

DISCUSSION PAPER SERIES

DP14666

THE VALUE OF TIME: EVIDENCE FROM AUCTIONED CAB RIDES

Nicholas Buchholz, Laura Doval, Jakub Kastl, Filip
Matejka and Tobias Salz

**INDUSTRIAL ORGANIZATION
INTERNATIONAL TRADE AND REGIONAL ECONOMICS
LABOUR ECONOMICS**



THE VALUE OF TIME: EVIDENCE FROM AUCTIONED CAB RIDES

Nicholas Buchholz, Laura Doval, Jakub Kastl, Filip Matejka and Tobias Salz

Discussion Paper DP14666

Published 27 April 2020

Submitted 24 April 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization
- International Trade and Regional Economics
- Labour Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Nicholas Buchholz, Laura Doval, Jakub Kastl, Filip Matejka and Tobias Salz

THE VALUE OF TIME: EVIDENCE FROM AUCTIONED CAB RIDES

Abstract

We estimate valuations of time using detailed consumer choice data from a large European ride hail platform, where drivers bid on trips and consumers choose between a set of potential rides with different prices and waiting times. We estimate consumer demand as a function of prices and waiting times. While demand is responsive to both, price elasticities are on average four times higher than waiting-time elasticities. We show how these estimates can be mapped into values of time that vary by place, person, and time of day. Regarding variation within a day, the value of time during non-work hours is 16% lower than during work hours. Regarding the spatial dimension, our value of time measures are highly correlated both with real estate prices and urban GPS travel flows. We apply our measures to quantify the opportunity cost of traffic congestion in Prague, which we estimate at \$483,000 per day.

JEL Classification: N/A

Keywords: N/A

Nicholas Buchholz - nbuchholz@princeton.edu
Princeton University

Laura Doval - laura@laura-doval.com
Caltech

Jakub Kastl - jkastl@princeton.edu
Princeton University and CEPR

Filip Matejka - filip.matejka@cerge-ei.cz
CERGE-EI and CEPR

Tobias Salz - tsalz@mit.edu
MIT

Acknowledgements

We thank Liftago for providing the data. We thank Stephen Redding and seminar participants at FTC, Maryland, MIT, Penn, Princeton, Stanford and UCL for useful comments. We are grateful for the financial support from the Transportation Economics in the 21st Century Initiative of the NBER and U.S. Department of Transportation. Kastl is grateful for the financial support of the NSF (SES-1352305). All remaining errors are ours.

The Value of Time: Evidence From Auctioned Cab Rides*

Nicholas Buchholz[†] Laura Doval[‡] Jakub Kastl[§]
Filip Matejka[¶] Tobias Salz^{||}

April 23, 2020

Abstract

We estimate valuations of time using detailed consumer choice data from a large European ride hail platform, where drivers bid on trips and consumers choose between a set of potential rides with different prices and waiting times. We estimate consumer demand as a function of prices and waiting times. While demand is responsive to both, price elasticities are on average four times higher than waiting-time elasticities. We show how these estimates can be mapped into values of time that vary by place, person, and time of day. Regarding variation within a day, the value of time during non-work hours is 16% lower than during work hours. Regarding the spatial dimension, our value of time measures are highly correlated both with real estate prices and urban GPS travel flows. A variance decomposition reveals that most of the substantial heterogeneity in the value of time is explained by individual differences as opposed to place or time of day. In contrast with other studies that focus on long run choices we do not find evidence of spatial sorting. We apply our measures to quantify the opportunity cost of traffic congestion in Prague, which we estimate at \$483,000 per day.

Keywords: Value of time, demand in transportation markets, ride hail.

JEL classification: C73; D83; L90; R12

*We thank Lyftago for providing the data. We thank Stephen Redding and seminar participants at FTC, Maryland, MIT, Penn, Princeton, Stanford and UCL for useful comments. We are grateful for the financial support from the Transportation Economics in the 21st Century Initiative of the NBER and U.S. Department of Transportation. Kastl is grateful for the financial support of the NSF (SES-1352305). All remaining errors are ours.

[†]Princeton University, 20 Washington Rd, Princeton, NJ 08540

[‡]California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125

[§]Princeton University, 20 Washington Rd, Princeton, NJ 08540

[¶]CERGE-EI, Politických vizòù 7, 111 21 Prague 1.

^{||}Massachusetts Institute of Technology, 50 Memorial Dr, Cambridge, MA 02142

1 Introduction

Allocating time is an important aspect of many economic decisions. Since [Becker \(1965\)](#), economists have used the labor-leisure trade-off to measure people’s value of time. Credible measures of the value of time are useful for a variety of important policy decisions. For instance, they are a critical input to infrastructure planning, particularly in the transportation sector. A unique empirical challenge in measuring the value of time is that in many economic settings — unlike in labor markets — time is not always directly priced. In this paper we use data from auctioned cab rides to overcome this empirical challenge, allowing us to estimate how the value of time is distributed across individuals, across places, and across time-of-day within a large urban area.

Transportation markets are generally revealing of the value of time: consider, for instance, the various decisions involved in the daily commute from home to work ([Domencich and McFadden, 1975](#)). In this context they might trade off cheaper-yet-slower off-peak trains against more expensive express trains; choose to leave earlier than their work requires to beat the rush hour; or choose between different service tiers in a ride-hailing platform, which involve different waiting times and prices. The choices made when faced with such tradeoffs are therefore informative about the value of time, or the opportunity cost of time spent traveling.

We use detailed consumer choice data from Liftago, a large European ride-hailing application. This platform uses a unique mechanism to allocate each ride through a rapid auction process in which nearby drivers bid on ride requests and requesting consumers choose between bids based on various characteristics. Most importantly, bids often involve tradeoffs between price and waiting time, or the time it would take the taxi to pick up the customer. Contrast this with platforms like Uber and Lyft that employ “surge” pricing to equilibrate demand and supply so that consumers do not get to directly express their preferences over prices and waiting times within the platform. We are able to observe both consumers’ individual choice sets as well as their ultimate selection for 1.9 million ride requests and 5.2 million bids.

The *first* contribution of this paper is to provide a direct and clean measurement of consumers’ willingness-to-pay to reduce waiting times. We use the variation in choice sets and choices to estimate a demand system that depends both on prices *and* waiting times. Such measures are of first-order importance for the provision of public transportation infrastructure as well as for the ride hail industry where price and waiting time are the two key variables on which firms compete. Our setting allows us to overcome some of the empirical challenges in measuring preferences over both prices and waiting-time.

Our *second* contribution, building on the work of [Small \(1982\)](#), is to provide a conceptual framework to interpret the disutility of waiting and to demonstrate how the willingness-to-pay for waiting-time reductions can be used to recover the value of time. When consumers choose a shorter wait time over a lower price, they reveal that the value of their time at a particular destination and time-of-day is greater than the value at the original location. Intuitively, the willingness to pay for lower wait times is simply the difference between the *value of time* at the destination and the *value of time* at the origin. In keeping with this interpretation we will refer to the willingness to pay for waiting time reductions as the *net value of time* (abbreviated as NVOT). We formalize this mapping, demonstrate that the spatial distribution of the value of time is identified and show how it can be recovered with a simple moment-based estimator. Our framework, therefore, leverages high-resolution choices in transportation markets to recover a summary measure of the value of economic activity across place, individuals, and time of day. The example of satellite data on night lights demonstrates that alternate measures of economic activity have often sparked fruitful new lines of research and filled in gaps where other data sources are not available or deficient ([Henderson et al., 2012](#)). The types of data required for our framework is increasingly made available by platforms in the transportation sector. Although our approach relies in part on unique features of our platform, we believe that it could be easily adapted to conduct a similar study within the US context, for example, with data from Uber or Lyft. Our measurement is complementary to other long-run measures of spatial economic activity such as real estate prices.

The demand results show that consumers respond substantially to changes in both price and waiting time. We find that the consumers average NVOT is \$10.80, meaning that on an hourly basis they would be willing to pay that much to reduce wait times. Price elasticities are about three times higher than waiting-time elasticities, however, there is large variation in these measures within the day and across space. This heterogeneity underscores the importance of the context in which price and time tradeoffs are made. From the overnight hours to the mid-day, the willingness to pay for lower wait times approximately doubles. Geographic differences are estimated to have an order of magnitude difference from one extreme to another. Because our data includes panel identifiers of both passengers and drivers we are also able to credibly identify the individual specific heterogeneity in both the elasticities and the implied NVOT. We recover heterogeneity across passengers both in their utility of income and dis-utility to wait times. We rank individuals by their relative sensitivity to prices and wait times and find that the top quartile have NVOT measures about 3.5 times higher than the bottom quartile.

We use our estimates to investigate the sources of variation in the value of time and determine how much is driven by differences across people, locations, and times of day. Our approach is similar to a branch of the labor literature which decomposes differences in wage into firm and worker specific variation (Abowd et al. (1999)). This exercise provides a number of important insights. First, we find that almost eighty percent of the variation in the value of time is driven by differences across people. Once individual-specific variation is taken into account, differences across places play a relatively small role. We also find that people who express a higher vOT are not necessarily doing so for the same places. In fact, the relationship between high value of time people and high value of time places is slightly negative. This finding is interesting in light of recent evidence of positive sorting of high earners in the long run residential housing market (e.g., Bayer et al. (2007)). A small survey of passengers reveals that the estimated value of time during the starting time of a typical work day is very close to the mean wage in the survey sample. This finding provides evidence that our results may be credibly extrapolated to other contexts.

When people are traveling they forgo valuable time at origin or destination. We can therefore use our value of time measures to quantify the opportunity cost of traffic congestion. Building on this, we use our estimates to quantify the cost of congestion for the population of riders on the app and, under additional assumptions, for the entire city of Prague, which has 1.3 million residents. We find that the cost of traffic congestion, counting only work days, is about \$0.5 million per day and \$75 annually for each vehicle driver on the road. This provides, to our knowledge, the first city-wide estimates of the opportunity cost of congestion directly derived from observed choices in a market setting.

Related literature The paper contributes mainly to four strands of literature, which we describe below.

The *first* strand is the literature on the opportunity cost of time. While the standard labor-leisure choice model implies that an extra hour of leisure should be valued at the shadow wage, the literature starting from the seminal work of Becker (1965) recognizes that an agent's time is also an input to other non-market activities. Recent papers in this literature have used a variety of widely available micro-data to study the trade-off between market goods and time (see, for instance, Aguiar and Hurst (2007); Aguiar et al. (2012); Nevo and Wong (2019)). Though these studies utilize rich and comprehensive datasets of consumption behavior (for example, household scanner data), they are only able to measure the opportunity cost of time indirectly through other market transactions. A related question is to what extent workers value flexible work schedules.

Mas and Pallais (2017) investigate preferences for such flexibility in an experiment with call-center workers, Bloom et al. (2015) study work-from-home preferences and performance differences among workers in a large travel agency, and Chen et al. (2017) study the value of work hours flexibility among Uber drivers.¹ Our work contributes to this literature in a number of ways. First, we extend the analysis of the opportunity cost of time outside of the workplace and study how this value varies across individuals, space, and time. Second, our data allows us to directly measure consumers' opportunity costs of time and to disaggregate these measures along various dimensions.

The *second* strand is the emerging literature that utilizes high resolution spatial data to shed light on urban sorting and segregation (Athey et al., 2019; Davis et al., 2017; Couture et al., 2019; Kreindler and Miyauchi, 2019; Almagro and Dominguez-Iino, 2019).² Kreindler and Miyauchi (2019) use people's commuting flows to estimate a gravity model where consumers choose where to work, as a function of where they live and wage values. The model yields estimates of the relative desirability of locations. They then show that the implied values from the estimation correlate well with the empirical distribution of wages and night lights in the city. In contrast, since we directly observe how people trade-off waiting time and monetary savings when choosing to move from one location to another, we can obtain a direct measure of their value of spending time at a specific location. Our approach allows us to perform a number of important quantifications that would be hard to do based on GPS data alone. For instance, we demonstrate in our application how our measures can be used to quantify the cost of traffic congestion.

The *third* strand is the literature in transportation economics and industrial organization, dating to the pioneering work of Daniel McFadden (McFadden, 1974; Domencich and McFadden, 1975), on the value of travel time savings. While the rich spatial nature of our data allows us to link consumers' willingness to pay for reductions in waiting time to the value of time across locations, the studies in this strand measure the benefits of travel time savings through surveys or revealed preference analysis based on mode choice. 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. Couture et al. (2018) study the determinants of driving speed and the deadweight loss of travel, where hours in traffic are valued at half the average wage. In contrast, our study directly uses our VOT measure as

¹The taxi industry has long provided a laboratory for empirical work on flexible work hours. See, e.g., Camerer et al. (1997), Farber (2005), Farber (2008), Crawford and Meng (2011), Thakral and Tô (2017).

²While not using high resolution spatial data, the work of Su (2018) is also related to this strand. The author studies the causal link between the value of time and gentrification. He argues that the increase in the value of time of high-wage workers led them to seek living areas with shorter commuting times, which then leads to gentrification of former poor, but close to downtown, areas.

the opportunity cost of lost time due to congestion. [Bento et al. \(2020\)](#) use commuter tollway choices to infer consumers' urgency from their willingness-to-pay for travel time savings. [Hall \(2018\)](#) analyzes the benefits of choice over toll and non-toll lanes.³ Finally, a recent literature studies equilibrium outcomes resulting from transportation infrastructure. This includes studies that analyze the impact of new roads on driving behavior ([Duranton and Turner, 2011](#)), the impact of transit on urban development and spatial sorting ([Heblich et al., 2018](#)), the welfare effect of transportation improvements ([Allen and Arkolakis, 2019](#)), and the design of optimal transportation networks ([Fajgelbaum and Schaal, 2017](#)). Our study complements this literature by showing how transportation options and choices are related to the opportunity cost of time, an important component of the overall welfare of transit. We also provide quantification of the baseline costs associated with transit externalities in a specific setting.

Fourth, our estimates are relevant to the literature that studies taxi and ride hail markets. Recent papers analyze the supply side in ride-hail markets. [Buchholz \(2018\)](#) quantifies the impact of uniform pricing regulation and search frictions on the spatial allocation of drivers and passengers in the NYC taxi market; [Frechette et al. \(2018\)](#) assess the effect of entry restrictions and market thickness on efficiency in the NYC taxi market.⁴ In these papers, the demand for taxis is estimated either as a function of prices ([Buchholz \(2018\)](#)) or waiting times ([Frechette et al. \(2018\)](#)). [Cohen et al. \(2016\)](#) and [Castillo \(2019\)](#) estimate demand for rides on the Uber platform, and in particular [Castillo \(2019\)](#) also estimates the demand for both waiting-time and price. His paper has a different focus and quantifies the benefit of surge pricing whereas we provide a conceptual framework to link the disutility of waiting-time to the spatial distribution of the value of time. In terms of the setting, ours has the benefit that we observe a direct and salient choice that includes a price waiting-time trade-off and also a panel-structure for riders, which provides a convincing way to estimate the population heterogeneity.

³There are a host of other studies. 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-Diaz 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. [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. [Kreindler \(2018\)](#) evaluates the welfare effect of congestion pricing using both travel behavior data and a field experiment.

⁴In a similar vein, [Liu et al. \(2019\)](#) and [Hall et al. \(2019\)](#) study various aspects of the design of DiDi and Uber, respectively. These papers, however, do not estimate demand.

