Measuring the Value of Places: A Geographic Decomposition of the Value of Time*

Nicholas Buchholz † Laura Doval‡ Jakub Kastl § Ranie Lin ¶

Tobias Salz ‖

July 30, 2025

Abstract

This paper develops a framework that decomposes estimates of consumers' value of time (VOT) from transportation settings into location-specific activity values. We show that when consumers choose between faster, more expensive options and slower, cheaper alternatives, they reveal information about how they value their time spent at both origin and destination locations. We illustrate our framework using recent VOT estimates from the Prague taxi market. These estimates significantly correlate with travel flows between locations, suggesting these measures capture meaningful information about location preferences. We estimate location-specific activity values and show they are positively correlated with real estate prices, demonstrating a link between these short-run measures of value and the long-run values implied by property markets. Furthermore, a variance decomposition shows that almost 80% of activity value variation is explained by individual differences rather than location-specific factors. We demonstrate the usefulness of measuring activity values through two applications examining optimal congestion pricing and highway procurement policy, where accounting for heterogeneity in these values can improve the pricing of time and congestion externalities.

Keywords: Value of time, demand in transportation markets, ride hail.
JEL classification: C73; D83; L90; R12

*We thank Liftago for providing the data. Portions of this work originally appeared in a manuscript entitled "The Value of Time: Evidence from Auctioned Cab Rides". That original manuscript has been revised to focus on platform pricing issues, and all materials on the microfoundation and decomposition of the value of time have instead been developed here.

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

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

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

¶Princeton University, 20 Washington Rd, Princeton, NJ 08540

‖Massachusetts Institute of Technology, 50 Memorial Dr, Cambridge, MA 02142

1 Introduction

Many goods and services are rationed by waiting times, either alongside or in place of prices. When markets allow consumers to directly trade off time and money, researchers can quantify a "value of time." However, because time can neither be created nor destroyed, these measured values inherently reflect *differences* in how time can be allocated; when consumers pay to reduce waiting time, they are choosing to spend less time in one state and more in another. This suggests that estimates of the value of time can be used to learn about more fundamental economic objects, such as how consumers value different activities and the places where they occur.

This paper develops a framework that decomposes estimates of consumers' value of time from transportation settings into location-specific activity values. We argue that when consumers choose between faster, more expensive options and slower, cheaper alternatives, they reveal information not only about how they value reductions in waiting time, but also about how they value time spent at both origin and destination locations. Understanding these location-specific values informs urban policy by shedding light on how places themselves are valued, which in many cases is difficult to measure directly, for example in non-priced spaces like public parks. Recent empirical work has made significant progress in measuring consumers' willingness to pay to reduce waiting time in transportation settings from taxi and ride-hail markets (Buchholz et al., 2024; Goldszmidt et al., 2020) to public transit (Almagro et al., 2024) to private vehicles (Kreindler, 2024). Our framework provides a systematic way to extract additional insight from such measurements by decomposing them into underlying valuations of time at specific locations.

Our first contribution is theoretical: we show how to map empirical measurements of waiting time valuations into the underlying value of activities at origin and destination. These measurements can be individual and time-of-day specific. This framework can be applied wherever researchers have estimated consumers' willingness to pay for waiting time reductions.

Our second contribution is empirical: we illustrate our decomposition approach using value of time estimates from Buchholz et al. (2024), which recovers detailed estimates from a ride-hailing platform in Prague that offers explicit trade-offs between price and waiting time. As shown in their work, consumers on this platform respond substantially to changes in both price and waiting time, with an average willingness to pay for one hour of time savings of $14.05. Leveraging these estimates and their underlying location- and time-specific heterogeneity, we apply our framework to decompose origin-destination value of time into location-specific activity values. This decomposition reveals four key insights: First, the measured value of time across locations, which represents consumers' willingness to pay for waiting time reductions, is significantly correlated with travel flows across origin-destination pairs, suggesting that value of time estimates encode preferences for spending time in different areas with different available activities. Second, the subsequently estimated activity values are significantly correlated with real estate prices, suggesting a link between short-run values and long-run values implied by real estate markets. Third, almost eighty percent of the variation in the value of time is driven by differences across people, with individual-specific variation dominating location and time-of-day effects. Fourth, people who express a higher average value of time are not necessarily doing so for the same places—in fact, the relationship between high value of time individuals and high value of time locations is slightly negative.

To demonstrate the applicability of our framework, we use our activity value estimates in two applications. The first examines congestion pricing following the classical "bottleneck" model of Vickrey (1969). Our analysis reveals that optimal tolls exhibit substantial spatial variation, ranging from zero on uncongested routes to over $5 on highly congested routes in the downtown urban core of Prague. This variation depends on both congestion levels and heterogeneity in location-specific activity values. Under mild congestion, roughly half of the routes require no toll, while under severe congestion, a multi-tiered optimal pricing structure emerges. We offer the empirical link to a key insight of Vickrey (1969), that congestion pricing depends not just on traffic volumes but fundamentally on how people value time at different locations.

In the second application, we study the use of time incentives in a road infrastructure procurement setting. This exercise builds on an existing literature studying the impact of time incentives in highway procurement, demonstrating the importance of accounting for heterogeneity in the underlying social costs of road closures (Lewis and Bajari, 2011). In the setting of Prague, we show that using a uniform value of time to determine optimal time incentives, a standard approach of policy makers, will lead cities to misprice the social cost of closures by up to ninety percent depending on the location and time of road closures.

Related literature Our analysis builds on and extends the work of Buchholz et al. (2024), which uses the same ride-hailing platform data to estimate demand for taxi services as a function of both price and wait time. While their focus is on studying the welfare implications of personalized pricing strategies, our paper takes the value of time estimates as inputs to address a fundamentally different question: what do transportation decisions reveal about the spatial distribution of activity values?

In doing so, we contribute to three strands of literature.

The first strand is the literature on the opportunity cost of time. While the standard labor-leisure choice model implies that an extra hour of leisure should be valued at the shadow wage, the literature starting from the seminal work of Becker (1965) recognizes that an agent's time is also an input to other non-market activities. Recent papers in this literature have used a variety of widely available micro-data to study the trade-off between market goods and time (see, for instance, Aguiar and Hurst (2007); Aguiar et al. (2012); Nevo and Wong (2019)). Though these studies utilize rich and comprehensive datasets of consumption behavior (for example, household scanner data), they are only able to indirectly measure the opportunity cost of time through other market transactions. Within transportation economics specifically, a rich theoretical literature has developed frameworks for understanding how individuals value time spent traveling. Vickrey (1969) pioneered the 'bottleneck model' of congestion, showing how queuing time creates welfare losses and demonstrating that optimal time-varying tolls could eliminate these inefficiencies. Small (1982) further developed this approach by modeling how commuters value schedule delays relative to preferred arrival times. Arnott et al. (1990, 1993) extended these models to incorporate heterogeneity in commuter preferences and analyze the distributional effects of congestion pricing. This work has recently been revisited in the context of road pricing when carpooling is feasible (Ostrovsky and Schwarz, 2023). Our framework offers a micro-foundation to the scheduling and congestion costs that arise in these models by linking them to preferences over activities in different locations.

