The Value of Time: A High-frequency Analysis of Ride-Hail Auctions*
(Preliminary, Do not circulate)

Nicholas Buchholz † Laura Doval ‡ Jakub Kastl § Filip Matejka ¶
Tobias Salz ‖

December 20, 2018

Abstract

We use detailed consumer choice data from a large ride-hailing application to study consumer valuations of time. This application offers a unique mechanism that allows taxi drivers in the Czech Republic to bid on trips, and allows consumers to choose between a set of characteristics of a ride: price, waiting time (based on the distance from consumers to each available taxi), car type and driver ratings. We leverage rich variation in bids and customer choices to directly measure consumer willingness-to-pay for time savings and decompose it further by time-of-day and location. Baseline estimates from Prague show the value of waiting to exceed median wages; from $0.40 per minute in peak weekday morning hours, to $0.05 per minute in off-peak weekend hours. Waiting times are important to explain customers' choices, on average accounting for 13.7 percent of total valuation. We use our preference estimates to quantify the welfare benefits of offering a menu of time and price options. Our results show that a flexible menu of waiting time and choices improves welfare by up to 48% over a more standard dispatch mechanism.

Key Words:

JEL classification: C73; D83; L90; R12

*We thank Liftago for providing the data and seminar participants at FTC, Maryland and Princeton for useful comments. Kastl is grateful for the financial support of the NSF (SES-1352305). All remaining errors are ours.

†Princeton University

‡California Institute of Technology

§Princeton University, NBER and CEPR

¶CERGE-EI, Charles University

‖Columbia University

1 Introduction

How do consumers value time outside of labor market transactions? Standard models of labor-leisure choice imply that an extra hour of leisure should be valued at the shadow wage available to a worker for that hour. Dispersion in worker productivity and wages therefore implies a commensurate dispersion in the opportunity cost of time across individuals. Moreover, the dispersed cost of time is of direct consequence to the functioning of markets that ration goods and services through some degree of waiting or queuing. Most markets with binding capacity constraints experience rationing through waiting times. Such markets are abundant, including restaurants, transportation, medical services and many others. It is not clear, however, that consumers value time in these settings according to their wage. There are many reasons for such a discrepancy. Most importantly, a large fraction of workers are not able to freely adjust their hours and the notion of a labor market wage as a shadow cost on time use is therefore misguided.

How much is time valued and to what extent does time rationing influence allocations?[1] Any empirical analysis of markets characterized by time rationing will necessarily require accounting for the time costs that manifest in addition to the standard monetary sources of welfare and profit.

While many markets are either organized around time costs (e.g., delivery services) or capitalize on them (e.g., "skip-the-line" services of airlines), the ability to measure these costs empirically is typically confounded by the bundled nature of time-saving services. For example, frequent flyers may access short check-in lines at the airport but they also enjoy other service amenities on flights. The bundled nature of the product therefore makes it difficult to identify how much the time-component of the service is worth to consumers.

In this paper we use detailed consumer choice data from a large taxi and ride-hailing application in the Czech Republic. This unique market allows taxi drivers to bid on trips, and allows consumers to choose between a set of prices and waiting times (depending on the distance from consumers to each taxi), in addition to driver ratings and model-year of the car. We observe the full menu of choices faced by consumers across 1.9M ride requests and 5.2M bids, as well as their ultimate selection. We leverage rich variation in bids and customer choices to estimate the demand system, enabling us to directly measure consumer willingness-to-pay for time savings. We further decompose these measures across a variety of sources of consumer heterogeneity.

To provide context to our approach, let us envision an ideal experiment aimed at measuring individuals' opportunity costs of time. In this experiment we would provide choices to our subjects in which we allow them to pay a (random) price in order to receive additional minutes of time to be freely allocated to any desired activity. Alternatively, we could offer a random price to eliminate time. In principle we could then recover our subjects' value of time directly via measuring differences

[1]On top of the direct allocative consequences of non-price rationing, here we might also envision dispersion in time costs as an additional dimension on which to segment customers and exercise market power.

in choices across menus. The problem, of course, is that we do not possess a time-machine. Unlike money, or any other material good, the experimenter cannot simply add or subtract time to a subject's time budget. What we can do, however, is ask how much a subject would value additional minutes of freely disposable time versus a more restricted and specific activity, where in our case this activity is waiting for a taxi. In workplace-related experiments, this tradeoff often involves choosing between free time and the dis-amenity value of work at a specific job. In contrast, waiting for a service is an activity that entails less *specific* dis-amenities. Given our interest in measuring the value of freely disposable time, we believe this setting gets us closer to the ideal experiment we envision above.

As early as Becker (1965), economic models of household decision-making have emphasized the importance of time as a resource. A more recent literature has used a variety of widely available micro-data to study the trade-off between market goods and time (Aguiar and Hurst (2007) and Nevo and Wong (2015)). Though these studies utilize rich and comprehensive datasets of consumption behavior (for example, household scanner data), they are only able to measure this tradeoff indirectly. In contrast, we are able to observe consumer selections over time and money trade-offs directly. The labor economics literature has also studied the value of flexible work arrangements, which are becoming an increasingly common phenomenon. Mas and Pallais (2017) uses experimental evidence to estimate workers' preference for flexibility, finding substantial willingness-to-pay for non-fixed work schedules. They find that most people are not willing to pay for flexibility and that workers have a strong distaste to work outside of the standard nine-to-five hours. While their results suggest that the average worker is willing to forgo only small amounts of income to gain flexibility, the tail of workers have very high willingness to pay. Our results mirror this finding. We find that most people are much more price elastic than waiting-time elastic. A small set of riders, however, have persistent and very high willingness to pay for reductions in waiting time. Leveraging the panel structure of our data, and perhaps not surprisingly, we also find substantial heterogeneity in these elasticities across time-of-day and day-of-week within individuals.

Chen et al. (2017) study the value of labor flexibility among Uber drivers and similarly find substantial welfare gains to flexibility. They also demonstrate that these benefits are due to the variability of opportunity costs across the day. Our work builds on this literature in a number of ways. First, we extend the analysis out of the workplace and measure consumers' opportunity costs of time. Our data allow us to disaggregate these measures along various dimensions. Second, we leverage the choice-setting to study counterfactuals in which the platform operator restricts choice and instead matches buyers and sellers according to different mechanisms. These exercises allow us to understand the impact of preference heterogeneity on market participation in settings where waiting time is used to ration goods and services.

Measurements of time cost are also an important ingredient to plan new and evaluate existing transportation infrastructure. For example, the US department of transportation distinguishes in its guidance on Valuation of Travel Time between 'on-the-clock' business travelers and personal travel (Belenky, 2011). For the former the valuation of travel time is assigned to be the nationwide median gross compensation based on the 2015 BLS National Occupational Employment and Wage Estimates. For personal travel, estimates are based on survey results from Miller (1989). Like Miller (1989) many studies that estimate time valuations are situated within the transportation literature, largely based on stated preference reports (see Abrantes and Wardman (2011) for the UK and Cirillo and Axhausen (2006) for Germany). Jara-Daz et al. (2008) combines detailed data (travel diaries and interviews) from Chile, Germany and Switzerland with a theoretical model in the spirit of Becker (1965) to estimate people's value of leisure time. They find that the marginal valuation of leisure is 65.9% of the average hourly wage in Chile, 119.8% in Germany and 87.8% in Switzerland. Small (2012) reviews the travel time literature and presents stylized facts suggesting that the value of personal travel time is about 50% of the gross wage rate and that the value of travel time increases less than proportionally with income/hourly wage - with elasticity estimates ranging from 0.5 to 0.9. Borjesson et al. (2012) study two identical surveys given to car commuters in Sweden in 1994 and 2007 and find that people with below median income have elasticity of travel time with respect to income indistinguishable from zero, and those with above median income have elasticity close to 1. Lam and Small (2001) use survey of California commuters on Route 91 which includes free lanes and tolled "express lanes." The value of time is estimated at $22.87 - that is at 72% of the average wage during their sample period. Fosgerau et al. (n.d.) study data collected from interviews with over 6,000 Danish people and obtain estimates of the value of time being about 67% of the mean after-tax wage. In a long study Significance Quantitative Research (2007) find that the value of time in the Netherlands is about €8.76, with business trips valued at €24 per hour.

Our approach has the advantage that it is directly derived from economically relevant choices. This is especially relevant as many public authorities are explicitly introducing more dynamic pricing mechanisms to manage congestion. Our measures capture both inter-temporal variation in time cost and also cross-sectional variation. We are able to use our estimates to assess how different parts of the population are differentially affected by certain pricing and incentive schemes. These estimates can then be used as inputs into diverse applications ranging from designing optimal reward/punishment schemes to incentivize on-time procurement project completion to optimal pricing of express lanes.

During work-shift hours (6AM to 3PM) we estimate average time valuations of around $20.96 per hour (35 cents per minute) and for more disposable afternoon and late evening time we estimate valuations of $13.92 (23 cents per minute). These averages mask stark heterogeneity within the population. Due to the panel-nature of our data, we obtain individual-specific estimates. Within work hours, individuals in the top quartile value an hour at $40.45 and $9.71 in the bottom quartile. Disposable time in the top quartile is valued at $31.47 and in the bottom quartile $1.22. To put these numbers into perspective the average wage in Prague is about $9.5.[2] The users of the ride hailing platform that is studied here are likely not a representative sample. We are, therefore, in the process of administering a survey in cooperation with Liftago about salaries and work hours of riders that we will be able to link to our individual specific estimates.