The rest of this paper proceeds as follows. [Section 2](#) describes the institutional setting and our data. [Section 3](#) describes the conceptual framework which motivates our analysis. [Section 4](#) lays out the model of consumer choice and the details of identification and estimation. We present the results of our estimation in [Section 5](#), counterfactuals in [Section 6](#) and conclude in [Section 7](#).

2 Setting and Data

2.1 The Platform: A Unique Approach to Matching and Price Discovery

Liftago is a ride-hail platform, which was founded in 2015 and services rides through licensed taxi drivers. In Prague, all taxis need to be operated by licensed drivers. Moreover, taxis need to be equipped with a separate meter, which captures the number of kilometers traveled in the “occupied” mode together with the billed amount.⁵ A licensed driver may find fares by searching for street-hail passengers or by choosing to participate in a dispatch service. Among dispatch options, there are traditional telephone-based dispatch services and, more recently, the app-based ride-hail platform that we study. Note that this regulatory environment is different from most U.S. municipalities, in which there is nearly free entry into the ride-hail market through firms such as Uber and Lyft. Uber has also been present in Prague since 2014, but its presence is not as large as in a typical US city of similar size, partially since it is still fighting several legal battles due to various licensing and taxation issues.⁶

Drivers pay a percentage fee for each ride that is booked through the platform. By tracking both the taxi’s GPS and the time of the trip, it provides an approximate fare both before the trip begins and after its completion. While the platform has a strong presence in the Czech Republic, it is still less well known in other countries. This has the advantage that few riders are tourists, making it easier to interpret our estimates in light of local economic quantities.⁷

Importantly, drivers and passengers are matched by a combination of a dispatch algorithm and an auction. Whenever a passenger requests a ride, the system looks for nearby available cars and sends requests to a certain number of them, typically four, to elicit an offer. A cab driver who receives a request observes the details of the trip —

⁵Licensing requires both a fee and an exam. Meters need to be certified every two years by a state agency. Each meter records the aggregate numbers of kilometers billed together with the revenues.

⁶Since the EU court’s decision from December 2017, Uber is viewed as a transportation company and hence its drivers need to be properly licensed.

⁷This is also reflected in the relatively small fraction of airport rides, which comprise about 2 percent of total trips.

the location of the passenger, the destination, passenger rating and payment via cash or credit. A taxi driver who is interested in performing the job submits a bid, which is chosen from a set of pre-programmed tariffs.⁸ A tariff consists of a flag fee, a per-minute waiting fee and a per-kilometer fee with a regulatory cap at CZK 36 (\approx \$1.41). The platform takes any tariff bids and combines them with a query to Waze, a real-time traffic mapping service and Google subsidiary, which provides estimated trip time and distance information. The tariff bids are then translated into a single expected price for a trip. The passenger then observes bids as final trip prices together with other bid-specific attributes: the waiting time until the taxi arrives (ETA), the make and model of the car, and the driver’s rating. Importantly, these non-price attributes are automatically attached to the bids; in the case of waiting time, the Waze query also determines the expected ETA of each bidder to reach the passenger. The passenger may select one of the bids, in which case the ride occurs, or else may decline all bids. When the ride is completed, the passenger pays the fare shown on the meter.⁹ Figure 14 shows the interface that riders see on the app before the request and after the request arrives.

A noteworthy consequence of Liftago’s mechanism is that it allows for variation in both prices and waiting times: a driver with a high ETA may submit a lower bid than a driver with a low ETA, and vice-versa. Contrast this with traditional taxi services, where prices are fixed so that the market clears through adjustments in waiting time, and with other ride-hail platforms, where prices are adjusted to keep waiting times stable.

2.2 Data

Our dataset covers 1.9 million trip requests and 1.1 million actual trips on the platform 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 identifier for each driver and passenger. There are 1,455 unique drivers and 113,916 unique passengers over the sample period.

For each ride request we complement the data with geospatial and public transportation data from Google Places and Transit Matrix APIs, based on the GPS addresses for each point of origin and destination in the Liftago data. The API data provide alternative public transit times and routes as well as a measure of over 90 types of businesses and

⁸When starting the ride, the driver selects a tariff among the options he has pre-programmed 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. Some drivers who specialize in Liftago trips have over 20 different tariffs. Note that neither the passengers nor the other drivers observe tariffs that were not chosen.

⁹The app and an email receipt from Liftago only display an estimated amount. However, if the actual amount diverges from the estimated amount, customers are encouraged to report the discrepancy. Drivers can be banned from the platform if they are found regularly overcharging.

Table 1: Bid, Order, and Daily Summary Statistics

Variable	P10	Mean	P90	S.D.
PANEL A: ORDER SUMMARY ($N = 1,874,409$)				
Price of Trip (USD, across-auction)	5.65	11.66	19.75	6.261
Wait Time (minutes, across-auction)	3.00	6.8	12.00	3.816
PANEL B: BID SUMMARY ($N = 5,229,724$)				
Number of Bids (within-auction)	1.00	2.79	4.00	1.087
Price of Trip (USD, within-auction)	5.04	9.85	16.16	1.059
Wait Time (minutes, within-auction)	3.50	6.85	11.25	1.856
PANEL C: DAILY SUMMARY ($N = 638$)				
Number of Requests	1961	2938.3	4098	956.67
Number Rides	1160	1786.2	2470	557.15
Number Drivers	411	499.23	585	97.372

NOTE: The table shows summary statistics at the bid level (Panel A), the auction level (Panel B) and for entire days (Panel C). P10 refers to the 10th percentile and P90 to the 90th percentile of the respective variable.

local amenities within both 100m and 1km radii around pickup and drop-off locations.¹⁰ Furthermore, we use hourly rainfall data to attach prevailing weather characteristics to each ride.¹¹ Finally, we use data on GPS-specific land values and land zoning types from GIS coded data.¹²

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.¹³

2.3 Preferences over Time and Money: Intra-daily Patterns

In this section we describe spatial and inter-temporal patterns in prices, waiting times, and choices. Those patterns show that there is large and interpretable heterogeneity in consumer choices on which our model estimates build. 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.


¹⁰For more information see https://developers.google.com/places/supported_types.

¹¹We obtain this from <https://www.noaa.gov/>.

¹²Those are available at <http://www.geoportalpraha.cz>.

¹³We discard auctions with more than four bids, representing only 0.33% of the sample.

Figure 1: Prices and Waiting Times by Hour and Day



(a) Average Prices

(b) Average Wait Times


Consumers in our data often face a non-trivial trade-off between price and waiting time when choosing among bids. A trade-off implies that there exist options such that one has a lower waiting-time but a higher price and vice-versa. Depending on the time of day about 58-70 percent of auctions involve a trade-off between waiting less and paying more.¹⁴

Figure 2 shows how consumers solve the trade-off between time and monetary costs at different times of the day. At all times of day, consumers are more likely to pick the minimum price option than the minimum waiting time option. The elasticities that we back out from our model are very much in line with this observation. Moreover, the magnitudes of these differences vary throughout the day. We see that during work hours, there is a significant dip in the likelihood to choose the lowest price option and an even larger and significant increase in the likelihood to choose the lowest wait option.¹⁵ This pattern can be attributed to some combination of preference heterogeneity across customers as well as within-customer heterogeneity across the day. Since we observe customer identifiers, our model leverages the variation across consumers in the timing of trips and the choices within trips.

¹⁴Figure 15 in Appendix A shows how the fraction of auctions with a trade-off varies by hour of day.

¹⁵Note that those two do not have to add up to one since a consumer might, for example, choose a driver with the highest rating and that driver offers neither the lowest price nor provides the lowest wait time.

Figure 2: Tradeoffs and Choices by Hour




NOTE: The graph shows the mean probability that a customer who faces trade-offs between price and waiting time chooses either the lowest price or lowest waiting time. 95% confidence intervals for each series are also shown.

We next show how choices differ by locations. Instead of using administrative boundaries to arrive at a partition into smaller locations we prefer a data driven approach. We employ a clustering approach using exact GPS locations of trip origins and destinations. The partitioning is done according to a simple *k-means* procedure on latitude and longitude with the number of locations set to $A = 30$. We chose the value $k = 30$ to balance modeling the richness of spatial preference heterogeneity against sacrifices due to sample size. The resulting map is shown on [Figure 17](#).

[Figure 3](#) compares choices over price and wait times by location. The figure shows, within each pickup location, the probability that customers choose the lowest price and/or wait time among all available bids, computed only within auctions where a trade-off between the choices is present. Locations are sorted by the probability of choosing the lowest price. Like [Figure 2](#), [Figure 3](#) demonstrates that consumers exhibit preferences for lower prices and waiting times. It shows that minimum prices are chosen about 2-3 times more often than minimum waiting times, but there is substantial heterogeneity across locations. Those differences will eventually allow us to infer the different values that riders assign to different locations. We will also decompose how much of this variation is coming from place-innate characteristics and how much is driven by differences across people who travel between different locations.

Figure 3: Choices and Bids by Pickup Location




NOTE: This figure shows the mean probability that a customer who faces trade-offs between price and waiting time chooses either the lowest price or lowest waiting time. The locations in each are sorted by the probability on price. 95% confidence intervals for each series are also shown.

3 Conceptual Framework

We now describe the conceptual framework that shows how the choices that we observe are related to the underlying value of time at different locations. A consumer’s day is characterized by an allocation of time to various activities in different locations. Those activities (e.g., leisure or family time at home versus production at work) have different productivities or different intensity of pleasure at different times. Comparing her options at any given point in time, a consumer thus decides whether to move to a different location and spend her time there. Moving between locations is costly, both in terms of money and time, and a consumer has a choice between various transportation options. Liftago’s auction mechanism allows us to observe a particularly clean set of decisions about these transitions.

Figure 4 illustrates how the decision to accept a particular bid affects the time allocation between the origin and the destination. A consumer requests a ride at time t to the destination. She receives two bids. Since both drivers are supposed to take the same optimal route, the time from the pickup to the destination is the same and given by Δ . The two bids differ in the estimated time of arrival of the driver. The second bid leads to a longer wait time w_2 . By accepting bid 1, the passenger decides to spend $w_2 - w_1$

Figure 4: Origin Destination Trade-Off in Waiting Time Choice



NOTE: This figure shows the trade-off between two rides and how the choice of longer wait implies a trade-off between where to spend time.

less time at the origin and $w_2 - w_1$ more time at the destination.

Assume that the value of time for one unit at the origin is given by VOT_a^o and the value of time at the destination is VOT_a^d . The subscripts denote an area $a \in \mathcal{A}$ and the superscripts d and o indicate whether this area serves as a destination or as an origin. Thus when comparing both trips, the consumer is comparing $\text{VOT}_a^o \cdot (w_2 - w_1)$ against $\text{VOT}_a^d \cdot (w_2 - w_1)$. Letting the difference between w_2 and w_1 be one minute, the willingness to pay for waiting time reductions, which we call the *net value of time*, can be expressed as:

$$\text{NVOT}_{a \rightarrow \hat{a}} = \text{VOT}_{\hat{a}}^d - \text{VOT}_a^o \quad (1)$$

Different people will assign different values to places. Going forward, all the objects that we describe above will have i -subscripts to reflect the fact that we recover a distribution of the *value of time* ($\text{VOT}_{\hat{a},t,i}^d$ and $\text{VOT}_{a,t,i}^o$) which will give rise to a distribution of the *net value of time* ($\text{NVOT}_{a \rightarrow \hat{a},t,i}$).¹⁶

Since our empirical setting entails choices similar to the one we describe above we are able to observe distributions of the *net value of time* but not directly the underlying value of time. In [Appendix C](#) we formally derive the conditions under which the distributions of $\text{VOT}_{a,t,i}^o$ and $\text{VOT}_{\hat{a},t,i}^d$ are identified from observed distributions of the *net value of time*. The result requires one location normalization either for the origin or destination. To achieve non-parametric identification the result relies on known deconvolution techniques and for parametric identification (normal distribution) on a straightforward rank condition.

¹⁶Our data would also allow us to estimate a non-linear relationship between the origin and destination value of time. For instance, arriving ten minutes late to an appointment could be more than twice as damaging than arriving five minutes late. We have experimented with such non-linear relationships but abandoned them for the conceptual clarity of a linear specification.

3.1 Model of Travel Choice

So far, $\text{NVOT}_{a \rightarrow \hat{a}}$ was treated as directly observed. This section describes how to recover $\text{NVOT}_{a \rightarrow \hat{a}}$ from a standard discrete choice framework. In the next section we describe how to decompose the NVOT into the location-specific VOTs.

There is a set of locations $\mathcal{A} = \{1, \dots, A\}$ indexed by a and a set of consumers $\mathcal{I} = \{1, \dots, I\}$, indexed by i . When presented with a menu of bids (or offers) for a ride between a and \hat{a} , the consumer makes a discrete choice between J options. Each alternative $j \in \mathcal{J} = \{1, \dots, J\}$ is characterized by a tuple consisting of price, wait time, route-characteristics such as distance, characteristics of the car (model, year, color), driver (ratings and name) and the stochastic part $\epsilon_{i,j,t}$. We capture observable trip differences with $\mathbf{x}_{i,j,t}$. We also have to account for an additional term, $\xi_{a,\hat{a},t}$, that captures unobserved conditions affecting demand on a particular route, such as big sporting events, holiday travel, etc. We discuss endogeneity concerns in the estimation section. With this setup, the indirect utility from option j can be written as:

$$u_{i,j,t} = \beta_{i,h_t,a,\hat{a}}^w \cdot w_{j,t} + \beta_{i,h_t}^p \cdot p_{j,t} + \beta_{i,h_t}^x \cdot \mathbf{x}_{i,j,t} + \xi_{a,\hat{a},t} + \epsilon_{i,j,t}, \quad (2)$$

where β_{i,h_t}^p and β_{i,h_t}^w reflect preferences over waiting time and price; the subscripts h_t , a , and \hat{a} in β^w indicate that we allow preferences over waiting time to vary with the hour of the day as well as by origin and destination. Finally, the coefficient β_{i,h_t}^x captures preferences for other ride- j -specific characteristics (distance, driver's rating, car type, etc.) as well as environmental conditions common to all j : hour-of-day, public transit availability, traffic speeds, trip distance and time, rainfall, origin and destination neighborhoods, and whether the order is placed on the street or in a building. Note that the latter set of variables allow us to richly condition on many determinants of the outside option.

We can then map the preference parameters in [Equation 2](#) into NVOTs for different locations and different times of day. These NVOT's are 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, on average, indifferent between paying or waiting more:

$$\beta_{i,h_t}^p \cdot p_{j,t} + \beta_{i,h_t}^w \cdot w_{j,t} = \beta_{i,h_t}^p \cdot p_{j',t} + \beta_{i,h_t}^w \cdot (w_{j,t} + 1) \quad (3)$$

This implies that a minute of time at destination \hat{a} relative to its value at the origin a

is valued as

$$\text{NVOT}_{a \rightarrow \hat{a}, h_t, i} = p_{j,t} - p_{j',t} = \frac{\beta_{i, h_t, a, \hat{a}}^w}{\beta_{i, h_t, a, \hat{a}}^p}. \quad (4)$$

Equation 4 demonstrates that individual estimates of the Net Value of Time can be recovered directly from the estimated demand model by taking a ratio of coefficients.