The second strand is the emerging literature that utilizes high resolution spatial data to shed light on urban economic activity (Athey et al., 2019; Davis et al., 2019; Couture et al., 2024; Kreindler and Miyauchi, 2021; Almagro and Domınguez-Iino, 2019). Kreindler and Miyauchi (2021) provides a particularly useful comparison as they use individual level commuting flows to estimate a gravity model where consumers choose where to work as a function of where they live and wage values. Their model also yields estimates of the relative desirability of locations. They then show that the implied values from the estimation correlate well with the empirical distribution of wages and night lights in the city. In contrast, since we directly observe how people trade off waiting time and monetary savings when choosing to move from one location to another, we can obtain a direct measure of their value of spending time at a specific location. Our approach allows us to not only quantify the intensity of preferences over locations in a way that cannot be accomplished with GPS data alone, but to also show that such quantity-based measurements are positively correlated with preference intensity.

The third strand is the literature studying the value of travel or waiting time savings in transportation markets (Small (2012), Hall (2018), Goldszmidt et al. (2020), Bento et al. (2024), Cook and Li (2023), Buchholz et al. (2024), Kreindler (2024), Castillo (2023), Rosaia (2025)).[1] A recent literature also studies equilibrium outcomes resulting from transportation infrastructure, such as the impact of new roads on driving behavior (Duranton and Turner, 2011), the impact of transit on urban development and spatial sorting (Heblich et al., 2020), the welfare effect of transportation improvements (Allen and Arkolakis, 2022), and the design of optimal transportation networks (Fajgelbaum and Schaal, 2020). Our study complements this literature by providing a conceptual framework linking consumers' willingness to pay for lower waiting times with their underlying valuations of time spent on activites at the destination and at the origin, an important component of the overall welfare of transit infrastructure. We then show how these measures can be used to more accurately set time incentives in a highway procurement setting.

Our approach differs from existing measures of spatial economic value in several key ways. While economists have extensively studied long-run effects of place of residence on individual outcomes (Chetty et al., 2018; Couture et al., 2024), residential location may imperfectly capture the benefits people derive from different places. Studies using GPS and consumption data suggest that actual spatial usage patterns are less segregated than residential patterns: Athey et al. (2019) find that time use across spatial regions is less segregated than residential locations, while Davis et al. (2019) show restaurant consumption is only half as segregated as residences. Our v measure captures valuations from individuals' full activity sets, providing a direct short-run complement to these longer-term measures.

The rest of this paper proceeds as follows. Section 2 describes the conceptual framework and model that motivates our analysis. Section 3 describes the institutional setting and data of our empirical application. Section 4 provides the details of identification and estimation of location-specific activity values using traditional value of time estimates. We present the results of our activity value estimation in Section 5. We present the applications to congestion pricing and highway procurement in Section 6 and conclude in Section 7.

[1] Taxi and ride-hail markets, in particular, have long provided a laboratory for empirical work on flexible work hours. See, e.g., Camerer et al. (1997), Farber (2005), Chen et al. (2019), Buchholz et al. (2025).

2 Model

2.1 Conceptual Framework

We now describe our conceptual framework, which shows how willingness to pay for waiting time reductions are related to the underlying values of time spent on activities in different locations. This framework maps closely to our empirical application of choices among ride-hail options, but can easily accommodate more general choices across modes. A consumer's day is characterized by the allocation of time to various activities. Those activities (e.g., leisure or family time at home versus production at work) have different productivities or different intensities of pleasure at different times and locations. Decisions on when to transition from one place to another implicitly reveal a change of value among activities at different locations. Comparing her options at any given point in time, a consumer thus decides whether or not to move to a different location and spend her time there. Moving between locations is costly, both in terms of money and time, and a consumer has a choice between various transportation options.

A consumer is defined by a pair of valuations $(v^o, v^d) : [0, T] \mapsto \mathbb{R}^2$, which denote their utility from spending time t at either the origin, O, or the destination, D. The following specification assumes that the value of the time spent on the ride is 0; alternatively, this means that v^{o} represents the value of spending time at the origin, net of the value of the ride, while v^{d} represents the value of spending time at the destination, net of the value of the ride.

Further, these functions will depend on the locations $a \in \mathcal{A}$ of the origin and the destination, which we denote by $v^o_a, v^d_{\hat{a}}$ if location a is the origin and $\hat{a}$ is the destination.

The consumer starts at time $t = 0$ at the origin. To move from the origin to the destination, a consumer has different travel options, each with different prices and waiting times, but with equal travel times Δ. Suppose, for sake of illustration, that there are two such options. If the consumer requests a ride at $q \geq 0$ with an estimated time of arrival of w and a price of p, her payoff is given by

$$u(q, w) - p = \int_0^{q+w} v^o(t)dt + \int_{q+w+\Delta}^{T} v^d(t)dt - p. \quad (1)$$

Figure 1 illustrates how the decision to choose a particular trip option affects the time allocation between the origin and the destination. The two trip options shown differ in their waiting times, w_1 and w_2, though both take the same amount of time, Δ, to travel from origin to destination once travel is underway. By choosing the option with shorter waiting time, w_1, the passenger spends $w_2 - w_1$ less time at the origin and $w_2 - w_1$ more

Figure 1: Origin Destination Trade-Off in Waiting Time Choice

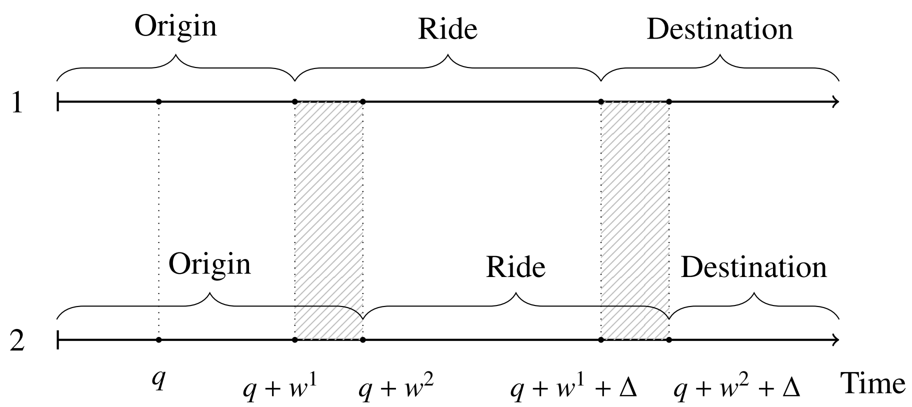

time at the destination.

