

Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides*

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

February 28, 2024

Abstract

We recover valuations of time using detailed data from a large ride-hail platform, where drivers bid on trips and consumers choose between a set of rides with different prices and wait times. Leveraging a consumer panel, we estimate demand as a function of both prices and wait times and use the resulting estimates to recover heterogeneity in the value of time across consumers. We study the welfare implications of personalized pricing and its effect on the platform, drivers, and consumers. Taking into account drivers' optimal reaction to the platform's pricing policy, personalized pricing lowers consumer surplus by 2.5% and increases overall surplus by 5.2%. Like the platform, drivers benefit from personalized pricing. ETA-based pricing—where different prices are set for various wait times and where extensive use of consumer data is not required—can capture a significant portion of the profits garnered from personalized pricing, while simultaneously benefiting consumers.

Keywords: *value of time, transportation markets, demand in transportation markets, ride hail, platforms, price discrimination*

JEL classification: C73; D83; L90; R12

*We thank the editor, Aviv Nevo, and four anonymous referees for feedback that has greatly improved this paper. We thank Juan Camilo Castillo, Gabriel Kreindler, Rob Porter, Stephen Redding, and seminar participants at FTC, Maryland, MIT, Ohio State, Penn, Princeton, Stanford, UCL, UCLA, Wash U., and IO² for useful comments. We thank Stefano Baratucce, Jacob Dorn, Ranie Lin, Thi Mai Anh Nguyen, Roi Orzach, and Yining Zhu for exceptional research assistance. 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

[‡]Columbia University, 665 West 130th Street, New York, NY 10027

[§]CERGE-EI, Politických vizdù 7, 111 21 Prague 1

[¶]Massachusetts Institute of Technology, 50 Memorial Dr, Cambridge, MA 02142

1 Introduction

The trade-off between time and money lies at the core of all transportation markets, where consumers often face the choice of paying a higher price for faster or more immediate travel. How consumers resolve this trade-off largely determines the demand for transportation services and therefore the benefits of policies, such as infrastructure investments, congestion taxes, and pricing. The rise of ride-hailing platforms has made this trade-off between time and money even more salient: platforms offer consumers an increasingly tailored set of options, making reducing the time spent waiting more or less expensive. Relying on the ability to quote prices to consumers on a smartphone app, these platforms can also fine-tune their prices to specific observables, such as time of day and location or aggregate market conditions. Recent work has established the efficiency gains of more sophisticated pricing policies relative to traditional taxi markets (see, e.g., Buchholz, 2022; Castillo, 2019; Rosaia, 2020).

At the same time, ride-hail platforms engage in frequent interactions with consumers, allowing them to learn about consumers' sensitivity to prices and wait times. Consequently, these platforms can assess consumers' willingness to pay for reduced wait times, which we refer to as their *value of time*. Although this ability enables platforms to provide consumers with better tailored options, it also allows them to engage in personalized pricing.

The welfare effect of price discrimination is known to be ambiguous and thus a matter of empirical measurement (Pigou, 1920). However, such measurement is challenging for a number of reasons. Obtaining credible estimates of a single consumer's preferences requires the analyst to have access to substantial amounts of observational data for any given consumer. Two-sided markets, such as ride-hailing platforms, exacerbate this challenge because any changes to consumer pricing necessarily impact the other side of the market. Thus, to measure the welfare implications of pricing policies, the analyst also needs to account for the supply side's incentives.

In this paper, we use detailed consumer choice data from a large European ride-hailing platform, Liftago, to measure the welfare implications of price discrimination on the basis of consumer heterogeneity in the value of time. We make two primary contributions. First, relying on the unique features of our setting, we obtain estimates of the heterogeneity in the value of time across consumers on the platform by estimating a demand system that depends on both prices *and* wait

times. Second, we use these estimates to quantify the welfare effects of platform pricing policies that exploit the heterogeneity in consumers' value of time to price discriminate. As part of this exercise, we estimate a model of driver bidding behavior, which allows us to infer drivers' opportunity costs of serving rides. This approach allows us to account for the payments that are necessary to incentivize drivers' participation under different counterfactual pricing policies.

To measure consumers' preferences over time and money, we exploit the unique features of Liftago's ride-allocation mechanism, which allows consumers to directly express preferences over price and wait time. Liftago allocates rides through a rapid auction process in which nearby drivers bid on ride requests. Requesting consumers then choose between bids based on various characteristics. Most importantly, bids often involve a trade-off between price and wait time, allowing us to observe how consumers resolve this trade-off. We observe consumers' individual choice sets as well as their ultimate selection for 1.9 million ride requests and 5.2 million bids. Because we observe the same consumers repeatedly interacting with the platform, we can recover persistent differences in the value of time across consumers.

Our demand results quantify how much consumers respond to changes in both price and wait time, as well as the underlying persistent heterogeneity in this response. Price elasticities are four to 10 times as large as wait-time elasticities. Expressed as an hourly quantity, we find that consumers' value of time (henceforth, VOT) is on average \$13.21. Although these measures vary widely within the day and across space, most of the variation is driven by latent differences across riders independent of observable sources of variation. Ranking individuals by their relative sensitivity to prices and wait times, we find the VOT of the top quartile is about 3.5 times higher than that of the bottom quartile.

We then use our demand estimates to study the welfare effects of platform pricing strategies that exploit the heterogeneity in consumers' preferences over time and money. We allow the platform to use consumers' historical data as an input to pricing and analyze the impact of different forms of personalized pricing on consumer welfare as well as driver and platform profits. Because of the increasing abundance of consumer data, the potential consequences of personalized pricing are the subject of recent policy debate.¹ In our counterfactual, we allow the

¹See, for example, White House Council of Economic Advisors (2015), OECD Directorate for financial and enterprise affairs, Competition Committee, Note by the United States (2016), and Bourreau and De Streel (2018).

platform to offer a menu of trips to the consumers and separate payments to drivers to ensure their participation. In contrast to the current platform policy, where the platform collects a percentage fee and directly implements drivers' bids as prices, we allow the platform to directly choose the price the rider faces.

An important piece of our counterfactual exercise is the estimation of a model of driver bidding. As discussed above, a unique aspect of platform price discrimination is that different forms of platform pricing have implications for drivers' earnings. Consequently, the platform's costs of procuring a ride are *endogenous* to its pricing policy. Thus, drivers' incentives act as a constraint on the platform's ability to benefit from price discrimination. As a result, how additional surplus is split between the platform and the drivers is not immediately clear. To the best of our knowledge, ours is the first study to consider how personalized pricing affects the supply side in platform markets. This is relevant for many other platform settings because platforms observe consumer data that suppliers do not.

Our model of driver bidding allows us to infer drivers' opportunity cost of serving a ride and evaluate drivers' incentives. In choosing their bids, drivers weigh the revenues from a trip against the value of their outside option, which is heterogeneous and privately observed by drivers. We adopt methods from the empirical auction literature (Guerre et al., 2000; Jofre-Bonet and Pesendorfer, 2003) that allow us to map from observed bids to costs that capture these opportunity costs.

Having estimated the platform's demand and supply of rides, we quantify the welfare effects and the platform's profits from personalized pricing. Relative to the baseline scenario in which offers are determined via a competitive auction and the platform only collects a 10% fee on the winning bid, personalized pricing leads to a threefold increase in platform profits but a decrease in overall welfare, most of which comes from reduced driver profits (-74%), increased ride prices (+7%), and fewer trips being completed on the platform (-20%).

Relative to the baseline, personalized pricing allows the platform to both exercise its market power—by setting prices on both sides of the market—and use its information about consumer preferences. To understand the effects of each of these changes separately, we consider an intermediate counterfactual, *uniform pricing*, in which the platform sets prices on both sides of the market but does not condition its pricing policy on individual consumer preferences. This counterfactual captures the platform's gains from centralizing pricing given the consumers' and drivers' outside options. We find uniform pricing explains most of the losses in

consumer welfare and driver profits relative to the baseline scenario. Consumer surplus decreases by 36% and driver profits decrease by 74%. The number of requests in which all trip offers are rejected increases from 36% to 49% because of higher prices.

Relative to uniform pricing, personalized pricing has a small but negative effect on average consumer welfare. This aggregate loss in consumer surplus, however, masks interesting distributional effects across consumers. Indeed, most consumers (62.5%) benefit from personalized pricing, but these gains are offset by the platform’s ability to increase prices for the most inelastic consumers. Relative to uniform pricing, average prices fall slightly under personalized pricing and the market expands by up to 7.6%. Furthermore, conditional on the platform exercising its market power, drivers benefit from the platform’s ability to incorporate consumer information into its pricing policy, as reflected by a 11% increase in profits relative to uniform pricing. Because both drivers and the platform benefit while consumer surplus decreases only modestly, personalized pricing increases welfare by up to 6.33% relative to uniform pricing.

In practice, platforms may be reluctant to use detailed consumer order histories for their pricing policies. Instead, they may offer menus of trips in which trips with shorter wait times have higher prices to take advantage of consumers’ heterogeneous preferences over wait times. Indeed, ETA-based pricing is becoming increasingly popular among the major ride-hailing platforms.² Motivated by this observation, our final counterfactual, *ETA-based pricing*, allows the platform to set different prices on rides with different wait times, but not to condition its pricing on individual consumer preferences. We compare what fraction of surplus ETA-based pricing can capture relative to the case in which the platform knows individual wait-time sensitivities. We find ETA-based pricing increases profits by 0.9% relative to uniform pricing, which is about two-thirds of the increase from pricing on individual wait-time sensitivity. Furthermore, overall welfare under ETA-based pricing is close to welfare when the platform knows individual wait-time sensitivities.

Our results highlight the nuanced welfare effects of incorporating detailed consumer information into pricing in two-sided markets. Relative to the competitive

²For example, Lyft recently introduced “wait and save,” which grants a discount for riders who are willing to wait longer, in turn serving faster rides to consumers who have more urgent requests (Helling, 2023). Uber followed with a similar feature called UberX Priority; see Uber (2023). A discussion of how Uber uses (personal) data to price is provided by Martin (2019).

baseline mechanism where prices are set via auctions, our uniform pricing counterfactual shows unexercised pricing power by the platform, which comes at a considerable welfare cost for consumers and drivers. However, conditional on the platform exercising this pricing power, the more information the platform incorporates into its pricing policy, the larger the overall welfare gains. Furthermore, our results highlight our welfare conclusions depend on what side of the market we look at: once we account for drivers' incentives, drivers command a substantial share of the surplus created by using more consumer information. Finally, our results suggest that ETA-based pricing strategies can command a substantial portion of the profits arising from personalized pricing.

Related literature As we describe below, the paper contributes to the literatures on price discrimination, taxi and ride-hailing markets, and the transportation literature that studies the value of time.

The empirical literature on price discrimination focuses on non-platform settings and measures the benefits of second-degree (Miravete, 1996; Hendel and Nevo, 2013; Luo, Perrigne, and Vuong, 2018) and third-degree (e.g., Hendel and Nevo, 2013; List, 2004; Bauner, 2015; Levitt, List, Neckermann, and Nelson, 2016) price discrimination, as well as different forms of nonlinear and personalized pricing (e.g., Rossi, McCulloch, and Allenby, 1996; Shiller, 2013; Nevo, Turner, and Williams, 2016; Dubé and Misra, 2023).³ Recent papers study price discrimination in the context of big data; see, for instance, Ali, Lewis, and Vasserman (2022) and Jin and Vasserman (2021) for the benefits of voluntary data disclosure, Kehoe, Larsen, and Pastorino (2018) for personalized pricing in the market of experience goods, Aridor, Che, and Salz (2023) for the effects of privacy-protection policies, and Doval and Skreta (Forthcoming) for how consumers' forward-looking behavior affects firms' incentives to collect data in the first place. Our paper contributes to this literature in at least two ways. First, we quantify the benefits of price discrimination in the context of a two-sided platform, which introduces challenges relative to price discrimination in one-sided markets (see [Section 6](#)). Second, our individual-level measurement of the value of time allows us to quantify the effects of personalized pricing on latent unobservables. This is in contrast to studies such as Dubé and Misra (2023), in which personalized pricing is only based on observable characteristics.

³The literature on the welfare effects of price discrimination goes back to the seminal work of Pigou (1920); see Aguirre, Cowan, and Vickers (2010) and Bergemann, Brooks, and Morris (2015) for recent studies.

Our paper contributes to the literatures on taxi and ride-hail markets. Some of these papers estimate demand for taxis or ride-hailing as a function of prices (Buchholz, 2022; Gaineddenova, 2021) or wait times (Frechette, Lizzeri, and Salz, 2019), but not both. More closely related are Castillo (2019), Rosaia (2020), and Goldszmidt, List, Metcalfe, Muir, Smith, and Wang (2020), which, like us, estimate demand in ride-hail markets as a function of both wait time and price, but with a different focus and data. Castillo (2019) quantifies the benefits of surge pricing, and Goldszmidt et al. (2020) measure the value of time through an experiment on Lyft. Rosaia (2020) studies platform competition and the role of platform pricing policies in determining the distribution of supply across the platforms.⁴ In contrast to these papers, we recover individual-level heterogeneity in demand, which we then use to study the welfare implications of the platform’s ability to personalize prices and steer consumers to different drivers.

Our data allow us to directly measure consumers’ willingness to pay for reductions in wait time based on choices on the platform, which is distinct from, but related to the value of travel-time savings, i.e., the value that people assign to shorter trips. For this reason, our paper contributes to the literature in transportation economics and industrial organization on the value of travel-time savings, dating back to the pioneering work of Daniel McFadden (McFadden, 1974; Domencich and McFadden, 1975). These studies measure the value of travel-time savings through surveys or revealed-preference analysis based on mode choice. Small (2012) provides an excellent review of this literature. Recent studies take advantage of more detailed micro data. Hall (2018) analyzes the benefits of choice over toll and non-toll lanes. Kreindler (2023) experimentally measures the value of peak-congestion pricing in Bangalore. Bento, Roth, and Waxman (2020) use commuter tollway choices to infer consumers’ urgency from their willingness to pay for travel-time savings.

Organization The rest of this paper proceeds as follows. [Section 2](#) describes the institutional setting and our data. [Section 3](#) describes the demand and supply models and [Section 4](#) their estimation. [Section 5](#) presents our estimation results. [Section 6](#) analyzes different forms of price discrimination on the platform. [Section 7](#) concludes.

⁴Gaineddenova (2021) analyzes the effects of platform pricing policies that do not allow riders to sort in terms of their willingness to pay for a ride.

2 Setting and data

2.1 A unique approach to matching and price discovery

Liftago is an app-based ride-hail platform that was founded in 2015 and services rides through licensed taxi drivers in many cities in Europe. We focus on Prague, where licensing requires both a fee and an exam. Moreover, taxis need to be equipped with a physical meter, which captures the number of kilometers traveled in the “occupied” mode and the billed amount. 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. A licensed driver may find rides by searching for street-hail consumers or by choosing to participate in a dispatch service. Among dispatch options are traditional telephone-based dispatch services and, more recently, Liftago. This regulatory environment is different from most US cities in which there is nearly free driver entry into the ride-hail market with firms such as Uber and Lyft. In the period that we study, Liftago is by far the dominant platform in Prague.⁵

Drivers pay a 10% fee for each ride booked through the platform. By tracking both the taxi’s GPS location and the time of the trip, the platform provides an approximate fare before the trip begins and a final fare after its completion. Because Liftago is not well known internationally, few riders are tourists, making our estimates easier to interpret in terms of local economic quantities. This is also reflected in the relatively small fraction of airport rides, which constitute about 2% of total trips.