4 Estimation

This section discusses the details of the estimation, which involves two steps. Based on observing the individual choices over bids on the app, we first estimate a likelihood based model. To capture time- and location-specific heterogeneity in time values, we leverage the panel nature of our data to compute random coefficients on the waiting-time using an MCMC procedure. We then use the coefficient estimates from the first step to estimate the NVOT, as we outline above, and a moment-based estimator to recover the distribution $\text{VOT}_{\hat{a}}^d$ (*the value of time*) separately. For clarity in exposition and the interpretation of work and non-work hours, the analysis hereafter utilizes weekday data only.

4.1 Mixed Logit Discrete Choice Model

Under the assumption that $\epsilon_{i,j,t}$ are independently and identically distributed according to a Type I extreme value distribution, choosing the maximum among J alternatives with utilities given by Equation 2 reduces to the standard logit. The probability that an alternative j will be chosen depends on its relative mean utility and is given by:

$$l(w_{j,t}, p_{j,t}, \mathbf{x}_{i,j,t}; \theta) = \frac{\exp(\beta_{i, h_t, a, \hat{a}}^w \cdot w_{j,t} + \beta_{i, h_t, a, \hat{a}}^p \cdot p_{j,t} + \beta_{i, h_t}^x \cdot \mathbf{x}_{i,j,t})}{\exp(-\xi_{a, \hat{a}, t}) + \sum_j \exp(\beta_{i, h_t}^w \cdot w_{j,t} + \beta_{i, h_t, a, \hat{a}}^p \cdot p_{j,t} + \beta_{i, h_t}^x \cdot \mathbf{x}_{i,j,t})}, \quad (5)$$

Endogeneity concerns: For bid-specific attributes we the econometrician are on almost equal footing with consumers because we observe all relevant bid attributes up to the drivers' names and photos. These unobserved features are therefore part of $\epsilon_{i,j,t}$. However, because drivers might condition their bids on $\xi_{a, \hat{a}, t}$, bids could still be correlated with unobservable demand conditions and thereby bias the price coefficients. To deal with this concern, we exploit persistent differences in bids among different drivers.¹⁷ These differences might, for example, come from pre-programmed bid increments in the meter. However, a straightforward GMM implementation is not feasible

¹⁷This approach is similar to a literature that exploits different leniency standards of judges and known as the *judge design*. See Waldfoegel (1995) for the first such strategy.

since our model relies on individual choice data and is likelihood-based. This prevents us from using standard inversion techniques to isolate $\xi_{a,\hat{a},t}$ and directly instrument. To get around this we, instead, concentrate $\xi_{a,\hat{a},t}$ out using a control function approach (Petrin and Train, 2010). This step consists of a simple regression of trip prices on a set of driver fixed effects, from which we recover a residual which then enters the likelihood as a control. In Figure 16 in Section Appendix A we show the resulting distribution of fixed effects, which demonstrates that there is large and persistent variation in driver bids. The interquartile range is \$1.9 or 20% of the average fare and the range from the 10th to the 90th percentile is \$4 or 43% of the average fare. Without the control function we find smaller price and waiting time elasticities resulting in an overall mean NVOT which is 19% lower than with the control function.

Control variables: We specify h_t as follows: the price coefficient is allowed to vary across work hours (defined to be 9am-6pm) and non-work hours, and the waiting time coefficient is allowed to vary across five blocks of time. The omitted category is the midnight hour, 12am to 1am. The remaining blocks are 1am to 5am, 6am to 9am, 10am to 3pm, 4pm to 6pm, and 7pm to 12am. We also control for drivers' quality ratings, car type, traffic speed and the distance of the trip. Further, we add controls rain. In Equation 2, β_{i,h_t}^x captures preferences for other ride- j -specific characteristics, which are the distance, the driver's rating, the car type, and environmental conditions common to all j .

Outside option: We allow the outside option to vary spatially at the level of each specific order in the data. This includes controlling for the environmental conditions (such as weather), and whether or not the trip was ordered from within a building or outside on the street. We also incorporate detailed public transit data including the time-of-day-specific presence of public transit availability within walking distance between each individual order's origin and destinations points. These characteristics impact the value of the outside option available to consumers, which is earned by choosing no alternative from \mathcal{J} . We normalize the value of the outside option to zero during 12pm and 1am at a location without any nearby public transit option, when it is not raining and when the order is place inside a building. Note that this specification allows for spatial (across different origin-destination pairs) and time-of-day variation in the outside option.

Estimation details: The estimation exploits the panel structure of our data to capture the full heterogeneity in time values in a tractable way. In particular, we opt for a hierarchical Bayes mixed-logit model to obtain individual specific estimates for the disutility of waiting 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 follow techniques described in [Rossi and Allenby \(2003\)](#), [Rossi et al. \(2005\)](#) and [Train \(2009\)](#) and now describe the particular version of the Gibbs sampler that we construct.

We assume that the waiting time and price coefficients for each individual are additive in time of day, location, and an individual specific shifter that is normally distributed.

$$\beta_{i,h_t,a,\hat{a}}^w = \beta_i^w + \beta_a^w + \beta_{\hat{a}}^w + \beta_{h_t} \quad (6)$$

$$\beta_{i,h_t,a,\hat{a}}^p = \beta_i^p + \beta_a^w + \beta_{\hat{a}}^w + \beta_{h_t} \quad (7)$$

Suppose that $\beta_i \sim N(\mu, \Sigma)$, where $\beta_i = (\beta_i^w, \beta_i^p)$, the vector of coefficients that vary at the individual level. The matrix Σ denotes the variance of the individual specific components of the coefficients as well as their covariance. The covariance $\sigma_{w,p}$ tells us whether or not people who are more elastic to waiting times are also more elastic to price. We would expect this to be true as the utility of income should be related to opportunity costs. However, one can also imagine segments of the population who are wealthy and yet have an abundance of disposable time, such as individuals who may be well-off and retired.

$$\Sigma = \begin{pmatrix} \sigma_w & \sigma_{w,p} \\ \sigma_{w,p} & \sigma_p \end{pmatrix} \quad (8)$$

We assume that $\mu \sim N(\mu_0, \Sigma_0)$, where Σ_0 is a diffuse prior (unboundedly large variance). We assume that the hyper-parameters of the variance are Inverse-Wishart, $\Sigma_0 \sim IW(v_0, S_0)$. One can then iteratively update the β_i -coefficient vector, 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 $\bar{\mu}^l$ be the sample mean of coefficients of iteration l in the chain and S^l be the sample variance of the Inverse-Wishart.

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,

we have

$$K(\beta_i|\mu^l, \Sigma^l, \mathbf{y}_i) \propto \prod_i^{T_i} l(y_{ii}; \beta_i) \cdot \phi(\beta_i|\mu^l, \Sigma^l) \quad (9)$$

where \mathbf{y}_i is the vector of choices and covariates observed for passenger i with T_i observations and $l(y_{ii}; \beta_i)$ is the likelihood contribution of a particular choice.

1. Draw a new posterior mean μ^l for the distribution of coefficients from $N(\bar{\mu}^{l-1}, \frac{W}{N})$.
2. Draw Σ^l from $IW(K + N, S^l)$, where $S^l = \frac{K \cdot I + N \cdot S_1}{K + N}$, and $S_1 = \frac{1}{N} \cdot \sum_i^N (\beta_i^{l-1} - \bar{\mu}^{l-1}) \cdot (\beta_i^{l-1} - \bar{\mu}^{l-1})'$.
3. For each i , draw β_i^l , according to the Metropolis Hastings algorithm, with new proposal β_i^{pl} starting from β_i^{l-1} using density $\phi(\beta_i|\mu^l, \Sigma^l)$.

Such an MCMC procedure is known to be slow for a large dimensional parameter space. To avoid slow convergence we, therefore, first estimate the model without the random coefficients using standard maximum likelihood and then employ the above Gibbs sampler to obtain the distribution of random coefficients separately, starting from the maximum likelihood estimates.

4.2 Moment Estimator to Recover Time-Location Valuations

In this section we describe how we recover the place flow values of different locations from the $\text{NVOT}_{i,t}$ estimates. We use a simple moment estimator that decomposes NVOT measures into location flow values. [Theorem 1](#) in [Appendix C](#) provides the conditions for identification on which this estimator is based. [Equation 10](#) relates the known NVOT measure that we back out from our logit-model estimates to a specific parametrization of location flow values.

$$\text{NVOT}_{i,h_t,a \rightarrow \hat{a}} = \underbrace{\text{VOT}_{i,\hat{a},h_t}^d}_{\text{value of destination}} - \underbrace{\text{VOT}_{i,a,h_t}^o}_{\text{value of origin}} \quad (10)$$

We build the following moments, which are the empirical counterparts to the simple difference equation above:

$$g_{i,h,a,\hat{a}}(\theta) = \frac{1}{N_{i,h,a,\hat{a}}} \sum_{i,t} \text{NVOT}_{i,t} - (\text{VOT}_{i,\hat{a},h_t}^d - \text{VOT}_{i,a,h_t}^o) \quad (11)$$

The optimization is then given by the following linear program with with inequality constraints:¹⁸

$$\min_{\theta} g_{i,h,a,\hat{a}}(\theta)' \cdot \hat{\Omega} \cdot g_{k,h,a,\hat{a}}(\theta) \quad (12)$$

$$s.t. \text{vOT}_{i',a',h'}^o = K \text{ and } \text{vOT}_{i,\hat{a},h}^d \geq 0 \quad \text{vOT}_{i,a,h}^o \geq 0 \quad \forall a, i, h \quad (13)$$

We maintain our previous decomposition of time of day into six bins and there are thirty different locations. We further allow the place values to depend on whether the car is ordered from inside or outside the building and by whether it rains or not. Together with the thirty locations we have more than twenty-thousand time dependent origin-destination pairs within which we could observe each passenger. This curse of dimensionality means that we can not perform the above decomposition for all passengers in our data separately. We therefore sort passengers along the ratio β_i^p / β_i^w and group them in five percent bins of this ratio. This grouping preserves most of the rich variation in the value of time across people and gives us enough data per bin to perform the above decomposition.

5 Results

We first present the results from the logit-demand model and the implied waiting time and price elasticities. We then present the results on the value of time.

5.1 Logit Model Results

Table 2 shows the coefficients and standard errors that we obtain from the demand model. The dis-utility of money is higher and almost identical across working and non-working hours. The intra-daily coefficients on waiting time vary over time-of-day, increasing in absolute magnitude (more negative) into the mid-day peak and declining into the evening. In addition to the time dimension, an important component of our subsequent analysis is to what extent the willingness to pay for waiting time reductions scales both with the origin and destination. In addition, there is a large amount of origin- and destination-specific heterogeneity in utility. Due to the large number of coefficients we show those separately in Figure 19.

The outside option is chosen about 33% of the time. Several coefficients measure an interaction effect between waiting time and additional factors related to the outside option: public transit availability, whether the trip is ordered on the street or not, and the

¹⁸We estimate the model based on Equation 12 as a linear program in JuMP (Dunning et al. (2017)).

presence of rain. Since the latter two of these interact with waiting time their net effects are not immediately apparent. We compute marginal effects to see the impact: rainfall confers a 0.59% increase in choosing the outside option. On-street ordering leads to a 0.16% increase. Within the auction, consumers prefer drivers with higher ratings and better cars. They prefer taxis for longer distance trips.

Table 2: Model Coefficient Estimates

DESCRIPTION	COEFFICIENT	STD ERROR
PRICE 6PM-6AM	-1.183	0.002
PRICE 6AM-6PM	-1.175	0.003
WAITING TIME 1AM-5AM	-0.018	0.011
WAITING TIME 6AM-9AM	-0.081	0.011
WAITING TIME 10AM-3PM	-0.089	0.01
WAITING TIME 4PM-6PM	-0.062	0.01
WAITING TIME 7PM-11PM	-0.031	0.01
WAITING × ON-STREET ORDER	-0.036	0.002
WAITING × RAINING	0.015	0.004
WAITING TIME SQUARED	-0.006	0.0
DRIVER RATING	11.063	0.097
CAR: MID QUALITY	0.248	0.005
CAR: HIGH QUALITY	0.708	0.009
TRIP SPEED	-0.051	0.002
ALT. TRANSIT AVAILABLE	0.044	0.007
ORDER ON STREET	0.211	0.014
RAIN	-0.148	0.034
TRIP DISTANCE	5.65	0.071
WAITING × PICKUP LOCATION FE 1-30		
WAITING × DROPOFF LOCATION FE 1-30		
PICKUP LOCATION FE 1-30	Omitted - See Figure 19	
DROPOFF LOCATION FE 1-30		
HOUR FE		

NOTE: This table shows coefficient estimates and standard errors of the logit demand model for each consumer type. The final 120 rows are omitted for exposition. These parameter estimates comprise outside option shifters and waiting time preference interactions with each of 30 pickup and dropoff locations as defined in [Section D.1](#). The omitted results are instead depicted graphically in [Figure 19](#).

[Table 3](#) show the elasticities of price and waiting time, computed as the percent change in selecting the bid with respect to a percent change in price and waiting time, respectively.

The table shows a set of bid-level elasticities, which measure the competitiveness of alternative bids, as well as a set of order-level elasticities, which measure the competitiveness of the outside option. We see a general pattern that consumers are much more

price elastic than waiting-time elastic: price elasticities range from four to eleven times higher, with starker differences in the evening compared to the day time.

In column 2 we use the estimated joint distribution of random coefficients to categorize four types of individuals as high (H) and low (L) sensitivity to price and waiting times. High price-sensitivity individuals have *below* median β_i^p , meaning they experience the highest disutility from price. Likewise, high waiting-time-sensitivity individuals have *below* median β_i^w . Consumers have highly heterogeneous elasticities: between the two extreme groups both price and waiting time elasticities differ by about a factor of four. We estimate a modest positive correlation between the sensitivity to price and waiting times: more price-sensitive passengers are also more waiting time sensitive, and vice-versa.

These elasticity estimates convey that both price and waiting time are important factors in the consumer decision and that waiting time elasticities vary throughout the day in ways that reflect the varying value of work and non-work related tasks.¹⁹

Table 3: Estimated Elasticities

Time of Day	Individual Type	Bid Level Elasticities		Order Level Elasticities	
		PRICE	WAITING TIME	PRICE	WAITING TIME
Daytime 6am-6pm	Overall	-4.37	-1.01	-3.9	-0.89
	H Price, H Wait Sensitivity	-8.59	-1.81	-7.36	-1.53
	H Price, L Wait Sensitivity	-2.86	-0.85	-2.8	-0.76
	L Price, H Wait Sensitivity	-5.1	-1.04	-4.47	-0.96
	L Price, L Wait Sensitivity	-2.03	-0.52	-2.06	-0.51
Evening 6pm-6am	Overall	-5.49	-0.5	-4.9	-0.49
	H Price, H Wait Sensitivity	-8.72	-0.8	-7.48	-0.75
	H Price, L Wait Sensitivity	-3.4	-0.37	-3.43	-0.37
	L Price, H Wait Sensitivity	-6.16	-0.52	-5.39	-0.52
	L Price, L Wait Sensitivity	-2.49	-0.22	-2.63	-0.24

NOTE: This table shows the demand elasticity of price and waiting time across daytime and evening hours and individual type groupings. We distinguish as *high (H) price sensitivity* individuals who have below median values for β_i^p and *low (L) price sensitivity* individuals as those with above median values for β_i^p , and similarly for waiting time sensitivity. The first two columns show these elasticities among competing bids, reflecting the change in demand due to a 1% change in price or waiting time on a single bid. The second two columns show them with respect to choosing the outside option, reflecting a change in demand due to a 1% change in price or waiting time on *all* bids.

