Online Appendix

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

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

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

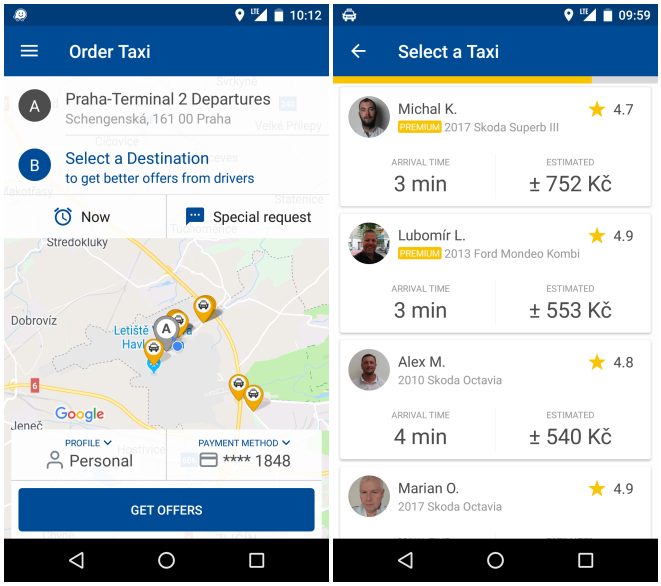

*Princeton University, 20 Washington Rd, Princeton, NJ 08540
†Columbia University, 665 West 130th Street, New York, NY 10027
‡CERGE-EI, Politickỳch vìzòù 7, 111 21 Prague 1
§Massachusetts Institute of Technology, 50 Memorial Dr, Cambridge, MA 02142

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

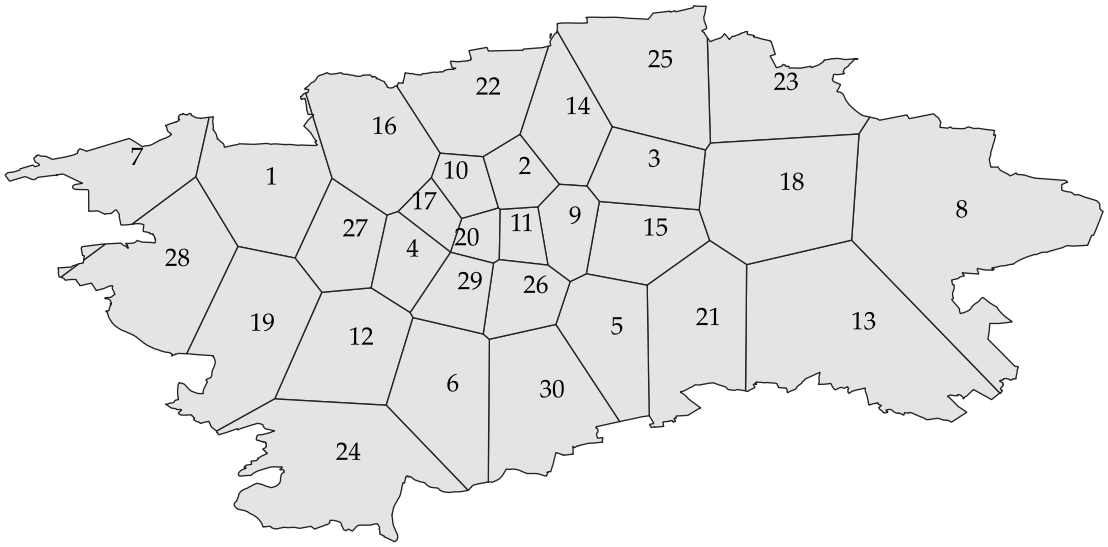

Note: This figure maps the boundaries of the city of Prague with locations defined by a k-means clustering procedure on GPS locations of trip origins and depicted as Voronoi cells that contain the clustered points. Displayed index values correspond to indices used in the paper. We define the city center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