Drivers and consumers are matched by a combination of a dispatch algorithm and an auction. Whenever a consumer requests a ride, the system looks for nearby available cars and sends requests to a number of them, typically four, to elicit an offer. A driver who receives a request observes the details of the trip —the location of the consumer, the destination, consumer rating, and payment via cash or credit. A driver who is interested in fulfilling the ride submits a bid from a set of pre-programmed tariffs.⁶ A tariff consists of a flag fee, a per-minute waiting fee, and a per-kilometer fee with a regulatory cap of CZK 36 (\approx \$1.41). The platform takes tariff bids and combines them with a query to Waze, a real-time traffic

⁵Uber has also been in Prague since 2014, but its presence is not as large as in a typical US city of similar size, partially because it is still fighting several legal battles, due to various licensing and taxation issues. Since an EU court’s decision in December 2017, Uber is viewed as a transportation company, and hence, its drivers also need to be properly licensed.

⁶Drivers who supply rides in Liftago typically have many pre-programmed bid increments.

mapping service, which provides estimates of the taxi arrival time (equivalently, the consumer’s wait time), the trip time, and trip distance. The tariff bids are then translated into a single expected price for a trip. The consumer then observes bids as final trip prices together with other bid-specific attributes: the wait time until the taxi arrives, the make and model of the car, and the driver’s rating. Importantly, these non-price attributes are automatically attached to each of the bids; in each auction, drivers only have control over the tariff. The consumer may select one of the bids, in which case the ride occurs, or may decline all bids. When the ride is completed, the consumer pays the fare shown on the meter. In [Figure A.1](#) in [Appendix A](#), we show the interface that riders see before making the request and after the offers arrive.

Liftago’s mechanism allows for variation in both prices and wait times: a driver with a high wait time may submit a lower bid than a driver with a short wait time, and vice versa. Contrast this market-clearing mechanism with traditional taxi services, in which prices are fixed and the market clears through adjustments in wait time ([Frechette et al., 2019](#)), and with other ride-hail platforms, in which prices adjust to keep wait times stable ([Castillo, 2019](#)).

2.2 Data

Our dataset covers 1.9 million trip requests and 1.1 million fulfilled trips on Liftago between September 30, 2016, and June 30, 2018. For each request, we observe the time of the request, the pick-up and drop-off location, trip price bids and estimated wait times from each driver, and which bid the consumer chose, if any. In addition, we observe a unique identifier for each driver and consumer. The sample period includes 1,455 unique drivers and 113,916 unique consumers. We complement the data for each ride request with public-transit availability based on the GPS addresses for each origin and destination in the Liftago data. Furthermore, we use data on hourly rainfall in Prague to attach prevailing weather characteristics.⁷

[Table 1](#) summarizes daily activity on the platform. About 3,000 trip requests are made each day, 61% of which become rides. The average bid is \$10.72 and the average wait time is seven 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 (less than 0.33%), no more than four bids

⁷Public data are available from the National Oceanic and Atmospheric Administration (<https://www.noaa.gov/>).

are made. We discard auctions with more than four bids.

Table 1: Bid, Order, and Daily Summary Statistics

Variable	Full Sample <i>N=1,872,362 orders</i>				Weekdays Only (M-F) <i>N=1,337,826 orders</i>			
	P25	Mean	P75	S.D.	P25	Mean	P75	S.D.
Panel A: Chosen Trip								
Price of Trip (USD)	6.16	9.31	11.23	4.63	6.13	9.24	11.11	4.58
Wait Time (minutes)	4.00	6.07	8.00	3.08	4.00	6.14	8.00	3.08
Panel B: All Bids								
Price of Trip (USD)	9.95	10.72	11.47	1.17	9.79	10.57	11.31	1.16
Wait Time (minutes)	5.81	7.01	8.22	1.85	5.88	7.11	8.36	1.88
Number of Bids	2.00	2.79	4.00	1.09	2.00	2.81	4.00	1.08
Panel C: Daily								
Requests per Day	2319	2934	3520	956	2376	2940	3486	947
Trips per Day	1410	1785	2149	557	1502	1831	2180	569
Drivers per Day	464	499	553	89	491	518	569	90

NOTE: This table shows summary statistics at the auction level (Panel A), the bid level (Panel B), and the daily level (Panel C). P25 refers to the 25th percentile and P75 to the 75th percentile of the respective variable. Panel B reports the average over each statistic computed within each auction. The data are based on 638 days of observations.

Table 2 summarizes drivers’ daily activity on the platform. On average, drivers participate in 16 auctions per day on the platform and win around 3.8 of these auctions per day. Although we do not observe drivers’ off-platform trips, we observe they spend around one hour driving on the platform. We also find the time spent between serving two platform rides is on average 112 minutes. The drivers’ activity on the platform is consistent with the drivers’ being licensed taxi drivers who supplement their street-hail business with on-platform rides. Taken together, these statistics suggest drivers rely on the street-hail business for the majority of their earnings.

2.3 Preferences over time and money: Intra-daily patterns

In this section, we provide descriptive evidence for patterns in prices, wait times, and choices. We document large and interpretable heterogeneity in consumer choices, which provides important identifying variation for our model. In Figure 1, we show the average prices (Figure 1a) and wait times (Figure 1b) of offered trips by day of the week and time of the day. Prices are lower during weekday afternoons and higher during weekends, whereas wait times tend to be substantially higher

Table 2: Daily Driver Statistics

Variable	Full Sample <i>N=318,458 driver-days</i>				Weekdays Only (M-F) <i>N=235,511 driver-days</i>			
	P25	Mean	P75	S.D.	P25	Mean	P75	S.D.
Bids Per Day/Driver	6.00	16.40	23.00	13.38	6.00	15.99	22.00	12.93
Trips Per Day/Driver	1.00	3.58	5.00	3.44	1.00	3.54	5.00	3.36
Win Probability	0.09	0.23	0.33	0.19	0.09	0.23	0.33	0.19
Daily Time on Trip (min.)	14.78	61.99	93.58	60.27	15.07	62.46	94.32	60.32

NOTE: This table describes the daily distribution of bids, trips, win probabilities, and minutes spent on trips for a driver. P25 refers to the 25th percentile and P75 to the 75th percentile of the respective variable.

during the day than overnight.

Consumers in our data often face a non-trivial trade-off between price and wait time when choosing among bids. A trade-off implies the existence of options such that one has a shorter wait time but a higher price, and vice versa. Depending on the time of day, about 58%–70% of auctions involve a trade-off between waiting less and paying more (see [Figure B.1](#) in [Appendix B](#) for additional detail.)

In [Figure 2a](#), we show 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 wait-time option. The elasticities we back out from our model are in line with this observation. Moreover, the magnitudes of these differences vary throughout the day. During work hours, the likelihood of choosing the lowest price option significantly dips, and the likelihood of choosing the shortest wait option increases by even more.⁸ This pattern can be attributed to some combination of preference heterogeneity across consumers as well as within-consumer heterogeneity throughout the day. Because we observe consumer identifiers, our model leverages the variation across consumers to identify individual-level preference heterogeneity.

In [Figure 2b](#), we compare choices over price and wait times by pickup location. We show in [Figure 2b](#) the probability of consumers choosing the lowest price or shortest wait time among all available bids in each pickup location, computed only within auctions that feature a trade-off between prices and wait times. Locations are sorted by the probability of choosing the lowest price. As with [Figure 2a](#), in [Figure 2b](#), we see consumers pick rides with lower prices and shorter wait

⁸These two likelihoods need not add up to 1, because a consumer may choose a driver with neither the lowest price nor the shortest wait time if, for instance, this driver has the highest rating.

Figure 1: Prices and wait Times by Hour and Day

Figure (a) Average Prices





Figure (b) Average Wait Times



NOTE: These figures show the average offer prices (Figure 1a) and average wait times (Figure 1b) across all bids submitted on different days of the week and at different times of day.

times. The figure shows minimum prices are chosen about two to three times more often than minimum wait times, but heterogeneity also exists across locations. While there are differences across locations, these differences are relatively small compared with the intertemporal variation in how consumers solve the price and wait-time trade-off.

Figure 2: Trade-Offs and Choices

Figure (a) By Hour





Figure (b) By Pickup Location



NOTE: These figures show the mean probability of a consumer who faces a trade-off between price and wait time choosing either the lowest price or shortest wait time. Probabilities are computed among orders in which one of the bids was chosen over the outside option. In [Figure 2b](#), we sort locations by the probability of choosing the lowest price.

3 Platform model

We describe the theoretical framework through which we study the consumers' and drivers' choices in the platform. A ridesharing platform connects consumers

and taxi drivers in exchange for a percentage fee on the ride fare. On a given date t and location a , the interaction unfolds as follows: First, a consumer in need of a ride from a requests a ride. Second, the platform sends out ride requests to the drivers based on their proximity to the passenger. Third, conditional on accepting the request, drivers submit a bid for the ride. Fourth, the consumer observes the bids selected by the platform’s algorithm and decides which to accept, if any. Finally, if the consumer accepts one of the bids, the platform collects a 10% fee from the driver. Below, we describe each of these events, working backwards from the consumer’s choice of which ride to accept.

3.1 Demand side and the value of time

Consider a consumer, indexed by i , who submits a ride request between two locations. We summarize the request r by its origin a , its destination \hat{a} , and the date and time it is submitted, t , denoted by $r = (t, a, \hat{a})$. The consumer is then presented with a menu with J_r offers. Each offer $j \in \{1, \dots, J_r\}$ is characterized by its price, b_j , wait time, w_j , and observable trip characteristics, x_j . The observable trip characteristics include request-dependent characteristics common to all drivers—for example, hour of day, public transit availability, traffic speeds, trip distance and time length, rainfall, origin and destination, and whether the order is placed on the street or in a building—as well as driver j ’s characteristics—for example, driver’s name and rating, car year, color, and model (basic, premium, or luxury).⁹ In what follows, we use J_r to indicate both the set of drivers associated with a request and the offer attributes $(b_j, w_j, x_j)_{j \in J_r}$ when it does not risk confusion.

The consumer’s preferences over the tuple (b_j, w_j, x_j) are summarized by (i) a vector of coefficients, $(\beta_{ir}^w, \beta_{ir}^p, \beta^x)$, (ii) a stochastic part ϵ_{ijr} , and (iii) an additional term, ξ_r , that captures unobserved conditions affecting demand on a particular route, such as large sporting events or transit delays.

Formally, consumer i ’s utility from option $j \in J_r$ can be written as

$$u_{ijr} = \beta_{ir}^w w_j + \beta_{ir}^p b_j + \beta^x x_j + \xi_r + \epsilon_{ijr}. \quad (1)$$

Importantly, the coefficients β_{ir}^w and β_{ir}^p are consumer specific. Our specification also includes a utility parameter on the square of the wait time, which we omit from

⁹Premium cars (56% of our sample) include brands such as Audi and Lexus. Luxury cars (10% of our sample) include brands such as Tesla and Ferrari.

Equation 1 for expositional clarity. As we discuss in Section 5.1, the parameter estimate is small and therefore does not play a major role in the analysis.

Consumer choice out of a menu The consumer’s outside option is to reject all the bids in a given request and perhaps take another form of transportation. We normalize the value of this outside option to 0. However, we do allow consumer utility from platform rides to shift according to several factors summarized in x_j , such as the availability of public transit. These factors allow us to control for the shifting value of the inside good relative to the outside option.

Under this normalization and the assumption that ϵ_{ijr} are independently and identically distributed according to a Type I extreme value distribution, the likelihood of consumer i choosing driver j out of menu J_r is given by

$$l_{J_r}(w_j, b_j, x_j, \xi_r; \beta) = \frac{\exp(\beta_{ir}^w w_j + \beta_{ir}^p b_j + \beta^x x_j + \xi_r)}{1 + \sum_{k \in J_r} \exp(\beta_{ir}^w w_k + \beta_{ir}^p b_k + \beta^x x_k + \xi_r)}. \quad (2)$$

Value of time The preference parameters in Equation 1 allow us to describe a consumer’s value of time (VOT) at different locations and different times of day. These VOT are obtained via the following equality, which compares the utility of offer j with the utility of a hypothetical option j' that adds a single minute to the wait time but is otherwise identical. In particular, the time from pickup to the destination is the same for both trips. The difference $b_j - b_{j'}$ that solves the equation reflects the additional units of money needed to make the consumers indifferent between paying more for j or wait more for j' :

$$\beta_{ir}^p b_j + \beta_{ir}^w w_j = \beta_{ir}^p b_{j'} + \beta_{ir}^w (w_j + 1). \quad (3)$$

When comparing trips j and j' , the consumer trades off spending one more minute at the origin (trip j') or at the destination (trip j). Equation 3 implies a minute of time at destination \hat{a} relative to its value at origin a is valued as

$$\text{VOT}_{ir} = b_j - b_{j'} = \frac{\beta_{ir}^w}{\beta_{ir}^p}, \quad (4)$$

where $r = (t, a, \hat{a})$. Equation 4 shows we can recover individual estimates of VOT directly from the estimated demand model by taking a ratio of coefficients.

3.2 Model of driver bidding

We now present the model of driver bidding behavior. We use this model to rationalize the observed bids as arising from drivers’ privately observed opportunity costs of serving a ride. Such costs inform us about the distribution of markups drivers are able to earn and, consequently, how welfare is distributed in the market. When we turn to our analysis of platform pricing, we require these costs as an input to the platform pricing problem and the associated counterfactual pricing and welfare analyses.

Whereas the model that follows is static, we show in [Appendix D.2](#) that it can be microfounded by a dynamic model, that accounts for drivers’ dynamic incentives in the platform, in the spirit Lagos (2000), Buchholz (2022), Brancaccio et al. (2020) and Brancaccio et al. (2023). The model in this section has the advantage of relying on fewer assumptions and, as we argue in [Appendix D.1](#), the costs recovered from this model are sufficient for the counterfactuals we consider.

Consider driver j who is near location a at time t and receives a request $r = (t, a, \hat{a})$. This request is associated with wait time w_j and observable trip characteristics x_j . The wait time w_j together with the trip’s length—included in x_j —determine the number of periods until the passenger is dropped off, $\tau(w_j, x_j)$, which may affect the driver’s cost of serving the ride.

When considering what bid to submit, the driver compares the expected benefit of winning the auction against the opportunity cost of successfully bidding for the trip. The opportunity cost of successfully bidding for the trip summarizes what the driver gives up when serving the trip: the value of remaining at location a at time t net of the value of being at location \hat{a} at time $t + \tau(w_j, x_j)$. The value of remaining at a at time t includes both the current foregone opportunities at that moment—for instance, serving a ride in the taxi market—and the continuation value of remaining at location a at time t . In other words, the opportunity cost of serving the ride is a combination of current and future foregone opportunities, and we refer to it as the driver’s inclusive cost of serving the ride.