The estimated model fits the data well. In Appendix [Section E.3](#), we show quality of our model fit on both aggregate moments and specific choices.

¹⁹We can also decompose elasticities by trip origins and destinations as with [Table 9](#) and [Table 10](#). Broadly similar patterns between demand types are revealed, though we see that there are large differences in each elasticity measure from one location to another. In general, price elasticities are more variable than waiting time elasticities.

5.2 The Net Value of Time

We now present results for the willingness-to-pay for waiting time reductions implied by our estimates, which we refer to as the *net value of time* or NVOT, scaled to USD per hour. NVOT represents the difference between the value of time attainable at a *destination* from the value of time attainable at the trip *origin*, given the activities and features of each location at each time for each person. They are computed using the above coefficients together with Equation 4, where we account for all trip-specific and environmental factors that affect valuations of the trip and the outside option, such as public transit and rainfall.

Table 4 summarizes the NVOT results. The overall mean value is \$13.47 per hour, an average NVOT across all trips and individuals. The most prominent source of heterogeneity is between individuals. We again report four groups of individuals, those with above- and below-median random coefficient estimates on both price preferences and waiting time preference. The low price sensitivity and high waiting time sensitivity group exhibits NVOT nearly twice the overall average at \$23.39 per hour, while individuals with low sensitivity to price and high sensitivity to waiting time have average NVOT of \$5.00. All groups have similar time-of-day patterns, with the highest values in the morning hours between 6am and 9am. Going forward, to be able to extrapolate from our specific context, we will interpret our VOT measure during this time of day as the wage rate. Not only is this the highest VOT measure during the day it is also most plausibly the value assigned to work related activities. Supporting evidence for this interpretation comes from a small survey that we conducted among passengers together with Liftago. We find that that wage rate among these survey participants is \$15.23, very close to our measurement of the VOT during work hours, which is \$15.44.

The individual-specific parameters reveal stark differences in valuation, but there is also substantial heterogeneity by time-of-day as well as place. We see that late morning to mid-day hours have estimated values that are approximately 50% higher compared to evening and overnight hours. Finally, we divide Prague's spatial regions into *core* and *non-core*, where core regions are the locations including and adjacent to regions 11 and 20 as depicted in Figure 17. In Table 4 we label as *urban core trips* any trip that involves one of these regions as origin or destination. We find that NVOT in core regions are on average 41% higher than non-core trips.

One important application of these results is that they provide insight into the interpretability of geographic (GPS) time-use data, such as cell phone location data. In several studies, for example Kreindler and Miyauchi (2019), authors use this type of data and interpret aggregate location decisions as a measure of place-specific amenities

Table 4: Net Value Of Time Estimates

SUBSAMPLE	Net Value of Time (NVOT)					
	12a-6a	6a-9a	10a-2p	3p-6p	7p-12a	All Hours
ALL TYPES	12.56	16.38	16.63	13.91	10.96	13.47
H PRICE, H WAIT SENSITIVITY	14.44	17.29	17.46	15.53	13.52	15.03
H PRICE, L WAIT SENSITIVITY	4.34	6.76	7.08	5.19	3.82	5.0
L PRICE, H WAIT SENSITIVITY	22.51	26.21	26.44	23.96	20.01	23.39
L PRICE, L WAIT SENSITIVITY	8.99	12.41	12.91	9.85	2.33	10.02
URBAN CORE TRIPS	12.96	17.1	17.31	14.37	11.4	13.97
NON-CORE TRIPS	9.93	11.78	9.61	10.57	7.6	9.94


NOTE: This table shows NVOT estimates implied by the logit demand model. All estimates are presented in US dollars.

or productive capacity.

Our analysis is, to our knowledge, the first to pair highly granular spatial data, that records where people go within a city, with direct NVOT measures, the latter of which capture the relative attractiveness of locations as per our illustration in [Figure 4](#). A natural question is, therefore, how close binary measures of time spent are to our direct measures of the relative attractiveness of locations. To answer this question we correlate our WTP for waiting time reductions for going from a to \hat{a} with the fraction of all trips that go from a to \hat{a} . The two different measures line up well but not perfectly. [Figure 5](#) shows a scatter plot of the nine hundred different directional data points as well as the binscatter points on top. As the plot illustrates, the two are highly significantly correlated with a correlation coefficient of 0.56 and t -statistic of 20.46. We can further condition this analysis on different times of day. We find that at 0.43 this correlation is highest during the middle of the day, which falls into work hours (10am-3pm), and lowest during midnight at 0.38. Both of these correlations are significant at all conventional levels.

These results convey a number of important insights. First, travel flows appear to be a good but not perfect measure of time valuations. The variation in those correlations throughout the day gives further guidance on when those travel flows align better with place-specific time valuations. Time valuations that are based on revealed preferences in transportation markets can therefore serve as an additional important quantification of the relative attractiveness of locations. In the next section, we explore this idea further by estimating the model outlined in [Section 3](#), which micro-founds the expressed willingness to pay to reduce waiting times through place specific values of time.

Figure 5: Relationship between Travel Flows and the NVOT



NOTE: This graph shows the scatter (transparent round dots) and binscatter (white diamonds) relationship between the NVOT for an origin-destination pair and the respective traffic flow (as a fraction of total).

5.3 The Value of Time

Economists have devoted considerable attention to measuring the long run effect of place of residence on individual outcomes (Chetty et al., 2018; Couture et al., 2019). However, place of residence might be an imperfect measure of the extent to which people benefit from the resources that different places have to offer. For example, a study of “experienced segregation” based on GPS data reveals that time use across different spatial regions is less segregated than residential locations (Athey et al., 2019). Another example is Davis et al. (2017) who use Yelp data to show that restaurant consumption is only half as segregated as residences.

We complement this literature by providing a direct measure of the monetary value that people assign to spending time at different places in the short-run. Our estimates further allow us to study whether or not different strata of society value different places differently. We exploit the fact that our data contain a panel of riders in which we can observe the same rider making decisions about time allocations across many different places. If we found that most of the variation in the value of time is driven by places and not people, this would suggest that place-specific factors accrue equally to different people. For example, there might be differences in the provision of public goods across places, which are typically equally available to all people. On the other hand, differences in the value of time might be predominantly driven by differences across people. This would be true, for example, if wealthy people with productive jobs enjoy activities that are equally exclusive across different places.

To address these questions we first provide several summary measures of our esti-

mates and then a full variance decomposition that separates time-, place-, and individual-specific variation in the value of time.²⁰

An important consideration for our measurement is that time at a destination might be more valuable per-se than time at an origin. Such differences might arise from planning and other coordinated activities. For example, people may travel to a particular destination location to be productive and then later exit the location when their productive task is completed, so that the location becomes an “origin”. While we are unable to directly observe plans, our decomposition allows us to estimate the value of time at a location both for when it serves as an origin and as a destination. We are able to do this because we observe riders going in both directions and can recover both $NVOT_{a \rightarrow \hat{a}}$ and $NVOT_{\hat{a} \rightarrow a}$. With this data we can construct a reduced form measure that captures differences in productivity of time use across origin and destination.²¹ To this end, let us rewrite the relationship between the NVOT and VOT as follows:

$$NVOT_{a \rightarrow \hat{a}} = VOT_{\hat{a}}^d - VOT_a^o = VOT_{\hat{a}}^d - \frac{VOT_a^d}{VOT_a^d} \cdot VOT_a^o = VOT_{\hat{a}}^d - \delta_a \cdot VOT_a^d = VOT_{\hat{a}} - \delta_a \cdot VOT_a \quad (14)$$

This measure is denoted as δ_a and given by the ratio of $VOT_{\hat{a}}^o$ to $VOT_{\hat{a}}^d$. Following [Equation 14](#), we interpret VOT_a^d as the inherent place value and δ_a as a depreciation factor due to less (or more) productive time use at origin. From now on we drop superscripts o and d .


[Figure 6](#) shows the unconditional distribution of the value of time v_a (USD/hour) that we back out. The histogram reveals large variation. The interquartile range goes from \$6.8/h to \$16.3/h. The tenth percentile of the distribution is \$3.7/h and the ninetieth percentile is \$21.3/h.

How do our value of time estimates compare to other quantities that capture spatial differences in economic activity? To answer this question we correlate our VOT measures with property prices. A binscatter plot of the relationship is shown in [Figure 7](#). There is a strong positive relationship between real estate prices and our VOT measures. However, the “returns” to higher land values are diminishing. From a regression of log-VOT on log land-values we measure an elasticity of 0.25 ($p < 0.001$). So for every 1% higher land value we measure a value of time that is 0.25% higher. On the one hand, this strong positive relationship serves as a validation of our measures since we have

²⁰The place specific estimates should not necessarily be interpreted as deep structural parameters that are inherent to a place. They could capture certain agglomeration effects, such as the taste for meeting certain people who typically can be found at a given place.

²¹In the empirical specification we use additional data that is informative about the extent of planning. We obtained the exact location of the customer when a ride was ordered. This allows us to address how much a trip was planned in advance, since we can distinguish, for example, if a customer is standing outside of the building at the time of placing the order and hence is likely to have only limited value for her extra time at the origin.


Figure 6: Histogram VOT



NOTE: This graph shows the unconditional distribution of the values of time that we back out.

not used land values at any point in the estimation. On the other hand, we find that with an R^2 of less than 0.02 that real estate prices can only explain a very small percentage of the variation in VOT. This re-emphasizes the point that the value of time is a useful complementary measure of short run economic activity across space.

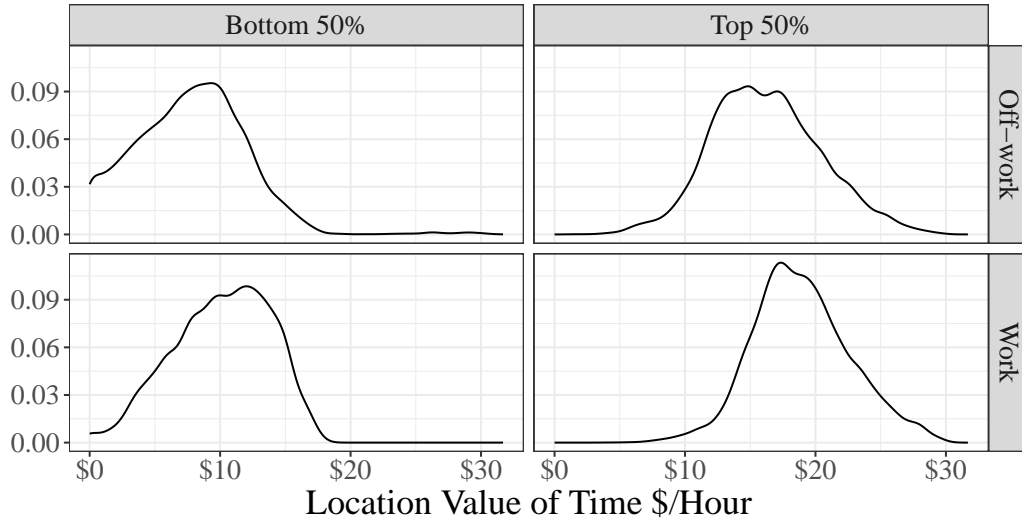
Figure 7: VOT by Group and Time



NOTE: The x-axis orders the thirty locations by their average value. The y-axis in the left panel shows the values of both types during work times and right panel during non-work times.

Figure 8 conditions the value of time estimates on work and non-work time as well as the bottom 50% and top 50% in terms of VOT. The histograms shows that both groups express a higher value of time during work hours.

Figure 8: vOT by Group and Time



NOTE: The x-axis orders the thirty locations by their average value. The y-axis in the left panel shows the values of both types during work times and right panel during non-work times.

Table 5: Summary Statistics of Value of Time

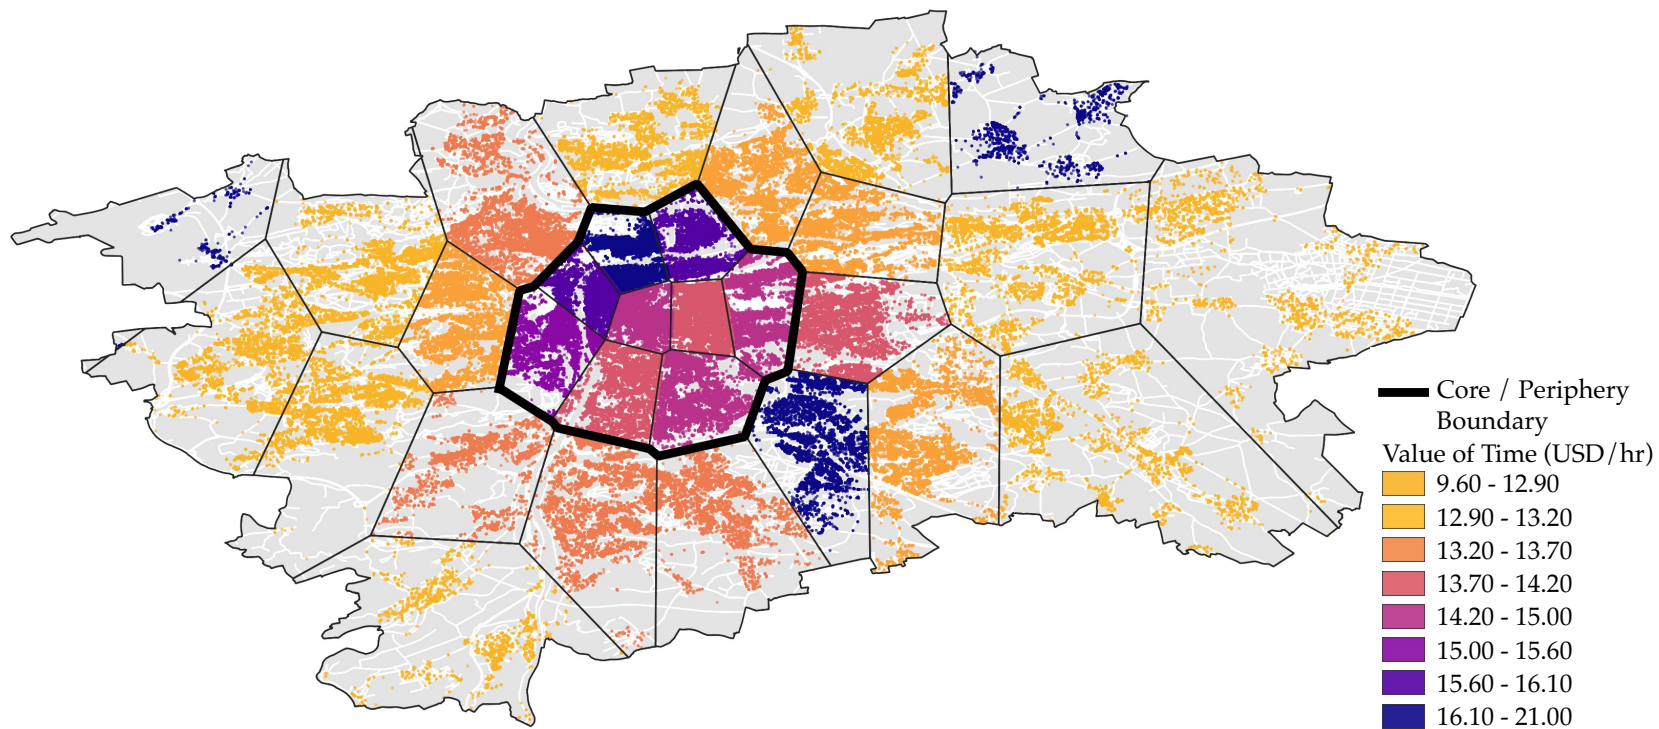
	Work Time (USD)		Non Work Time (USD)		Non Work Time vOT
	MEAN	STD	MEAN	STD	Work Time vOT (%)
Location Values (v_{i,a,h_t})					
All	15.44	9.71	13.18	10.11	0.87
vOT < median(vOT)	20.73	10.68	18.06	10.6	0.81
vOT ≥ median(vOT)	10.0	3.97	8.14	6.46	0.85

NOTE: This table shows summary statistics on the estimates that capture how much less productive time at the origin is relative to the destination separate by type. The last column shows the relative magnitude of the non work-time vOT over the work-time vOT in percent.