While the main focus of this paper is mostly on the demand side of the market, i.e., preferences of passengers, particularly with respect to the trade-off between money and time, our interest certainly lies in analyzing the whole platform. In particular, as hinted upon earlier in the introduction, our goal is to evaluate the design of the platform subject to the constraints provided by the available data. In a companion paper, we are working on conducting a full equilibrium analysis of the platform in order to evaluate possible changes in the design - such as evaluate the costs and benefits of destination-based pricing, evaluate the inefficiency caused by per cent fee structure or to evaluate to cost of price caps typically imposed by cities on cab fares which prevent prices for certain rides to equate supply and demand. We are also working with the platform to conduct field experiments in order to learn about how to best incentivize drivers to accept short rides even when their opportunity cost might be high, which according to the company's executives is the most pressing problem and for which the price caps are typically binding.

There are several recent papers that focus on analyzing the supply side in ride-hail markets. Cook et al. (2018) study gender differentials in earnings among Uber drivers. While the compensation scheme is by definition the same for both women and men, and thus leaves no room for the direct gender-based discrimination in wages, Cook et al. (2018) document that women still earn considerably less on average, even though these differences disappear for highly experienced female drivers. Buchholz (2018) studies dynamic search and spatial equilibrium in the New York taxicab market and quantifies the impact of uniform pricing regulation and search frictions on the spatial allocation of drivers and passengers. Frechette et al. (2018) provide a model of dynamic labor supply and search in the New York taxicab market to asses the effect of entry restrictions and market thickness on efficiency.

The rest of this paper proceeds as follows. We begin by describing the institutional setting and our data in Section 2. Section 3 lays out the model of consumer choice, the drivers and the platform. Section 4 discusses our estimation approach. We present the results of our estimation in Section 5, counterfactuals in Section 6 and conclude in Section 7.

[2]This is computed from an average monthly wage of $1,529.35 and an average of 40 work hours, https://www.jobspin.cz/2016/04/salary-levels-in-prague-brno-and-other-czech-cities-and-various-industries-sectors/

2 Institutional Setting and Data

2.1 Institutional Setting

The landscape of cabs and ride-sharing companies operating in Prague during the period of our study includes the usual suspects. The most important players are the legacy taxi cabs. These cars need to be operated by licensed drivers who in addition to paying the licensing fee need to pass a non-trivial exam. Moreover, these cars need to be equipped by an old-fashioned meter, which captures the number of kilometers traveled in the "occupied" mode together with the billed amount. This meter needs to be certified every two years by a state agency, which then for every meter records the aggregate numbers of kilometers billed and the revenues. A legacy cab driver has several options how to operate: he/she can look for jobs independently (by looking for street hail) and/or subscribe to one of the several phone dispatchers who typically charge a flat monthly fee on the order of $500 and/or operate for Liftago, the platform we study in this paper and whose matching algorithm we will describe below.

Platform Details

Liftago was founded in 2015 with the intention of combining current riding services (street cabs) with the upcoming ride-sharing services. Licensed cab drivers can register with the platform and use its services without any monthly fee. The platform charges a per cent fee per ride instead. The drivers can multihome - and operate using both Liftago and the old-fashioned street-hail or phone dispatch. There are two main services that the Liftago platform provides. First, it matches drivers with nearby passengers seeking a ride, and it does so much faster than a phone dispatch. Second, by tracking both the GPS of the cab and the time of the trip it can provide the approximate fare both before the trip begins and after its completion - and it can even facilitate the payment. This second feature is especially valuable in markets plagued by fare hikes and crooked drivers - and Prague certainly is one of such markets. While Liftago has a strong presence in the Czech Republic, it is still less well known in other countries, even though there are plans for expansions to other European countries in the near future. For us this has the advantage that there are very few tourists in the population of riders (even though the app itself is available in English as well). This makes it easier for us to interpret the magnitudes that we recover from our demand model.[3]

Matching algorithm

The matching of drivers and passengers is implemented by a combination of an optimization algorithm and an auction. Whenever a passenger requests a ride, the system looks for nearby available

[3]This is also reflected in the relatively small fraction of airport rides, which comprise about 2 percent of total trips.

cars and pings a certain number of them (typically four) to elicit an offer. A cab driver who is pinged observes the details of the requested trip - the location of the passenger, the destination, passenger rating and type of payment (cash, credit card, payment through the app etc.). A cab driver who is interested in performing the job submits a bid by selecting one of the tariffs associated with his profile in the app.[4] A tariff consists of a per kilometer fare (capped by the law at CZK 28 ≈ $1.40), a flagging fee and a per minute waiting fee. The passenger observes the bids and selects the preferred one (or the outside option) and the ride is then performed. When starting the ride, the driver selects the appropriate fare among the ones he has available on the meter. This is also why the bidding is not completely unrestricted, a typical driver has only about 5 fare combinations on his meter, but there are notable exceptions, typically drivers who specialize in riding for Liftago customers, who have over 20 different fares. When the ride is completed the customer uses his preselected mode of payment and pays the fare shown on the meter (and can get a tax receipt), the app and an email receipt from Liftago only display an estimated amount.[5]

One noteworthy feature, which separates Liftago both from traditional taxi markets and other well-known ride hailing platforms, is that both waiting times and prices are allowed to vary flexibly. Variation in both dimensions can therefore give rise to a market equilibrium. Traditional taxi markets mostly exhibit rigid fares, which means that the market clears through time rationing. Established ride hail platforms, on the other hand, typically aim to keep waiting times stable and adjust prices to achieve this goal. However, since willingness to pay in both dimensions - time and money - is potentially heterogeneous within the population of riders neither of these two extreme mechanisms is necessarily optimal.

2.2 Data

Our dataset covers 1.9 million trip requests and 1.1 million actual trips on Liftago between September 30 2016, and June 30 2018. For each request, we observe the time of the request, pick-up and drop-off location, trip price bids and estimated waiting times from each driver, and which bid the passenger chose, if any. In addition, we observe a unique id for each driver and passenger. There are 1,455 unique drivers and 113,916 unique consumers over the sample period.

For each ride request we complement the data by precise timestamp-specific trip time estimate by employing Google Maps API based on the GPS addresses of the point of origin and of the destination. Furthermore, we use hourly rainfall data from https://www.noaa.gov/ and daily temperature and snowcover data from the Czech Hydrometeorological Institute.[6] to attach prevailing

[4]Note that neither the passengers nor the other drivers observe these alternative tariffs.

[5]However, if the actual amount diverges from the estimated amount it is extremely easy (and encouraged) for the customer to report the discrepancy and drivers can be banned from the platform if they are found regularly overcharging.

[6]Source: http://portal.chmi.cz/historicka-data/pocasi/denni-data.

Panel A: Bid and Order Summary					
Variable	**Obs**	**10%ile**	**Mean**	**90%ile**	**S.D.**
Price of Trip ($, per bid)	5,229,724	5.65	11.66	19.75	6.261
Wait Time (minutes, per bid)	5,229,724	3.00	6.8	12.00	3.816
Number of Bids (per order)	1,874,407	1.00	2.79	4.00	1.087
Panel B: Daily Summary					
Variable	**Obs**	**10%ile**	**Mean**	**90%ile**	**S.D.**
Number of Requests	638	1961	2938.3	4098	956.67
Number Rides	638	1160	1786.2	2470	557.15
Number Drivers	638	411	499.23	585	97.372

Table 1: Bid, Order, and Daily Summary Statistics

weather characteristics to each ride. Finally, we use data on GPS-specific land values from GIS coded data available at http://www.geoportalpraha.cz.[7]

Data Summary

Table 1 summarizes the daily activity of the platform. There are on average about 3,000 trip requests each day, of which 61% become rides. The average bid is $11.66 and the average waiting time is 7 minutes. In addition, about one third of the drivers in the sample were active each day. The average number of drivers bidding in each auction is 2.8, and except in rare cases there are no more than four bids.[8]

Figure 1 Panels (a) and (b) show the average trip price and wait times by day of the week and time of the day. We see that prices are lower during the weekday afternoons and higher during the weekends, while wait times tend to be substantially higher during the day hours compared with overnight hours, across both weekday and weekend.

Figure 1 Panels (c) and (d) summarize the week-over-week trends in prices (panel (c)). Beginning 2017 the ridership and drivers stop to grow and enter a relatively stationary period, although there are some large swings in the number of passengers towards the end of 2017. It shows the relationship between labor supply and price during the sample period. As expected, price decreases as supply increases during the winter holiday seasons and increases as supply falls during the summer. The second panel shows how average waiting time and average price evolve during the sample period.

[7] We are currently in the process of running a survey, which should provide us with further passenger covariates: education, income, commuting patterns etc.

[8] We discard auctions with more than four bids, representing only 0.33% of the sample.