2.2 Decomposition

Let a and $\hat{a}$ denote the origin and destination's locations, respectively. When considering the two trip options, the consumer compares $v^o_a \cdot (w_2 - w_1)$ against $v^d_{\hat{a}} \cdot (w_2 - w_1)$, up to an approximation. Letting the difference between w_2 and w_1 be one unit of time (for example, one minute), the willingness to pay for waiting time reductions, which we call the *value of time*, can be expressed as

$$\text{VOT}_{a \to \hat{a}} = v^d_{\hat{a}} - v^o_a. \tag{2}$$

Different people may assign different values to places, and they may assign different values to the same places at different times. Going forward, the objects that we describe above have $\hat{a}$, i and t subscripts to reflect that we recover a location-, individual-, and time-of-day-dependent distribution of the *activity value* ($v^d_{\hat{a},t,i}$ and $v^o_{a,t,i}$), which gives rise to a distribution of the *value of time* ($\text{VOT}_{a \to \hat{a},t,i}$).

Time at a destination might be more valuable per-se than time at an origin. Such differences might arise from planning and other coordinated activities.To account for the fact that destination time is more valuable than origin time, we introduce a depreciation factor and rewrite the relationship between the VOT and v as:

$$\text{VOT}_{a \to \hat{a}} = v^d_{\hat{a}} - v^o_a = v^d_{\hat{a}} - \frac{v^d_a}{v^d_a} \cdot v^o_a = v^d_{\hat{a}} - \delta_a \cdot v^d_a = v_{\hat{a}} - \delta_a \cdot v_a. \tag{3}$$

Note that riders may simultaneously travel the same in both directions, so that one person's origin is another person's destination. Following Equation 3, we interpret v^d_a

as the inherent location value and δ_a as a depreciation factor due to less (or more) productive time use at the origin. Hereafter we drop superscripts o and d.

2.3 A Useful Special Case: Ideal Arrival Times

While our general framework does not require assumptions about optimal arrival times, a useful special case emerges when we consider settings where consumers have ideal arrival times at their destinations. This special case closely relates to canonical models in transportation economics, and provides a concrete interpretation for the depreciation factors we estimate in Section 5.3.2.

In transportation economics, a central concept is consumers' ideal arrival time (Vickrey, 1969; Arnott et al., 1993). This arrival time, t^*, represents when a consumer would most prefer to reach their destination, such as the start of a meeting, work, or an appointment. The values that consumers place on time spent at origins versus destinations often change around this ideal arrival time. For instance, the value of time at a destination typically increases substantially at t^*, while the value of time at the origin might decrease as activities there conclude.

In this special case, we treat the underlying value of time functions $v_a(t)$ and $v_{\hat{a}}(t)$ as constant in neighborhoods around t^*. This allows us to work with the simplified notation where v_a and $v_{\hat{a}}$ represent the (constant) per-unit-time values in these neighborhoods. More formally, for small time differences $(w_2 - w_1)$, we have:

$$\int_{t_1}^{t_2} v_a(t)dt \approx v_a \cdot (t_2 - t_1) \tag{4}$$

where v_a is the constant value of time at the origin in the relevant time interval.

To illustrate how trip options allow us to infer changes in activity values around ideal arrival times, consider the example in Figure 2. The value of time at the origin is constant, whereas the value at the destination jumps up at t^* (from $\delta_{\hat{a}} v_{\hat{a}}^D$ to $v_{\hat{a}}^D$ with $\delta_{\hat{a}} < 1$), reflecting that the destination becomes more valuable at t^*.

If the consumer requests a trip at time q, she observes at that time the *realized* waiting times and prices for each of the available trip options, as well as the trip duration, Δ. In this example, the consumer has a choice between two trips. Trip 1 has waiting time w_1, so that the consumer departs at $q + w_1 \equiv t_O^1$ and arrives at $q + w_1 + \Delta \equiv t_D^1$, whereas trip 2 has a longer waiting time of $w_2 > w_1$. Given the realized waiting times and prices, the consumer chooses the best available option, the utility of which is evaluated using Equation 1.

Ideal arrival times give depreciation factors a particularly clear interpretation. Before

Figure 2: Activity Values Around an Ideal Arrival Time

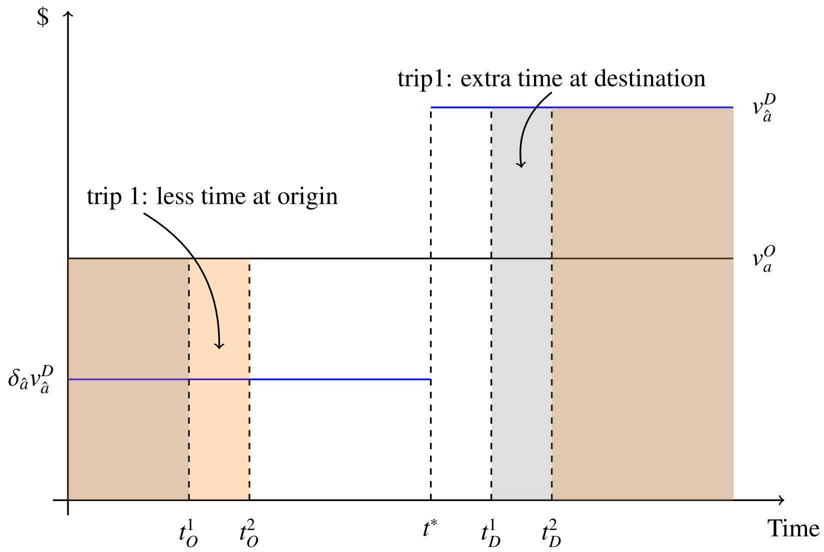

NOTE: This graph shows the differences in total activity values that two trips create. It also shows how the destination value changes after t^*.

t^*, time at the destination is less valuable because the main productive activities have not yet begun. After t^*, the full value of the destination is realized. This creates sharp changes in the willingness to pay for waiting time reductions around t^*.

When we turn to applications in Section 6, we make use of this special case to map our estimated activity values to the schedule delay parameters in canonical congestion pricing models and to further understand how construction delays affect travelers differently depending on their ability to adjust departure times.

3 Empirical Framework: Auctioned Cab Rides

Our analysis uses ride-level data from Liftago, a European ride-hailing platform operating in Prague. The dataset covers 1.9 million trip requests and 1.1 million fulfilled trips between September 2016 and June 2018. A unique feature of this platform is its auction mechanism: when a consumer requests a ride, nearby drivers submit bids, and the consumer then chooses between multiple offers that vary in both price and waiting time. Unlike traditional taxi markets where prices are fixed and the market clears through adjustments in wait time (Frechette et al. (2019), Buchholz (2022)), or modern ride-hailing platforms that use surge pricing (Castillo (2023)), Liftago's auction mechanism generates direct evidence about how consumers value the trade-off between time and money.