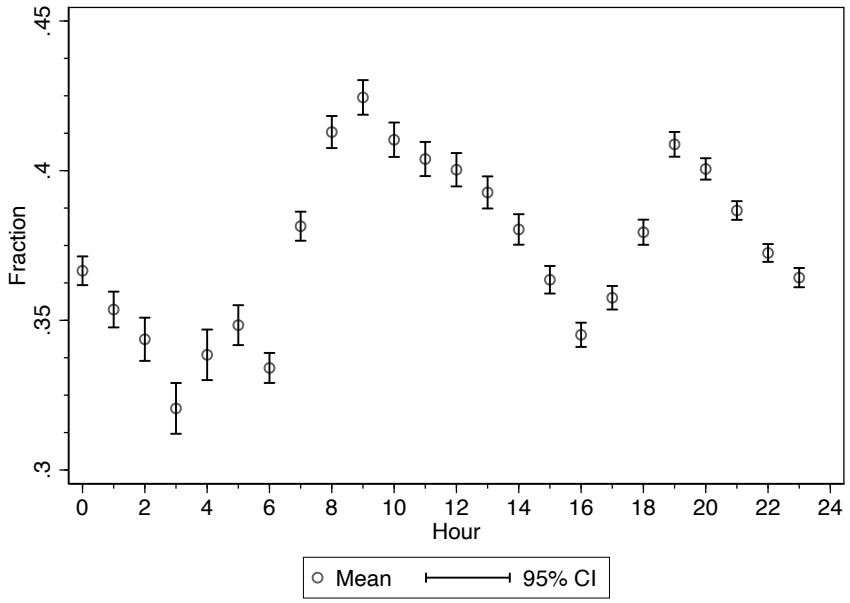

Linearity of choice probability in wait time In Figure B.2, we provide further detail about the probability of choosing to take a trip on the platform. In Figure B.2a, we show the probability of choosing a trip as a function of the wait time, residualized after taking into account time of day, origin 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