Our model of driver bidding behavior captures in reduced form the driver’s inclusive cost of serving a ride. We assume that conditional on winning the auction for a trip request r with characteristics (w_j, x_j) , driver j incurs a cost c_{jr} . In the dynamic microfoundation in [Appendix D.2](#), c_{jr} is a combination of foregone opportunities in location a at time t and foregone continuation values (see [Equa-](#)

tion D.5). As we explain below, c_{jr} alone is enough to determine the driver's optimal bid.

Driver j 's optimal bid then solves

$$\max_b \gamma(b|w_j, x_j, \xi_r) (0.9 b - c_{jr}), \quad (5)$$

where $\gamma(b|w_j, x_j, \xi_r)$ denotes the probability driver j wins the ride when he submits bid b . It follows that driver j 's optimal bid satisfies the first-order condition

$$c_{jr} = 0.9 \left(b + \frac{\gamma(b|w_j, x_j, \xi_r)}{\gamma'(b|w_j, x_j, \xi_r)} \right). \quad (6)$$

The win-probability γ depends on the consumer's preferences over the submitted bids, as well as how many other drivers bid for the ride and their wait times and characteristics. Drivers do not observe anything about the other bidders when submitting their bid. The bids and quality attributes of other drivers, summarized by (c_{-j}, w_{-j}, x_{-j}) , the number of competing drivers, and consumer's preferences are therefore all stochastic from driver j 's perspective. Taking into account the distribution of the number of competing drivers, their bids and wait times, and the consumer's preferences, driver j 's winning probability is given by

$$\gamma(b|w_j, x_j, \xi_r) = \mathbb{E} \left[l_{J_r}(w_j, b_j, x_j, \xi_r; \beta) | j \in J_r \right]. \quad (7)$$

The expectation is therefore taken over (i) the number of competing drivers, (ii) the competing drivers' bids, $b_{j'}$, their costs $c_{j'r}$, and their characteristics $(w_{j'}, x_{j'})$, and (iii) the consumer's preferences as summarized by β and the shocks ϵ .

4 Estimation

4.1 Demand estimation

We now discuss demand estimation details. We estimate a likelihood model based on the individual choices over bids on the app. The model captures two types of heterogeneity in VOT: time- and location-specific heterogeneity, and individual-specific heterogeneity.

To capture common elements of time- and location-specific heterogeneity in VOT, we introduce time and location heterogeneity in price and wait-time coefficients. Specifically, the price coefficient β_{ir}^p can vary between work (9am–6pm) and non-

work hours. Instead, the wait-time coefficient β_{ir}^w can vary across five blocks of time: 1am–5am, 6am–9am, 10am–3pm, 4pm–6pm, and 7pm–12am. We denote by h_t the cell of the partition to which a given date and time t belongs.

To capture individual-specific heterogeneity, we leverage the panel structure of our data to compute random coefficients on both price and wait time using an MCMC procedure. To allow for better interpretation of this heterogeneity, the analysis hereafter only utilizes weekday data.

More specifically, we assume a consumer’s wait-time preference coefficient is additive in time of day, location, and an individual-specific shifter:

$$\beta_{ir}^w = \beta_i^w + \beta_a^w + \beta_{\hat{a}}^w + \beta_{h_t}^w, \quad (8)$$

whenever $r = (t, a, \hat{a})$. Instead, we only allow for minimal variation in price coefficients within an individual consumer. We assume price coefficients are additive in an individual-specific and a time-of-day-specific shifter:¹⁰

$$\beta_{ir}^p = \beta_i^p + \beta_{h_t}^p. \quad (9)$$

We assume the individual-specific shifters (β_i^w, β_i^p) are normally distributed, with mean μ and variance-covariance matrix Σ . The covariance of the individual-specific components captures whether people who are more elastic to wait times are also more elastic to price. Because income utility should be related to the opportunity cost of time, a positive covariance is expected.

Random coefficients logit estimated via MCMC To capture the full heterogeneity in consumer preferences, we exploit the panel structure of our data. We adopt a hierarchical Bayes mixed-logit model to obtain individual-specific estimates for both wait-time and money preferences. We use an MCMC method using data augmentation of latent variables as in Tanner and Wong (1987). In this approach, the unobserved random coefficients are simulated at each iteration. This method sidesteps the need to evaluate multidimensional integrals, by instead sampling from a truncated normal distribution.

Following techniques described in Rossi et al. (2005) and Train (2009), we construct a Gibbs sampler, the details of which are in [Appendix B.2](#). Such an MCMC

¹⁰Allowing the price coefficients to vary across day and night hours allows us to capture, among other observations, that daytime business trips may be reimbursed.

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

Control function Because drivers might condition their bids on the unobserved demand conditions ξ_r , the price coefficients may be biased. Because our model is likelihood based, we use a control function approach to address this issue (Petrin and Train, 2010).

Based on the first-order conditions of the drivers' bidding problem in Equation 6, we can approximate their bids as a function of a driver-specific cost component \bar{c}_j , an order-driver-specific deviation from this average denoted Δc_{jr} , and a function of the demand conditions $g(\xi_r)$:¹¹

$$b_{jr} = \frac{1}{0.9} (\bar{c}_j + \Delta c_{jr}) + g(\xi_r). \quad (10)$$

Our control function approach exploits the variation in persistent cost differences \bar{c}_j across drivers, which we found to be large in the data (see Figure B.3 in Appendix B.2.1).¹² Because the driver selection process is determined by physical proximity, the assignment of drivers to consumers is quasi random, and \bar{c}_j therefore provides the random identifying variation. To implement this approach, we first regress offered bids on a set of driver fixed effects. From this regression, we take the residual and average it within each order to predict $g(\xi_r)$. We then add this predicted value as a control for ξ_r in the consumer's indirect utility function.

We provide in Appendix B.2.1 additional detail on the design of our control function approach. Figure B.3 depicts the resulting distribution of fixed effects, which shows average driver bids have large and persistent variation relative to the overall mean. The interquartile range is \$1.90, or 20% of the average fare, and the range from the 10th to the 90th percentile is \$4.00, or 43% of the average fare. In Table B.1, we also provide results from a Monte Carlo exercise that demonstrates, given our model assumptions, the control function's ability to recover unbiased parameter estimates in the presence of unobserved demand shocks.

¹¹This approximation is based on a Taylor expansion; see Appendix B.2.1.

¹²This approach is similar to that in the literature that exploits different leniency standards of judges, known as the judge design. See, e.g., Waldfogel (1995).

4.2 Driver costs: Identification and estimation details

We now provide details of the estimation of the drivers' costs. First, recall that the problem faced by the drivers can be mapped into a static auction (Jofre-Bonet and Pesendorfer, 2003), except that the winning probabilities are determined by the consumer's choices. Second, given our demand estimates, we show we can use techniques similar to those in Guerre et al. (2000) to back out the drivers' inclusive costs that rationalize the observed bids.

Recall that driver j 's problem when faced with request $r = (t, a, \hat{a})$ with characteristics (w_j, x_j) delivers the following first-order condition¹³:

$$c_{jr} = 0.9 \left(b + \frac{\gamma(b|w_j, x_j, \xi_r)}{\gamma'(b|w_j, x_j, \xi_r)} \right). \quad (11)$$

Equation 11 implies the distribution of inclusive costs is non-parametrically identified from the demand estimates and the drivers' bids. Indeed, we directly observe the bid on the right-hand side of Equation 11. Furthermore, we can compute the sample analogues of the win probability $\gamma(\cdot|w_j, x_j, \xi_r)$ and its derivative $\gamma'(\cdot|w_j, x_j, \xi_r)$ with estimates from the demand system (cf. Equation 7), so that all the objects on the right-hand side are observed.

Using Equation 11, we can simply back out driver j 's inclusive cost bid by bid. However, for each order, we only observe one realization of competing bids. To compute $\gamma(\cdot|w_j, x_j)$, which includes drivers' time- and location-dependent expectations about competitors and consumers, we sample from the observed distribution of bids and consumers (Hortaçsu and McAdams, 2010). For each bid observed in the data, we simulate 50 requests. For each request, we draw a consumer-preference vector β_i and a set of competing drivers. To simulate the correct conditional expectation, the bids and attributes of competing drivers are drawn conditional on each location and time period, as well as the trip's length. Given these simulation draws, we construct the sample analogues of $\gamma(\cdot|w_j, x_j, \xi_r)$ and $\gamma'(\cdot|w_j, x_j, \xi_r)$.

Remark 1 (From inclusive costs to flow costs and continuation values). *As we explain in Section 6, the identification of the drivers' inclusive costs, which aggregate the drivers' on- and off-platform opportunity costs over time, is enough*

¹³We assume a connected bid space. This is an approximation because drivers in reality bid in small increments of several cents.

for our counterfactual exercises. In many applications, separately identifying the on- and off-platform components of the costs and further decomposing them into their per-period components may be useful. For this reason, after introducing our dynamic microfoundation for the drivers' model in [Appendix D.2](#), we show both that the model is identified in [Appendix D.3](#) and that all driver primitives can be recovered through a simple regression of the inclusive costs on a set of time- and location-specific dummies. In particular, we show in [Proposition 1](#) that the driver's current foregone opportunities and all location- and time-dependent expectations that give rise to c_{jr} are separately identified.

5 Results

We first present the results from our demand estimation and the implied elasticities for both wait time and price. We then show the VOT results implied by the demand estimates. Lastly, we present results on the estimated costs of drivers.

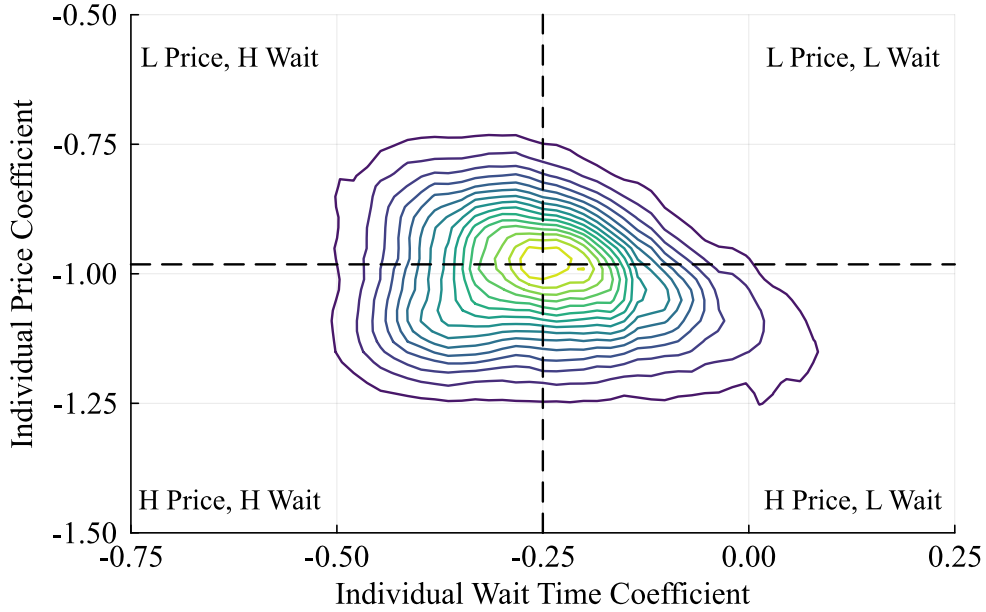
5.1 Demand results

Our results include a set of estimated utility parameters for price, wait time, and additional shifters that are common to all consumers, as well as individual-specific preference estimates for price and wait time. We start by reporting the preference parameters that are individual-specific, and then report those parameters that are common.

[Figure 3](#) shows a contour plot of the joint distribution of individual consumer preferences over prices, β_{ir}^p , and wait times, β_{ir}^w . Preferences are heterogenous along both dimensions. A slight negative dependence also exists between price and wait-time coefficients. For instance, the average price coefficient of consumers with wait-time coefficients below the median is -1.11 and -1.17 for wait-time coefficients above the median. The graph also shows our classification into four types of individuals, depending on whether they have high (H) or low (L) sensitivity to price and wait time. To do so, we split consumers along the median of the respective distributions of coefficients. We use this categorization to present elasticity comparisons below.

In [Table 3](#), we report the coefficients that are common to all consumers, with standard errors. We find the common part of the price coefficient varies minimally across working and non-working hours. The intra-daily variation in wait-time

Figure 3: Individual-Specific Preference Estimates



NOTE: This figure provides a contour plot of the kernel density of individual estimates of β_{ir}^w and β_{ir}^p . Each quadrant defines a high (H) and low (L) relative sensitivity to wait time and price, used in analysis below.

coefficients shows people are the most wait-time sensitive at night and in the early morning hours, and the least wait-time sensitive during the middle of the day. We find the coefficient on wait-time squared is close to 0, which is consistent with what we see in the data: in [Figure B.2](#) in [Appendix B.1](#), we show that the likelihood of picking a particular trip is close to a linear function of the wait time and of the minimum wait time. We also find variation in wait-time sensitivity across locations, although these differences are much smaller than the aforementioned variation across times of day. Due to the large number of location coefficients, we summarize the full set of location- and time-of-day-specific estimates separately in [Figure B.4](#) and [Figure B.5](#).

Additional coefficients measure an interaction effect between wait time and other indicator variables that impact the value of a trip compared with the outside option: public-transit availability, whether the trip is ordered on the street, and the presence of rain in the hour the trip was ordered. These environmental factors are relatively small but significant, with the marginal effects of each implying less than a 1% decrease in the probability of choosing a ride on the platform.

Beyond recovering heterogeneity in wait-time and price preferences, we also es-

timate rich substitution patterns dependent on driver-specific factors. Both the driver rating and the car type significantly affect consumer choices. Our estimates suggest consumers value each rating point by \$0.10. Consumers have an average willingness-to-pay of \$0.64 for luxury cars and \$0.22 for premium cars, compared with basic cars.

Table 3: Common-Preference Estimates

Description	Coefficient	Std Error
Price 6pm-6am	-1.18	0.002
Price 6am-6pm	-1.167	0.003
Wait Time 1am-5am	-0.027	0.011
Wait Time 6am-9am	-0.088	0.011
Wait Time 10am-3pm	-0.098	0.01
Wait Time 4pm-6pm	-0.069	0.01
Wait Time 7pm-11pm	-0.036	0.01
Wait \times On-Street Order	-0.035	0.002
Wait \times Raining	0.012	0.004
Wait Time Squared	-0.006	0.001
Driver Rating Points	11.097	0.097
Car: Mid Quality	0.245	0.005
Car: High Quality	0.714	0.009
Trip Speed	-0.052	0.002
Alt. Transit Available	0.038	0.007
Order on Street	0.204	0.014
Rain	-0.106	0.034
Trip Distance	5.482	0.071
Waiting \times Pickup Location FE 1-30	✓	✓
Waiting \times Dropoff Location FE 1-30	✓	✓
Pickup Location FE 1-30	✓	✓
Dropoff Location FE 1-30	✓	✓
Hour FE	✓	✓

NOTE: This table provides coefficient estimates and standard errors from the logit demand model for each consumer type. Estimates are conditioned on 60 additional wait-time interactions and 66 additional fixed effects. Additional details for these values are in [Figure B.4](#). These parameter estimates comprise outside option shifters and wait-time preference interactions with each of 30 pickup and dropoff locations as defined in [Appendix A.2](#). The omitted results are instead depicted graphically in [Figure B.4](#).

