired.com/story/uber-changes-rules-drivers-adjust-strategies/amp. [Online; accessed 17-August-2020].

Murota, K. 2003. *Discrete convex analysis.* SIAM.

Paul, K. 2018. There's a quiet battle of wills between uber drivers and customers over who cancels first. https://www.marketwatch.com/story/new-york-citys-new-uber-rules-could-make-those-5-cancellation-fees-go-away-2018-08-16. [Online; accessed 17-September-2018].

Rayle, L., Shaheen, S., Chan, N., Dai, D., and Cervero, R. 2014. App-based, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in san francisco. Tech. rep., University of California Transportation Center (UCTC).

Uber. 2017. Drive with uber - make money on your schedule. https://www.uber.com/us/en/drive/. [Online; accessed 12-December-2017].

Uber. 2019. To california drivers: keeping you in the driver's seat. https://www.uber.com/blog/california/keeping-you-in-the-drivers-seat-1/. [Online; accessed 17-August-2020].