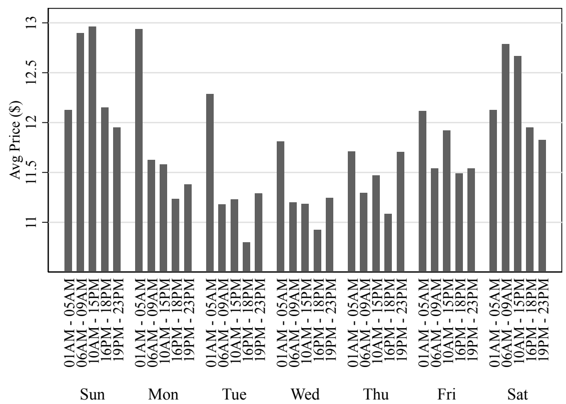

(a) Average Prices

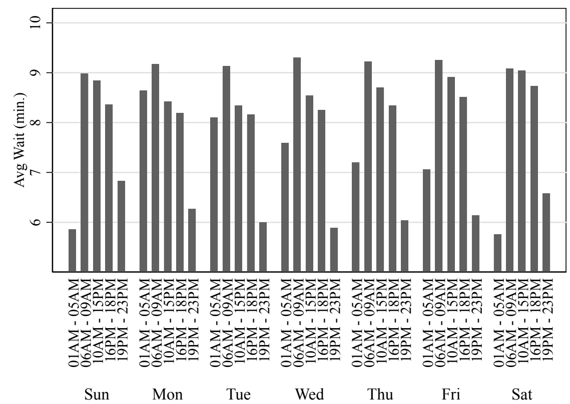

(b) Average Wait Times

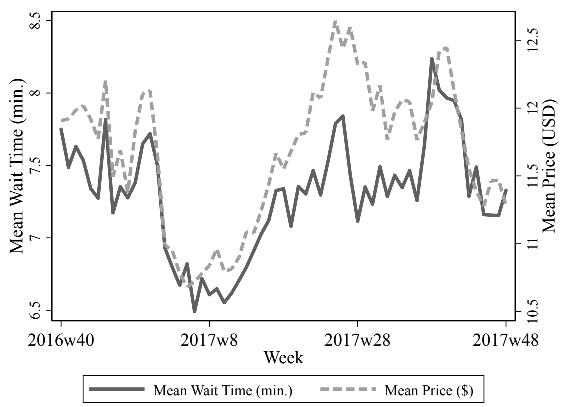

(c) Average Prices

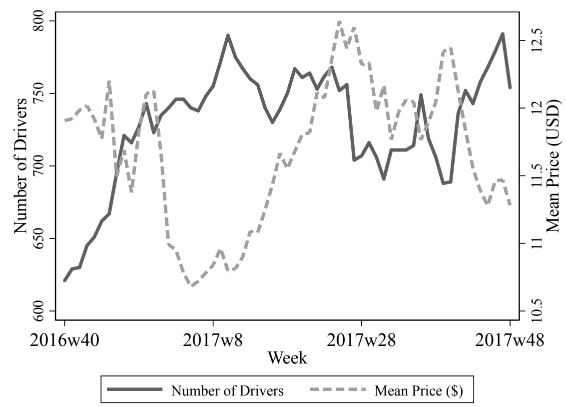

(d) Price and Driver Count

Figure 1: Day and Time Statistics

2.3 Preferences for Time and Money

Figure 3 shows the discrepancy between bids and selections across price and waiting time. Just as prices clear markets in traditional settings, both prices and wait times are equilibrium objects in this setting, dictated by the interaction of supply and demand as well as the preferences of customers over each dimension. What we see in the figure is that preferences seem to matter, as chosen prices and wait times are both systematically lower than the bids. Moreover, the magnitudes of these differences vary throughout the day, which is attributable to some combination of preference heterogeneity across customers as well as within-customer heterogeneity across the day. The most volatile swings across these two dimensions of choice characteristics come from waiting times, of which the average daily high is larger than the average daily low by around 60-80%.

By observing trade-offs in the choice set available to each customer, the data allow us to estimate

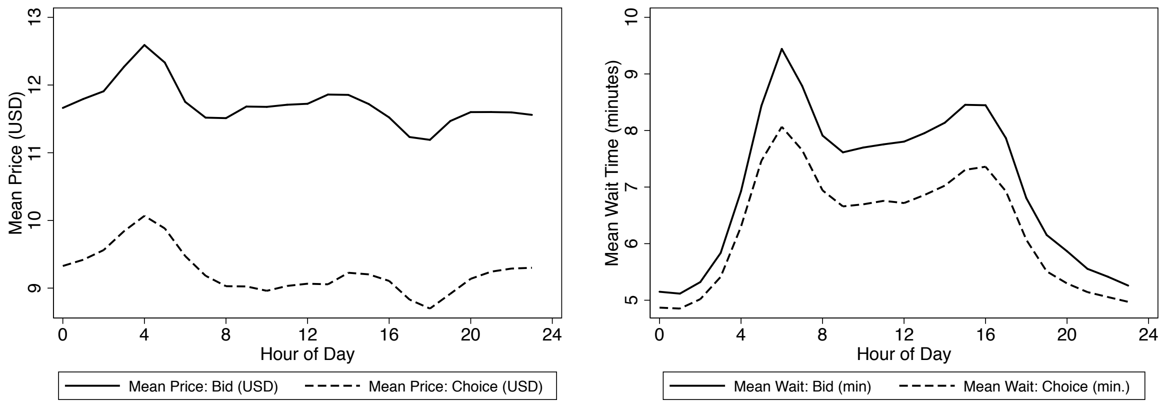

(a) Number of Trips (b) Number of Drivers

Figure 3: Average Bid and Chosen Price and Waiting Times by Hour

such heterogeneity. To illustrate this, the first panel of Figure 5 shows the ratio of rides with no tradeoff across the day. A ride has no tradeoff if the driver with the minimum bid also has the shortest waiting time, excluding rides where there's only a single bid. The second panel shows how passenger's preference changes within a day. Specifically, in rides where there is a tradeoff between price and waiting time, passengers are more likely to choose minimum price at night and minimum waiting time during the day. These two lines intersect at 5 am and 8 pm.

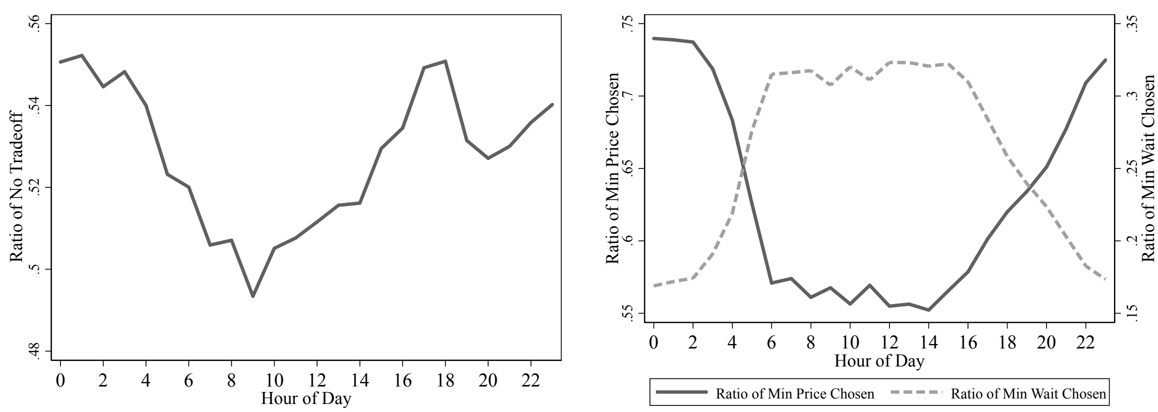

(a) No Tradeoff Ratio (b) Min Price vs. Wait Chosen

Figure 5: Ratio of Rides w/o tradeoff by hour

The goal of our demand model will be two-fold. First, we want to estimate preferences for time and money in order to measure the value of time implied by customer behavior. Second, we want to understand how participation in the market is related to customer heterogeneity by estimating distributions of preferences and their relation to the availability of choices offered by the platform at different times.

3 Model

We split our discussion of the model into two parts. We begin by the demand side which focuses on passengers and their choice. The goal of the model is to recover the underlying preferences over the various characteristics of cab rides. The second part will focus on drivers, i.e., the supply side of the market. The goal of the supply side of the model is to recover the distribution of outside options that the drivers face when competing for the rides. These outside options are essentially coming from a value function associated with the dynamic program each driver is (implictly) solving. We will then put everything together in order to evaluate various counterfactual scenarios.

3.1 Demand Side

Passenger side: A passenger i located in location g requesting a ride to location h in time t faces a standard discrete choice problem of choosing between $J \equiv |\mathcal{J}|$ alternatives in the index set $\mathcal{J}$ (including the outside option, the utility of which is normalized to zero). To keep the notation in check we designate a route index $r \in \mathcal{R} = \mathcal{G} \times \mathcal{H}$.