The data include detailed information about each request, including the pick-up and drop-off locations, the set of price bids and estimated wait times from each driver, which bid the consumer chose (if any), and unique identifiers for both drivers and consumers. Over our sample period, we observe 1,455 unique drivers and 113,916 unique consumers.

For the analysis of location-specific activity values, we augment the core ride data with several spatial datasets. First, we partition Prague into distinct locations using a spatial k-means clustering algorithm applied to the GPS coordinates of trip origins, giving 30 locations with approximately equal numbers of originating trips. This data-driven approach ensures our spatial units reflect actual travel patterns rather than administrative boundaries. Figure 6 shows the resulting partition overlaid on a map of Prague together with empirical results.

We complement these ride data with location characteristics from several sources: Land values from municipal GIS data, providing a measure of long-run location value, public transit accessibility metrics from the Google Transit Matrix API, hourly Prague weather data from NOAA weather stations.

Our empirical strategy leverages two key features of these data. First, the panel structure allows us to estimate individual-specific preferences. Second, the spatial richness of observing consumers traveling between all possible origin-destination pairs enables us to decompose aggregate willingness-to-pay measures into location-specific values.

3.1 Discrete Choice Framework

The estimation of location-specific values of time requires two key ingredients: (1) estimates of consumers' willingness to pay for waiting time reductions across different origin-destination pairs, and (2) a framework for decomposing these estimates into location-specific values. The literature has taken various approaches to measuring consumers' willingness to pay for waiting time reductions. Some papers exploit variation in tolled versus untolled lanes (e.g., Hall (2018)), while others use experimental variation in prices (e.g., Goldszmidt et al. (2020)) or analyze how consumers trade off different arrival times (e.g., Kreindler (2024)). Our empirical analysis builds on the discrete choice framework developed in Buchholz et al. (2024), which leverages auction-based variation in prices and waiting times. While our decomposition framework can accommodate different approaches to estimating the willingness to pay for waiting time reductions, we briefly summarize their model here for completeness.

The basic model analyzes consumers' (indexed by i) choices over options for taxi rides. Each ride is indexed by $r = (t, o, d)$, a tuple of origin o, destination d, and the date and time the request is submitted, t. A request then generates J_r offers from nearby drivers, where each offer $j \in \{1, ..., J_r\}$ is characterized by a price (or bid) b_j, a wait time w_j, and other observable characteristics that shift riders' utility (such as car quality or driver ratings), summarized in the vector x_j. Consumer i's utility from option j can be written as:

$$u_{ijr} = \beta^w_{ir} w_j + \beta^p_{ir} b_j + \beta^x_{ir} x_j + \xi_r + \varepsilon_{ijr} \tag{5}$$

where β^w_{ir} and β^p_{ir} capture individual-specific preferences over wait time and price, ξ_r represents unobserved demand conditions for request r, and ε_{ijr} is an idiosyncratic error term. The panel structure of the data, with repeated observations for each consumer, allows us to estimate these individual-specific preferences. With demand preferences, we can recover $\textsc{vot}_{a \to \hat{a},t,i}$ as the ratio $-\beta^w_{ir}/\beta^p_{ir}$.[2] Our focus in this paper is on using these estimates as an input to recover location-specific values of time, as outlined in our conceptual framework above.

4 Estimation

We employ a two-stage estimation approach to estimate location-specific activity values. In the first stage, we follow Buchholz et al. (2024) to estimate the demand model specified above using maximum likelihood with a control function to address price endogeneity. We estimate random coefficients that capture consumer-specific heterogeneity in preferences using an MCMC procedure to recover the distribution of individual-specific coefficients (β^w_i, β^p_i). The control function, based on persistent driver-specific pricing differences, addresses potential endogeneity in prices. This approach provides us with individual-specific estimates of VOT that account for persistent differences in how consumers trade off time and money.

In the second stage, we use the first stage estimates and decompose them into location-specific activity values, While our framework allows for estimates at arbitrary temporal granularity, for computational tractability, we group times of day t into coarser blocks denoted by h_t. The notation h_t denotes hourly blocks from 1am–5am, 6am–9am, 10am–3pm, 4pm–6pm, and 7pm–12am. We aggregate locations into thirty neighborhoods shown in Figure 6. This partitioning allows us to characterize preference heterogeneity across both locations and meaningful time periods while maintaining sufficient statistical power for estimation.

[2] We refer readers to Buchholz et al. (2024) for complete details on identification and estimation of these demand parameters.

To decompose VOT into location-time activity values, we use a moment estimator based on the relationship between the observed $\text{VOT}_{i,h_t,a\to\hat{a}}$ and the unknown pairs of $v_{i,\hat{a},h_t}$ and $\delta_{i,a,h_t} \cdot v_{i,a,h_t}$:

$$\text{VOT}_{i,h_t,a\to\hat{a}} = v_{i,\hat{a},h_t} - \delta_{i,a,h_t} \cdot v_{i,a,h_t} + \eta_{i,h_t,a\to\hat{a}}, \tag{6}$$

where we interpret $\eta_{i,h_t,a\to\hat{a}}$ as orthogonal measurement error. For each i and h_t there are $2 \cdot N_a - 1$ parameters after normalizing one origin value. For each pair of (i, h_t) we observe N_a^2 equations.

Denoting the collection of v's as $\overline{v}$ and requiring that all activity values are (weakly) positive, we can specify an optimization problem as

$$\begin{aligned} \underset{\overline{v}}{\text{minimize}} \quad & \sum \eta^2_{i,h_t,a\to\hat{a}} \\ \text{subject to} \quad & 0 \leq v_{i,a,h_t}, \ \forall i, h_t, a \ \text{ and } \ \delta_{i,1,h_1} = 0 \ \forall i. \end{aligned}$$

Appendix A provides conditions for identification and shows that we need one location normalization. We normalize the origin v of a single location-time to zero. We select the reference point as $h_t = 12\text{am}$ at location $a = 1$. With this normalization, the magnitude of activity values are interpreted relative to the value of the normalized origin location and time. Recall locations become origins when consumers are departing them, which we further interpret as a state in which activities are finished. To estimate the model we use JuMP (Dunning et al. (2017)).

5 Results

In this section, we first present the demand estimates that capture consumers' willingness to pay for waiting time reductions across different origin-destination pairs. The underlying preference estimates come from Buchholz et al. (2024), which we transform in to the VOT estimates that serve as inputs to our spatial decomposition. We demonstrate that VOT estimates embed meaningful information about underlying location preferences by examining their correlation with observed travel flows between locations. Next, we move to our main contribution: the decomposition of these estimates into location-specific activity values. We then use this decomposition to convey four key insights about activity values: (1) the spatial distribution of activity values across Prague, (2) the correlation between our activity values and real estate prices, (3) a variance decomposition that isolates the contributions of individual, location, and time-of-day heterogeneity, and (4) the productivity differential between the beginning and end of activities.