We now turn to discuss what these estimates imply for the sensitivity of demand to changes in price and wait-time. In [Table 4](#), we show price and wait-time elasticities respectively, and also a set of order-level elasticities, which measure only substitution to the outside option when either all prices or all wait times change

Table 4: Estimated Elasticities

Time of Day	Type Partition	Bid-Level Elasticities		Order-Level Elasticities	
		Price	Wait Time	Price	Wait Time
Daytime 6am-6pm	Overall	-4.35	-1.02	-3.98	-0.92
	H Price, H Wait	-7.68	-1.71	-6.68	-1.46
	H Price, L Wait	-3.04	-0.87	-3.06	-0.79
	L Price, H Wait	-4.93	-1.02	-4.3	-0.95
	L Price, L Wait	-1.96	-0.5	-2.04	-0.5
Evening 6pm-6am	Overall	-4.92	-0.46	-4.53	-0.46
	H Price, H Wait	-8.4	-0.79	-7.26	-0.75
	H Price, L Wait	-2.93	-0.32	-3.12	-0.33
	L Price, H Wait	-6.16	-0.52	-5.35	-0.52
	L Price, L Wait	-2.37	-0.21	-2.54	-0.23

NOTE: This table provides the demand elasticity of price and wait 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 wait-time sensitivity. The first two columns show these elasticities among competing offers, reflecting the change in demand of a particular offer due to a 1% change in that offer's price or wait time. The second two columns show the elasticities with respect to choosing the outside option, reflecting a change in demand for all offers due to a 1% change in price or wait time on *all* offers.

by a small amount. We see a general pattern: consumers are much more price elastic than wait-time elastic. Price elasticities range from four to 11 times higher than wait-time elasticities, with starker differences in the evening. Consumers have highly heterogeneous elasticities: between the two extreme groups (i.e., H Price, H Wait and L Price, L Wait), both price and wait-time elasticities differ by about a factor of four.

These elasticity estimates convey that both price and wait time are important factors in the consumers' decisions, and that wait-time elasticities vary throughout the day in ways that reflect the patterns of consumer choices we see in Figure 2.¹⁴

5.2 Value of time results

We now present results on the VOT implied by our estimates, scaled to USD per hour. We compute VOT using the coefficients in Table 3 together with Equation 4. We summarize the results in Table 5. The overall mean VOT across all trips,

¹⁴We can also decompose elasticities by trip origins and destinations as we have done in Table B.3. Broadly similar patterns between demand types are revealed, though each elasticity measure varies from one location to another. In general, price elasticities are more variable than wait-time elasticities.

expressed as an hourly quantity, is \$13.21. We find significant individual-, time-of-day-, and location-specific heterogeneity underlying this average. The most prominent of the three is the heterogeneity across consumers. As before, we report four groups of individuals: those with above- and below-median random coefficient estimates on both price and wait-time preferences. The low-price-sensitivity and high-wait-time sensitivity group exhibits VOT nearly twice the overall average at \$22.08 per hour, whereas individuals with high sensitivity to price and low sensitivity to wait time have an average VOT of \$4.95 per hour. All groups have similar time-of-day patterns, with the highest values in the morning between 6am and 9am. We find VOT estimates are higher in the late morning and mid-day hours than in the evening and overnight.

Finally, we report results separately for Prague’s city center and the city periphery. We provide definitions of these regions in [Figure A.2](#). We find the VOT for trips in the city center is higher than for trips in the city periphery. Although variation in VOT is meaningful both spatially and intertemporally, variation in the individual component of the VOT explains by far the largest share of its overall variation.

In contrast to earnings data from the Czech Statistical Office, our estimates of VOT are larger than the average wage in Prague, which is approximately \$9.50 per hour during the sample period. This finding is perhaps not surprising, because taxis are a relatively expensive mode of transit, and thus, taxi riders are likely positively selected on income.¹⁵ Even for this selected set of riders, we find large heterogeneity in their VOT. As we illustrate in the next section, this heterogeneity has important implications for the pricing counterfactuals.

Because in our discrete-choice model, preference parameters are individual but not trip-specific, our VOT estimates for a given consumer are obtained by averaging over all the trips the consumer takes. However, a given consumer may have a higher or lower VOT depending on the circumstances surrounding a ride (e.g., more or less urgent trips), and our VOT estimates average across all these circumstances. In [Appendix B.5](#), we demonstrate our data can be used to recover different VOT estimates between trips with a drop-off time close to the start of a new hour, which are more likely to involve deadlines, and those at different times. This analysis provides an example of how we might recover different types of trip-specific

¹⁵A platform-conducted survey about riders’ wage rates shows the average wage among the respondents is \$15.23, which is higher than the average wage of \$9.15 in Prague at the time. This finding confirms that consumers on the platform are positively selected in terms of income, as they are on other major ride-hail platforms.

Table 5: Value-of-Time Estimates

Subsample	Value of Time (VOT)					
	12a–6a	6a–9a	10a–2p	3p–6p	7p–12a	All Hours
All Types	12.39 (0.47)	15.03 (0.43)	15.69 (0.36)	13.36 (0.38)	11.58 (0.52)	13.21 (0.4)
H Price, H Wait	15.93 (0.09)	18.11 (0.08)	18.8 (0.08)	16.73 (0.07)	15.19 (0.07)	16.61 (0.05)
H Price, L Wait	4.33 (0.06)	6.59 (0.05)	7.07 (0.06)	5.0 (0.05)	3.54 (0.04)	4.95 (0.04)
L Price, H Wait	21.19 (0.51)	24.22 (0.43)	25.05 (0.44)	22.27 (0.53)	20.09 (0.52)	22.08 (0.49)
L Price, L Wait	8.07 (0.9)	10.89 (0.8)	11.61 (0.74)	9.18 (0.9)	7.4 (1.09)	9.04 (0.82)
City-Center Trips	12.79 (0.48)	15.56 (0.42)	16.23 (0.39)	13.79 (0.37)	12.03 (0.51)	13.66 (0.41)
Non-City-Center Trips	9.79 (0.62)	11.75 (0.53)	12.15 (0.36)	10.15 (0.39)	8.27 (0.56)	10.05 (0.5)

NOTE: This table provides VOT estimates implied by the logit demand model. All estimates are presented in USD. We report bootstrap standard errors in parentheses, based on 100 bootstrap iterations.

heterogeneity that is otherwise averaged in the results we report in Table 5.

5.3 Supply results

In Table 6, we summarize our estimates of the drivers' inclusive cost. We find drivers earn rents over their inclusive costs as indicated by markups of around 30%. Winning drivers bid about \$1.36 lower on average than remaining drivers and have costs that are on average \$0.75 lower than the remaining drivers. We also find drivers with shorter wait times bid higher and, conditional on winning, earn higher markups than drivers with higher wait times. This finding suggests drivers are aware that short wait times are a quality attribute that makes their overall bid more competitive. The markup of drivers with the shortest wait time is more than 14% higher than the markup of the remaining drivers in that order. Within an order, the bid of the shortest-wait-time driver has a higher markup than that of the driver with the highest rating. Interestingly, drivers with the highest rating bid 58 cents less, which is about 6.1% lower. These results suggest the wait-time premium is substantial and that high ratings are maintained in part by offering lower-than-average prices. This finding contrasts with more common

Table 6: Driver Cost Estimates

	Mean	Median	P25	P75	S.D.
Bid (\$), winner	8.75	8.0	6.35	10.56	3.26
Bid (\$), shortest wait	10.21	10.08	8.36	11.88	2.68
Bid (\$), highest rating	9.53	9.43	7.8	11.09	2.49
Bid (\$), all bidders	10.11	9.29	7.08	12.37	4.02
Cost (\$), winner	6.28	6.09	4.86	7.47	1.94
Cost (\$), all bidders	7.03	6.81	5.37	8.43	2.26
Markups winner	0.37	0.32	0.21	0.48	0.21
Markups all bidders	0.29	0.23	0.15	0.37	0.19
Markups shortest wait	0.33	0.28	0.18	0.44	0.21
Markups highest rating	0.33	0.27	0.17	0.43	0.21

NOTE: This table provides a breakdown of markups and cost estimates that we recover from the driver model. The table also shows summary statistics of prices for each of the breakdowns.

services such as Uber, where drivers have no discretion over pricing.

The costs and markups that we recover exhibit large variance. This finding suggests the platform’s bidding mechanism plays an important role in discovering the lowest-cost drivers. We break down driver costs by hour and by location in [Figure 4](#), where we show some systematic variation in the average cost. However, these averages by location and hour mask large variations across drivers. The gray shaded area represents one standard deviation above and below the average cost. We find that idiosyncratic cost variation is large relative to predictable variation in costs due to either time-of-day or location. In [Figure C.1](#) in [Appendix C.1](#), we show that markups are highest during the day, especially in the early morning hours.

Figure 4: Cost Heterogeneity

Figure (a) By Hour


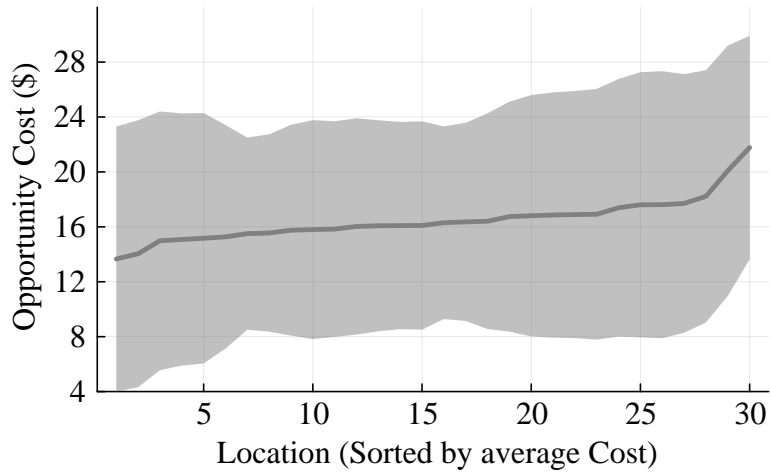


Figure (b) By Location



NOTE: In these figures, the gray lines denote the average of estimated driver costs by hour (Figure 4a) and location (Figure 4b). The gray shaded area represents one standard deviation above and below the average cost. In Figure 4b, we order locations by average driver cost to better highlight the cost heterogeneity.

6 Pricing the value of time

Our estimates demonstrate large heterogeneity in both consumer VOT and driver costs. In this section, we evaluate the welfare effects of counterfactual platform pricing strategies that take advantage of the latent individual-specific variation in consumer preferences, while also respecting drivers' participation constraints.

We focus our analysis on studying *personalized pricing*, including its effects on consumers, drivers, and the platform. This focus leverages the panel dimension of our consumer data and complements prior work on livery vehicles, which has investigated the pricing of aggregate or observable sources of demand variation (Buchholz, 2022; Castillo, 2019; Rosaia, 2020). Understanding this type of pricing is increasingly relevant as ride-hail and other platforms turn to more sophisticated pricing strategies, building on a wealth of accumulated consumer data.

In what follows, we first describe how we set up the platform’s problem. In the counterfactuals, we allow the platform to *decouple* prices on the drivers’ and consumers’ side, which is in contrast to the baseline in which Liftago charges the winning driver a fixed 10% fee and implements that driver’s bid as a price the consumer must pay. We now allow the platform to choose a price for each option the consumer faces and separately choose driver payments for each option, so it is able to influence not only how much the consumer pays and the driver receives, but also whom the consumer matches with.

Direct mechanism approach We model the platform’s pricing and matching policy as arising from a direct mechanism that takes as inputs the drivers’ costs and outputs a menu from which the consumer chooses. This approach has the advantage that the consumers’ and drivers’ behavior are captured by a set of (mostly linear) constraints that the platform’s mechanism must satisfy, and therefore does not require the computation of equilibria.¹⁶ We assume the platform offers the same drivers to the consumer as we observe in the data.

Formally, for a given trip request r and for a given set of drivers J_r , the platform chooses a pair of transfers $t_j^r(c_1, \dots, c_{J_r})$ and prices $p_j^r(c_1, \dots, c_{J_r})$ for each driver $j \in J_r$ and each profile of drivers’ costs $\bar{c}_r \equiv (c_1, \dots, c_{J_r})$. Tariff $t_j^r(\bar{c}_r)$ represents the payment driver j receives and $p_j^r(\bar{c}_r)$ represents the price the consumer pays for a ride with driver j . Whereas the transfers and prices may depend on observable trip characteristics such as the drivers’ wait times, we suppress this dependence from the notation for ease of exposition.

Because the platform does not know the consumer’s logit shocks, the platform’s menu offer $(p_j^r(\bar{c}_r), w_j)_{j \in J_r}$ leads to a set of choice probabilities. As in [Section 3.1](#), these probabilities are given by the likelihood in [Equation 1](#), replacing the drivers’ bids by the platform-chosen prices.

¹⁶For this reason, this approach can be exploited to characterize the platform’s pricing policy that maximizes a combination of platform profits and consumer and driver welfare.

Given consumers' prices and drivers' transfers, the platform's profits are

$$\Pi(t^r, p^r; \beta, \bar{c}_r) = \sum_{j=1}^{J_r} (p_j^r(c_j, c_{-j}) l_{J_r}(p_j(c_j, c_{-j}), \cdot; \beta) - t_j^r(c_j, c_{-j})).$$

That is, the platform collects price p_j^r from the consumer if the consumer selects driver $j \in J^r$ and makes an expected payment of t_j^r to driver j , where the expectation is relative to the probability of serving the ride.¹⁷

Finally, we specify the information the platform has about consumers and drivers when choosing its pricing policy. On the consumer side, our counterfactuals consider different scenarios for what the platform knows about the individual-specific components of β . Below, we denote the platform's information about consumers by \mathcal{I} . On the drivers' side, we assume the platform knows the drivers' cost distribution, but not the drivers' realized costs. Thus, when choosing its pricing policy, the platform needs to satisfy each driver's participation and incentive compatibility constraint. Because driver j knows their own cost and the trip request, but not how many other drivers' the platform requests and their costs, driver j faces an expected transfer $T_j(\hat{c}_j)$ and an expected probability of being assigned to the consumer $L_j(\hat{c}_j)$ when their submitted costs are \hat{c}_j . Thus, the platform's policy must satisfy the following constraints for each driver j , cost c_j , and reported cost \hat{c}_j :

$$\begin{aligned} T_j(c_j) - c_j L_j(c_j) &\geq 0, & (\text{PC}_j(c_j)) \\ T_j(c_j) - c_j L_j(c_j) &\geq T_j(\hat{c}_j) - c_j L_j(\hat{c}_j). & (\text{IC}_j(c_j, \hat{c}_j)) \end{aligned}$$

The first constraint, $\text{PC}_j(c_j)$, implies the platform must cover each driver's inclusive cost of serving the ride, which accounts for the foregone current and future opportunities both on and off the platform. The second constraint, $\text{IC}_j(c_j, \hat{c}_j)$, captures that the platform does not know the drivers' opportunity costs and must respect drivers' incentives when selecting the transfers.