Each alternative $j \in \mathcal{J}$ is characterized by a tuple consisting of price, time to wait and characteristics of the car (model, year, color), driver (ratings and name) and the stochastic part $\epsilon_{i,j,r,t}$. His (indirect) utility from option j can be written as:

$$v_{i,j,r,t} = \beta^{w}_{i,r,h_t} \cdot w_{i,j,t} + \beta^{p}_{i,r,t} \cdot p_{i,j,r,t} + \beta^{x}_{i,r,h_t} \cdot X_{i,j,r,t} + \epsilon_{i,j,r,t} \qquad (1)$$

Under the assumption that $\epsilon_{i,j,r,t}$ are independently and identically distributed according to a Type I extreme value distribution, choosing the maximum among J alternatives with utilities given by (1) reduces to the standard logit. The (ex-ante) probability that an alternative j will be chosen depends on its relative mean utility.

We have a panel of passengers who repeatedly face the above described discrete choice with rich variation in choice sets and characteristics of the alternatives. Since this is a well-understood model, there is no need for us to dwell on identification or estimation issues. In particular, standard endogeneity concerns are much less severe in our case since we observe passenger choice sets along with all attributes that should determine the relative attractiveness of the observed options. This is, however, not true for the outside option and one could construe that choice attributes interact with unobserved aggregate shifters. Note that some of the obvious candidates are observable (rain, snowfall etc.) and can thus be controlled for. Without random coefficients, the model can be estimated by maximum likelihood. When allowing for random coefficients, we estimate the model by Markov Chain Monte Carlo (MCMC) Gibbs Sampling in order to avoid the need for calculating multidimensional complicated integrals.

3.2 Supply Side

On the supply side of the market, the drivers face a decision problem which is inherently dynamic. Whenever they get asked to submit a bid for a ride requested via the Liftago app, they compare the value of participating in the ensuing auction against their outside option, which is, however, unobserved to the econometrician. Nevertheless, the model will allow us to link the observed bids with these unobserved outside options.

Denote ω a per-period value that drivers draw while not servicing trips on the platform, where a period is defined as a small time interval. We can interpret ω as representing the (per period) expected fare that drivers can receive for street pickups or any other alternative use of their time. It is drawn from some time- and location-dependent distribution. This set of conditional distributions is the key primitive that we want to learn from the model. We assume that location transitions are exogenous, and hence we do not model optimal location search a la Buchholz (2018).[9] Let $\delta(a_t)$ denote the probability of a Liftago ping during one time period. We make a specific timing assumption: in every period, a driver might receive a fare draw with Liftago ping and thus face a choice between ω or participate in the auction, or collect ω directly without a Liftago ping. This specific timing assumption allows us to learn about outside option draws from the bidding behavior in the auction. Denote by $\mathcal{H}^t(a_t, a_{t+\tau}, \omega)$ the value of receiving a Liftago ping requesting a ride from origin a_t to destination $\hat{a}_{t+\tau}$ taking τ time periods. We can then define $\mathcal{S}^t(a_t, \omega)$ as the value of being in area a with a fare draw ω:

$$\mathcal{S}^t(a_t, \omega) = (1 - \delta(a_t)) \cdot \Big[\omega + \mathbb{E}_\omega[\mathcal{S}^{t+1}|a_t]\Big] + \delta(a_t) \cdot E_{\hat{a},\tau}[\mathcal{H}^t(a_t, \hat{a}_{t+\tau}, \omega)|a_t] \tag{2}$$

where:

$$E_{\hat{a}_{t+\tau}}[\mathcal{H}^t(a_t, \hat{a}_{t+\tau}, \omega)|a_t] = \int_{\hat{a},\tau} \mathcal{H}^t(a_t, \hat{a}_{t+\tau}, \omega) dF(\hat{a}, \tau|a_t) \tag{3}$$

starts at a_t and ends at $\hat{a}_{t+\tau}$. This essentially integrates out the possible ways to get to all possible location-outside option pairs along all paths taking τ periods. This means that the expectation is a double integral both over areas as well as the travel times to get there.

[9] This is fine for the purposes of this paper as all counterfactuals regard the platforms pricing and matching strategy as fixed, and given our maintained assumption of a "small Liftago" it shouldn't affect the conclusions that we draw from counterfactuals. However, if we were interested in changing the design of the platform, this could become an issue.

Focusing on the value of getting a Liftago ping with a requested trip to area $a_{t+\tau}$ (assuming trip duration $\tau \geq 1$), let $p(b)$ denote the probability that the driver wins the auction with bid b. We thus obtain:

$$\mathcal{H}^t(a_t, \hat{a}_{t+\tau}, \omega) = \max_b \left\{ p(b) \cdot \left(b + \beta^\tau \cdot \mathbb{E}_{\hat{\omega}} \left[\mathcal{S}^{t+\tau}(\hat{a}_{t+\tau}, \hat{\omega}) \right] \right) \right. \\ \left. + (1 - p(b)) \cdot \left(\omega + \beta \cdot \mathbb{E}_{\hat{\omega}}[\mathcal{S}^{t+1}(a_t, \hat{\omega})] \right) \right\}$$

Note that we do not model separately a driver's choice to reject a ping. This can be subsumed in this equation by assuming that the driver chooses a price of infinity. Alternatively, one could introduce an additional parameter, such as some fixed cost of bidding. While perhaps more realistic, we do not expect this to have any significant impact on our objects of interests: that is the estimates of preferences with respect to price and waiting time.

Since the choice of b does not impact the expectation defined in (3), we can rewrite the maximization problem in a way that directly maps into a static auction setup:

$$\mathcal{H}^t(a_t, \hat{a}_{t+\tau}, \omega) = \max_b p(b) \cdot \left(b + \beta^\tau \cdot \mathbb{E}_{\hat{\omega}} \left[\mathcal{S}^{t+\tau}(\hat{a}_{t+\tau}, \hat{\omega}) \right] - \mathbb{E}_{\hat{\omega}} \left[\mathcal{S}^{t+1}(a_t, \hat{\omega}) \right] - \omega \right) + \\ \omega + \beta \cdot \mathbb{E}_{\hat{\omega}} \left[\mathcal{S}^{t+1}(a_t, \hat{\omega}) \right] \quad (4)$$

Let:

$$v(a_t, a_{t+\tau}, \omega, t, \tau) \equiv - \left(\beta^\tau \cdot \mathbb{E}\left[\mathcal{S}^{t+\tau} | a_t, a_{t+\tau} \right] - \mathbb{E}\left[\mathcal{S}^{t+1} | a_t \right] - \omega \right) \quad (5)$$

We can, therefore, rewrite the bidder's problem as:

$$\max_b p(b) \cdot (b - v(a_t, a_{t+\tau}, \omega, t, \tau)) \quad (6)$$

Assuming the bid space is connected, we can directly appeal to (the procurement version of the) GPV and back out $v(a_t, a_{t+\tau}, \omega, t, \tau)$ by assuming optimal bidding behavior and inverting the bids. It should also be clear that since $\mathbb{E}[\mathcal{S}^{t+\tau}|a_t, a_{t+\tau}]$ and $\mathbb{E}[\mathcal{S}^{t+1}|a_t]$ enter additively, we can identify these values separately from the random part. Moreover, we can simply estimate these values by running a fixed effects regression, and obtain ω as the resulting residual.

If the bid space is discrete, we can use a bounds approach described below. Although it might, in practice, be hard to separately identify $\mathbb{E}[\mathcal{S}^{t+1}|a_t]$ from the distribution of ω, both of which depend on a_t.

To get at the bounds approach, it will be useful to approach the optimal bidding problem by utilizing our demand side estimates. Recall that the optimal bid for a trip from location a_t to location a_{t+tau} given opportunity cost $v\left(a_t, a_{t+\tau}, \omega, t, \tau\right)$ solves:

$$b \in \arg\max_b \mathbb{P}\Big[\text{Driver } j \text{ chosen}|b, x_{i,j}\Big] \cdot \left(b - v\left(a_t, a_{t+\tau}, \omega, t, \tau\right)\right) \tag{7}$$

where $x_{i,j}$ is a vector of characteristics of passenger i and driver j. Note that the probability that j is chosen of course implicitly depends on rivals' actions.