Table 5 gives an overview of means and standard deviations for the top and bottom 50% in the distribution of vOT. During both work and non-work time, the top 50% exhibit a vOT that is almost twice that of the bottom part of the population. Within type, however, the difference between work-time and non-work-time is modest. For the top 50% the value of time is 15% higher during work than during non work times. For the bottom 50% the discrepancy is larger, their value of time is 19% higher during work hours.

Figure 9 shows vOT heterogeneity by location. Going from the lowest-value location to the highest-value location implies a difference of \$11.58 in the hourly value of time. This corresponds to a 54.5% increase in percentage terms. However, the interquartile range is already substantially smaller with a difference of \$2.51, which in percentage terms is 16.5%. This suggests a relatively small contribution of places to the variance

Figure 9: Map of vOT Estimates in Prague



NOTE: This figure depicts each GPS destination point in the Liftago data shaded by the average value of time in that destination, or vOT_a^d across places $a \in \{1, \dots, 30\}$ within the city limits of Prague. White lines depict the city's street map. The bold black line delineates the boundaries which we specify as the urban core in presenting results.

in the value of time. However, both the location breakdown and the person-specific breakdown mask the variation from the respective other source. For example, to the extent that they are correlated one may mistake variation in the value of time in the spatial dimension for person-specific variation and vice versa. To separate out the respective sources of variation and investigate their relationship in more detail we now turn to a variance decomposition. A variance decomposition will allow us to investigate whether people with high average vOT spent more time in places with high average vOT. The extent of such sorting effects has been of long standing interest in the economics of marriage markets (Becker (1973)), labor markets, (Abowd et al., 1999; Eeckhout and Kircher, 2018; Shimer and Smith, 2000) and housing markets (Couture et al. (2019)). Our approach is similar to the labor literature that tries to decompose wages into firm specific components and worker specific components, most importantly Abowd et al. (1999). For the variance decomposition we derive an accounting equation from the following regression model:

$$\text{VOT}_{i,a,h_t} = \alpha_i + \eta_{h_t} + \gamma_a + \epsilon_{i,t}, \quad (15)$$

where α_i is a person fixed effect η_{h_t} is a time-period fixed effect, γ_a is a place fixed effect, and $\epsilon_{i,t}$ is a residual. Due to the curse of dimensionality, we run this specification by replacing individual values with the percentile means, so that i becomes an indicator of the percentile bin.²² This regression then gives rise to the following variance decomposition.

$$\begin{aligned} \text{Var}(\text{VOT}) = & \text{Var}(\alpha) + \text{Var}(\eta) + \text{Var}(\gamma) + \\ & 2 \cdot \text{Cov}(\alpha, \eta) + 2 \cdot \text{Cov}(\alpha, \gamma) + 2 \cdot \text{Cov}(\gamma, \eta) + \text{Var}(\epsilon) \end{aligned} \quad (16)$$

Table 6 shows the results from this exercise. At 78%, by far the largest contributor to the observed variance of the vOT are differences across people. In comparison, place-specific variance is small and accounts for 10% of the total. Intra-daily changes in the vOT contribute another 9%.

Moreover, the covariances show that high vOT people do not spend more time in high vOT places. With a covariance of -0.4 we measure a slightly negative sorting effect. Similarly, there is a negative correlation between people who express a high value of time the times of day with higher average vOT.²³

²²Our panel is not long enough to observe a given individual in all possible combinations of origin, destination and time of day.

²³We have also repeated this same exercise with the nvOT as opposed to vOT. Location fixed effects in this exercise are defined over directional pairs since the net value of time is defined over those. Under this alternative exercise, we come to a very similar conclusion. Most of the variation is driven by differences across people instead of differences across places. This should dissuade any concern that our results are driven by specific aspects of the decomposition that we perform.

Table 6: Variance Decomposition of the Value of Time

$Var(\alpha_i)$	$Var(\eta_{h_i})$	$Var(\gamma_a)$	$2 \cdot Cov(\alpha_i, \gamma_a)$	$2 \cdot Cov(\alpha_i, \gamma_{h_i})$	$2 \cdot Cov(\gamma_a, \eta_{h_i})$	$Var(\epsilon_{i,t})$
27.7	3.25	3.5	-0.4	-0.1	0.3	1.6
0.78	0.09	0.1	-0.01	-0.002	0.008	0.04

NOTE: The first row of this table shows the variances of each of the components of the model above. The second row shows the variance of each of the components divided by the total variance, which can loosely be interpreted as fractions of the total variance.

We next discuss our estimates of differentially effective time use at origin and destination, as measured by δ_{i,a,h_t} . Table 7 gives an overview. Across the entire population, unplanned time at the origin is only 31% as valuable as time at the destination during work hours while it is 36% as valuable during non-work time. Holding the location fixed, meetings and other types of work-required planning therefore lead to a larger discrepancy in the value of planned and unplanned time. Moreover, we find that there is a large difference in δ within the population. Not surprisingly, people that we have labeled as more time sensitive ($\beta_i^w > median(\beta^w)$) have a larger discrepancy in the value of planned and unplanned time. This shows that higher willingness to pay for lower waiting times is driven both by a larger inherent value of time at any given location and also a bigger discrepancy between planned and unplanned time.

Table 7: Summary Statistics for δ_{i,a,h_t}

	Work Time		Non Work Time	
	MEAN	STD	MEAN	STD
Location Values (vOT_{i,a,h_t}) in USD				
All	0.19	0.2	0.22	0.31
$\beta_i^w > median(\beta^w)$	0.12	0.1	0.11	0.1
$\beta_i^w \leq median(\beta^w)$	0.29	0.34	0.44	0.74

NOTE: This table shows summary statistics on the estimates that capture how much less productive time at the origin is relative to the destination separate by type.

To summarize our results, we find large differences in the value of time across individuals and also large differences in differentially productive time use across origin and destination. Most of the variation in the value of time comes from person-specific heterogeneity as opposed to place specific heterogeneity. Interestingly, this suggests that, in the short run, different places do not confer benefits that are equally valued by everyone. These results are an important short-run complement to the recent debate around the importance of place-specific factors for long-run outcomes. Another interesting implication lies in the finding that the value of time of high types can, in part,

be explained by a larger necessity for planning as measured by the ratio between origin and destination flow values for the same area. This suggests that high VOT types depend to a larger extent on complementary inputs with whom a need for coordination arises. For example, high skilled work places might require more meetings and coordination with others to use time productively.

6 Application: The Time Cost of Traffic Congestion

Traffic congestion is a growing phenomenon across the globe. The 2019 Urban Mobility report estimates that the average urban commuter in the United States spent an extra 54 hours of travel time on roads as a result of congestion.²⁴ INRIX reports a similar global measure taken across 220 cities among 38 countries to estimate that congestion increases traffic time on average by 108 hours in 2018, a 1% increase over the previous year.²⁵ This problem has the attention of regulators. Large metropolitan cities are experimenting with a variety of policy instruments including congestion pricing, variable tolling and stricter limits to parking. While congestion is a salient problem for commuters and businesses alike, the costs are normally difficult to quantify precisely because commuters' opportunity costs in traffic are hard to measure. In this section, we use our framework to quantify the opportunity cost of congestion in Prague.

Our approach has two ingredients. First, we use time and distance data from the 1.9 million trips in our sample to measure average traffic speed by time of day and route. We find the hour of day at which average traffic speeds are highest and denote this as the free-flowing or un-congested traffic speed. By comparing each observed trip time against the counterfactual trip time under un-congested conditions, we compute a measure of extra travel time due to congestion. Figure 10 panel (a) plots average excess travel times by hour across Prague. During peak times, trips take an extra 5-6 minutes or about 25% longer. In evening and overnight hours, congestion falls substantially as average traffic speeds become close to the free-flowing speeds.


The second ingredient is our measure of time cost. Our time of day and place-specific VOT estimates serve as these measures and can be interpreted as a summary of the value of activities at certain places throughout the day. When we estimate a positive WTP for transiting from one place to another, this reveals that the destinations must be higher valued than the origins *at that time*. Thus, the time cost of congestion will depend on whether congestion reduces time spent at a productive destination ($VOT_{a,h,i}$) or a less productive origin ($\delta_{a,h,i} \cdot VOT_{a,h,i}$). One feature of our value of time estimates is

²⁴See Schrank, Lomax and Eisele (2012).

²⁵See <https://inrix.com/scorecard/>.

that we can distinguish between anticipated and unanticipated congestion cost. Anticipated congestion costs, such as during an hour in which a road is typically congested, will induce travelers to leave early and thereby sacrifice time at the origin instead of the destination. Conversely, unanticipated congestion costs will cost travelers time at their destinations, which are most often higher valued. To compute congestion costs, we simply multiply the number of excess minutes in traffic by the value of time at origin or destination, depending on whether congestion is anticipated or unanticipated. We assume that the value of time spent in a vehicle is zero, consistent with our normalization of time spend waiting at an airport. Figure 10 panel (b) combines trip-specific excess travel times and their associated VOT measures to plot the cost of congestion. We find congestion cost patterns that mirror the daily commute. If this were completely anticipated by travelers, then an average rush-hour trip would impose around \$0.25 – \$0.30 of time cost as passengers would choose to replace less valuable time at origins with time in congested traffic. When delays are unanticipated, however, the average per-trip cost to travelers is between \$1.00 – \$1.30 during rush-hour.

Figure 10: Excess Time in Traffic and Estimated Congestion Costs



The above figures make clear that congestion costs are borne most heavily in the rush-hour periods, and, when traffic accidents and other shocks to commute time impose unanticipated delays, the costs can grow five-fold. We now explore congestion cost heterogeneity across space. Higher congestion costs in core sections of urban areas has led some municipalities to implement zone-based congestion pricing, as London did in 2003.²⁶ Figure 11 and Figure 12 separate trips by *urban core* or *non-core*. Urban core designates routes that begin or end in the high density center of Prague (Locations including and adjacent to 11 and 20 and in Figure 17).

²⁶Similar programs are being planned for New York City. See <http://www.mta.info/press-release/bridges-tunnels/mta-announces-selection-transcore-build-nation-leading-central>.

Figure 11: Excess Travel Time in Center and Periphery

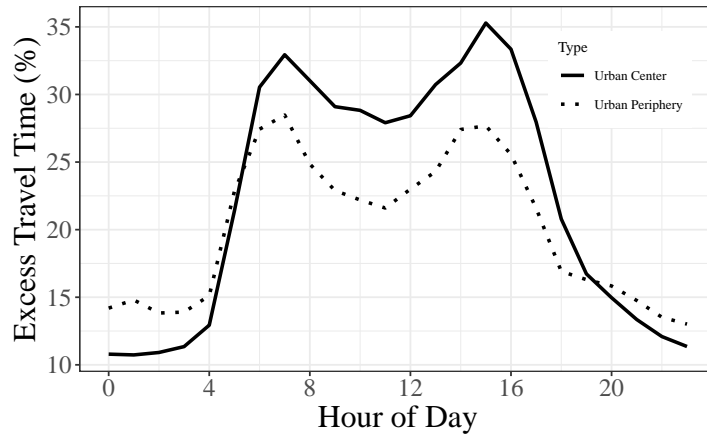



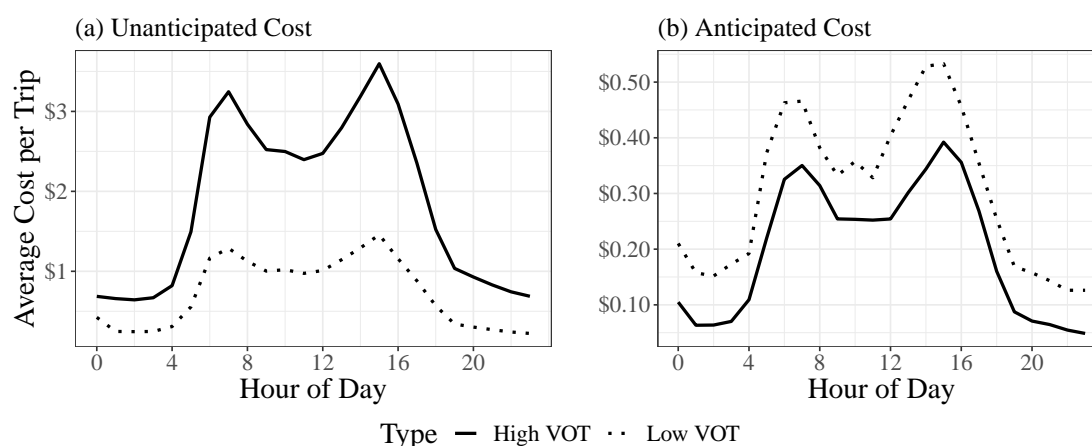
Figure 11 shows that, on average, the extra time spent driving relative to time in free-flowing traffic is larger by a factor of 2-3 throughout the day on routes connecting to the urban core. Trips in the city periphery exhibit similar time of day patterns, but with less extreme differences across night and day. In Figure 12 we show congestion costs across core and non-core routes. Unexpected congestion costs (Panel (a)) in the urban core are mostly higher during the day time, but evening congestion is more costly outside of the urban core. This is consistent with our finding that VOT at peripheral destinations, which are largely residential, are higher valued than core origins in the evening hours. This is because anticipated Panel (b) shows anticipated congestion costs in and out of the core over the day. Since peripheral residential areas outside of the urban core tend to have higher values as origin-locations than non-residential areas, we see that anticipated costs, which impose costs on origin time rather than destination time, are higher on average outside the core.

Figure 12: Planned and Unplanned Congestion Cost in Center and Periphery



Individual heterogeneity in VOT can induce disparate impacts of congestion. We denote high-VOT individuals, denoted as individuals with above-median WTP for waiting time reductions, as *time-sensitive* types. The remaining individuals are denoted *time-insensitive* types. Figure 12 Panel (b) shows that unanticipated congestion costs for time-sensitive individuals are slightly more than double that of time-insensitive individuals, with average rush-hour costs up to \$2.02 per trip. Anticipated congestion costs both types nearly the same amount, which suggests that travelers' time costs in unproductive states (i.e. periods in which anticipated time costs can be borne more cheaply) is much more equal than in productive states. These differences are relevant because many solutions to urban congestion involve pricing through tolls, zone-pricing, or automotive taxes. Nevertheless, this distinction suggests that price-based solutions to congestion will result in time-sensitive drivers and passengers choosing to pay while time-insensitive drivers exit the road.