The platform then chooses transfers and prices (t_j, p_j) to solve

$$\begin{aligned} \max_{p^r, t^r} \mathbb{E}_{\beta, \bar{c}} [\Pi(t^r, p^r; \beta, \bar{c}_r) | \mathcal{I}] & & (\Pi(\mathcal{I})) \\ \text{s.t. } \text{PC}_j(c_j), \text{IC}_j(c_j, \hat{c}_j) & \text{ for all } j \in J_r, c_j, \text{ and } \hat{c}_j. \end{aligned}$$

¹⁷That is, $t_j(\bar{c}_r) = l_{J_r}(p_j(\bar{c}_r), w_j, x_j; \beta_i) \cdot \tilde{t}_j(\bar{c}_r)$, where \tilde{t}_j are the payments conditional on serving the ride and $l_{J_r}(\cdot)$ is the likelihood the driver is chosen out of the menu.

Given the above formulation, one remark is in order about the model of platform pricing. The constraints the platform faces in terms of drivers' behavior are expressed in terms of the drivers' inclusive costs, which aggregate their foregone on- and off-platform opportunities at the baseline. Whereas changing the platform pricing may not affect the value of the off-platform opportunities, it does affect the value of the on-platform ones and the drivers' choices between on- and off-platform activities, both of which affect their inclusive costs. As we show in Table 2 and discuss further in Appendix D.4, the data suggest on-platform activities represent only a very small portion of drivers' overall daily earnings, which has two implications. First, the inclusive costs we recover are largely driven by the foregone off-platform opportunities. Second, changes to the platform's pricing policy will have small effects on the drivers' inclusive costs. It follows that in our setting, ignoring the changes in the drivers' opportunity costs is not a first-order concern. Nevertheless, our identification and estimation results for the dynamic microfoundation of the drivers' model provide the necessary machinery to re-estimate the drivers' costs as the platform's policy changes at the cost of more assumptions than those needed to recover the costs c_j . In Appendix D.4, we utilize this machinery to illustrate our results are robust to changes in the drivers' continuation values as a result of the platform's pricing policy change.

Counterfactual description Our personalized pricing counterfactuals distinguish between the case in which the platform knows β_i^w and the case in which it knows both β_i^w and β_i^p . In each case, the platform sets prices conditional on both the drivers' costs and its knowledge of individual consumer preferences. We contrast the results of the personalized pricing counterfactuals with two other scenarios:

First, we consider a *uniform pricing* counterfactual, in which the platform sets prices on both sides of the market conditional on drivers' costs, but not on individual consumer preferences. The comparison between the uniform pricing and baseline results explains what portion of the welfare and profit effects of personalized pricing are a consequence of the platform's ability to set prices on both sides of the market. Instead, the comparison between the uniform and personalized pricing results explains what portion of the welfare and profit effects of personalized pricing are a consequence of the platform using information about consumer preferences in its pricing.

Second, we examine *ETA-based pricing*, which is another non-personalized pricing

policy in which the platform conditions prices both on drivers’ wait times and costs. That is, under ETA-based pricing, the platform may assign different prices to drivers with the same reported costs if they have different wait times. ETA-based pricing is a useful benchmark both because Uber and Lyft are already engaging in this type of pricing and because platforms may be reluctant to use consumer order histories to estimate β_i^w and β_i^p due to consumer backlash. Wait-time pricing may then give consumers incentives to reveal their high wait-time sensitivity by picking a driver with shorter wait time and higher price.

Given our focus on unobserved sources of demand variation at the individual level, we conduct our counterfactuals for a subset of trips that are observationally similar. Specifically, we select trips from the modal origin to the modal destination between the highest-volume hours of 7pm and 12am.¹⁸

6.1 Counterfactual results

We now discuss our counterfactual results, which we summarize in [Table 7](#). Relative to the baseline, personalized pricing on both wait-time and price coefficients leads to a three-fold increase in the platform’s profits, a substantial part of which is due to an almost four-fold decrease in drivers’ profits. Consumers face higher prices relative to the baseline and hence, the share of rides in the platform decreases by 20%. Higher prices also lead to a fall in consumer welfare of 40%. The effects of personalized pricing on the wait-time coefficients are more modest, but still substantial. Notably, because the platform can now segment the market on less variables, the share of rides in the platform decreases by 7% and drivers’ profits fall by 9% relative to personalized pricing on both coefficients. Interestingly, consumer welfare is higher when the platform conditions its pricing policy on the wait-time coefficient alone, despite average fares being higher than under personalized pricing on both coefficients.

The welfare and profit consequences of personalized pricing combine both the platform’s ability to exert its price-setting power *and* its use of consumer information. Comparing the baseline welfare and profits to those under uniform pricing reveal how much of the changes in welfare and profits are due to the platform’s market power. We document that the value to the platform of exercising its market power and directly setting prices on both sides is substantial. Relative to the baseline,

¹⁸The implications of our computations are qualitatively unchanged if we focus on different subsets of trips.

Table 7: Platform Pricing: Welfare Results

	Baseline	Platform Pricing			
	10% flat fee	uniform pricing	wait coef. only	wait, price coef.	ETA pricing
Inside Option Share	0.632	0.472	0.477	0.508	0.485
<i>Δ% from uniform</i>	+34.03%	–	+1.08%	+7.64%	+2.79%
Average Fare	\$6.144	\$6.646	\$6.635	\$6.574	\$6.438
<i>Δ% from uniform</i>	-7.55%	–	-0.18%	-1.08%	-3.13%
Platform Profit	\$0.419M	\$1.175M	\$1.189M	\$1.279M	\$1.185M
<i>Δ% from uniform</i>	-64.31%	–	+1.12%	+8.86%	+0.83%
Platform Revenue	\$2.123M	\$2.138M	\$2.17M	\$2.327M	\$2.166M
<i>Δ% from uniform</i>	-0.71%	–	+1.46%	+8.82%	+1.27%
Driver Profit	\$1.08M	\$0.256M	\$0.26M	\$0.284M	\$0.23M
<i>Δ% from uniform</i>	+321.59%	–	+1.52%	+10.98%	-10.12%
Cons. Surplus	\$1.139M	\$0.744M	\$0.743M	\$0.725M	\$0.776M
<i>Δ% from uniform</i>	+53.07%	–	-0.15%	-2.53%	+4.3%
Total Welfare	\$2.639M	\$2.176M	\$2.192M	\$2.289M	\$2.192M
<i>Δ% from uniform</i>	+21.28%	–	+0.73%	+5.21%	+0.73%

NOTE: This table summarizes the results of the pricing counterfactuals. All values estimate the welfare achieved among the subset of rides of the most common origin, destination and time of day in our sample. The first column, Baseline, denotes the welfare effects of the status-quo 10% fee structure. The second column, uniform pricing, shows the effects of optimal platform pricing conditional on knowledge of the distribution of consumer preferences and driver costs alone. The third and fourth columns show the effects of optimal platform pricing conditional on, respectively, knowledge of only consumers' individual preference for wait time, and full knowledge of individual preferences for wait time and price. The last column, ETA-based pricing, shows the effects of optimal platform pricing similar to uniform pricing but now also setting prices according to drivers' ETA and costs. Percentage changes are shown relative to the uniform pricing counterfactual.

the platform raises prices for consumers while lowering the transfers to drivers, thereby appropriating surplus from both market sides. Driver profits relative to their outside option fall by 78%, and consumer surplus falls by 34%. The platform's pricing leads to a significant quantity distortion with 52% of consumers choosing the outside option compared with 36% in the baseline. One major take-away is that the platform leaves substantial profits on the table by operating passively through the 10% fee rule. By exercising market power over both consumers and drivers in the form of centralizing and decoupling prices, the platform

profits increase even without engaging in price discrimination.¹⁹ Relative to uniform pricing, when the platform just uses knowledge of the individual wait-time coefficient, profits increase by about 1% and by more than 8% when it conditions prices on both wait-time and price coefficients.

Another important takeaway from the comparison between the baseline and uniform pricing is that the losses in consumer surplus from personalized pricing are mostly due to the platform’s ability to centralize pricing and not its knowledge of consumer preferences. Relative to uniform pricing, the loss of consumer surplus due to both forms of personalized pricing are modest. Furthermore, personalized pricing on the wait-time coefficient expands the inside market share by 1.1%, whereas personalized pricing on both coefficients expands the inside market share by 7.6%. In both cases, average prices are slightly lower than uniform pricing. As we detail further below, the most price and wait time-sensitive consumers tend to benefit through lower prices, whereas the least sensitive consumers face higher prices.

A final takeaway from the comparison between the baseline and uniform pricing relates to the effects of personalized pricing on drivers’ profits. This question is an important one that is relevant in many e-commerce settings, in which platforms typically have much more information about consumers than suppliers. Relative to uniform pricing, we document that in percentage terms, drivers benefit even more from the platform itself from the platform’s knowledge of consumer preferences. When the platform only conditions on the wait-time coefficient, driver profits are 1.5% higher than under uniform pricing. When the platform conditions on both the wait-time and price coefficients, driver profits increase by more than 11%. Overall, personalized pricing increases surplus relative to uniform pricing: when the platform uses knowledge of both price and wait-time coefficients, total surplus goes up by more than 5%.

In practice, platforms may choose not to use individual order histories to learn consumers’ preferences. In this case, offering different prices for different wait times can be used as an alternative to give consumers incentives to reveal and act on their preferences. We therefore consider ETA-based pricing. We find ETA-based pricing leads to a 0.82% increase of profits relative to uniform pricing, which is about two-thirds of the profit increase from pricing on wait-time coefficients.

¹⁹A potential explanation for the low pricing we observe is that it serves as a kind of investment in market share and long-run adoption, a facet of the platform optimization problem we do not address.

This finding suggests ETA-based pricing strategies can capture a large share of the profits of pricing strategies that directly rely on consumer-preference data. We also see overall welfare under ETA-based pricing is close to the welfare of personalized pricing on the wait-time coefficient. Conversely, such menu pricing does not attain the platform and driver profits found under full personalized pricing, because it cannot effectively capture differences in underlying price elasticities.

Table 8 summarizes the distribution of prices under each counterfactual pricing regime. Moving to any platform-based pricing raises overall prices almost everywhere in the distribution. Relative to uniform pricing, sales prices become much more dispersed under personalized pricing, and increasingly so with full personalization (i.e., on both consumer coefficients). Under fully personalized prices, the 10th percentile of the price distribution is slightly lower than the baseline. We offer additional insight on the prevailing price and wait-time distributions in Appendix D.5 and show personalized pricing does little to sort wait time-sensitive consumers to lower wait-time rides, because these preferences are instead internalized by the platform through differential pricing.

Table 8: Exploring the Price Distribution

	Q10	Q25	Q50	Q75	Q90
<i>Baseline (10% Flat Fee)</i>	\$5.62	\$5.74	\$6.24	\$6.47	\$6.83
<i>Uniform Pricing</i>	\$6.01	\$6.05	\$6.18	\$6.96	\$7.59
<i>Personalized (β_i^w only)</i>	\$5.88	\$6.05	\$6.35	\$7.13	\$7.71
<i>Personalized (β_i^w and β_i^p)</i>	\$5.46	\$5.86	\$6.6	\$7.34	\$7.9
<i>ETA Pricing</i>	\$5.87	\$5.9	\$6.14	\$7.14	\$7.41

NOTE: This table summarizes the distribution of sales prices across pricing counterfactuals. Price variation comes from differences in winning bids in the baseline and differences in transaction prices in each counterfactual according to each pricing regime. All estimates correspond to the welfare analysis in Table 7.

Distributional effects Personalized pricing has interesting distributional implications. In Figure 5a we show the consumer-surplus effects of personalized pricing relative to uniform pricing across different consumer groups. We show the extent to which preference heterogeneity determines the “winners and losers” of personalized prices. Although consumer surplus falls, we find that the majority (62.5%) of consumers benefit from personalization. For the most sensitive consumers in the

Figure 5: Distributional Effects of Policies

Figure (a) Consumer Surplus Effects by Individual Preference Type





Figure (b) Consumer Surplus Effects by Consumer VOT type



NOTE: These figures show the consumer-surplus effects of moving from platform pricing with no personalization to personalized platform pricing. In Panel 5a, we show these effects across the joint distribution of consumers' individual preferences for price and wait time expressed in quartiles of the absolute value of preferences. Thus, larger numbers indicate higher disutility. In Panel 5b, we repeat the exercise by collapsing individual preferences to individual VOT distributions.

upper-right cell, surplus increases nearly three-fold once prices are personalized. The reason is that the platform lowers prices enough to grow participation among the more elastic consumers. At the same time, for the least sensitive consumers in the bottom-left cell, surplus is cut in half because the platform can now isolate

and mark up rides to only these consumers.

In [Figure 5b](#), we explore the role of VOT heterogeneity. Because VOT is defined as the ratio $\beta_{ri}^w/\beta_{ri}^p$, a high-VOT individual may have a relatively high sensitivity to wait time and an average price sensitivity, or an average wait-time sensitivity and a relatively low price sensitivity. As a result, the surplus patterns are not monotone in VOT.

7 Conclusions

The demand for transportation services depends on how consumers trade off time and money. Ride-hailing platforms commonly offer tailored options for different wait times, indicating an interest in using data to implement some degree of price discrimination. Given the vast and growing stores of consumer data, we ask how such platforms and both sides of the market, consumer and driver, would be affected by setting personalized prices that are tailored according to the platform’s estimates of individual preferences for time and money.

We use panel data from a large European ride-share platform that offers menus with explicit trade-offs between time and money. This unique feature allows us to estimate a demand model based on choices from these menus and recover consumers’ preferences over time and money, as well as their implied willingness to pay to reduce wait time. By observing the same consumers over time, we are able to recover individual-level heterogeneity in these estimates.

Our demand model results reveal noteworthy patterns in how individuals value time and money. Consumers are substantially more price elastic than wait-time elastic. Large variation exists in the willingness to pay for wait-time reductions in the population of riders. We show how these differences vary throughout the day; for instance, we show the willingness to pay for lower wait times is higher during work hours. Most of this variation is driven by latent demand characteristics across consumers as opposed to observable differences.

We also exploit the auction format to estimate drivers’ opportunity costs. We demonstrate that potentially complex and dynamic bidding decisions can be transformed into a static auction, allowing us to recover driver opportunity costs through an estimator in the spirit of [Guerre et al. \(2000\)](#). On the driver side, we show large variation in both costs and markups. Thus, in our counterfactuals, it is important that the platform is able to account for these varied opportunity

costs when setting prices directly.