5.1 Demand Estimates

Table 1: Value-of-Time Estimates

Subsample	Value of Time (VOT)					
	12a–6a	6a–9a	10a–2p	3p–6p	7p–12a	All Hours
Panel A: Location-Time Heterogeneity						
All O-D Pairs (mean)	9.76	11.38	11.87	10.16	8.05	9.9
25th pct.	6.75	8.64	9.18	6.74	5.43	7.14
50th pct.	9.53	11.97	12.35	10.17	8.11	9.78
75th pct.	12.15	15.27	15.39	12.64	10.15	12.23
City-Center O-D Pairs (mean)	12.49	14.91	14.53	12.22	10.23	12.55
25th pct.	8.54	11.81	11.76	9.93	8.03	9.82
50th pct.	10.93	14.18	14.35	11.88	9.54	11.56
75th pct.	13.13	16.68	16.84	14.05	11.43	14.05
Non-City-Center O-D Pairs (mean)	6.67	7.49	8.98	7.94	5.67	7.14
25th pct.	4.87	6.48	7.21	4.55	3.56	5.44
50th pct.	7.65	9.29	10.02	7.52	6.09	7.56
75th pct.	10.15	11.91	12.55	10.31	8.06	9.59
Panel B: Individual-Time Heterogeneity						
All Types (mean)	14.14	15.89	16.09	13.9	12.5	14.05
H Price, H Wait	13.67	16.84	17.17	14.92	13.36	14.66
H Price, L Wait	4.08	6.67	7.03	5.21	3.68	4.91
L Price, H Wait	21.3	25.39	25.91	23.11	19.88	22.76
L Price, L Wait	19.27	12.04	12.15	11.07	15.71	13.85

NOTE: This table provides VOT estimates from Buchholz et al. (2024). All estimates are presented in USD per hour.

Table 1 presents the VOT results implied by our estimated demand model. The average implied value of time across all trips is $14.05 per hour, with substantial heterogeneity across consumer types, locations, and times of day. Panel A shows variation in VOT across different origin-destination-time groups, with the mean taken across individuals who request trips in that group. This panel shows that demand estimates for trips to and from the city center exhibit higher VOT. A similar pattern applies to trips in the middle of the day. Panel B characterizes demand heterogeneity across individuals for sub-categories of consumers with above- and below-median price and waiting time sensitivities. This panel demonstrates the wide range of individual-level preference heterogeneity. For example, consumers with low price-sensitivity and high wait-time sensitivity exhibit an estimated VOT that is over 50% larger than the overall average. Table 1 emphasizes the spatial VOT variation we use in our decomposition exercise; however, a summary of all underlying coefficient estimates and additional VOT-specific analysis is provided in Buchholz et al. (2024).

5.2 Comparison with Travel Flow Measures

Before moving to a decomposition of VOT, we first demonstrate that VOT estimates themselves capture meaningful information about how people value different locations. To do this, we compare our estimates to travel flows, a measure of how many travelers move from one point to another, a commonly used and widely-available alternative measure. For example, Kreindler and Miyauchi (2021) use GPS-based commuting patterns to infer relative wages across locations, assuming that areas with higher wages disproportionately attract workers.

While such approaches have provided valuable insights, they typically require making structural assumptions about specific amenities (like wages) that make locations attractive. Our approach offers a direct monetary quantification of time value through revealed preferences in transportation choices. If such VOT estimates are capturing meaningful location values, they should correlate with travel flows—people should generally travel more frequently to destinations they value highly.

Figure 3 shows this relationship across the 900 different directional origin-destination pairs in our data. As the plot illustrates, the two measures are highly significantly correlated with a correlation coefficient of 0.56 and a t-statistic of 20.46. This correlation is highest during the middle of the day and lowest at midnight.

These results imply that travel flows provide a reasonable but imperfect proxy for time valuations. The strong correlation validates our VOT estimates as capturing meaningful information about location preferences. However, the variation in correlation throughout the day indicates that the relationship between flows and values is not constant. Moreover, the fact that there are some high flow routes that have low values provides one rationale for a congestion surcharge that disincentivizes low value trips. In the next section, we use our decomposition approach to extract more precise location-specific activity values from these VOT estimates, which will allow us to examine the underlying drivers of these patterns in greater detail.

5.3 The Value of Activities: Quantification and Interpretation

We now turn to results surrounding location-specific activity values, obtained via our decomposition of underlying VOT estimates. We present and interpret several stylized

Figure 3: Relationship between Travel Flows and the VOT

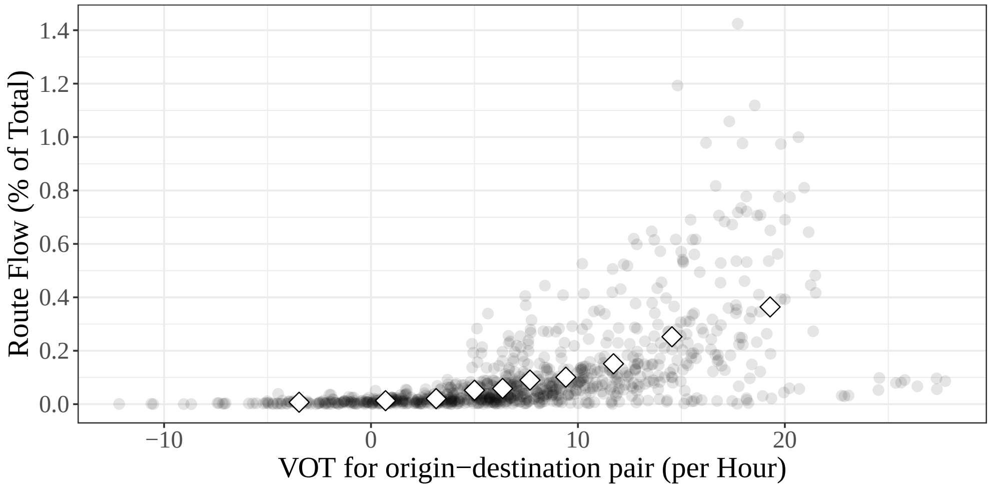

NOTE: This graph shows the scatter (transparent round dots) and binscatter (white diamonds) relationship between the VOT for an origin-destination pair and the respective traffic flow (as a fraction of the total).

facts about the distribution of activity values across locations, time, and individuals, and the depreciation of value between the beginning and end of activities.