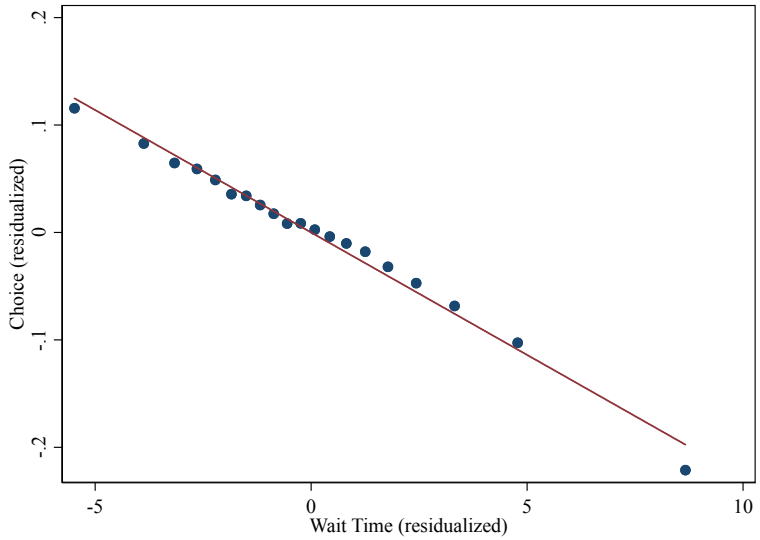

Figure (b) Trip Choice as Function of Minimum Wait Time

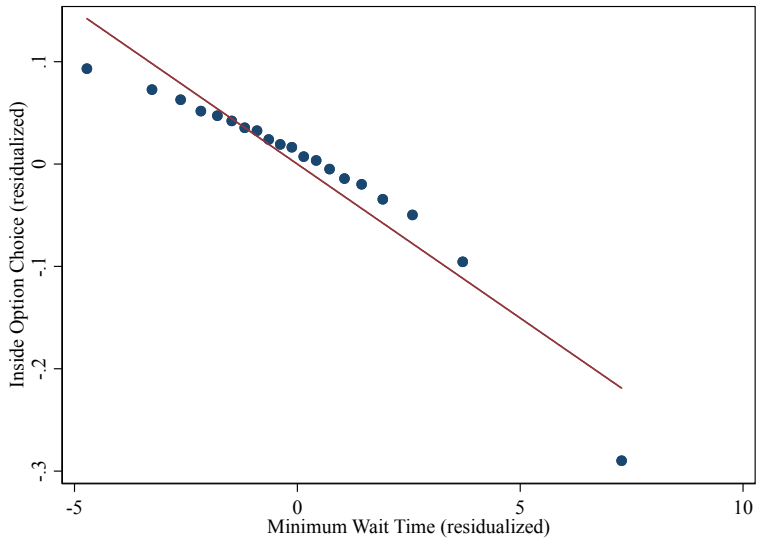

B.2 Demand estimation

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

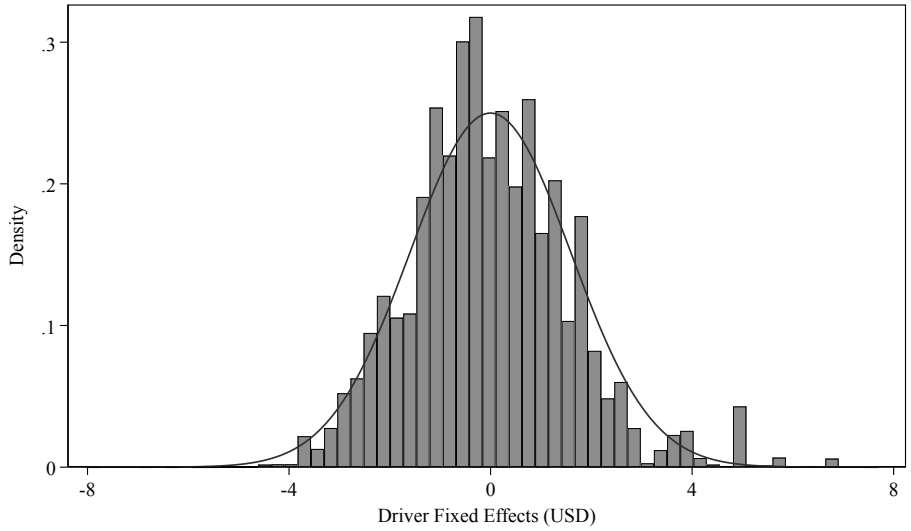

Monte Carlo study for control function In this Monte Carlo study, we assume all consumers choose between two ride offers and the outside option. Similar to the specification in the main text, the utility of the outside option is normalized to 0, whereas the utility of a given offer is given by[1]

$$u_{rj} = \beta^p \, b_j + \beta^w \, w_j + \beta^x \, x_j + \xi_r + \epsilon_{rj}. \tag{B.1}$$

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.

[1] Our Monte Carlo study normalizes the coefficient on the square of the wait time, β^{sq}, to 0 because it is close to 0 in our estimation.

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

$$b_{jr} = \overline{c}_j + g(\xi_r) + \Delta c_{jr}. \tag{B.2}$$

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	0.178	0.02	0.73	0.01	0.178	-0.09	0.73
$SD(\hat{\theta})$	0.044	0.017	0.06	0.05	0.02	0.005	0.018	0.01
$\frac{1}{S}\sum_{s=1}^{S} abs(\hat{\theta}_s - \theta)$	0.48	0.023	0.38	0.07	0.49	0.022	0.49	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.

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

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

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	0.200	0.42	0.81	0.52	0.20	0.43	0.801
$SD(\hat{\theta})$	0.046	0.017	0.07	0.03	0.02	0.005	0.02	0.01
$\frac{1}{S}\sum_{s=1}^{S} abs(\hat{\theta}_s - \theta)$	0.038	0.013	0.06	0.026	0.025	0.004	0.03	0.012

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.

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:

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

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

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

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 $\overline{\mu}^l$ 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 $\mathcal{N}(\overline{\mu}^{l-1}, \frac{W}{N})$.
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 $S_1^l = \frac{1}{N} \cdot \sum_i^N (\beta_i^{l-1} - \overline{\mu}^{l-1}) \cdot (\beta_i^{l-1} - \overline{\mu}^{l-1})'$.
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.