Given our estimated variation in individual consumer preferences and driver costs, we conduct a counterfactual analysis to quantify the welfare implications of personalized pricing. We model the platform’s pricing and matching policy as arising from a direct mechanism that takes as inputs drivers’ costs and outputs a menu from which the consumer chooses. This approach allows us to study the platform’s optimal pricing policies in a tractable way that avoids the need to compute equilibria among the two sides of the market.

Our counterfactual results show the platform does not exercise the full extent of its market power. Conditional on the platform setting prices on both sides of the market, we show that pricing policies such as personalized pricing that leverage consumer data are welfare-improving relative to those that do not. Indeed, personalized pricing leads to increased welfare relative to uniform pricing due to lower average prices and a small market expansion. Under personalized pricing, the platform, the drivers, and most consumers benefit at the expense of the least elastic consumers. Nevertheless, the net beneficiaries are the platform and drivers, because consumers incur small overall welfare losses. These results highlight potentially interesting distributional implications of pricing policies that rely heavily on consumer data. Moreover, we show that non-personalized approaches, such as ETA-based pricing, may achieve some of these welfare gains.

References

- AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): “Monopoly price discrimination and demand curvature,” *American Economic Review*, 100, 1601–15.
- ALI, S. N., G. LEWIS, AND S. VASSERMAN (2022): “Voluntary Disclosure and Personalized Pricing,” *The Review of Economic Studies*, 90, 538–571.
- ARIDOR, G., Y.-K. CHE, AND T. SALZ (2023): “The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR,” *RAND Journal of Economics*, 54, 695–730.
- BAUNER, C. (2015): “Mechanism choice and the buy-it-now auction: A structural model of competing buyers and sellers,” *International Journal of Industrial Organization*, 38, 19–31.
- BENTO, A., K. ROTH, AND A. R. WAXMAN (2020): “Avoiding traffic congestion externalities? The value of urgency,” *National Bureau of Economic Research*.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The limits of price discrimination,” *American Economic Review*, 105, 921–57.
- BOURREAU, M. AND A. DE STREEL (2018): “The regulation of personalised pricing in the digital era,” .
- BRANCACCIO, G., M. KALOUPTSIDI, AND T. PAPAGEORGIU (2020): “Geography, transportation, and endogenous trade costs,” *Econometrica*, 88, 657–691.
- BRANCACCIO, G., M. KALOUPTSIDI, T. PAPAGEORGIU, AND N. ROSAIA (2023): “Search frictions and efficiency in decentralized transport markets,” *The Quarterly Journal of Economics*, 138, 2451–2503.
- BUCHHOLZ, N. (2022): “Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry,” *The Review of Economic Studies*, 89, 556–591.
- CASTILLO, J. C. (2019): “Who Benefits from Surge Pricing?” *Available at SSRN 3245533*.
- DOMENCICH, T. A. AND D. MCFADDEN (1975): “Urban travel demand—a behavioral analysis,” Tech. rep.
- DOVAL, L. AND V. SKRETA (Forthcoming): “Purchase history and product personalization,” *RAND Journal of Economics*.
- DUBÉ, J.-P. AND S. MISRA (2023): “Personalized pricing and consumer welfare,” *Journal of Political Economy*, 131, 131–189.
- FRECHETTE, G. R., A. LIZZERI, AND T. SALZ (2019): “Frictions in a competitive, regulated market: Evidence from taxis,” *American Economic Review*, 109, 2954–92.

- GAINEDDENOVA, R. (2021): “Pricing and Efficiency in a Decentralized Ride-Hailing Platform,” Tech. rep., University of Wisconsin, Madison.
- GOLDSZMIDT, A., J. A. LIST, R. D. METCALFE, I. MUIR, V. K. SMITH, AND J. WANG (2020): “The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments,” Tech. rep., National Bureau of Economic Research.
- GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): “Optimal nonparametric estimation of first-price auctions,” *Econometrica*, 68, 525–574.
- HALL, J. D. (2018): “Pareto improvements from Lexus Lanes: The effects of pricing a portion of the lanes on congested highways,” *Journal of Public Economics*, 158, 113–125.
- HELLING, B. (2023): “Lyft Wait And Save: Cheaper Rides For Waiting Longer,” <https://www.ridester.com/lyft-wait-and-save/>.
- HENDEL, I. AND A. NEVO (2013): “Intertemporal price discrimination in storable goods markets,” *American Economic Review*, 103, 2722–51.
- HORTAÇSU, A. AND D. MCADAMS (2010): “Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market,” *Journal of Political Economy*, 118, 833–865.
- JIN, Y. AND S. VASSERMAN (2021): “Buying data from consumers: The impact of monitoring programs in US auto insurance,” Tech. rep., National Bureau of Economic Research.
- JOFRE-BONET, M. AND M. PESENDORFER (2003): “Estimation of a dynamic auction game,” *Econometrica*, 71, 1443–1489.
- KEHOE, P. J., B. J. LARSEN, AND E. PASTORINO (2018): “Dynamic competition in the era of big data,” Tech. rep., Stanford University.
- KREINDLER, G. (2023): “Peak-hour road congestion pricing: Experimental evidence and equilibrium implications,” Tech. rep., National Bureau of Economic Research.
- LAGOS, R. (2000): “An alternative approach to search frictions,” *Journal of Political Economy*, 108, 851–873.
- LEVITT, S. D., J. A. LIST, S. NECKERMANN, AND D. NELSON (2016): “Quantity discounts on a virtual good: The results of a massive pricing experiment at King Digital Entertainment,” *Proceedings of the National Academy of Sciences*, 113, 7323–7328.
- LIST, J. A. (2004): “The nature and extent of discrimination in the marketplace: Evidence from the field,” *Quarterly Journal of Economics*, 119, 49–89.

- LUO, Y., I. PERRIGNE, AND Q. VUONG (2018): “Structural analysis of nonlinear pricing,” *Journal of Political Economy*, 126, 2523–2568.
- MARTIN, N. (2019): “Uber Charges More If They Think You’re Willing To Pay More,” <https://www.forbes.com/sites/nicolemartin1/2019/03/30/uber-charges-more-if-they-think-youre-willing-to-pay-more/?sh=3ae3409f7365>.
- McFADDEN, D. (1974): “The measurement of urban travel demand,” *Journal of public economics*, 3, 303–328.
- MIRAVETE, E. J. (1996): “Screening consumers through alternative pricing mechanisms,” *Journal of Regulatory Economics*, 9, 111–132.
- NEVO, A., J. L. TURNER, AND J. W. WILLIAMS (2016): “Usage-based pricing and demand for residential broadband,” *Econometrica*, 84, 411–443.
- OECD DIRECTORATE FOR FINANCIAL AND ENTERPRISE AFFAIRS, COMPETITION COMMITTEE, NOTE BY THE UNITED STATES (2016): “Round table on ‘price discrimination’,” Tech. rep.
- PETRIN, A. AND K. TRAIN (2010): “A control function approach to endogeneity in consumer choice models,” *Journal of marketing research*, 47, 3–13.
- PIGOU, A. C. (1920): *The Economics of Welfare*, Macmillan.
- ROSAIA, N. (2020): “Competing platforms and transport equilibrium: Evidence from New York City,” Tech. rep., Harvard University.
- ROSSI, P. E., G. M. ALLENBY, AND R. McCULLOCH (2005): *Bayesian Statistics and Marketing*, Wiley.
- ROSSI, P. E., R. E. McCULLOCH, AND G. M. ALLENBY (1996): “The value of purchase history data in target marketing,” *Marketing Science*, 15, 321–340.
- SHILLER, B. R. (2013): “First degree price discrimination using big data,” Tech. rep.
- SMALL, K. A. (2012): “Valuation of travel time,” *Economics of Transportation*, 1, 2 – 14.
- TANNER, M. A. AND W. H. WONG (1987): “The Calculation of Posterior Distributions by Data Augmentation,” *Journal of the American Statistical Association*, 82, 528–540.
- TRAIN, K. E. (2009): *Discrete Choice Methods with Simulation*, Cambridge University Press.
- UBER, U. B. (2023): “Uber Charges More If They Think You’re Willing To Pay More,” <https://www.uber.com/en-CA/blog/uberx-priorirty-faq/>.

WALDFOGEL, J. (1995): “The selection hypothesis and the relationship between trial and plaintiff victory,” *Journal of Political Economy*, 103, 229–260.

WHITE HOUSE COUNCIL OF ECONOMIC ADVISORS (2015): “Big Data and Differential Pricing,” Tech. rep., Executive Office of the President of the United States.


Online Appendix

A Platform and data details

A.1 Interface

In Figure A.1, we show Liftago’s app interface.


Figure A.1: Liftago app’s interface



A.2 Locations

Using the exact GPS points of trip origin, we partition our data into 30 locations. In Figure A.2, we show these locations together with an index value for comparing the 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 spatial point of a request and the weighed center of all points within the same cluster, with the constraint that each cluster has an equal number of requests. The depicted locations are close approximations of the k -means clustering procedure, 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 A.2: Locations in Prague




Note: This figure maps the boundaries of the city of Prague with locations defined by a k -means 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. We define the city center as locations including and adjacent to regions 11 and 20.

B Demand model: Omitted details and figures

B.1 Choices and trade-offs

In Figure B.1, we show the proportion of trips that involve a trade-off between spending less and waiting less by hour.

Figure B.1: Trips with Price - Wait Time Trade-Off by Hour



Linearity of choice probability in wait time In [Figure B.2](#), we provide further detail about the probability of choosing to take a trip on the platform. In [Figure B.2a](#), we show the probability of choosing a trip as a function of the wait time, residualized after taking into account time of day, origin and destination location, weather conditions, and the respective other prices and wait times. In [Figure B.2b](#), we show the probability of choosing any trip over the outside option as a function of the minimum wait time, residualized for time of day, origin and destination locations, and weather conditions.

Figure B.2: Choices as a Function of Wait Time

Figure (a) Trip Choice as Function of Wait Time





Figure (b) Trip Choice as Function of Minimum Wait Time



B.2 Demand estimation


We expand on the details of the demand estimation in this section. In particular, we explain the control function approach in estimating the model to instrument for the drivers' bids and the Gibbs sampler.

B.2.1 Control function approach

We present here the results related to our control function approach to demand estimation discussed in Section 4.1, and also how our v estimates would change without it.

Driver fixed effects for control function approach Recall that to obtain an estimate of the unobservable demand conditions, we regress drivers' bids on a set of driver fixed effects. In Figure B.3, we depict the distribution of estimated driver fixed effects, which we use to construct the control function.

Figure B.3: Driver Fixed Effects



Monte Carlo study for control function In this Monte Carlo study, we assume all consumers choose between two ride offers and the outside option. As in the specification in the main text, the utility of the outside option is normalized to 0, whereas the utility of a given offer is given by

$$u_{rj} = \beta^p b_j + \beta^w w_j + \beta^x x_j + \xi_r + \epsilon_{rj}. \quad (\text{B.1})$$

As in the main text, we abuse notation slightly and let j index both a specific offer within a request and the driver who is making this offer. Furthermore, we equate requests to consumers; that is, each consumer corresponds to one request and we drop index i from the notation.

For the Monte Carlo, we assume all rides have 0 wait time, and we abstract away from the dependence on time of day and location. We assume the observable trip characteristics x_j are distributed $N(0, 3)$ and the unobservable demand conditions

ξ_r , which we allow to be order specific, are distributed $N(0, 1)$. Finally, the errors ϵ_{rj} are independently and identically distributed, T1EV.

As in Equation 10, drivers' bids are assumed to be determined as

$$b_{jr} = \bar{c}_j + g(\xi_r) + \Delta c_{jr}. \quad (\text{B.2})$$

We conduct 50 iterations, the results of which we report in Table B.2. Table B.2 shows we obtain unbiased estimates of the parameter values in the indirect utility function under our control function approach. For comparison, Table B.1 shows the same table without the control. We can see that without control for the unobservable, the price coefficients are close to zero or negative and the non-price coefficient is also downwards-biased.

Table B.1: Monte Carlo Results When the Control Function Is Excluded

	$N = 1000$				$N = 10000$			
	β^p	β^x	β^p	β^x	β^p	β^x	β^p	β^x
True value θ	0.5	0.2	0.4	0.8	0.5	0.2	0.4	0.8
$\frac{1}{S} \sum_{s=1}^S \hat{\theta}_s$	0.02	0.177	0.035	0.73	0.02	0.177	-0.08	0.73
$SD(\hat{\theta})$	0.046	0.016	0.05	0.04	0.01	0.005	0.017	0.01
$\frac{1}{S} \sum_{s=1}^S \text{abs}(\hat{\theta}_s - \theta)$	0.47	0.024	0.36	0.07	0.48	0.022	0.48	0.07

NOTE: This table shows results of the Monte Carlo study for our control function approach *without* the control function included. It presents four different scenarios: two different sample-size scenarios and two different sets of parameters. The table presents the average estimate, the standard deviation of the estimates, and the average absolute deviation from the true parameter.

Table B.2: Monte Carlo Results When the Control Function Is Included

	$N = 1000$				$N = 10000$			
	β^p	β^x	β^p	β^x	β^p	β^x	β^p	β^x
True value θ	0.5	0.2	0.4	0.8	0.5	0.2	0.4	0.8
$\frac{1}{S} \sum_{s=1}^S \hat{\theta}_s$	0.52	0.199	0.42	0.80	0.53	0.201	0.42	0.799
$SD(\hat{\theta})$	0.045	0.021	0.05	0.03	0.01	0.004	0.019	0.01
$\frac{1}{S} \sum_{s=1}^S \text{abs}(\hat{\theta}_s - \theta)$	0.038	0.048	0.05	0.025	0.028	0.003	0.025	0.01

NOTE: This table shows results of the Monte Carlo study for our control function approach *with* the control function included. It presents four different scenarios: two different sample-size scenarios and two different sets of parameters. The table presents the average estimate, the standard deviation of the estimates, and the average absolute deviation from the true parameter.

Approximation of drivers' bids as in Equation 10 We provide here the details to derive the regression of drivers' bids on their costs and the unobservable trip characteristics that we use in our control function approach. Recall that driver j 's optimal bid for trip request r satisfies Equation 6, which we reproduce here for ease of reference

$$\frac{1}{0.9}c_{jr} - \left(b - \frac{\gamma(b|w_j, x_j, \xi_r)}{\gamma'(b|w_j, x_j, \xi_r)} \right) = 0.$$

We can rewrite the above as

$$\frac{1}{0.9}c_{jr} - H_{w_j, x_j}(b, \xi_r) = 0.$$

Assuming the above expression is invertible in b , this defines a function $b_{jr} = F_{w_j, x_j}(c_{jr}, \xi_r) = F_{w_j, x_j}(\bar{c}_j + \Delta c_{jr}, \xi_r)$. A Taylor expansion around $(0, 0)$ then delivers:

$$b_{jr} = \frac{\partial}{\partial \bar{c}_j} F_{w_j, x_j}(0, 0) (\bar{c}_j + \Delta c_{jr}) + \frac{\partial}{\partial \xi_r} F_{w_j, x_j}(0, 0) \xi_r + R(\bar{c}_j + \Delta c_{jr}, \xi_r), \quad (\text{B.3})$$

where the latter is the remainder term.