Figure 4: Histogram of v

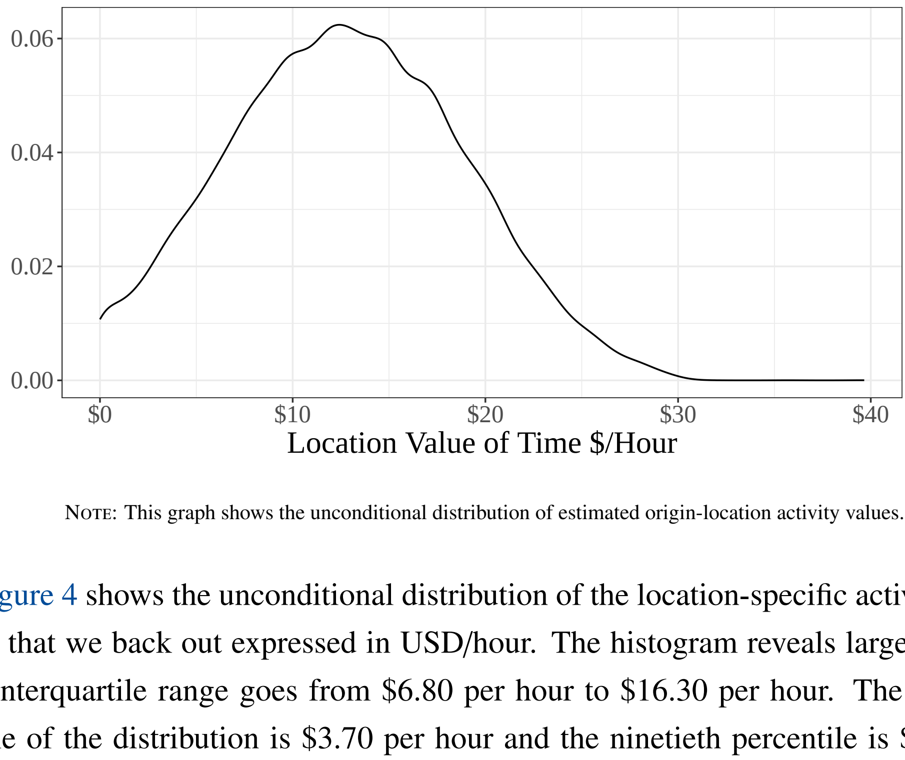

NOTE: This graph shows the unconditional distribution of estimated origin-location activity values.

Figure 4 shows the unconditional distribution of the location-specific activity values v_{i,a,h_t}, that we back out expressed in USD/hour. The histogram reveals large variation. The interquartile range goes from $6.80 per hour to $16.30 per hour. The tenth percentile of the distribution is $3.70 per hour and the ninetieth percentile is $21.30 per hour.

Spatial Heterogeneity Figure 6 shows heterogeneity by location, measured as the average across people and time of day of all v estimates for this location. Going from the lowest-value location to the highest-value location implies a difference of \$10.83 in the hourly value of time, or an increase of 51.96%. However, the interquartile range is already substantially smaller with a difference of \$2.89, or 17.78%. This relatively compressed distribution suggests that fixed location characteristics are not the largest driver of variation in activity values.

These results provide direct measures of short-run monetary valuations for spending time at different locations. Unlike real-estate pricing measures, our v estimates encompass valuations from individuals' complete activity patterns across all locations in our sample.

To explore how our value of time measures compare to other quantities that capture spatial differences in economic activity, we correlate our v estimates with real estate values. We use a binscatter plot to show this relationship in Figure 5. There is a strong positive relationship between real estate prices and our v measures. However, the "returns" to higher land values are diminishing. From a regression of log-v on log land-values we measure an elasticity of 0.25 ($p < 0.001$). So for every 1% higher land value we measure a value of time that is 0.25% higher. On the one hand, this strong positive relationship validates our measures since we have not used land values at any point in the estimation. On the other hand, we find that with an R^2 of less than 0.02, real estate prices can only explain a very small percentage of the variation in v. This finding highlights how our value of time measures serve as a unique short-run analog to other widely available long-run measures of economic value across space. Such measures are important when policies target short-run decision making. We illustrate this in our application on procurement incentives.