This is essentially an asymmetric first price auction with competition on price b and quality (characteristics x_j such as car type or wait time), where the latter is exogenous by the time the request arrives. The probability of "winning" the auction is determined by customer's preferences for both the fare as well as other characteristics of the options in his/her choice set and hence it can be obtained using our demand system. Let $w_{i,j}$ be the implied waiting time for passenger i until the pick-up by driver j and let r be the trip requested characterized by the location of the passenger, or, equivalently, the point of origin of the ride and the destination. We can then simulate, the following integral:

$$\mathbb{P}\Big[\text{Driver } j \text{ chosen}|b, X_j, r\Big] = \int \frac{exp(\beta^w_{s,r,h_t} \cdot w_{i,j} + \beta^f_{s,r,h_t} \cdot b_j + \beta^x_{s,r,h_t} \cdot X_{s,j})}{1 + \sum_{k \neq j} exp(\beta^w_{s,r,h_t} \cdot w_{i,k} + \beta^f_{s,r,h_t} \cdot f_k + \beta^x_{s,r,h_t} \cdot X_{s,k})} d\mathcal{F}(s, \boldsymbol{w}_{-j}, \boldsymbol{b}_{-j}) \tag{8}$$

The above expression shows that we can simulate the winning probability under the customer "scoring rule" once we have the demand model by sampling from the customer population s (customer attributes other than his/her location and destination are unknown to the driver at the time of the bid) as well as competitors bids and distances. We will next show that we can use this known object along with observed bids under a moment inequality approach to learn about the underlying distributions of valuations. For notational simplicity let $\mathbb{P}\Big[\text{Driver } j \text{ chosen}|b, X_j, r\Big] = p[b_j]$. Furthermore, let $\mathcal{B}_j = \{b_{1,j}, ..., b_{n,j}\}$ be the support of possible bids that driver j can choose in his meter. Let c (for chosen) indicate the index of the optimal bid, i.e., b_c is in the set defined in (7). We then have:

$$p[b_c] \cdot \left(b_c - v\left(a_t, a_{t+\tau}, \omega, t, \tau\right)\right) \geq p[b_{c-1}] \cdot \left(b_{c-1} - v\left(a_t, a_{t+\tau}, \omega, t, \tau\right)\right)$$

which after rearranging gives us the lower bound:

$$v\left(a_t, a_{t+\tau}, \omega, t, \tau\right) \geq \frac{p[b_{c-1}] \cdot b_{c-1} - p[b_c] \cdot b_c}{p[b_{c-1}] - p[b_c]} \equiv L_{c,t}$$

And similarly:

$$p[b_c] \cdot (b_c - v(a_t, a_{t+\tau}, \omega, t, \tau)) \geq p[b_{c+1}] \cdot (b_{c+1} - v(a_t, a_{t+\tau}, \omega, t, \tau))$$

get us the upper bound:

$$U_{c,t} \equiv \frac{p[b_c] \cdot b_c - p[b_{c+1}] \cdot b_{c+1}}{p[b_c] - p[b_{c+1}]} \geq v(a_t, a_{t+\tau}, \omega, t, \tau)$$

where b_{c-1} and b_{c+1} are the bids adjacent to the chosen one, b_c, among those that driver j can choose from.

With these continuation values at hand, we can begin by perform a similar descriptive decomposition as with the value of time. In particular, we will be interested in learning about patterns of opportunity cost during different hour of the day, days of the week or across different locations.

4 Estimation

We begin by laying out our estimation procedure for the parameters of the demand model and will present the details of the estimation of the supply side later.

4.1 Estimation Details of the Demand Side

4.1.1 Simple Logit

To illustrate how willingness to pay varies throughout the day we first estimate a simplified variant of (1) where the price and time coefficient are allowed to vary by hour, i.e. $h_t \in \{0, ..., 23\}$ and that does not include individual specific coefficients. Figure 7 and Figure 8 in the following section are based on this specification. We then move to our preferred specification of the model with random coefficients but where we partition the day into five blocks, by which we allow willingness to pay to vary. The omitted category is the midnight hour (12AM to 12:59AM) and the remaining blocks are 1AM to 5AM, 6AM to 9AM, 10AM to 3PM, 4PM to 6PM, and 7PM to 11:59AM. We also control for drivers' quality ratings, car type, traffic speed and the distance of the trip. Furthermore, we add controls for weather (dummy for rain measured at hourly frequency)

The coefficient estimates allow us to construct time valuations for different individuals and different times of the day. This value is obtained via the following equality, which compares the utility of choice j with the utility of some hypothetical option j' that adds a single minute to waiting time, but otherwise has the same characteristics. The price difference $p_{j,t} - p_{j',t}$ that solves the equation reflects the additional units of money needed to make consumers indifferent between paying more or waiting more (on average).

$$\beta^p_{i,r,h_t} p_{j,t} + \beta^w_{i,r,h_t} w_{j,t} + \beta^x_{i,r,h_t} X_{i,j,r,t} = \beta^p_{i,r,h_t} p_{j',t} + \beta^w_{i,r,h_t}(w_{j,t}+1) + \beta^x_{i,r,h_t} X_{i,j,r,t} \tag{9}$$

This implies that a minute of time is valued as

$$\mathbf{WTP}(\$) = p_{j,t} - p_{j',t} = \frac{\beta^w_{i,r,h_t}}{\beta^p_{i,r,h_t}}. \tag{10}$$

4.1.2 Random Coefficients Logit

To capture rider heterogeneity we allow for individual specific coefficients in Equation 1. Note that due to individual level panel data and the absence of endogeneity problems we do not need to transform this into a GMM problem. We instead opt for a hierarchical bayes mixed logit procedure to obtain individual specific estimates for the dis-utility of waiting, which we estimate via an MCMC method using data augmentation of latent variables as in Tanner and Wong (1987). In this approach the unobserved random coefficients are simulated and then these simulations are treated as data, which sidesteps the need to evaluate multidimensional integrals by sampling from a truncated Normal distribution instead. We closely follow Train (2009). For clarity of exposition we now describe the univariate version of the Gibbs sampler that we construct, whereas for the full estimation we will later allow a vector of individual specific coefficients. We posit that the waiting time coefficient is additive in time of day, place, and an individual specific shifter that is normally distributed. Let $v_t \in \{v_1, ..., v_K\}$ be an indicator that points to a specific time interval that we obtain from a partition of day. We will allow for individual specific coefficients for each of these intervals, $\boldsymbol{\beta} = (\beta_{i,v_1}, ..., \beta_{i,v_K})'$, and also allow these to be correlated with each other.

$$\beta_{i,r,t} = \beta_{i,v_t} + \beta_r + \beta_{h_t}, \quad \beta_i \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{11}$$

We assume that $\mu \sim \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\sigma}_0)$, where $\boldsymbol{\sigma}_0$ is a diffuse prior (unboundedly large variance). Assume that $\Sigma \sim IW(v_0, s_0)$. One can then iteratively update the β_i-coefficients, the mean of the coefficients as well as the standard deviations. The specific assumptions on the priors lead to conjugate distributions where the posterior mean of β_i is itself normal and the variance again in the family of inverse gamma distributions. To describe the updating algorithm, let $\overline{\beta}^l$ be the sample mean of coefficients of iteration l in the chain and $\overline{s}^l$ be the sample variance.

The key simplification exploited in the Gibbs sampler is that one does not have to obtain an analytical expression for the posterior distribution of the β_i's, which instead only needs a proportionality factor that can be easily computed at each step. In particular:

$$K(\beta_i|\boldsymbol{\mu}^l, \boldsymbol{\Sigma}^l, \mathbf{y}_i) \propto \Pi_t L(y_{ti};\beta_i) \cdot \phi(\beta_i|\boldsymbol{\mu}^l, \boldsymbol{\Sigma}^l) \tag{12}$$

where $\mathbf{y}_i$ is the vector of choices observed for passenger i and $L(y_{ti};\beta_i)$ is the likelihood contribution of a particular choice.

1. Draw a new posterior mean μ^l for the distribution of coefficients from $\mathcal{N}(\overline{\beta}^{l-1}, \frac{W}{N})$.
2. Draw $\boldsymbol{\Sigma}^l$ from $IW(K+N, s^l)$, where $s^l = \frac{K \cdot I + N \cdot S_1}{K+N}$, and $S_1 = \frac{1}{N} \cdot \sum_n (\boldsymbol{\beta}_{l-1} - b)(\boldsymbol{\beta}_{l-1} - b)'$.
3. For each i, draw $\boldsymbol{\beta}_i^l$, according to the metropolis hastings algorithm, with new proposal $\boldsymbol{\beta}_i^{pl}$ starting from $\boldsymbol{\beta}_i^{l-1}$ using density $\phi(\boldsymbol{\beta}_i|\boldsymbol{\mu}^l, \boldsymbol{\Sigma}^l)$.[10]

However, since such an MCMC procedure is known to be slow for a large dimensional parameter space we iterate on the MCMC procedure to get the β_i's and a standard likelihood procedure for the remaining parameters.

4.2 More Details on Estimation of the Supply Side

Estimation of the model is straightforward. In the first stage we simulate the probabilities of winning, for which we use the demand model.[11] This allows us to directly obtain bounds on the opportunity cost. For estimation it will be useful to translate these bounds into valuations for an entire hour. If the units of t and τ are set in minutes, we can scale the above bounds by $60/\tau$. We also assume that the distribution of valuations is fixed in an hour so that we can index a trip that falls within an hour in terms of that hour or in terms of the two adjacent hours if the trip jumps the hour mark. We can therefore express:

$$\frac{60 \cdot L_{c,t}}{\tau} \leq 60 \cdot \frac{\omega_c(a, h_t)}{\tau} \leq \frac{60 \cdot U_{c,t}}{\tau} \tag{13}$$

As on the demand side, estimation is easy in this case and can be conducted by MCMC + data augmentation techniques, which operates on the following three steps. (I) guess a distribution for the structural error terms of the valuation that we want to explain. (II) Take a draw from the distribution that takes into account the truncation of the dependent variable, which essenatially becomes the valuation. This draw allows us to get an observed left hand side. (III) Run a regression of valuation on covariates to recover the estimates of the other coefficients of the linear index.

[10] Additional details on these steps and the Metropolis Hastings algorithm are provided in **HERE EVENTUALLY REFERENCE TO CORRECT APPENDIX CHAPTER**.

[11] One could also follow the usual approach in empirical analysis of auction data and estimate the (conditional) bid distributions and integrate over the number and locations of participating drivers and use that to evaluate the probability of winning.