Figure B.4: Location-Specific Coefficient Estimates

Figure (a) Wait-Time Coefficients

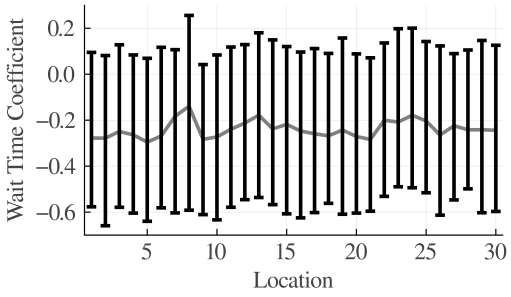

Figure (b) Price Coefficients

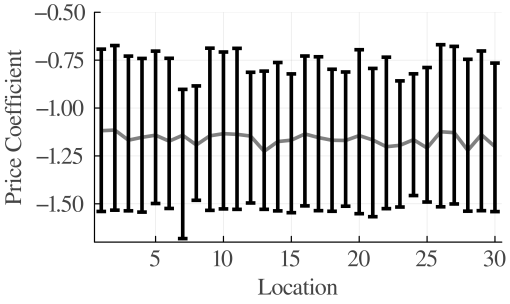

Note: These figures summarize all location-specific wait-time and price coefficient estimates.

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 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: Time-Specific Coefficient Estimates

Figure (a) Wait-Time Coefficients

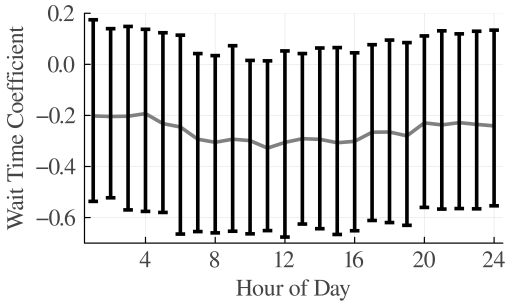

Figure (b) Price Coefficients

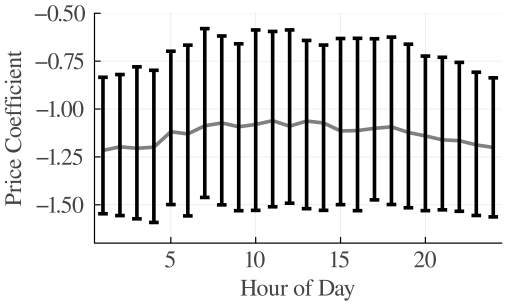

Note: These figures summarize all hour-specific wait-time and price coefficient estimates.

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

	Origin Locations		**Destination Locations**	
Location Index	Price	Wait time	Price	Wait time
1	-7.09	-0.83	-6.10	-0.65
2	-3.97	-0.62	-3.38	-0.81
3	-5.69	-0.94	-5.49	-0.67
4	-4.84	-0.71	-3.84	-0.83
5	-6.21	-0.95	-5.93	-0.64
6	-7.46	-1.10	-6.99	-0.61
7	-13.89	-0.67	-12.73	-0.74
8	-13.88	-1.37	-11.40	-0.63
9	-4.03	-0.77	-3.99	-0.72
10	-4.12	-0.59	-3.53	-0.76
11	-3.59	-0.51	-3.28	-0.72
12	-6.94	-1.02	-6.41	-0.56
13	-10.60	-1.15	-9.72	-0.55
14	-5.78	-0.85	-5.39	-0.66
15	-4.86	-0.85	-4.92	-0.68
16	-5.54	-0.83	-4.58	-0.75
17	-4.69	-0.64	-3.75	-0.88
18	-9.45	-1.46	-8.44	-0.68
19	-7.90	-1.07	-7.08	-0.58
20	-4.02	-0.54	-3.60	-0.88
21	-7.49	-1.09	-7.34	-0.61
22	-7.44	-1.16	-6.27	-0.62
23	-10.96	-1.27	-8.06	-0.46
24	-13.80	-1.38	-10.93	-0.60
25	-8.80	-1.14	-7.88	-0.62
26	-4.26	-0.67	-4.00	-0.73
27	-5.27	-0.90	-4.65	-0.70
28	-9.40	-1.12	-8.16	-0.58
29	-4.36	-0.59	-3.95	-0.70
30	-7.30	-1.02	-6.70	-0.56