B.2.2 Gibbs sampler

We now explain the specific version of the Gibbs sampler that we construct. Our exposition closely follows Chapter 12 in Train (2009), adapted to our notation. Recall that the vector of coefficients β follows a normal distribution with mean μ and variance-covariance matrix Σ . We assume μ is normally distributed with mean μ_0 and variance-covariance matrix Σ_0 , where Σ_0 is a diffuse prior (unboundedly large variance). We assume the hyper-parameters of the variance are Inverse-Wishart, $\Sigma_0 \sim \text{IW}(v_0, S_0)$.

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 requires a proportionality factor that can be easily computed at each step. In particular, we have

$$K(\beta_i | \mu^l, \Sigma^l, \mathbf{y}_i) \propto \prod_{t=1}^{T_i} l(w_j, b_j, x_j; \beta) \cdot \phi(\beta_i | \mu^l, \Sigma^l), \quad (\text{B.4})$$

where \mathbf{y}_i is the vector of choices and covariates observed for consumer i with T_i

observations and $l(\cdot; \beta_i)$ is the likelihood contribution of a particular choice. The specific assumptions we make about the priors lead to conjugate distributions where the posterior mean of β_i is itself normal and the variance is again in the family of inverse gamma distributions.

One can then iteratively update the coefficient vector, β_i , the mean of the coefficients, as well as their standard deviations. To describe the updating algorithm, let $\bar{\mu}^l$ be the sample mean of coefficients of iteration l in the chain, and let S^l be the sample variance of the Inverse-Wishart. The iterative updating is then given by the following steps:

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

B.3 Location and time-of-day parameter estimates

In this section, we provide further details on location and time-of-day parameter estimates.

Location In [Figure B.4](#), we summarize all location-specific wait-time and price-coefficient estimates omitted from [Table 3](#). Because each location contains multiple consumers and each consumer has individual preference estimates, we report each location-specific estimate as a bar. The dark line shows the median estimate and the vertical bars around each point display the 10th and 90th percentiles of all individual estimates within the location as indicated on the horizontal axis. Location indices may be cross-referenced with [Figure A.2](#). In [Figure B.4a](#), we display the wait-time coefficients, and in [Figure B.4b](#), we display the price coefficients.

Time of day In [Figure B.5](#), we summarize all time-of-day-specific wait-time and price coefficient estimates omitted from [Table 3](#). Similar to the analysis of the location-specific parameter estimates, we report each time-of-day-specific estimate as a bar. The dark line shows the median estimate and the vertical bars around each point display the 10th and 90th percentiles of all individual estimates within

Figure B.4: Location-Specific and Time-Specific Coefficient Estimates

Figure (a) Wait-Time Coefficients





Figure (b) Price Coefficients



Note: These figures summarize all location-specific wait-time and price coefficient estimates.

the hour indicated. In [Figure B.4a](#), we report the wait-time coefficients, and in [Figure B.4b](#), we report the price coefficients.

Figure B.5: Location-Specific and Time-Specific Coefficient Estimates

Figure (a) Wait-Time Coefficients





Figure (b) Price Coefficients



Note: These figures summarize all hour-specific wait-time and price coefficient estimates.

B.4 Elasticities by location

In [Table B.3](#), we present the price- and wait-time elasticity estimates by origin and destination locations. For example, location 1's origin elasticities are the price and wait-time elasticity associated with all trips that depart from location 1. Similarly, location 1's destination elasticities obtain from all trips that arrive to location 1. Recall that location indices may be found in [Figure A.2](#).

Table B.3: Bid Level Elasticities by Origin and Destination Location

LOCATION INDEX	Origin Locations		Destination Locations	
	PRICE	WAIT TIME	PRICE	WAIT TIME
1	-6.47	-0.87	-5.71	-0.63
2	-3.94	-0.64	-3.41	-0.82
3	-5.47	-0.89	-5.05	-0.64
4	-4.44	-0.66	-3.79	-0.81
5	-5.99	-0.93	-5.48	-0.62
6	-6.6	-1.04	-6.39	-0.54
7	-11.69	-0.44	-11.34	-0.96
8	-12.67	-1.25	-10.0	-0.49
9	-3.77	-0.75	-3.78	-0.69
10	-4.0	-0.61	-3.57	-0.78
11	-3.44	-0.49	-3.22	-0.72
12	-7.13	-1.15	-5.99	-0.49
13	-9.66	-1.04	-8.49	-0.48
14	-5.4	-0.81	-5.0	-0.64
15	-4.87	-0.92	-4.58	-0.65
16	-5.09	-0.8	-4.3	-0.72
17	-4.38	-0.62	-3.74	-0.87
18	-8.11	-1.39	-7.28	-0.64
19	-7.19	-0.97	-6.4	-0.61
20	-3.77	-0.51	-3.56	-0.88
21	-6.99	-1.02	-6.6	-0.57
22	-6.69	-0.98	-5.59	-0.54
23	-9.27	-1.35	-7.3	-0.41
24	-11.76	-1.21	-9.45	-0.55
25	-8.04	-1.05	-7.05	-0.61
26	-4.15	-0.68	-4.0	-0.73
27	-5.16	-0.91	-4.39	-0.67
28	-8.28	-0.93	-7.37	-0.51
29	-4.16	-0.6	-3.81	-0.69
30	-6.69	-0.94	-5.96	-0.57

NOTE: This table provides price and wait time elasticities across the Prague's thirty locations, both as origins and as destinations.

B.5 Trip-specific heterogeneity results

With the baseline results in [Table 5](#), we show the average VOT across different hours of the day. Within each time of day and within each individual's *type*, however, additional heterogeneity exists due to the fact that some trips are more or less time sensitive. For example, a given person may express higher VOT if she is late for an appointment.

To analyze this type of heterogeneity, we define time-sensitive trips as the subset

of trips in which having requested a ride, a consumer faces a set of bids in which (1) the arrival time falls around a rounded hour increment such as 9:00am or 2:00pm, (2) only one bid provides a trip that arrives before the hour, whereas all others would provide a trip that arrives after the hour, and (3) a trip occurs on the platform. Because the VOT estimated on trips generated on this subset is inherently selected due to point (3), we compare this against a similar subset of trips around a placebo clock time. We thus define placebo trips by selecting orders where the arrival time falls around a clock time ending in :23 or :53, such as 8:23am or 2:53pm. We then apply criteria (2) and (3) to these trips. The difference in VOT between time-sensitive and placebo trips reveals the relevant heterogeneity in time sensitivity.

To augment our baseline results, in [Table B.4](#), we report results for time-sensitive trips and shows the VOT for this subset is about 58% greater than the comparable measure in the placebo group. In the first column, we report VOT for all trips, as in [Table 5](#), as a comparison. Our baseline VOT results, therefore, can be interpreted as averages across this type of trip-specific heterogeneity.

Table B.4: Trip-Specific Heterogeneity in VOT

	All Trips	Time-Selected Trips	
		Placebo	Time-Sensitive
VOT (\$/hour)	\$13.21	\$17.92	\$28.34
<i>N</i> Trips	1,021,007	86,786	62,229
<i>N</i> Individuals	80,161	32,438	21,651

NOTE: This table provides baseline mean VOT estimates in the first column. The second column reports mean VOT estimates for trips in which exactly one bid occurs before the 23rd and 53rd minute of each hour, and additional bids fall beyond this time. The third column similarly reports VOT estimates for trips in which exactly one bid occurs before the first minute of each hour. All estimates are presented in US dollars.

C Driver supply side: Omitted results

C.1 Markups

In [Figure C.1](#), we show average markups across hours of the day ([Figure C.1a](#)) and location ([Figure C.1b](#)).

Figure C.1: Average Markups

Figure (a) By Hour





Figure (b) By Location



NOTE: These figures show the average of estimated driver markups by hour (Figure C.1a) and location (Figure C.1b). In Figure C.1b, we index locations in increasing order of average driver markup.

D Counterfactual computations

D.1 Implementation details

We offer here additional detail on the counterfactual computation. Recall from [Section 6](#) that all counterfactuals are obtained by computing the platform’s profit-maximizing direct mechanism consisting of a tuple of driver transfers, t_j , and consumer prices, p_j , which in turn affect the probability that a consumer matches with a specific driver. The platform may condition consumer prices and driver transfers on the drivers’ wait times and other trip observable characteristics, all of which the platform observes.

With this setup, the matching and pricing design problem has two constraints. First, the platform does not have full information about the consumers’ preferences over the different options, summarized by the parameters β and the logit shocks ϵ . Second, when we consider the drivers’ adjustment to the platform’s new policy, the platform does not know the drivers’ inclusive costs of serving the ride, which, as discussed in [Section 5](#), are heterogeneous.

To make results comparable to our baseline, we impose two restrictions on the counterfactual platform policy. First, the platform offers the same number of drivers to the passenger in the counterfactuals as in the baseline. Second, the probability of the platform assigning driver j to rider i is determined by the logit probabilities in [Equation 2](#), where the platform’s prices for the ride replace the drivers’ bids. Formally, the probability that driver j is chosen is given by

$$L_j(c_j, x_j, w_j) = \mathbb{E} \left[l_{J_r}(w_j, p_j(c_j, \cdot), x_j, \xi_r; \beta) | j \in J_r \right]. \quad (\text{D.1})$$

Thus, the platform’s policy must satisfy the participation and incentive compatibility constraints for each driver j , cost c_j , and reported cost \hat{c}_j :

$$\begin{aligned} T_j(c_j, w_j, x_j) - c_j L_j(c_j, x_j, w_j) &\geq 0, && (\text{PC}_j(c_j, x_j, w_j)) \\ T_j(c_j, w_j, x_j) - c_j L_j(c_j, x_j, w_j) &\geq T_j(\hat{c}_j, w_j, x_j) - c_j L_j(\hat{c}_j, x_j, w_j), && (\text{IC}_j(c_j, \hat{c}_j, x_j, w_j)) \end{aligned}$$

where $T_j(c_j, w_j, x_j)$ is driver j ’s expected payment from the platform when driver j ’s cost and wait time are c_j and w_j , and trip characteristics are x_j , under the assumption that other competing drivers are truthfully reporting their costs.

The platform then chooses the transfers and prices (t_j, p_j) to solve

$$\begin{aligned} \max_{p^r, t^r} \mathbb{E}_{\beta, \bar{c}} [\Pi(t^r, p^r; \beta, \bar{c}_r) | \mathcal{I}] & \quad (\Pi(\mathcal{I})) \\ \text{s.t. } \text{PC}_j(c_j, x_j, w_j), \text{IC}_j(c_j, \hat{c}_j, x_j, w_j) & \text{ for all } j \in J_r, c_j, \text{ and } \hat{c}_j. \end{aligned}$$

We implement this maximization problem as a maximization with inequality constraints, in direct analogy to the way we mathematically pose this problem above. Because computationally implementing this problem for a continuum of types is not feasible, we discretize the drivers' costs, the demand conditions x_j , and the distributions of the consumers' wait time and price coefficients. We divide each of these variables into terciles. We experimented with finer partitions and our results are not too sensitive to increasing the number of partitions.

D.2 A dynamic model of driver behavior

We provide here a dynamic model of driver behavior, which provides a microfoundation for the drivers' problem in Section 3.2 (cf. Equation 5). In the dynamic model, a day is partitioned into discrete periods, labeled by n . Thus, whereas in the main text time t refers to a ride's calendar date and time, the model below keeps track of time within a day, but not calendar date.

In each period n and each location a , timing is as follows. First, each available driver j observes the value of his outside option, κ_{jn} , which describes the flow value of not serving a trip on the platform. The outside option κ_{jn} is drawn from a possibly time- and location-dependent distribution, $G_n(\cdot | a)$. Second, each driver either receives a request from the platform ($\rho = 1$) or not ($\rho = 0$). From the driver j 's perspective, this event is random (and potentially location- and time-dependent). Below, we abuse notation and let $\rho_n(a)$ denote the total probability that the driver receives a request from the platform at location a in period n . A driver who receives no requests from the platform earns the outside option κ_{jn} . Instead, a driver who receives a request from the platform observes the trip's characteristics, (w_{jn}, x_{jn}) , and submits a bid for that trip simultaneously with the other drivers who received the same request. If the consumer selects driver j , driver j serves the ride. Instead, if the rider does not select driver j , the driver earns the outside option κ_{jn} . At the end of period n , drivers transition out of a only if they win the auction. No independent location choice is made.

Driver j 's state in period n then consists of the value of the outside option κ_{jn} and

the location a . Thus, driver j 's value at state (κ_{jn}, a) , $S_{jn}(\kappa_{jn}, a)$, is given by²⁰

$$S_{jn}(\kappa_{jn}, a) = (1 - \rho) \left[\kappa_{jn} + \delta \mathbb{E}_{G_{n+1}(\cdot|a)}[S_{jn+1}(\cdot, a)] \right] + \rho \mathbb{E}[V_{jn}(\kappa_{jn}, a, x_{jn})], \quad (\text{D.2})$$

where for simplicity we omit the dependence of the ping probability ρ on n and a .

We now unpack the value V_{jn} of receiving a request from the platform. Suppose driver j receives a request for a trip to location \hat{a} with wait time w_j and trip characteristics x_{jn} . Recall that $\tau(w_{jn}, x_{jn})$ denotes the number of periods until the consumer is dropped off at \hat{a} . That is, conditional on winning the auction, driver j will be at location \hat{a} $\tau(w_{jn}, x_{jn})$ periods from now. Thus, we can write driver j 's payoff from bidding for this trip as follows:

$$V_{jn}(\kappa_{jn}, a, x_{jn}) = \max_b \left\{ \gamma(b|w_j, x_{jn}) \left(0.9 b + \delta^{\tau(w_{jn}, x_{jn})} \mathbb{E}_{G_{n+\tau}(\cdot|\hat{a})} [S_{j, n+\tau}(\cdot, \hat{a})] \right) + \right. \\ \left. (1 - \gamma(b|w_{jn}, x_{jn})) \left(\kappa_{jn} + \delta \mathbb{E}_{G_{n+1}(\cdot|a)}[S_{jn+1}(\cdot, a)] \right) \right\}. \quad (\text{D.3})$$

In words, upon receiving a request, the drivers essentially participate in an asymmetric auction, where they compete on the price b under exogenous quality characteristics. With probability $\gamma(b|w_{jn}, x_{jn})$, driver j wins the auction, serves the ride at a price b , pays Liftago 10% of the trip's final price, and transitions to a new location in τ periods, where τ takes into account the wait time and the trip length. With probability $1 - \gamma(b|w_{jn}, x_{jn})$, the driver loses the auction and obtains the outside option and the opportunity to serve a trip in that location in period $n + 1$. We do not separately model a driver's choice to reject a request, because drivers can always submit a very high bid.