5 Results

5.1 Demand Estimation Results

Table 2 Panel A shows the coefficients that we obtain from the demand model along with standard errors. While the disutility of money is relatively constant through the day, the intra-daily coefficients on waiting time vary substantially from the day-time hours to the evening hours, with coefficients differing by up to a factor of 2.5 (from early morning to the morning rush-hours). Customers are more price sensitive but less waiting-time sensitive during the weekend, and more sensitive to waiting during periods of rain and while in the center of the city. Customers are less sensitive to waiting when getting dropped of at the airport, and while getting picked up near residential addresses. We also control for a variety of covariates associated with the market, trip, and bid. Within each bid, luxury cars and trucks yield higher utility, while motorcycles ("Car Type: Tiny") yield much less. Higher ratings are unsurprisingly more valued, though the variance is low. Characteristics of the trip also matter: longerrides, airport pickups, residential drop-offs and central city pickups all yield higher utility and thus higher likelihood of selecting a trip against the outside option. Market variables such as weekends, faster traffic conditions, and rain all confer lower utility.

Panel B shows the elasticities of price and waiting time, computed as the percent change in selecting the inside option with respect to a percent change in price and waiting time, respectively. Since the current specification only estimates random coefficients on waiting time, we show the elasticities against various quantiles of waiting preference (as measured by the estimated distribution of β_i^w's). Thus by construction the price elasticities are not varying with this. A noteworthy point is that even the top decile of waiting preferences are substantially less elastic than that of prices.

5.2 Simple Logit Results

First we present results from the model with a granular breakdown of hours but without heterogeneity. Figure 7 shows the pattern of disutility for price and waiting time (broken down by Weekend and Weekday) over each hour of the day. What emerges is a pattern in which customers have lower dis-utility of price and higher dis-utility of waiting during the daytime hours of around 5am to 7pm.

With these hourly coefficient estimates, we can turn to equation 10 to estimate the implied value of saving one minute. Figure 8 plots these values. Two stark patterns can be seen. First, weekday and weekend customers have very similar intra-day patterns of time valuation, though weekday customers' valuations are shifted up by approximately 8 cents per minute throughout the day. The value of time is estimated between 5 cents per minute in the weekend overnight hours up to 40 cents in the peak morning weekday hours.

Panel A: Parameter Estimates		
	Coefficient	Std Error
Price 1AM-5AM	-0.453	0.004
Price 6AM-9AM	-0.465	0.004
Price 10AM-3PM	-0.479	0.004
Price 4PM-6PM	-0.484	0.004
Price 7PM-11PM	-0.492	0.004
Price × Weekday	0.006	0.002
Waiting Time 1AM-5AM	-0.05	0.004
Waiting Time 6AM-9AM	-0.125	0.004
Waiting Time 10AM-3PM	-0.127	0.004
Waiting Time 4PM-6PM	-0.102	0.004
Waiting Time 7PM-11PM	-0.052	0.003
Waiting Time × Weekday	-0.035	0.003
Waiting Time × Airport Pick Up	-0.01	0.022
Waiting Time × Airport Drop Off	0.063	0.008
Waiting Time × Rain	-0.001	0.003
Waiting Time × Center	-0.038	0.011
Waiting Time × Res. Pickup	0.03	0.004
Waiting Time × Res. Dropoff	-0.01	0.004
Car Type: Luxury Low	0.092	0.008
Car Type: Luxury High	0.24	0.013
Car Type: Truck	0.684	0.076
Car Type: Tiny	-1.787	0.183
Trip Distance	10.89	0.081
Rating	3.903	0.09
Week Day	0.158	0.031
Trip Speed	-0.07	0.003
Airport Pick Up	0.202	0.189
Airport Drop Off	-0.807	0.072
Rain	-0.046	0.026
Center Trip	0.274	0.082
Residential Pickup	-0.179	0.031
Residential Dropoff	0.168	0.029

Panel B: Implied Elasticities		
	Price Elasticity	Waiting Elasticity
Waiting Preference: Mean	-2.689	-0.55
Waiting Preference: 10%ile	-2.689	-0.113
Waiting Preference: 25%ile	-2.689	-0.3
Waiting Preference: 50%ile	-2.689	-0.589
Waiting Preference: 75%ile	-2.689	-0.883
Waiting Preference: 90%ile	-2.689	-1.101

Table 2: Model Coefficients Estimates

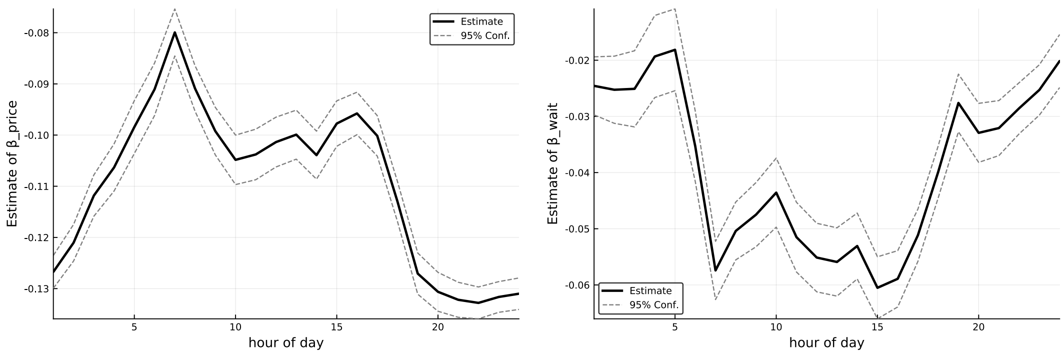

Figure 7: Parameter Estimates Example: Weekend Consumers

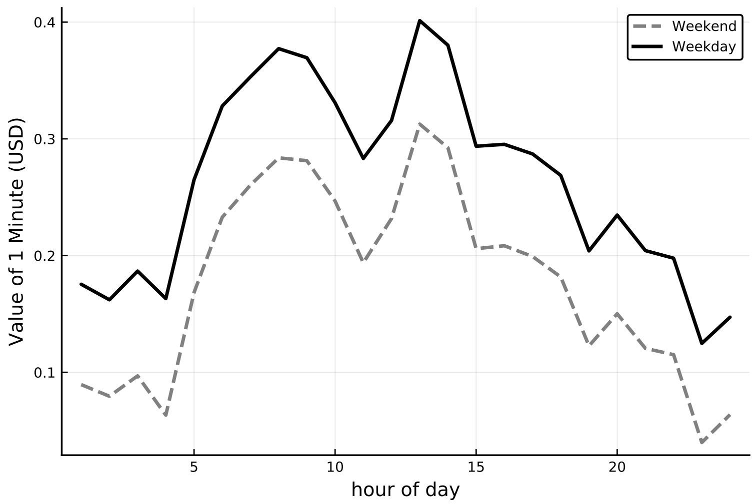

Figure 8: Value of One Minute by Group

Note that these inter-temporal patterns are both due to intra-personal differences of time valuations during the day but also due to cross-sectional differences across riders that takes trips at certain times of the day more than others. If we were to relax this assumption and accept that there is both within- and across-individual preference heterogeneity, then our estimates in Figure 8 do not allow us to distinguish the extent to which our results are driven by different times of day that select on different customers, versus a consistent pattern of within-individual preference heterogeneity across times of day.

5.3 Heterogeneity Estimates

We now present estimates of the MCMC estimation. Though we are now estimating individual heterogeneity, we continue to partition the estimated results into groups representing times of day and weekday/weekend. We estimate the distribution of parameters β^w_{i,r,h_t} within each group and compute the implied distribution of the value of time.

Figure 9 compares the estimates β^w_i as well as the implied value of time distributions across times-of-day and weekday-weekend. We estimate some consumers to have what appears as negative values - i.e., they exhibit a willingness to pay for *waiting* as opposed to saving time. One interpretation of this is that some consumers may utilize the choice set to select service according to some schedule (e.g., "I need a ride 20 minutes from now"), and therefore prefer to wait longer. Most consumers appear to assign positive values to saving time, however. Panel I shows that Morning and Mid-day customers tend to assign higher value to time savings, possibly commensurate with business and deadline-related travel. Panel II shows a similar pattern for weekdays compared with weekends.