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 81% 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)	$14.05	$12.84	$23.24
N Trips	500,000	89,045	70,470
N Individuals	66,583	30,882	27,645

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

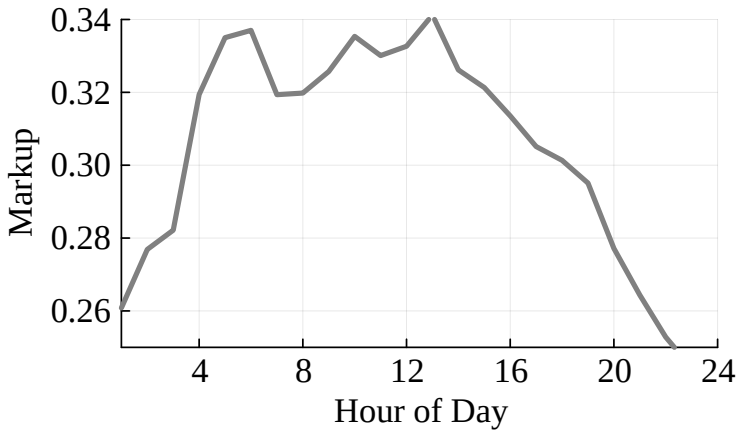

Figure (b) By Location

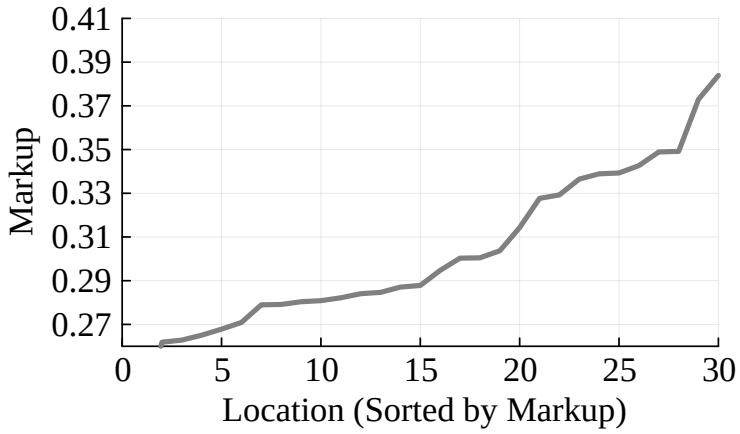

NOTE: These figures show the average of estimated driver markups by hour (Figure C.1a) and location (Figure C.1b). In Figure C.1b, we index locations in increasing order of average driver markup.

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

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

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

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

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 origin location o, 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, $G_n(\cdot|o)$. Second, each driver either receives a request from the platform ($\rho = 1$) or not ($\rho = 0$). From driver j's perspective, this event is random (and potentially location- and time-dependent). Below, we abuse notation and let $\rho_n(o)$ denote the total probability that the driver receives a request from the platform at location o 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 o 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 o. Thus, driver j's value at state (κ_{jn}, o), $S_{jn}(\kappa_{jn}, o)$, is given by[2]

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

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

We now unpack the value V_{jn} of receiving a request from the platform. Suppose driver j receives a request for a trip between locations o and d 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 d. That is, conditional on winning the auction, driver j will be at destination d $\tau(w_{jn}, x_{jn})$ periods from now. Thus, we can write driver j's payoff from bidding for this trip as follows:

$$\begin{aligned} V_{jn}(\kappa_{jn}, o, x_{jn}) = \max_b \Big\{ &\gamma(b|w_{jn}, x_{jn})\ \left(0.9\, b + \delta^{\tau(w_{jn}, x_{jn})}\, \mathbb{E}_{G_{n+\tau}(\cdot|d)}\left[S_{jn+\tau(\cdot)}(\cdot, d)\right]\right) + \\ &(1 - \gamma(b|w_{jn}, x_{jn}))\ \left(\kappa_{jn} + \delta\, \mathbb{E}_{G_{n+1}(\cdot|o)}[S_{jn+1}(\cdot, o)]\right) \Big\}. \quad \text{(D.3)} \end{aligned}$$