Figure 13: Estimated Hourly Congestion Costs by Individual Differences



Total Cost of Congestion in Prague

Finally, we ask what our congestion cost measures imply for traffic in Prague as a whole. The 2016 Prague Transportation Yearbook uses survey data to estimate that 1, 646, 600 automobile trips occur during each workday.²⁷ This study also provides estimates of the volume of traffic across each workday hour, each day of week and each month as well as average vehicle occupancy rates.

To determine the total costs of congestion, we first note that the population within Prague from which our estimates are drawn (i.e., taxi passengers) is likely to be selected on high-VOT individuals compared with the average Prague commuter. In order

²⁷See <http://www.tsk-praha.cz/static/udi-rocenka-2016-en.pdf>.

to extrapolate from our hourly congestion cost estimates, we first adjust these costs to reflect opportunity cost of car commuters who are *not* on the Liftago platform. To do this we make the assumption that at 9:00am, a typical workday starting time, our VOT estimate is equal to users' average wage. This value is equal to \$13.78 per-hour. The mean wage among Prague residents, however, is \$9.15 per-hour. Using this, we specify an adjustment factor $\alpha = 9.15/13.78 = 0.67$ and scale all VOT results by α . A second assumption we make is that the spatial distribution of trips in each hour is the same between Liftago users and car commuters. Combining these assumptions allows us to infer, for example, that at 5pm when Liftago riders face an unanticipated congestion cost of \$1.50 per trip, the average car commuter would face a cost of \$1.00 per trip.

We can now combine our hourly congestion costs with the hourly distribution of all traffic provided by the Prague Transportation Yearbook. We assign the annual number of car trips in Prague to each hour given by the yearbook, and multiply the number of trips per hour by the α -scaled average congestion costs. We again scale these measures by the average occupancy rate in Prague of 1.30. Finally, we determine how much congestion is anticipated by computing the mean traffic speeds by day of week and hour and treated any trip in which actual trip time is at or below the time implied by this expectation as a trip faced with anticipated congestion. Any additional trip time is treated as unanticipated.

Table 8 summarizes the results. Panel A compares the distributions of Liftago rides and all Prague car trips and shows that most car trips occur in the same hours in which our speed data imply the greatest congestion, as shown in Panel B. The next two panels compare congestion costs that can be planned ahead (Panel C) with those that are unplanned (Panel D). The last row of panel D shows that between 11 and 15 percent of congestion time is unplanned. The last row, Panel E, combines the data above to estimate total congestion costs of \$482,955 during each workday in Prague. At 254 work days per year, this implies that Prague drivers face \$123 million in annual congestion costs or about \$75 per driver.²⁸

Our congestion cost estimates only incorporate the cost of time. We do not quantify the added cost of pollution, noise, accidents or other costs. This estimate also averages over potentially important heterogeneity; a driver who only commutes into the city core at 9:00am and leaves at 6:00pm on each workday would face average annual costs of about \$281 over 46.7 lost hours. If we were to restrict attention to high congestion routes only, this number would likely grow further.

The ability to draw insights from market behavior to quantifying the cost of congestion can be valuable for local policy makers and regulators in a variety of contexts. For

²⁸Statistics on work-days per quarter come from the European Commission's Eurostat data for the Czech Republic.

Table 8: Estimated Congestion Costs

	<i>Time-of-Day</i>								<i>Aggregate</i>
	3:00am	6:00am	9:00am	12:00pm	3:00pm	6:00pm	9:00pm	12:00am	Total
A. Trip Volume Density*									
Liftago	0.08	0.08	0.1	0.08	0.09	0.15	0.17	0.27	1.0
All Vehicles	0.02	0.02	0.18	0.18	0.18	0.21	0.15	0.06	1.0
B. Mean Excess Travel Time/Trip (min.)									
All	0.9	0.9	5.43	4.31	5.0	5.62	2.1	1.16	3.29
Core-Locations	0.77	0.77	5.35	4.33	4.95	5.69	2.01	1.06	3.22
Non-Core	1.95	1.95	5.89	4.24	5.27	5.14	2.73	1.98	3.77
C. Anticipated Congestion Costs/Trip (USD)									
All	0.04	0.04	0.23	0.16	0.19	0.23	0.07	0.04	0.13
Core-Locations	0.54	0.54	1.26	0.95	1.18	1.07	0.5	0.33	0.12
Non-Core	0.13	0.13	0.32	0.23	0.3	0.29	0.14	0.1	0.21
High-VOT Types	0.03	0.03	0.22	0.15	0.18	0.22	0.07	0.03	0.12
Low-VOT Types	0.04	0.04	0.2	0.15	0.17	0.21	0.07	0.03	0.12
D. Unanticipated Congestion Costs/Trip (USD)									
All	0.22	0.22	1.24	1.0	1.16	1.23	0.39	0.19	0.72
Core-Locations	0.21	0.21	1.27	1.05	1.2	1.29	0.39	0.18	0.74
Non-Core	0.03	0.03	0.22	0.15	0.18	0.23	0.07	0.03	0.8
High-VOT Types	0.23	0.23	1.67	1.35	1.56	1.68	0.54	0.27	0.97
Low-VOT Types	0.16	0.16	0.8	0.65	0.74	0.76	0.21	0.09	0.45
Unanticipated Delays Fraction	0.11	0.11	0.14	0.15	0.15	0.15	0.12	0.11	0.13
E. Daily Estimated Congestion Costs (\$,000)*									
All Prague Traffic	3.14	3.14	103.36	87.18	101.81	129.66	39.35	11.18	482.95

NOTE: This tables reports a summary of our hourly congestion cost estimates. Each column aggregates results into three-hour bins so that, for example, 3:00pm refers to the period between 12:01am-3:00am. The final column reports means across all hours, except for panels indicated with an asterisk (*), in which the final column reports the sum over all hours.

example, suppose a city needs to write a procurement contract for road repairs which specifies bonus payments to contractors for timely completion. In order to determine what incentives to offer, our method could be used to measure the costs of congestion due to each additional day of construction, as long as anticipated delay time estimates for the particular project were available. Another use for congestion cost estimates is in determining the value of new infrastructure itself, such as the benefits to adding a new lane to a road, upgrading an intersection, or adding public transportation services. In each of these cases we would need to have information about the specific project to explicitly characterize its impact on traffic speeds and any spillover effects to nearby roads, but such an analysis remains feasible with our approach.

7 Conclusion

The trade-off between time and money is at the heart of many important economic decisions. In this project we use data from a large European ride share platform that offers menus with explicit trade-offs between time and money. We use this unique

feature to estimate a demand model based on choices from these menus. This allows us to recover riders' preferences and demand elasticities over time and money as well as their implied willingness-to-pay to reduce waiting time. Building on ideas in [Small \(2012\)](#) we then provide a framework to translate the estimates from this demand model into the implied *value of time* across different locations, times and types of individuals.

Our demand model reveals noteworthy patterns in how individuals value time and money. Consumers are about four times more price elastic than waiting time elastic. There are also large differences in willingness to pay for waiting time reductions in the population of riders. The top fifty percent of riders are willing to pay on average \$17.14 per hour of waiting time reduction, compared to \$4.54 for the bottom fifty percent. We show how these differences vary over the day, and we find that work hours lead to higher willingness-to-pay compared to evening hours. We also find that the origin-destination-specific WTP measures are highly correlated with the fraction of travel flows across the same routes, lending insight into studies using travel flows as a proxy for place-specific values.

We decompose the demand model estimates to quantify the value of time. We estimate the average value of time in Prague at \$14.48 per hour during work hours and \$12.55 per hour during non-work hours. We again find substantial heterogeneity by place, time and individuals. A variance decomposition reveals that 85% of the overall heterogeneity is driven by individual differences. Only a small fraction of the variation is due to place and time of day, once individual differences are accounted for. The model also allows us to distinguish between the value of time in planned and unplanned activities. We find that there are large differences in the population of riders in terms of the ratio of planned and unplanned time values and these differences are again increasing during work time. This finding suggests that more time-elastic demand might be driven by the need to coordinate with coworkers and other complementary production factors at work.

There are several qualitative insights that arise from our model. First, it provides a measure of the short-run value of places due to their set of features, amenities and activities, as it is based on directly observed decisions about same-day time allocation. Using other common methods to measure the value of places, such as housing prices, will risk confounding the short-run valuations with the long-run expectations that make up real estate demand. In contrast with literature on residential sorting, our results suggest that people with a higher value of time do not sort into places with a higher value of time. This suggests that the use of places may be more egalitarian than residential choices. Finally, our approach shows that data generated by ride hail companies and other transportation services can be exploited to better understand urban infrastructure

and potentially be used to detect and evaluate missing transport links.

To demonstrate how our estimates can be used in a policy context, we quantify the cost of traffic congestion in Prague. This is done by combining our high-resolution traffic speed data with our measures of planned and unplanned time costs. By recovering per-trip time costs across each hour of the day, we estimate that congestion imposes approximately \$483 thousand per day in time costs.

Bibliography

- Abowd, John M, Francis Kramarz, and David N Margolis**, “High wage workers and high wage firms,” *Econometrica*, 1999, 67 (2), 251–333.
- 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.
- , — , and **Loukas Karabarbounis**, “Recent developments in the economics of time use,” *Annu. Rev. Econ.*, 2012, 4 (1), 373–397.
- Allen, Treb and Costas Arkolakis**, “The welfare effects of transportation infrastructure improvements,” Technical Report, National Bureau of Economic Research 2019.
- Almagro, Milena and Tomás Dominguez-Iino**, “Location Sorting and Endogenous Amenities: Evidence from Amsterdam,” 2019.
- Athey, Susan, Billy Ferguson, Matthew Gentzkow, and Tobias Schmidt**, “Experienced Segregation,” 2019.
- Bayer, Patrick, Fernando Ferreira, and Robert McMillan**, “A unified framework for measuring preferences for schools and neighborhoods,” *Journal of political economy*, 2007, 115 (4), 588–638.
- Becker, Gary S**, “A Theory of the Allocation of Time,” *The Economic Journal*, 1965, pp. 493–517.
- , “A theory of marriage: Part I,” *Journal of Political economy*, 1973, 81 (4), 813–846.
- Belenky, Peter**, “Revised departmental guidance on valuation of travel time in economic analysis,” *US Department of Transportation. Washington, DC*, 2011.
- Bento, Antonio, Kevin Roth, and Andrew R Waxman**, “Avoiding traffic congestion externalities? the value of urgency,” Technical Report, National Bureau of Economic Research 2020.
- Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying**, “Does working from home work? Evidence from a Chinese experiment,” *The Quarterly Journal of Economics*, 2015, 130 (1), 165–218.

- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa**, “Discretizing unobserved heterogeneity,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2017, (2019-16).
- 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.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler**, “Labor supply of New York City cabdrivers: One day at a time,” *The Quarterly Journal of Economics*, 1997, pp. 407–441.
- Castillo, Juan Camilo**, “Who Benefits from Surge Pricing?,” *Available at SSRN 3245533*, 2019.
- 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.
- Chetty, Raj, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter**, “The opportunity atlas: Mapping the childhood roots of social mobility,” Technical Report, National Bureau of Economic Research 2018.
- 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.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe**, “Using big data to estimate consumer surplus: The case of uber,” Technical Report, National Bureau of Economic Research 2016.
- Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst**, “Income growth and the distributional effects of urban spatial sorting,” Technical Report, National Bureau of Economic Research 2019.
- , **Gilles Duranton, and Matthew A Turner**, “Speed,” *Review of Economics and Statistics*, 2018, 100 (4), 725–739.

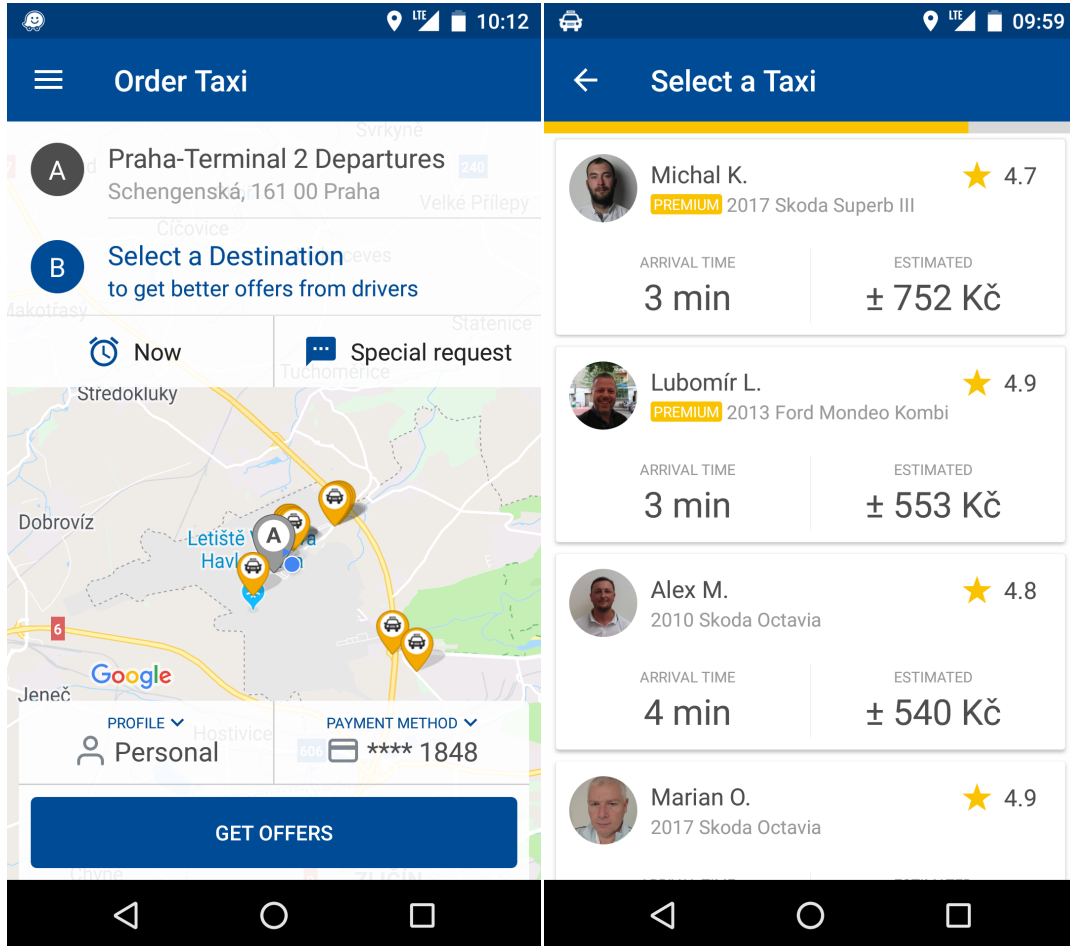
- Crawford, Vincent P and Juanjuan Meng**, “New york city cab drivers’ labor supply revisited: Reference-dependent preferences with rationally expected targets for hours and income,” *The American Economic Review*, 2011, *101* (5), 1912–1932.
- Davis, Donald R, Jonathan I Dingel, Joan Monras, and Eduardo Morales**, “How Segregated is Urban Consumption?,” 2017.
- Domencich, Thomas A and Daniel McFadden**, “Urban travel demand—a behavioral analysis,” Technical Report 1975.
- Dunning, Iain, Joey Huchette, and Miles Lubin**, “JuMP: A modeling language for mathematical optimization,” *SIAM Review*, 2017, *59* (2), 295–320.
- Duranton, Gilles and Matthew A Turner**, “The fundamental law of road congestion: Evidence from US cities,” *American Economic Review*, 2011, *101* (6), 2616–52.
- Eeckhout, Jan and Philipp Kircher**, “Assortative matching with large firms,” *Econometrica*, 2018, *86* (1), 85–132.
- Fajgelbaum, Pablo D and Edouard Schaal**, “Optimal transport networks in spatial equilibrium,” Technical Report, National Bureau of Economic Research 2017.
- Farber, Henry S.**, “Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers,” *The Journal of Political Economy*, 2005, *113* (1), 46–82.
- Farber, Henry S.**, “Reference-dependent preferences and labor supply: The case of New York City taxi drivers,” *The American Economic Review*, 2008, *98* (3), 1069–1082.
- Fosgerau, Mogens, Katrine Hjorth, and Stéphanie Vincent Lyk-Jensen**, “The Danish Value of Time Study,” Technical Report, Danmarks Transportforskning, Report 5, 2007.
- Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz**, “Frictions in a competitive, regulated market: Evidence from taxis,” Technical Report, National Bureau of Economic Research 2018.
- Hall, Jonathan D**, “Pareto improvements from Lexus Lanes: The effects of pricing a portion of the lanes on congested highways,” *Journal of Public Economics*, 2018, *158*, 113–125.
- Hall, Jonathan V, John J Horton, and Daniel T Knoepfle**, “Pricing Efficiently in Designed Markets: The Case of Ride-Sharing,” 2019.