Drivers' inclusive costs We now show how the driver's problem described in Equation D.3 delivers the problem in Equation 5 as a function of the driver's inclusive cost. Note we can rewrite the maximization problem in Equation D.3 in a way that directly maps into a static auction setup:

$$V_{jn}(\kappa_{jn}, a, x_{jn}) = \max_b \left\{ \kappa_{jn} + \delta \mathbb{E}_{G_{n+1}(\cdot|a)}[S_{jn+1}(\cdot, a)] + \right. \\ \left. \gamma(b|w_{jn}, x_{jn}) \left(0.9 b + \delta^{\tau(w_{jn}, x_{jn})} \mathbb{E}_{G_{n+\tau}(\cdot|\hat{a})}[S_{jn+\tau}(\cdot, \hat{a})] - \kappa_{jn} - \delta \mathbb{E}_{G_{n+1}(\cdot|a)}[S_{jn+1}(\cdot, a)] \right) \right\}. \quad (\text{D.4})$$

²⁰Implicit in our formulation is the assumption that a driver's bidding behavior does not depend on how they arrived at location a in period n : besides their individual characteristics (e.g., car make, rating), only the current time period, location, and outside option realization matter to determine their bidding behavior.

We can define the inclusive cost, $c_{jn}(\kappa_{jn}, w_{jn}, x_{jn})$, as follows:

$$c_{jn}(\kappa_{jn}, w_{jn}, x_{jn}) \equiv \kappa_{jn} + \delta \mathbb{E}_{G_{n+1}(\cdot|a)}[S_{jn}(\kappa_{n+1}, a)] - \delta^{\tau(\cdot)} \mathbb{E}_{G_{n+\tau}(\cdot|\hat{a})}[S_{jn+1}(\kappa_{n+\tau}, \hat{a})]. \quad (\text{D.5})$$

Thus, the problem in Equation 5 corresponds to that in Equation D.4, justifying our formulation in Section 3.2. In what follows, note that the inclusive cost can be written as the sum of two components: the flow payoff κ_{jn} and the difference in continuation values:

$$e_{jn}^{n+\tau}(w_{jn}, x_{jn}) \equiv \delta \mathbb{E}_{G_{n+1}(\cdot|a)}[S_{jn}(\kappa_{n+1}, a)] - \delta^{\tau} \mathbb{E}_{G_{n+\tau}(\cdot|\hat{a})}[S_{jn+1}(\kappa_{n+\tau}, \hat{a})]. \quad (\text{D.6})$$

Under our assumptions, κ_{jn} is exogenous to the platform's policies, but $e_{jn}^{n+\tau}$ is not.

D.3 Continuation values: Identification and estimation

We now discuss further identification results of the supply-side primitives, that is, the time and location-dependent distributions of outside options, $\{\kappa \sim G_n(\cdot|a) : a \in \{1, \dots, A\}, n \in \{1, \dots, N\}\}$. Instead, we observe

- (a) the driver request probabilities ρ ,
- (b) the fee f collected by the platform,
- (c) the probability of winning the auction with a bid b , denoted by γ , and
- (d) the conditional distribution of bids by drivers F .

Furthermore, we assume the discount factor δ is known. We then have the following result:

Proposition 1. *Assume the observables are as listed in items (a)-(d), and the discount factor, δ , is known. Furthermore, assume either of the following conditions hold:*

- (i) *For each $a \in \{1, \dots, A\}, n \in \{1, \dots, N\}$, there are trips to a with $\tau = 0$.*
- (ii) *For each $a \in \{1, \dots, A\}, n \in \{1, \dots, N\}$, a location $\tilde{a} \in \{1, \dots, A\}$ and period $\tilde{n} \in \{1, \dots, N\}$ exist such that there are both trips with $\tau = n - \tilde{n}$ and $\tau = n + 1 - \tilde{n}$ from \tilde{a} to a .*

Then, the set of conditional distributions of outside options $\{\kappa \sim G_n(\cdot|a) : a \in \{1, \dots, A\}, n \in \{1, \dots, N\}\}$ is non-parametrically identified.

Proof of Proposition 1. The argument in Section 4.2 implies the distribution of $c_{jn}(\kappa_{jn}, w_{jn}, x_{jn})$ is identified, so that

$$c_{jn}(\kappa_{jn}, w_j, x_j) = \kappa_{jn} + \delta \mathbb{E}_{G_{n+1}(\cdot|a)} [S_{jn+1}(\cdot, a)] - \delta^{\tau(\cdot)} \mathbb{E}_{G_{n+\tau}(\cdot|\hat{a})} [S_{jn+\tau}(\cdot, \hat{a})] \quad (\text{D.7})$$

is known. Under condition (i), for $\tau = 0$, the last two terms drop out and the distribution of outside options κ is therefore directly identified from the cost $c_{jn}(\kappa_{jn}, w_{jn}, x_{jn})$.

Dropping the dependence on the driver's index, j , and taking expectations on both sides of Equation D.7 using the distribution of κ conditional on (n, a) , $G_n(\cdot|a)$ we obtain that for $r = (t, a, \hat{a})$

$$\begin{aligned} \mathbb{E}_{G_n(\cdot|a_n)} [c_{nr}^{n+\tau}(\kappa, w_{jn}, x_{jn})] &= \mathbb{E}_{G_n(\cdot|a)} [\kappa_n] + \\ &\delta \mathbb{E}_{G_{n+1}(\cdot|a)} [S_{n+1}(\cdot, a)] - \delta^{\tau(\cdot)} \mathbb{E}_{G_{n+\tau}(\cdot|\hat{a})} [S_{n+\tau}(\cdot, \hat{a})]. \end{aligned} \quad (\text{D.8})$$

Under condition (ii), for each $l \geq 1$, $n \geq 0$, and location a , a location \tilde{a}_n and a time period m_n exist such that the following hold:

$$\begin{aligned} \mathbb{E}_{G_{m_n}(\cdot|\tilde{a}_n)} [c_{m_n}^{l+n}(\cdot)] &= \mathbb{E}_{G_{m_n}(\cdot|\tilde{a}_n)} [\kappa_{m_n}] + \\ &\delta \mathbb{E}_{G_{m_n+1}(\cdot|\tilde{a}_n)} [S_{m_n+1}(\cdot, \tilde{a}_n)] - \delta^{l+n} \mathbb{E}_{G_{l+n}(\cdot|a)} [S_{l+n}(\cdot, a)] \end{aligned} \quad (\text{D.9})$$

and

$$\begin{aligned} \mathbb{E}_{G_{m_n}(\cdot|\tilde{a}_n)} [c_{m_n}^{l+n+1}(\cdot)] &= \mathbb{E}_{G_{m_n}(\cdot|\tilde{a}_n)} [\kappa_{m_n}] + \\ &\delta \mathbb{E}_{G_{m_n+1}(\cdot|\tilde{a}_n)} [S_{m_n+1}(\cdot, \tilde{a}_n)] - \delta^{l+n+1} \mathbb{E}_{G_{l+n+1}(\cdot|a)} [S_{l+n+1}(\cdot, a)]. \end{aligned} \quad (\text{D.10})$$

Taking the difference between Equations D.9 and D.10, we obtain:

$$\begin{aligned} \Delta_l \equiv \frac{\mathbb{E}_{G_{m_n}(\cdot|\tilde{a}_n)} [c_{m_n}^{l+n}(\cdot)] - \mathbb{E}_{G_{m_n}(\cdot|\tilde{a}_n)} [c_{m_n}^{l+n+1}(\cdot)]}{\delta^{l+n}} &= \\ &\delta \mathbb{E}_{G_{l+n+1}(\cdot|a)} [S_{l+n+1}(\cdot, a)] - \mathbb{E}_{G_{l+n}(\cdot|a)} [S_{l+n}(\cdot, a)]. \end{aligned}$$

Furthermore, note that

$$\Delta_0 + \sum_{l=1}^{N-1} \delta \Delta_l = \delta \mathbb{E}_{G_{n+N}(\cdot|a)}[S_N(\cdot, a)] - \mathbb{E}_{G_n(\cdot|a)}[S_t(\cdot, a)] = (\delta - 1) \mathbb{E}_{G_n(\cdot|a)}[S_n(\cdot, a)],$$

where the last equality follows from the stationarity assumption. This identifies $\mathbb{E}_{G_n(\cdot|a)}[S_n(\cdot, a)]$ and therefore all other unknown expectations for a . We can repeat this procedure for all $a \in \{1, \dots, A\}$. Once the expectations are known, the conditional distributions of κ are directly identified from the inclusive costs, $c_n^{n+\tau}(\cdot)$. \square

Estimation: It follows from the proof of [Proposition 1](#) that all driver primitives can be recovered through a simple regression of the inclusive costs on a set of time- and location-specific dummies. To be precise, given the discount factor and the inferred costs from the GPV inversion at hand, we can run the following regressions following [Equations D.5 and D.6](#):

$$c_{jn}^{n+\tau}(\kappa, \cdot) = \sum_{\tilde{a} \in \{1, \dots, A\}} \sum_{l \in \{1, \dots, N\}} (\mathbb{1}_{\tilde{a}=a, l=n+1} \cdot \delta \alpha_{\tilde{a}l} - \mathbb{1}_{\tilde{a}=\hat{a}, l=n+\tau} \cdot \delta^\tau \cdot \alpha_{\tilde{a}l}) + \kappa. \quad (\text{D.11})$$

From this, we can back out the continuation values as $\mathbb{E}_{G_n(\cdot|a)}[S_n(\cdot, a)] = \alpha_{an}$.

Moreover, the residuals from this regression provide an estimate of the conditional distributions of outside options. So, $\hat{\kappa}_{jn}$ is an estimate of $G_n(\cdot|a)$. With this, we can then perform the decomposition of the inclusive costs into a primitive component and a remainder that depends on future cost draws, bids, and locations, $e_{jn}^{n+\tau}(a, \hat{a}) = c_{jn}^{n+\tau}(\cdot) - \kappa_{jn}$.

D.4 On-platform earnings and continuation values

In this section, we present evidence that drivers' continuation values are being driven mostly by off-platform activities (e.g., serving street-hail rides) and less by activities on the platform. Specifically, we investigate whether drivers' on-platform business is a substantial fraction of their total business and whether trips occur with enough frequency that drivers should care about future on-platform payoffs differing among different trip opportunities (and how those opportunities change in the counterfactual). Drivers receive on average around 16 bid requests in a day, but their average probability of winning any auction is around 23%, which leads

to an average of 3.8 on-platform trips served per driver per day. The average time spent between on-platform rides served is about 112 minutes, or nearly two hours. Note the average trip time is 17 minutes with a standard deviation of 9 minutes. Recall that Liftago drivers in Prague are all licensed to serve the traditional street-hail business *off-platform*. Despite our best efforts, we do not have data on drivers' (largely cash-based) street-hail business, but we can infer from other professional taxi driver markets that Liftago drivers' on-platform earnings are a relatively small fraction of the total. For example, in New York City, where we have access to comprehensive data, drivers work on average for 8.3 hours and serve around 23 trips per day. As a result, we think a reasonable assumption is that κ dominates the driver's perception of the outside option and any continuation values, because drivers' trip shocks and earnings shocks from the street-hail business are likely to dominate those due to changes to on-platform pricing. We essentially take this to imply that the counterfactual changes we consider would not meaningfully impact the spatial distribution of drivers.

We also tested the robustness of our counterfactuals to changes in on-platform earnings opportunities. To do this, we implement the decomposition strategy in the previous section to separately recover the distribution of off-platform outside options κ and of on-platform continuation values $e(\cdot)$ based on Proposition 1. With the estimated distribution of κ , we re-construct drivers' continuation values and inclusive costs by forward-simulation. We conducted this procedure under the baseline counterfactual, in which we take the κ distribution along with observed prices and driver win probabilities and simulate many future paths of auctions and possible driver outcomes, along with draws from the outside-option distribution. For each time- and location- state, we simulate a sequence of 40 periods (or 10 hours) on the platform 100 times. Even though the counterfactual payments to drivers under uniform pricing change by up to 75%, the results of this procedure suggest that the drivers' inclusive costs only change by about 1.5%.

We tested this change as a one-step iteration in the uniform pricing counterfactual, by first adjusting driver costs given the change in inclusive costs and then re-running the optimal platform pricing problem. Consistent with on-platform earnings being a small fraction of drivers' overall earnings, we find the counterfactual results also change only minimally and are qualitatively unchanged.

Table D.1: Price and Wait Time Distributions by Preference Type

	Baseline 10% flat fee	uniform pricing	wait coeff. only	wait, price coeff.	ETA pricing
Panel A: Average Wait Time By Wait Time Coefficient					
Average Wait Time					
<i>lowest wait-time coeff.</i>	4.75	4.57	4.58	4.57	4.65
	4.6	4.42	4.43	4.41	4.52
	4.44	4.26	4.27	4.24	4.37
<i>highest wait-time coeff.</i>	4.25	4.07	4.1	4.06	4.19
Panel B: Average Purchase Price By Wait Time Coefficient					
Average Purchase Price					
<i>lowest wait coeff.</i>	\$6.15	\$6.66	\$6.92	\$6.93	\$6.51
	\$6.14	\$6.62	\$6.73	\$6.67	\$6.48
	\$6.13	\$6.59	\$6.48	\$6.41	\$6.45
<i>highest wait coeff.</i>	\$6.12	\$6.56	\$6.28	\$6.17	\$6.43
Panel C: Average Purchase Price By Price Coefficient					
Average Purchase Price					
<i>lowest price coeff.</i>	\$6.18	\$6.75	\$6.74	\$7.6	\$6.62
	\$6.14	\$6.6	\$6.59	\$6.7	\$6.46
	\$6.11	\$6.51	\$6.48	\$6.2	\$6.36
<i>highest price coeff.</i>	\$6.07	\$6.41	\$6.39	\$5.69	\$6.26

NOTE: This table summarizes the average wait times and sales prices across pricing counterfactuals according to consumers' individual preferences for wait times and prices. All estimates correspond to the welfare analysis in Table 7.

D.5 Omitted counterfactual results

Table D.1 reports the average wait time and consumer prices incurred by different consumer types under each pricing policy. Treating wait time as a measure of product quality, this table represents a measure of allocative efficiency on preferences for quality. Customer types are sorted from *lowest* to *highest* wait time disutility (Panels A and B) and similarly in Panel C for price disutility. In the baseline there is a small degree of sorting, different preferences cannot be exploited by drivers, because they do not have access to consumers' types when bidding. As a result, sorting only occurs as a result of consumers' selecting different menu options as well as the outside option. By contrast, the platform does observe consumer types. As a result, we see in Panel A that high-sensitivity consumer types (i.e., those with a high wait-time disutility) tend to select lower ETA trips than they would under the baseline. Panel B shows that this tendency is in part due to the platform internalizing their preferences and pricing these trips lower than in the baseline. A similar result holds for price preferences. Panel C shows greater dispersion of pricing as prices become more personalized.

These results suggest the platform's optimal pricing does not induce sorting among consumers on their quality preferences. Instead, the platform changes the price distribution for wait-time-sensitive consumers to generate similar participation at lower prices. As a result, the most sensitive consumers face lower prices and the least sensitive face higher prices. This observation in part explains the market expansion we see under personalized pricing.