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:

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

[2]Implicit in our formulation is the assumption that a driver's bidding behavior does not depend on how they arrived at location o 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:

$$c_{jnr}(\kappa_{jn}, w_{jn}, x_{jn}) \equiv \kappa_{jn} + \delta \, \mathbb{E}_{G_{n+1}(\cdot|o)}[S_{jn}(\kappa_{n+1}, o)] - \delta^{\tau(\cdot)} \, \mathbb{E}_{G_{n+\tau}(\cdot|d)} \left[S_{jn+1}(\kappa_{n+\tau}, d)\right]. \tag{D.5}$$

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:

$$e^{n+\tau}_{jnr}(w_{jn}, x_{jn}) \equiv \delta \, \mathbb{E}_{G_{n+1}(\cdot|o)}[S_{jn}(\kappa_{n+1}, o)] - \delta^{\tau} \, \mathbb{E}_{G_{n+\tau}(\cdot|d)} \left[S_{jn+1}(\kappa_{n+\tau}, d)\right]. \tag{D.6}$$

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

D.3 Continuation values: Identification and estimation

We now discuss further identification results of the supply-side primitives. Because in this section the same location sometimes appears as an origin and as a destination, we denote by ℓ a generic location and by $\mathcal{L}$ the total number of locations. In what follows, we show the time and location-dependent distributions of outside options, $\{\kappa \sim G_n(\cdot|\ell) : \ell \in \{1, \ldots, \mathcal{L}\}, n \in \{1, \ldots, N\}\}$, are identified. Recall that 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 $\ell \in \{1, \ldots, \mathcal{L}\}, n \in \{1, \ldots, N\}$, there are trips to ℓ with $\tau = 0$.

(ii) For each $\ell \in \{1, \ldots, \mathcal{L}\}, n \in \{1, \ldots, N\}$, a location $\tilde{\ell} \in \{1, \ldots, \mathcal{L}\}$ and

period $\tilde{n} \in \{1, \ldots, N\}$ *exist such that there are both trips with* $\tau = n - \tilde{n}$ *and* $\tau = n + 1 - \tilde{n}$ *from* $\tilde{\ell}$ *to* ℓ.

Then, the set of conditional distributions of outside options $\{\kappa \sim G_n(\cdot|o) : o \in \{1, \ldots, A\}, n \in \{1, \ldots, N\}\}$ *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 for any trip between origin ℓ and destination ℓ'

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

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, o), $G_n(\cdot|o)$ we obtain that for $r = (t, o, d)$

$$\mathbb{E}_{G_n(\cdot|o)} \left[c_{jn}^{n+\tau}(\kappa, w_{jn}, x_{jn}) \right] = \mathbb{E}_{G_n(\cdot|o)} \left[\kappa_n \right] + \delta \, \mathbb{E}_{G_{n+1}(\cdot|o)} \left[S_{n+1}(\cdot, o) \right] - \delta^{\tau(\cdot)} \, \mathbb{E}_{G_{n+\tau}(\cdot|d)} [S_{n+\tau}(\cdot, d)]. \quad \text{(D.8)}$$

Below we use this equation for $d = \ell$ as we vary over origins $\tilde{\ell}_n$ and time periods. Under condition (ii), for each $k \geq 1$, $n \geq 0$, and destination ℓ, a location $\tilde{\ell}_n$ and a time period m_n exist such that the following hold:

$$\mathbb{E}_{G_{m_n}(\cdot|\tilde{\ell}_n)} \left[c_{m_n}^{k+n}(\cdot) \right] = \mathbb{E}_{G_{m_n}(\cdot|\tilde{\ell}_n)} \left[\kappa_{m_n} \right] + \delta \, \mathbb{E}_{G_{m_n+1}(\cdot|\tilde{\ell}_n)} \left[S_{m_n+1}(\cdot, \tilde{\ell}_n) \right] - \delta^{k+n} \, \mathbb{E}_{G_{k+n}(\cdot|o)} [S_{k+n}(\cdot, o)] \quad \text{(D.9)}$$