Figure 6: Map of v Estimates in Prague

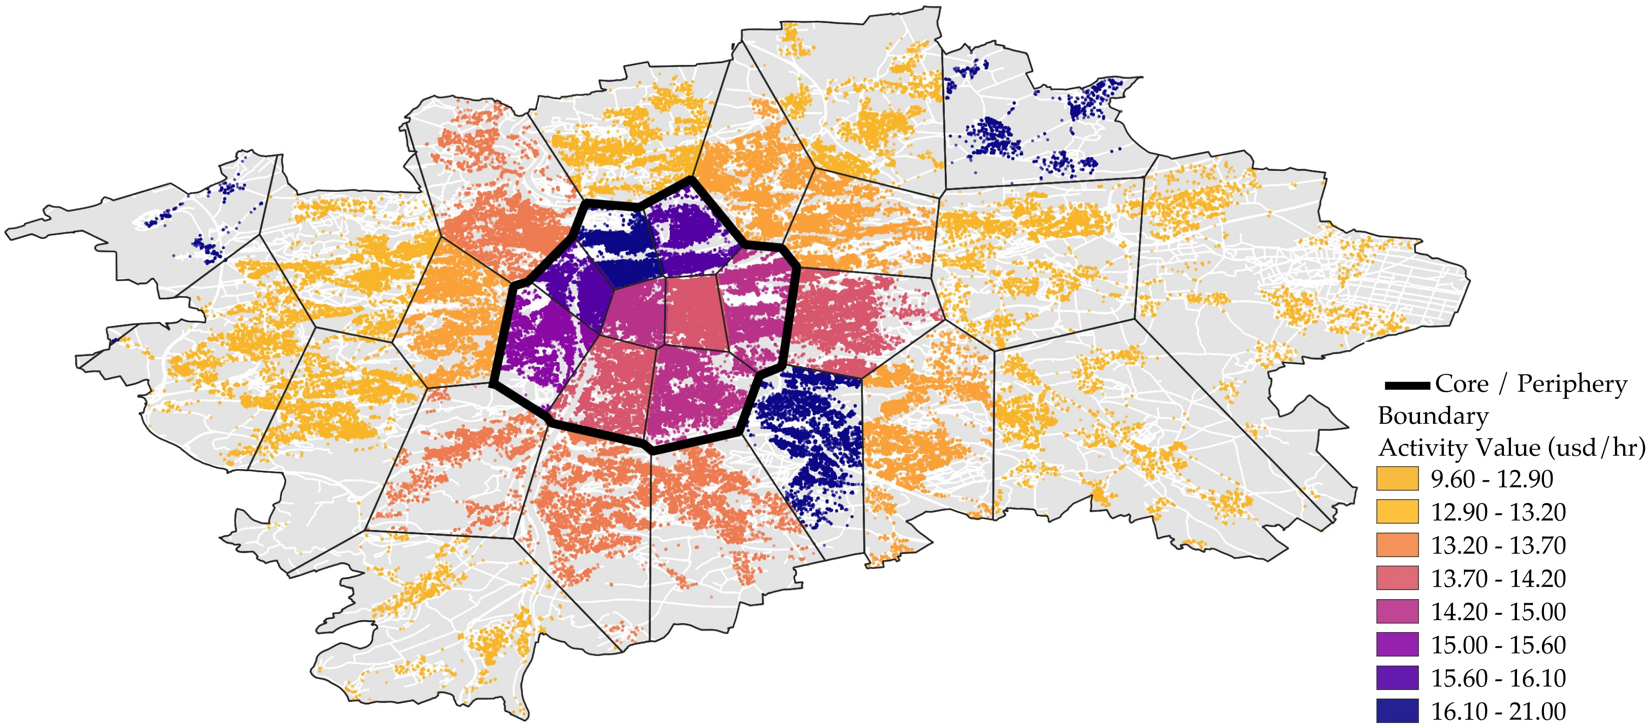

NOTE: This figure depicts each GPS destination point in the Liftago data shaded by the average value of activities in that destination, or v_a^d across places $a \in \{1, \ldots, 30\}$ within the city limits of Prague. White lines depict the city's street map. The bold black line delineates the boundaries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

Figure 5: Location-specific VOT and Land Values

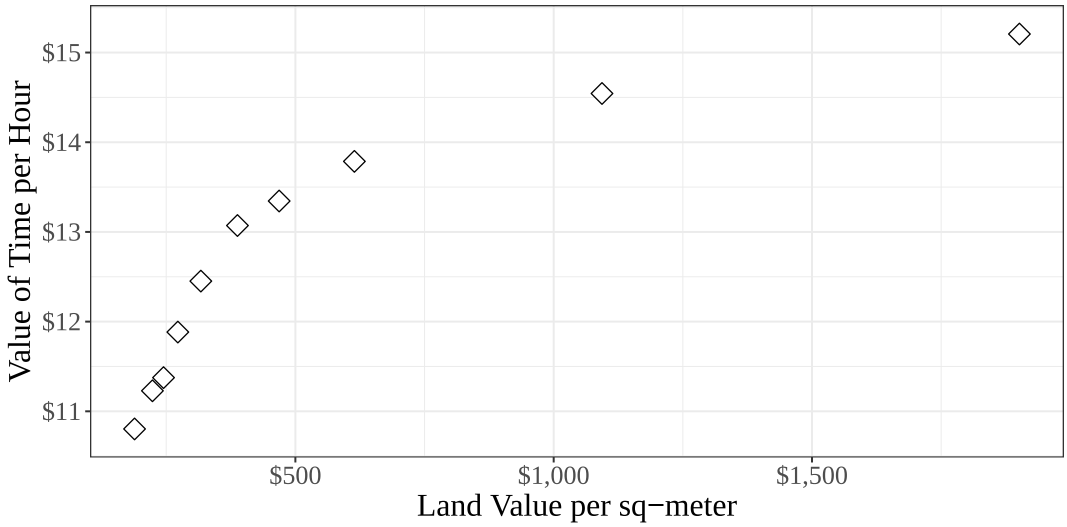

NOTE: This figure presents a binscatter plot of the value of time in each location compared to the average land values within each location. The x-axis orders the thirty locations by their average land value. The y-axis shows the average value of time in each of these locations.

Time-of-day and Individual Heterogeneity For this decomposition, we group time of day into two categories – work-hours, defined as the hours 6am–6pm, and non-work-hours. We group individuals by common preferences, just as in Table 1, according to whether the disutility of money is high (above-median, denoted H) or low (below-median, denoted L), and similarly with waiting time. Table 2 gives an overview of means and standard deviations of the v estimates across these times of day for the four different groups of passengers. During both work and non-work hours, the (L Price, H Wait) group expresses a v that is almost three times that of the (H Price, L Wait) group. Within type, however, the difference between work-hours and non-work-hours is modest. Across all groups the ratio of non-work-hours to work-hours v ranges from 74% to 87%. Those ratios are ordered in terms of the average v of these groups with (L Price, H Wait) types having the highest and (H Price, L Wait) types the lowest.

5.3.1 Variance Decomposition

Our activity estimates are indexed by time, location and person. Each of these categories contribute to the variation in activity values. In this section, we ask whether high-value individuals (i.e., those who place a higher value on average across all locations) tend to spend time in locations that have high value (i.e., regions that have systematically higher values for everyone else). If most of the variation in the value of time is driven by locations and not people, this would suggest that location-specific factors provide equal

Table 2: Summary Statistics of Value of Time

	Work Time (USD)		**Non Work Time** (USD)		**Non Work Time** v /
	Mean	STD	Mean	STD	**Work Time** v (%)
Location Values (v_{i,a,h_t})					
All	17.15	10.29	14.02	10.39	0.82
H Price, H Wait Sensitivity	19.18	7.0	15.95	7.05	0.83
H Price, L Wait Sensitivity	9.79	4.82	7.25	6.24	0.74
L Price, H Wait Sensitivity	27.05	12.43	23.45	12.73	0.87
L Price, L Wait Sensitivity	12.66	4.64	9.83	5.85	0.78

Note: This table shows summary statistics on the estimated value of time overall and across four groups of individuals. We separate estimates by the work hours from 6am–6pm and non work hours from 6pm–6am. The last column shows value of non-work time as a percentage of the value of work time.

benefit to different people. For example, there might be differences in the provision of public goods across locations, which are equally available to all people. On the other hand, differences in the value of time in different locations might be predominantly driven by differences across individuals. This would be true, for example, if high-value individuals derive utility from places solely because their workplaces are there, and otherwise derive no amenity value from public services.

We perform a variance decomposition to separate these two sources of variation and investigate their relationship in more detail. We take the approach of Abowd et al. (1999), which has previously been used in the labor literature to decompose wages into firm-specific and worker-specific components. For the variance decomposition we derive an accounting equation from the following regression model:

$$v_{i,a,h_t} = \alpha_i + \rho_{h_t} + \gamma_a + \epsilon_{i,t}, \tag{7}$$

where α_i is a person fixed effect, ρ_{h_t} is a time-period fixed effect, γ_a is a location fixed effect, and $\epsilon_{i,t}$ is a residual.[3] The regression gives rise to the following variance decomposition:

$$\begin{aligned} Var(v) = Var(\alpha) + Var(\rho) + Var(\gamma) + \\ 2 \cdot Cov(\alpha, \rho) + 2 \cdot Cov(\alpha, \gamma) + 2 \cdot Cov(\gamma, \rho) + Var(\epsilon). \end{aligned} \tag{8}$$

Table 3 shows the results from this exercise. At 78%, by far the largest contributor to the observed variance of the v are differences across people. In comparison, location-

[3] Our panel is not long enough to observe a given individual in all possible combinations of origin, destination and time of day. To deal with this curse of dimensionality, we replace individual values with their respective percentile mean (out of one hundred percentiles), so i becomes an indicator of each percentile bin.

specific variance is small and accounts for 10% of the total. Intra-daily changes in the v contribute another 9%.