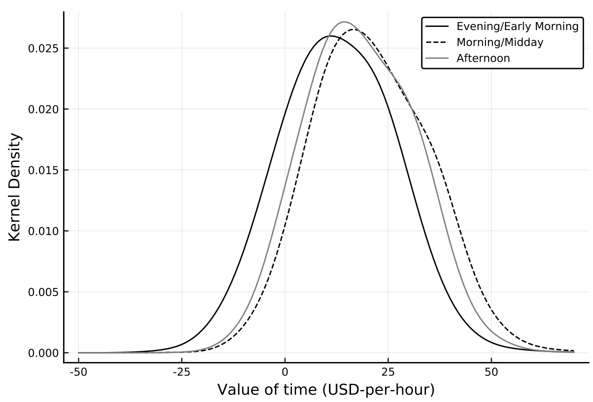

(a) Panel I: Time-of-Day

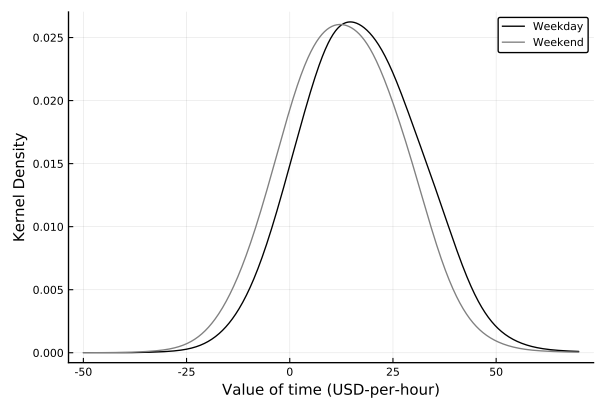

(b) Panel II: Weekday/Weekend

Figure 9: Value of Time Distributions

5.4 Income Comparison

One implication of our estimates is that the opportunity cost of time distribution that we estimate is somewhat higher than the local income distribution. Figure 11 compares the cumulative distribution of our estimates with the minute-level earnings deciles which we impute from Prague household income deciles.[12] This figure shows that the 80th percentile of income in Prague is $0.20,

[12]This calculation is based on data for the household income deciles in the Czech Republic, which we then scale such that median incomes match those in Prague.

which corresponds roughly to the 60th percentile in the value of time. Our results therefore suggest that opportunity costs may be higher than those implied by incomes.

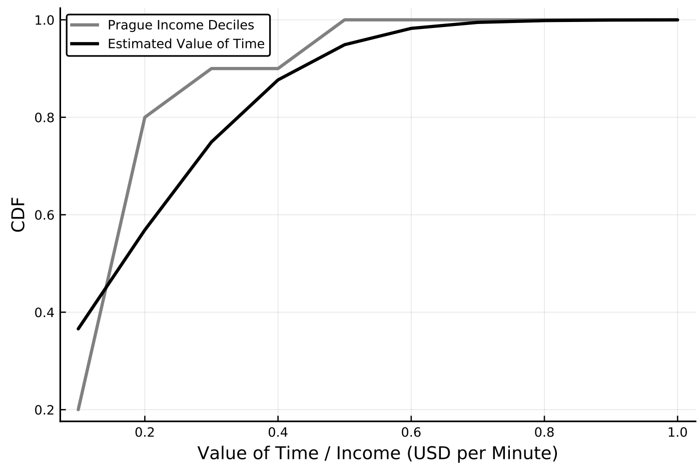

Figure 11: Distribution of Incomes and Time-Values

5.5 Comparing Work Time and Disposable Time

Previous literature has been able to estimate the value of flexibility in the context of a particular job, implying that opportunity costs of time might vary across the day. One of the key advantages of our setting and estimation approach is that we can measure each consumer's value of time at different times of the day, across work and non-work hours. Given these individual measures, we can look at several interesting quantities.

We first note that there is very high intra-daily correlation in time valuation. Across our sample, the correlation between time valuation in work and non-work hours is 94%. Figure 12 Panel I shows how time is valued across work and non-work hours across different quantiles of the time-value distribution. It shows that all value-types exhibit lower time valuation during disposable-time hours compared with work hours, with an average decline of around 34 percent.

We also examine the frequency that different consumer types use the service. Figure 12 Panel II shows the relative frequency of consumers, by quintile, that appear in our data. We see that during the day the top 20% of users by time value constitute 40-45% of trips on the platform. These users are less frequent in the evening, but still over-represented. By comparing this analysis against the within-person correlations, we learn how much observed variation in the data comes from compositional or selection variance versus within-person variance. [Insert the comparison here]

Finally we also measure the relative dispersion of opportunity costs at different times of day. By comparing the variance of time-valuation across our working- and non-working times, we see a remarkably consistent degree of dispersion with almost equal variances around 0.055. However, because the mean valuations are lower in the evening, the percent variance is therefore higher: 0.47% in the work hours versus 1.03% in the non-work period.

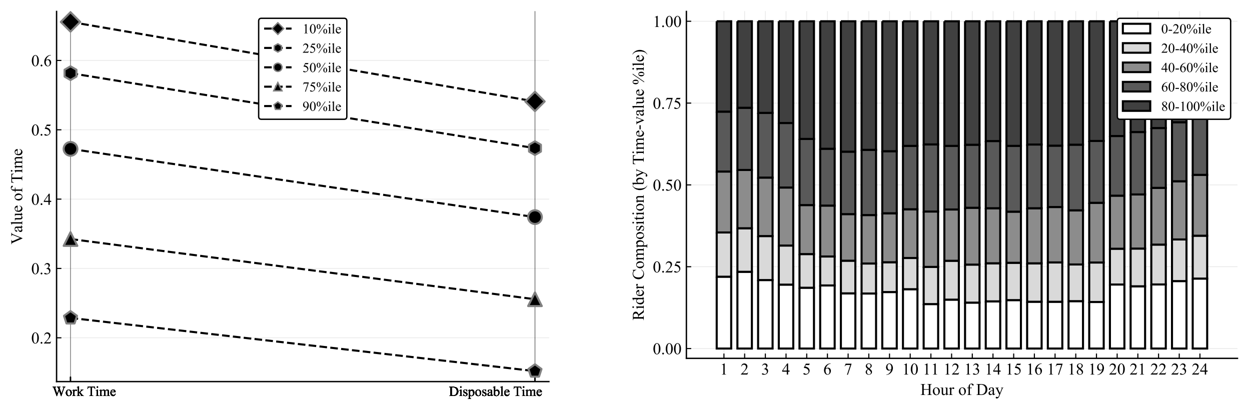

(a) Panel I: Within-person Shifts (b) Panel II: Hourly Composition of Riders

Figure 12: Intra-daily Patterns: Preference Shifts vs. Selection

6 Counterfactuals

We first ask how consumer welfare is impacted by three centralized matching mechanisms. Instead of permitting customers to make a choice from the menu of bids, each of our counterfactual mechanisms will instead assign a choice, taking the menu actually observed in the data as given. The first mechanism assigns the taxi with the lowest available waiting time regardless of price and other attributes. The second mechanism similarly assigns the taxi with the lowest available price. The third mechanism assigns the one with highest mean utility given our demand side estimates, without accounting for the individual's unobserved type (ascribed by β_i^w). To compute the counterfactuals, we take all observed orders in which there was more than a single choice offered. Using parameter estimates, we simulate outcomes under each mechanism by re-assigning each set of bids customers, assigning a trip from this set, and allowing consumers to take-it-or-leave-it by allowing them to choose the outside option. We aim to both measure the effect of removing choice on total trips and welfare, as well as to understand how market participation is influenced given preference heterogeneity in the population.

Our counterfactuals currently do not consider whether all allocations are also dynamically feasible. To give an example how this might matter: in a counterfactual we might match a passenger with a cab that, in reality, went on to serve a different passenger a moment after the request. If in our counterfactual that subsequent passenger is still assigned to the same cab, we violate the resource constraint that is borne out of the dynamic supply problem of rides. The demand model that we estimate here will be one of the building blocks for solving this dynamic supply problem.

Table 3 shows summary results for all trips in which more than one choice was present. It shows that imposing minimum wait time or prices can reduce market participation by up to 44 percent and reduce total surplus by up to 48 percent. Participation drops when consumers who would have selected a trip under the choice mechanism were instead assigned a different trip, for which the outside option was more attractive. These effects are driven by trips in which competing choices are actually present. Limiting results to these trade-off relevant trips alone (as in Panel 2) shows that it is exactly these choice sets that drive the overall effects. Panels 3 and 4 repeat the exercise, only by setting individual heterogeneity to zero. These results show that the counterfactual effects would look more pronounced when ignoring individual heterogeneity.

Notably, the "Best Observable Mechanism," which assigns the best trip on the basis of population mean preferences, reveals outcomes that are very close to the those of the choice mechanism. The trip quantity and welfare losses are around 1%. This suggests that assignment may be an attractive option for the platform if any components of the choice mechanism impose additional costs (e.g., cognitive costs of selection, uncertainty, confusion). However, the value of the selection on unobservables (not accounted for in the displayed results) would potentially make the choice mechanism much more valuable still.

Table 4 shows the benefit, in terms of both trip volume and per-consumer welfare, of moving from each counterfactual to the current choice mechanism. Column 1 shows the basic result that the regime in which minimum waiting time is automatically assigned is worse for consumers, as the choice mechanism delivers 29% more trips and 36% more welfare. This is corroborated by the relatively low waiting elasticities compared with price elasticities shown in Table 2; most consumers would not prefer to wait as little as possible and pay the higher prices associated with that extreme choice. In contrast, moving from a minimum price mechanism to a choice mechanism yields about 5% more trips and consumer welfare, suggesting this mechanism is not far off from matching customers to drivers that they would have chosen. Finally, the best-observable mechanism, which matches customers to the "best" option as defined by the mean utility estimates, but ignoring persistent unobservable heterogeneity, can get very close, to within 1% of the choice mechanism in terms of both trips and welfare. This outcome shows that prediction based matching could be implemented with little loss, allowing the platform to then leverage more control over matches, dynamically or otherwise, to achieve greater overall efficiency.