and

$$\mathbb{E}_{G_{m_n}(\cdot|\tilde{\ell}_n)} \left[c_{m_n}^{k+n+1}(\cdot) \right] = \mathbb{E}_{G_{m_n}(\cdot|\tilde{\ell}_n)} \left[\kappa_{m_n} \right] + \delta \, \mathbb{E}_{G_{m_n+1}(\cdot|\tilde{\ell}_n)} \left[S_{m_n+1}(\cdot, \tilde{\ell}_n) \right] - \delta^{k+n+1} \, \mathbb{E}_{G_{k+n+1}(\cdot|o)} [S_{k+n+1}(\cdot, o)]. \quad \text{(D.10)}$$

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

$$\Delta_k \equiv \frac{\mathbb{E}_{G_{m_n}(\cdot|\tilde{\ell}_n)}\left[c_{m_n}^{k+n}(\cdot)\right] - \mathbb{E}_{G_{m_n}(\cdot|\tilde{\ell}_n)}\left[c_{m_n}^{k+n+1}(\cdot)\right]}{\delta^{k+n}} =$$
$$\delta\ \mathbb{E}_{G_{k+n+1}(\cdot|o)}[S_{l+n+1}(\cdot,\ell)] - \mathbb{E}_{G_{k+n}(\cdot|o)}[S_{k+n}(\cdot,\ell)].$$

Furthermore, note that

$$\Delta_0 + \sum_{k=1}^{N-1} \delta\ \Delta_k = \delta\ \mathbb{E}_{G_{n+N}(\cdot|\ell)}[S_N(\cdot,\ell)] - \mathbb{E}_{G_n(\cdot|\ell)}\left[S_t(\cdot,\ell)\right] = (\delta - 1)\ \mathbb{E}_{G_n(\cdot|\ell)}[S_n(\cdot,\ell)],$$

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

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:

$$c_{jn}^{n+\tau}(\kappa,\cdot) = \sum_{\tilde{\ell}=1}^{\mathcal{L}} \sum_{m=1}^{N} \left(\mathbb{1}_{\tilde{\ell}=o,m=n+1} \cdot \delta\ \alpha_{\tilde{\ell}m} - \mathbb{1}_{\tilde{\ell}=d,m=n+\tau} \cdot \delta^{\tau} \cdot \alpha_{\tilde{\ell}m}\right) + \kappa. \qquad \text{(D.11)}$$

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

Moreover, the residuals from this regression provide an estimate of the conditional distributions of outside options. So, $\hat{\kappa}_{jn}$ is an estimate of $G_n(\cdot|o)$. 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, $e_{jn}^{n+\tau}(o,d) = c_{jn}^{n+\tau}(\cdot) - \kappa_{jn}$.

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.

D.5 Omitted counterfactual results

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

	Baseline	**Platform Pricing**			
	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					
lowest wait-time coeff.	4.77	4.57	4.57	4.57	4.77
	4.62	4.42	4.42	4.40	4.63
	4.45	4.25	4.26	4.24	4.46
highest wait-time coeff.	4.31	4.10	4.13	4.10	4.31
Panel B: Average Purchase Price By Wait Time Coefficient					
Average Purchase Price					
lowest wait coeff.	\$6.17	\$6.64	\$6.89	\$6.91	\$6.57
	\$6.16	\$6.60	\$6.66	\$6.64	\$6.54
	\$6.16	\$6.56	\$6.45	\$6.39	\$6.52
highest wait coeff.	\$6.15	\$6.54	\$6.28	\$6.19	\$6.51
Panel C: Average Purchase Price By Price Coefficient					
Average Purchase Price					
lowest price coeff.	\$6.20	\$6.73	\$6.71	\$7.52	\$6.67
	\$6.17	\$6.58	\$6.55	\$6.68	\$6.53
	\$6.14	\$6.48	\$6.45	\$6.22	\$6.44
highest price coeff.	\$6.10	\$6.38	\$6.35	\$5.71	\$6.34

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.

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.