- Heblich, Stephan, Stephen J Redding, and Daniel M Sturm**, “The making of the modern metropolis: evidence from London,” Technical Report, National Bureau of Economic Research 2018.
- Henderson, J Vernon, Adam Storeygard, and David N Weil**, “Measuring economic growth from outer space,” *American economic review*, 2012, 102 (2), 994–1028.
- Jara-Diaz, 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.
- Kreindler, Gabriel E**, “The welfare effect of road congestion pricing: Experimental evidence and equilibrium implications,” *Unpublished paper*, 2018.
- Kreindler, GE and Y Miyauchi**, “Measuring commuting and economic activity inside cities with cell phone records,” Technical Report, Working paper 2019.
- 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.
- Liu, Tracy, Zhixi Wan, and Chenyu Yang**, “The efficiency of a dynamic decentralized two-sided matching market,” *Available at SSRN 3339394*, 2019.
- Mas, Alexandre and Amanda Pallais**, “Valuing alternative work arrangements,” *American Economic Review*, 2017, 107 (12), 3722–59.
- McFadden, Daniel**, “The measurement of urban travel demand,” *Journal of public economics*, 1974, 3 (4), 303–328.
- 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,” *International Economic Review*, 2019, 60 (1), 25–51.
- Petrin, Amil and Kenneth Train**, “A control function approach to endogeneity in consumer choice models,” *Journal of marketing research*, 2010, 47 (1), 3–13.
- Research, John Bates Services Significance Quantitative**, “Values of time and reliability in passenger and freight transport in The Netherlands,” Technical Report 2007.

- Rossi, Peter E. and Greg M. Allenby**, “Bayesian Statistics and Marketing,” *Marketing Science*, 2003, 22, pp. 304–329.
- , —, and **Rob McCulloch**, *Bayesian Statistics and Marketing*, Wiley, 2005.
- Schrank, David, Tim Lomax, and Bill Eisele**, “2019 Urban Mobility Report,” *Texas Transportation Institute*, [ONLINE]. Available: <http://mobility.tamu.edu/ums/report>, 2012.
- Shimer, Robert and Lones Smith**, “Assortative matching and search,” *Econometrica*, 2000, 68 (2), 343–369.
- Small, Kenneth A.**, “The scheduling of consumer activities: work trips,” *The American Economic Review*, 1982, 72 (3), 467–479.
- Small, Kenneth A.**, “Valuation of travel time,” *Economics of Transportation*, 2012, 1 (1), 2 – 14.
- Su, Yichen**, “The Rising Value of Time and the Origin of Urban Gentrification,” Technical Report, Stanford University 2018.
- 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.
- Thakral, Neil and Linh T Tô**, “Daily Labor Supply and Adaptive Reference Points,” 2017.
- Train, Kenneth E.**, *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009.
- Tseng, Yin-Yen and Erik T Verhoef**, “Value of time by time of day: A stated-preference study,” *Transportation Research Part B: Methodological*, 2008, 42 (7-8), 607–618.
- Vickrey, William S.**, “Congestion theory and transport investment,” *The American Economic Review*, 1969, 59 (2), 251–260.
- Waldfogel, Joel**, “The selection hypothesis and the relationship between trial and plaintiff victory,” *Journal of Political Economy*, 1995, 103 (2), 229–260.


A Omitted Graphics

Figure 14: Picture of the Liftago App




NOTE: description here.

Figure 15: Trips with Price/Waiting Tradeoff



The graph shows the proportion of trips that involve an option to spend more and wait less.

Figure 16: Driver Fixed Effects



B Relation to the model in [Small \(1982\)](#)

We explain in this section how the model we use in the main body relates to the model in [Small \(1982\)](#), which is inspired by the work of [Vickrey \(1969\)](#), and in particular, to its microfoundation by [Tseng and Verhoef \(2008\)](#).

In Tseng and Verhoef (2008), a rider is defined by the value he assigns to being at the origin at time t , $\tilde{v}_O(t)$, the value he assigns to being at the destination at time t , $\tilde{v}_D(t)$, and the value he assigns to being on a ride at time t , $\tilde{v}_R(t)$. This, in turn, define a net value of being at the origin and at the destination at time t , $\text{vOT}_O(t) = \tilde{v}_O(t) - \tilde{v}_R(t)$ and $\text{vOT}_D(t) = \tilde{v}_D(t) - \tilde{v}_R(t)$, respectively.

The rider has an ideal arrival time at the destination, which is the time t^* at which $\text{vOT}_O(t^*) = \text{vOT}_D(t^*)$. However, moving from the origin to the destination involves some travel time and the time waiting for the ride. Letting Δ denote the length of the trip²⁹, t_O the time at which the ride is requested, and w the waiting time, the rider arrives at his destination at time $t_O + w + \Delta$.

Thus, the time cost of the trip is given by:

$$\begin{aligned} c_T(t_O + w, t_O + w + \Delta) &= \int_{t_O + w}^{t_O + w + \Delta} \text{vOT}_O(t) dt \\ &+ \mathbb{1}[t_O + w + \Delta < t^*] \int_{t_O + w + \Delta}^{t^*} (\text{vOT}_O(t) - \text{vOT}_D(t)) dt \\ &+ \mathbb{1}[t_O + w + \Delta > t^*] \int_{t^*}^{t_O + w + \Delta} (\text{vOT}_D(t) - \text{vOT}_O(t)) dt, \end{aligned} \quad (17)$$

The cost of time consists of three components: (i) the time on the ride could have been spent at the origin so that its cost is the foregone value at the origin, (ii) if the rider arrives before her ideal arrival time, then she foregoes the value of being at the origin (which is higher since $t_O + w + \Delta < t^*$), and (iii) if the rider arrives after her ideal arrival time, then she foregoes the value of being at the destination (which is higher since $t_O + w + \Delta > t^*$).

Now suppose a rider compares two different trips associated with two different waiting times with expected trip length Δ . It is clear that the first term drops out of the expression.

$$c_T(t_O + w_1, t_O + w_1 + \Delta) - c_T(t_O + w_2, t_O + w_2 + \Delta) \quad (18)$$

Since the disutility of travel plays no role in the decisions that we study, we are dropping this term from our expression. Since we are holding Δ constant, any early arrival relative to t^* implies an additional unit at the destination and one unit less at the origin relative to the plan. Similarly, any late arrival implies one unit less at the destination and one unit more at the origin relative to the plan. This means that the cost of early arrivals is given by:

²⁹Recall that we assume that it is independent of the starting time.

$$\int_{t_0+w}^{t^*} \text{VOT}_D(t) - \text{VOT}_O(t) dt \quad \text{if } t^* > t_0 + w$$

and the cost of late arrivals by:

$$\int_{t^*}^{t_0+w} \text{VOT}_O(t) - \text{VOT}_D(t) dt \quad \text{if } t^* < t_0 + w$$

[Small \(1982\)](#) considers a special case of the above model where $\text{VOT}_O, \text{VOT}_D$ are constant in time, so that [Equation 17](#) can be written as:

$$c_T(t_0 + w, t_0 + w + \Delta) = \alpha \Delta + \beta \max\{t^* - t_0 - w - \Delta, 0\} + \gamma \max\{t_0 + w + \Delta - t^*, 0\}, \quad (19)$$

where $\alpha = \text{VOT}_O, \beta = \text{VOT}_O - \text{VOT}_D,$ and $\gamma = \text{VOT}_D - \text{VOT}_O.$ This version of the model, where over the relevant time intervals, VOT_O and VOT_D are constant, is the one to which the model in the main body of the paper maps to. Whereas [Small \(1982\)](#), assumes that $\text{VOT}_O, \text{VOT}_D$ are constant everywhere, it is important to note that this is not the case in the exercise that we carry out. Indeed, we only assume that this holds over $[t_0, t_0 + w + \Delta];$ however, we allow these values to vary across different times of day, different locations, and different types of riders.

C Identification of Location Values

The following are the three assumptions that will be used in our main theorem below.

Assumption 1. (Independence across locations) $F(V_a^{i,k}, V_{a'}^{i,k'}) = F_{i,a,k}(V) \cdot F_{i,a',k'}(V)$
 $\forall (a, a') \in \mathcal{J}^2 \text{ and } (k, k') \in \{o, d\}^2$

The following assumption imposes the existence of at least one location where the mean of the values is known either when this location is the destination or it is the origin. For example, the mean (local) wage per minute might be a good approximation of the mean value per minute for trips originating in a typical business area.

Assumption 2. (Location Normalization) $\exists a \in \mathcal{A}, k \in \{o, d\} : \mathbb{E}[v_{i,a}^k] = \mu_0.$

While we will establish non-parametric identification of the value distributions, for the empirical application it will be useful to consider a special case of normal distributions.

Assumption 3. (Normal Distributions) $V_{i,a}^k \sim N(\mu_a^k, \sigma_a^k) \forall i, a \in \mathcal{A}, k \in \{o, d\},$ iid across $(i, a, k).$

With these assumptions we can state the following identification result. The first part is straightforward, but requires the knowledge of a subset of trips, such that one of the two values in the difference $V_{i,a}^d - V_{i,a'}^o$ is a known constant. The second part imposes a potentially weaker requirement: only that we know the mean value at one location (whether it is the origin or the destination) and that one part of difference can be held constant across a subset of observations. The last part is a special case, which imposes more parametric structure.

Theorem 1. (Sufficient Conditions for Identification)

1. Under Assumption 1, if $\exists (X^o, X^d, a) \subset X \times X \times \mathcal{A} : \Pr[V_{i,a}^k = \hat{v}_k | X^k, a] = 1$ for $k \in \{o, d\}$, then $F_{a,k}(v)$ is identified $\forall a \in \mathcal{A}, k \in \{o, d\}$.
2. Under Assumptions 1 and 2, when either $\exists k \in \{o, d\} : v_{i,a,t}^k = v_{i',a,t}^k \forall (i, i', a, t)$ or $\exists k \in \{o, d\} : v_{i,a,t}^k = v_{i,a,t'}^k \forall (i, a, t, t')$, then $F_{a,k}(v)$ is identified $\forall a \in \mathcal{A}$ and $k \in \{o, d\}$ whenever $|\mathcal{A}| \geq 3$.
3. Under Assumptions 2 and 3, (μ_a^k, σ_a^k) is identified $\forall a \in \mathcal{A}, k \in \{o, d\}$ whenever $|\mathcal{A}| \geq 3$.

Proof: See [Appendix C](#). The theorem says that we can identify the distributions of valuations non-parametrically whenever (1) values are independent across locations and we can isolate cases for which the value either at the origin or at the destination is known to be equal to some constant, (2) when values are independent and are location- and time-, but not individual-specific either at the origin or at the destination (e.g., everyone has the same value of spending a minute in a residential area) or values are location and individual, but not time-specific (e.g., an individual has always the same value of a minute of being at his business location at 8am every Monday), we know the mean value in at least one location and there are at least 3 locations. Given the previous result, we can of course also identify the values parametrically. However, this approach is particularly easy when the values are iid normal, when there are at least 3 locations and we have one location normalization for both mean and variance. (1) follows from restricting attention to cases where one part of the difference is known to be a constant, and hence the distribution can be fully recovered from the distribution of the differences. (2) follows from a deconvolution argument as in Li and Vuong (1998) since under the hypothesis of the theorem we can construct a sample in which we hold one part of the difference fixed. By applying the Kotlarski Lemma, one can then recover the distributions of both pieces of the difference (up to the location). The additional location normalization pins down the means separately. (3) follows since adding an additional location to the existing set of L locations comes with 4 new parameters (two

means and two standard deviations) and generates $2 \cdot (L - 1)$ additional equations and hence one needs $L \geq 3$ together with the two normalizations.

Proof. (1) Restricting attention to X^o and trips originating in a going to destination a' identifies $F_{D,a'}(v)$ trivially from the observed distribution of the differences by shifting the distribution of the data by the relevant constant, \hat{v}_O . Similarly, $F_{O,a'}(v)$ is identified from using only data from X^d and trips that originated in a and went to a' .

(2) Consider $|\mathcal{A}| = 3$ and the case $v_{i,a,t}^o = v_{i',a,t}^o \forall (i, i', a, t)$, i.e., that values at a given origin and time are the same for all individuals i . Since values at origin are the same, the deconvolution argument identifies (using variation over time) the distribution of values at origin and the distribution of the values at each destination (using variation across individuals going to each destination) - up to a location normalization, i.e., the two means cannot be identified separately directly.³⁰ With 3 locations we can setup the following system of 6 equations in 6 unknowns (in the means of values at origin and at destination), where the objects on the LHS are recovered from the deconvolution argument: $\mu_{12} = -\mu_1^o + \mu_2^d, \mu_{13} = -\mu_1^o + \mu_3^d, \mu_{21} = -\mu_2^o + \mu_1^d, \mu_{23} = -\mu_2^o + \mu_3^d, \mu_{31} = -\mu_3^o + \mu_1^d, \mu_{32} = -\mu_3^o + \mu_2^d$. This system is not identified since the matrix of coefficients has a rank of 5. Substituting equation from Assumption 2 for one of the equations including that mean restores the full rank of the coefficient matrix. For any additional location, 2 new parameters and $L - 1$ new equations are introduced, but the rank is still $(2L-1)$. Substituting equation from Assumption 2 brings the rank to $2L$, and we have $L \cdot (L - 1)$ equations. Hence we can simply keep any subset with rank $2L$ as at true parameters all of these equations will hold simultaneously.