6.1 Counterfactuals to be done

- Conduct field experiments to guide and constrain the model below.
- Dynamic mechanism design: designing the platform to provide dynamic incentives to drivers (link current payoff to continuation values)
- Analyze the dynamically optimal centralized mechanism: if platform sets all prices and makes matching assignments, could it induce better outcomes than the decentralized market.
  - Evaluate the inefficiency due to the 10% fee structure which currently operates essentially as a (per cent) tax on drivers.
  - How to best match supply and demand taking into account both prices and waiting times: what observables would be most valuable to collect
  - Allow for cross-subsidies in the two-sided price

Matching Regime	Cons. Surplus	Trip Profit	Total Trips	Total Surplus
1. All Orders				
Choice Mechanism	\$1.38	\$7.95	96831 (1.0)	\$0.9 M (1.0)
Min. Wait Mechanism	\$1.29	\$5.87	75226 (0.78)	\$0.67 M (0.74)
Min. Price Mechanism	\$1.38	\$7.55	92265 (0.95)	\$0.86 M (0.95)
Best Observable Mechanism	\$1.37	\$7.94	96215 (0.99)	\$0.9 M (1.0)
2. Orders with price vs. waiting tradeoffs				
Choice Mechanism	\$1.37	\$8.62	48552 (1.0)	\$0.48 M (1.0)
Min. Wait Mechanism	\$1.1	\$8.07	27215 (0.56)	\$0.25 M (0.52)
Min. Price Mechanism	\$1.36	\$8.63	44254 (0.91)	\$0.44 M (0.91)
Best Observable Mechanism	\$1.35	\$8.62	48015 (0.99)	\$0.48 M (0.99)
3. All Orders (Setting all $\beta_i^w = 0$)				
Choice Mechanism	\$1.08	\$7.8	106996 (1.0)	\$0.95 M (1.0)
Min. Wait Mechanism	\$1.04	\$7.35	79043 (0.74)	\$0.66 M (0.69)
Min. Price Mechanism	\$1.06	\$7.76	101401 (0.95)	\$0.89 M (0.94)
Best Observable Mechanism	\$1.08	\$7.8	106996 (1.0)	\$0.95 M (1.0)
4. Orders with price vs. waiting tradeoffs (Setting all $\beta_i^w = 0$)				
Choice Mechanism	\$1.05	\$8.43	56103 (1.0)	\$0.53 M (1.0)
Min. Wait Mechanism	\$0.9	\$7.76	28449 (0.51)	\$0.25 M (0.47)
Min. Price Mechanism	\$1.0	\$8.4	50807 (0.91)	\$0.48 M (0.9)
Best Observable Mechanism	\$1.05	\$8.43	56103 (1.0)	\$0.53 M (1.0)

Table 3: Counterfactual Summary

This table shows the counterfactual welfare results and trip counts under each matching mechanism. Total surplus sums individual consumer surplus as well as driver revenues across all completed trips. Panel 1 accounts for all trips in which there was more than 1 choice. Panel 2 accounts for all trips in which there was a tradeoff present among choices, one with a lowest waiting time and another with a lowest price. Panels 3 and 4 duplicate the first two, but using estimates which do not incorporate any heterogeneity, such that $\beta_i^w = 0$ for all i. Consumer surplus and driver profits are computed as the total consumer utility in dollars divided by the total number of orders.

Counterfactual	**All Consumers**	**Top 25% β_i^w**	**Bottom 25% β_i^w**	**No Heterogen.**
1. Trips by consumer				
Min. Wait Mechanism	28.72%	22.82%	31.44%	35.36%
Min. Price Mechanism	4.95%	1.14%	17.39%	5.52%
Best Observable Mechanism	0.64%	0.7%	2.34%	0.0%
2. Welfare by consumer				
Min. Wait Mechanism	35.82%	30.73%	36.96%	43.4%
Min. Price Mechanism	5.31%	1.69%	17.69%	6.34%
Best Observable Mechanism	0.82%	0.94%	2.86%	0.0%

Table 4: Trip and Welfare Gains from Choice Mechanism

This table shows the gains in total trips and welfare achieved by the consumer-choice mechanism compared with each counterfactual mechanism. Panel 1 shows that total number of trips gained over each counterfactual, and Panel 2 shows the consumer welfare gained. Columns 2 and 3 display the results among the top and bottom quartiles of the distribution of β^w_{i,h_t}, so that the top 25% are the least sensitive to waiting and the bottom 25% are the most sensitive to waiting.

7 Conclusion

In this project we exploit the unique features of a large European ride share platform to measure the time costs of a large population of riders. Measures of such costs are of great interest to both macroeconomists and labor economists. The trade-off between time and pecuniary resources is at the heart of many important economic decisions, and yet it is hard to directly measure time costs because those decisions are often empirically confounded by other attributes. In this paper we observe such a trade-off directly (and without confounding factors) and leverage it to estimate preferences over time and money directly. We find that time is valued between \$0.05 per-minute during the night to \$0.40 during the day. While a mechanism where passengers are making a choice from a menu generated via an auction with drivers submitting bids generates a trade-off between higher consumer surplus due to more choices to choose from and higher information processing requirements - since the passenger has to evaluate the choice itself. We evaluate this trade-off using a simple counterfactual exercise where each passenger, rather than being presented with a choice, is matched with a particular driver based on observable characteristics (i.e., with a choice that maximizes the mean utility of an average passenger). Our results suggest that such a mechanism gets to within 1% of surplus from the choice-based mechanism.

In future iterations of this study we will be able to correlate individual measures with results of a survey that we are preparing together with Liftago. We are currently working on increasing the dimensionality of individual specific betas which will allow us to correlate WTP for time during core work hours with WTP during more disposable time and other passenger characteristics. One of the goals of this projet is to feed into a paper that focuses on the mechanism design question itself: how to match supply and demand in order to maximize welfare in the context of dynamic spatial competition. One important aspect of this analysis is to provide the platform with tools for how to dynamically incentivize the drivers: link their current performance (such as accepting short and seemingly "low-value" rides) to continuation values by promising more lucrative rides in the future. While this is theoretically perhaps not too complicated an exercise - for any given objective function - it is much harder when real-world constraints enter into the picture.

Bibliography

Abrantes, Pedro AL and Mark R Wardman, "Meta-analysis of UK values of travel time: An update," *Transportation Research Part A: Policy and Practice*, 2011, *45* (1), 1–17.

Aguiar, Mark and Erik Hurst, "Life-cycle prices and production," *American Economic Review*, 2007, *97* (5), 1533–1559.

Becker, Gary S, "A Theory of the Allocation of Time," *The Economic Journal*, 1965, pp. 493–517.

Belenky, Peter, "Revised departmental guidance on valuation of travel time in economic analysis," *US Department of Transportation. Washington, DC*, 2011.

Borjesson, Maria, Mogens Fosgerau, and Staffan Algers, "On the income elasticity of the value of travel time," *Transportation Research Part A: Policy and Practice*, 2012, *46* (2), 368 – 377.

Buchholz, Nicholas, "Spatial Equilibrium, Search Frictions and Dynamic Efficiency in the Taxi Industry," 2018.

Chen, M Keith, Judith A Chevalier, Peter E Rossi, and Emily Oehlsen, "The value of flexible work: Evidence from uber drivers," Technical Report, National Bureau of Economic Research 2017.

Cirillo, Cinzia and Kay W Axhausen, "Evidence on the distribution of values of travel time savings from a six-week diary," *Transportation Research Part A: policy and practice*, 2006, *40* (5), 444–457.

Cook, Cody, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer, "The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers," Technical Report 2018.

Fosgerau, Mogens, Katrine Hjorth, and Stéphanie Vincent Lyk-Jensen, "The Danish Value of Time Study," Technical Report, Danmarks Transportforskning, Report 5, 2007.

Frechette, Guillaume, Alessandro Lizzeri, and Tobias Salz, "Frictions in a Competitive, Regulated Market: Evidence from Taxis," Technical Report 2018.

Jara-Daz, Sergio R., Marcela A. Munizaga, Paulina Greeven, Reinaldo Guerra, and Kay Axhausen, "Estimating the value of leisure from a time allocation model," *Transportation Research Part B: Methodological*, 2008, *42* (10), 946 – 957.

Lam, Terence C. and Kenneth A. Small, "The value of time and reliability: measurement from a value pricing experiment," *Transportation Research Part E: Logistics and Transportation Review*, 2001, *37* (2), 231 – 251. Advances in the Valuation of Travel Time Savings.

Mas, Alexandre and Amanda Pallais, "Valuing alternative work arrangements," *American Economic Review*, 2017, *107* (12), 3722–59.

Miller, Ted R, *The value of time and the benefit of time saving*, Urban Institute, 1989.

Nevo, Aviv and Arlene Wong, "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession," Technical Report, National Bureau of Economic Research 2015.

Research, John Bates Services Significance Quantitative, "Values of time and reliability in passenger and freight transport in The Netherlands," Technical Report 2007.

Small, Kenneth A., "Valuation of travel time," *Economics of Transportation*, 2012, *1* (1), 2 – 14.

Tanner, Martin A. and Wing Hung Wong, "The Calculation of Posterior Distributions by Data Augmentation," *Journal of the American Statistical Association*, 1987, *82* (398), 528–540.

Train, Kenneth E, *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009.