The case $v_{i,a,t}^o = v_{i,a,t'}^o \forall (i, a, t, t')$, is analogous, except variation over individuals is used for the deconvolution argument and the distribution of values at origin and destination is identified using variation within individuals as they go to different destinations. The means are then recovered as above.

(3) Since sum of two normally distributed random variables is also normally distributed with mean being the sum of means and the variance being the sum of the variances, the means can be recovered directly by using the argument in (2). However, the system of equations involving the variances is also less than full rank - and hence to identify these directly, one needs another location normalization (such as the knowledge of the variance at some location) or one needs to appeal to case (2), which establishes that the variances are identified non-parametrically.

□

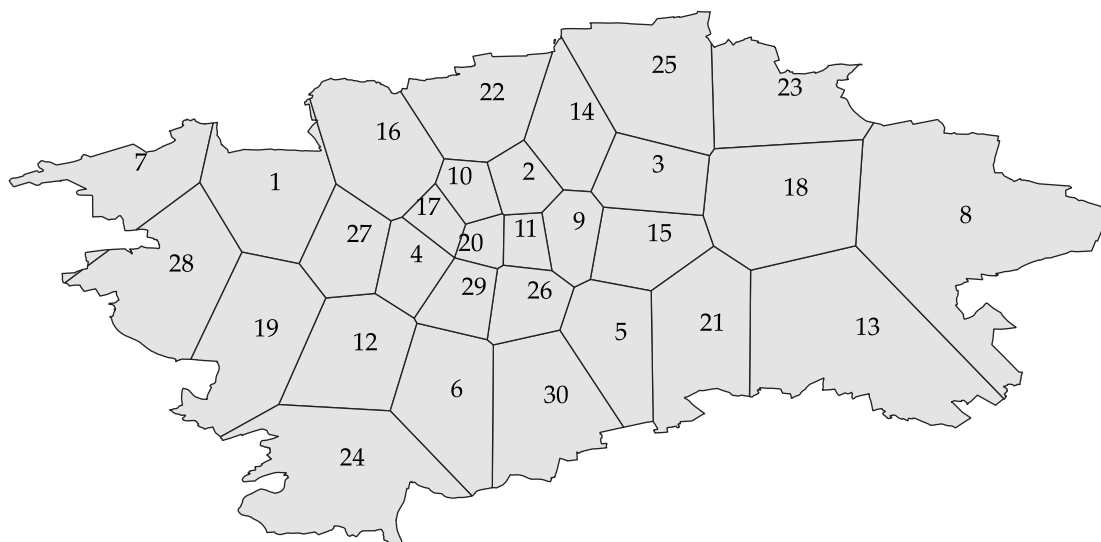
³⁰In the measurement literature setup $Y_1 = X + \varepsilon_1, Y_2 = X + \varepsilon_2$, the distributions $F(X), G(\varepsilon)$ are identified from $H(Y_1, Y_2)$ whenever the usual iid assumption holds, when the characteristic functions of F and G are non-vanishing everywhere and when $E(\varepsilon) = 0$.

D Data Details

D.1 Locations

Using the exact GPS points of trip origin, we partition our data into 30 locations. [Figure 17](#) shows these locations together with an index value for comparing results in [Section 5](#). The partitioning is done according to a simple *k-means* clustering procedure on the requested pickup locations with $k = 30$. This procedure minimizes the straight-line distance between each point and the weighed center of all points within the same cluster, with the constraint that each cluster has an equal number of points. The depicted locations are close approximations, displayed as Voronoi cells that contain the clustered points. This process allows location definitions to be independent of any political boundaries and better representative of places in which demand is concentrated.

Figure 17: Locations in Prague



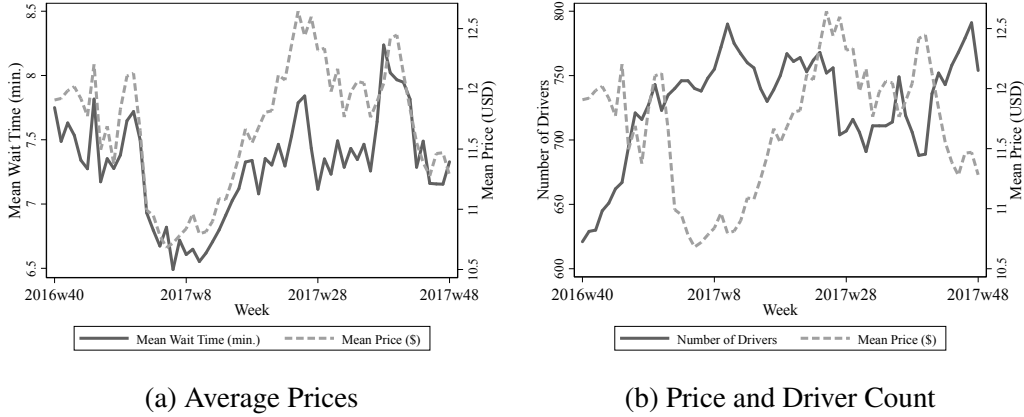
Note: This figure maps the boundaries of the city of Prague with locations defined by a kmeans-clustering procedure on GPS-locations of trip origins and depicted as Voronoi cells that contain the clustered points. Displayed index values correspond to indices used in the paper.

D.2 Market Time Series Summary

[Figure 18](#), panels (a) and (b) summarize the week-over-week trends in prices. 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.

Figure 18: Weekly Waiting Times and Number of Drivers Compared to Prices



E Estimation Details

E.1 Recovering Latent Types: Conditional *K-Means-Clustering* Procedure

The clustering procedure allocates passengers to initial latent classes using a “*k-means++*” algorithm and then reallocates passenger types to best explain individual choices after controlling for observable features of the environment. This approach is based on [Bonhomme et al. \(2017\)](#) and uses Julia’s *Clustering* package implementation of *k-means++*. Our discussion here is mainly based on the two.

E.2 Location Parameter Estimates

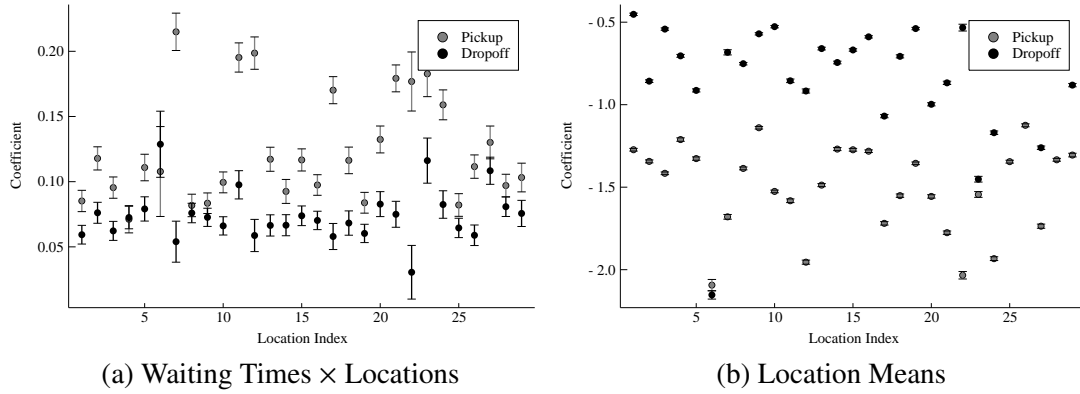
Here we offer additional details on location-specific parameter estimates. [Figure 19](#) plots these estimates across origins and destinations.

E.3 Model Fit

[Figure 20](#) illustrates the model’s ability to fit the observed choices in the data. For each trip and corresponding driver bids, we use our estimates to predict whether each customer will pick the outside-option and plot the average prediction for each week.


[Figure 21](#) Demonstrates the model’s ability to correctly predict specific choices across auctions. We do this in two ways: first we use our estimates to predict which

Figure 19: Location-Specific Coefficient Estimates



Note: This figure shows coefficient estimates for 120 location-specific parameters omitted from Table 2. Standard errors are depicted as bars around each point. Location indices may be cross-references with Figure 17. Panel (a) displays results for the coefficients on the interaction between waiting time and location, and Panel (b) displays results for the coefficients on location indicators, which can be interpreted as location-specific utility shifters relative to location index 1.

Figure 20: Model Fit On Outside Option Selection



option each customer will pick inclusive of simulated draws of $\epsilon_{i,j,t}$. Because the $\epsilon_{i,j,t}$ represent unobservable attributes associated with each choice, this simulation will tend to add additional noise to our choice predictions. We also make the same predictions but set $\epsilon_{i,j,t} = 0$ for all i, j, t , which improves our predictions by about 10%.

E.4 Elasticities by Area

See Table 9 and Table 10 below.

Table 9: Bid Level Elasticities by Origin Area

ORIGIN AREA	High VOT-Type Elasticities		Low VOT-Type Elasticities	
	PRICE	WAITING TIME	PRICE	WAITING TIME
1	-4.82	-0.98	-5.22	-0.93
2	-2.71	-0.37	-3.66	-0.3
3	-3.98	-0.56	-4.26	-0.46
4	-2.95	-0.38	-3.79	-0.29
5	-4.47	-0.37	-4.53	-0.38
6	-5.05	-0.59	-5.37	-0.57
7	-9.64	-0.27	-9.7	-0.22
8	-10.0	-0.55	-9.75	-0.48
9	-2.78	-0.46	-3.38	-0.41
10	-2.62	-0.33	-3.7	-0.25
11	-2.41	-0.36	-3.19	-0.26
12	-5.26	-0.47	-5.55	-0.43
13	-7.81	-0.47	-7.58	-0.44
14	-3.95	-0.51	-4.34	-0.44
15	-3.45	-0.49	-3.69	-0.44
16	-3.69	-0.38	-4.54	-0.31
17	-2.79	-0.34	-3.89	-0.24
18	-6.75	-0.78	-6.45	-0.7
19	-5.63	-0.56	-5.93	-0.53
20	-2.52	-0.31	-3.43	-0.22
21	-5.48	-0.52	-5.41	-0.47
22	-5.06	-0.7	-5.45	-0.6
23	-6.9	-0.64	-7.09	-0.6
24	-10.31	-0.44	-9.6	-0.38
25	-5.82	-0.44	-6.17	-0.41
26	-2.98	-0.39	-3.59	-0.34
27	-3.89	-0.56	-4.33	-0.5
28	-6.66	-0.43	-6.56	-0.39
29	-2.84	-0.36	-3.64	-0.29
30	-4.98	-0.57	-5.2	-0.51

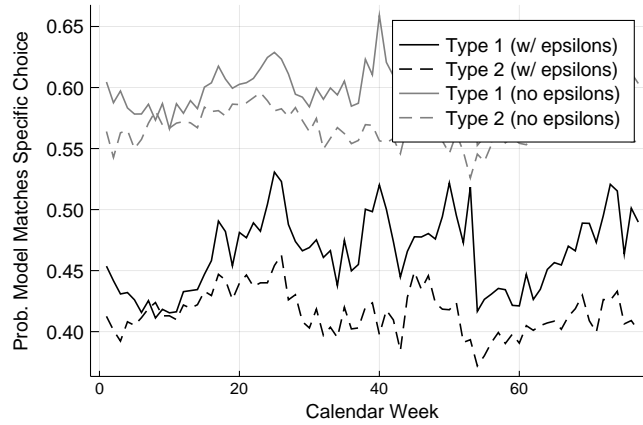
NOTE: This table shows price and waiting time elasticities across the thirty different origin places separated by high and low VOT types. [Figure 17](#) shows location indices on a map of Prague.

Table 10: Bid Level Elasticities by Destination Area

DESTINATION AREA	High VOT-Type Elasticities		Low VOT-Type Elasticities	
	PRICE	WAITING TIME	PRICE	WAITING TIME
1	-4.3	-0.47	-4.44	-0.42
2	-2.45	-0.47	-3.09	-0.45
3	-3.6	-0.32	-3.88	-0.26
4	-2.67	-0.45	-3.31	-0.47
5	-4.16	-0.43	-4.39	-0.32
6	-4.75	-0.28	-4.96	-0.21
7	-9.35	-0.38	-8.33	-0.39
8	-8.75	-0.32	-8.48	-0.28
9	-2.64	-0.37	-3.16	-0.36
10	-2.52	-0.44	-3.29	-0.45
11	-2.32	-0.35	-2.87	-0.35
12	-4.81	-0.34	-4.69	-0.24
13	-7.02	-0.26	-7.06	-0.21
14	-3.55	-0.32	-3.93	-0.27
15	-3.18	-0.32	-3.51	-0.3
16	-3.18	-0.41	-3.75	-0.37
17	-2.57	-0.45	-3.41	-0.46
18	-5.86	-0.3	-5.98	-0.24
19	-5.24	-0.33	-5.13	-0.26
20	-2.42	-0.48	-3.26	-0.5
21	-5.04	-0.33	-5.3	-0.28
22	-4.27	-0.31	-4.37	-0.23
23	-6.14	-0.19	-5.77	-0.17
24	-8.21	-0.38	-7.69	-0.31
25	-5.42	-0.36	-5.51	-0.28
26	-2.84	-0.41	-3.24	-0.37
27	-3.35	-0.35	-3.46	-0.29
28	-6.0	-0.37	-5.84	-0.27
29	-2.72	-0.36	-3.19	-0.34
30	-4.58	-0.33	-4.67	-0.26

NOTE: This table shows price and waiting time elasticities across the thirty different destination places separated by high and low VOT types. [Figure 17](#) shows location indices on a map of Prague.

Figure 21: Model Fit On Outside Option Selection



E.5 Estimation Results Omitting Control Function

Table 11 below presents the price and waiting time elasticity estimates in which the control function is omitted.

Table 11: Estimated Elasticities

Time of Day	Individual Type	Bid Level Elasticities		Order Level Elasticities	
		PRICE	WAITING TIME	PRICE	WAITING TIME
<i>Daytime</i> 6am-6pm	Overall	-1.83	-0.63	-1.74	-0.64
	H Price, H Wait Sensitivity	-4.08	-1.18	-3.83	-1.09
	H Price, L Wait Sensitivity	-0.95	-0.51	-1.04	-0.53
	L Price, H Wait Sensitivity	-2.15	-0.6	-2.05	-0.68
	L Price, L Wait Sensitivity	-0.65	-0.3	-0.75	-0.38
<i>Evening</i> 6pm-6am	Overall	-2.58	-0.25	-2.4	-0.28
	H Price, H Wait Sensitivity	-4.32	-0.39	-4.12	-0.4
	H Price, L Wait Sensitivity	-1.25	-0.2	-1.33	-0.22
	L Price, H Wait Sensitivity	-3.09	-0.26	-2.87	-0.31
	L Price, L Wait Sensitivity	-1.04	-0.15	-1.12	-0.18

NOTE: This table shows the demand elasticity of price and waiting time across daytime and evening hours and individual type groupings. This table replicates Table 3 except that the model is estimated with no control function. We distinguish as *high (H) price sensitivity* individuals who have below median values for β_i^p and *low (L) price sensitivity* individuals as those with above median values for β_i^p , and similarly for waiting time sensitivity. The first two columns show these elasticities among competing bids, reflecting the change in demand due to a 1% change in price or waiting time on a single bid. The second two columns show them with respect to choosing the outside option, reflecting a change in demand due to a 1% change in price or waiting time on *all* bids.