Table 3: Variance Decomposition of the Value of Time

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

NOTE: The first row of this table shows the variances of each of the components of the model in Equation 13. The second row shows the variance of each of the components divided by the total variance, which can loosely be interpreted as fractions of the total variance.

Interestingly, with a covariance of −0.4 the sorting effect is slightly negative, implying that people who express a high value of activities do not spend more time in places with activity values that are high for everyone.

5.3.2 How Much More Valuable are Destinations Compared to Origins?

We next examine the differential time values estimated in a given location, when that location is an origin compared with when it is a destination, as measured by δ_{i,a,h_t}. As discussed in Section 2.3, these depreciation factors can be interpreted as capturing the reduced productivity at origins relative to destinations, particularly around ideal arrival times. Table 4 gives an overview. Averaging across the entire sample, time at a location *as an origin* is only 18% as valuable as time at the same location as a destination during work hours, while it is 20% as valuable during non-work hours. Holding the location fixed, meetings and other types of work-related activities lead to a larger discrepancy in the value of time when activities are beginning versus when they have ended. Moreover, there is a large heterogeneity in δ within the population. Not surprisingly, people that are more waiting time sensitive (H Wait Sensitivity or above median disutility in w) have a larger discrepancy in the value of time between the beginning and end of activities. This shows that higher willingness to pay for lower waiting times is driven both by a larger inherent value of time at any given location and also a bigger discrepancy between the value of time at the start and end of activities.

To summarize our results, we find large differences in the value of time across individuals and also large differences in differentially productive time use across locations. Most of the variation in the value of time comes from person-specific heterogeneity as opposed to location-specific heterogeneity. Interestingly, these results suggest that in the short run, different places do not confer benefits that are equally valued by ev-

Table 4: Summary Statistics for δ_{i,a,h_t}

	Work Time		Non Work Time	
	MEAN	STD	MEAN	STD
Origin Depreciation Factors (δ_{i,a,h_t}) in USD				
All	0.18	0.17	0.2	0.24
H Price, H Wait Sensitivity	0.14	0.11	0.12	0.11
H Price, L Wait Sensitivity	0.32	0.34	0.5	0.79
L Price, H Wait Sensitivity	0.14	0.12	0.14	0.12
L Price, L Wait Sensitivity	0.18	0.18	0.21	0.25

NOTE: This table shows summary statistics for our estimates that capture how much less productive time at the origin is relative to the destination.

eryone. These results are also an important short-run complement to the recent debate around the importance of place-specific factors for long-run outcomes. Another interesting implication lies in the finding that the value of time for people who are highly sensitive to waiting time, in part, be explained by a greater need for planning as measured by the origin depreciation factor δ. This suggests that these individuals depend to a larger extent on complementary inputs that require coordination. For example, high skilled work places might require more meetings and coordination with others to use time productively.

6 Applications

The value of time spent in different activities and locations fundamentally shapes transportation decisions and policy. While the transportation economics literature has long recognized that people value time differently at home versus work (Vickrey, 1969), these activity values have typically been assumed or calibrated rather than directly estimated from revealed preferences. In this section, we apply our location-specific activity value estimates to two important transportation policy problems: congestion pricing and optimal time incentives in highway procurement. Our method to decompose VOT estimates into underlying activity values enables recovering the heterogeneous valuations of time at origins and destinations that are essential for both applications.

6.1 Congestion Pricing

Our first application examines congestion pricing, a policy approach that uses tolls to manage traffic congestion, particularly during peak times. While congestion pricing has long been advocated by economists, its implementation has been limited in practice. Our activity value estimates can help design efficient tolling systems that account for the heterogeneity in how people value time across different locations and times of day.

For this application, we adopt the special case framework from Section 2.3, which maps directly to the classical "bottleneck" model of Vickrey (1969), further developed by Arnott et al. (1993) and now studied in modern theoretical applications (e.g., Ostrovsky and Schwarz (2023)). An important feature of this framework is that optimal congestion pricing does not require knowledge of the value of travel time itself (commonly denoted as α in the literature). As we review below, the optimal tolls depend only on the schedule delay costs: the disutility of arriving early (β) or late (γ), which our decomposition can directly identify from the relative activity values at origins versus destinations. This makes our estimates particularly well-suited for congestion pricing applications.

In the model, N identical commuters with ideal arrival times uniformly distributed on a time interval $[t_1, t_2]$ travel from origin a to destination $\hat{a}$ with a "bottleneck" in between. The bottleneck has a capacity constraint: only $n_o < \frac{N}{t_2 - t_1}$ commuters can pass through per time increment, which gives rise to congestion. As congestion forms, consumers must wait to pass through the bottleneck. The waiting time to pass through the bottleneck at time t (or equivalently, arrive at destination $\hat{a}$ at time t) is denoted as $q(t)$, or time spent waiting in queue.

A commuter who arrives at t with ideal arrival time t^* incurs costs:

$$C(t) = \alpha q(t) + \beta(t^* - t)\mathbb{1}[t < t^*] + \gamma(t - t^*)\mathbb{1}[t \geq t^*]. \tag{9}$$

As we depict in Figure 2, the value of each minute of late arrival (that is, after t^*) can be derived as the difference in spending an additional minute of *unproductive* time in an origin location a and one minute less of *productive* time in the destination location $\hat{a}$. The value of early arrival is defined analogously, but now time at origin is productive while time at destination is not, before reaching t^*. Since origins and destination values are linked through equation (3) we can directly express the early and late arrival parameters (β and γ) from the theoretical literature in terms of our estimated objects:

$$\beta = v_a - \delta_{\hat{a}} v_{\hat{a}} \tag{10}$$

$$\gamma = v_{\hat{a}} - \delta_a v_a \tag{11}$$

where v_a and $v_{\hat{a}}$ are the values of time at the origin and destination, and δ_a and $\delta_{\hat{a}}$ are the depreciation factors at those locations. This specification captures our finding that commuters value time differently at origins versus destinations, and that this valuation changes with activity status. In what follows, we describe key results from the literature, closely following Vickrey (1969), Arnott et al. (1993), and Ostrovsky and Schwarz (2023) to show how we compute optimal tolling.

6.1.1 No-toll Equilibrium

In the no-toll equilibrium, each commuter chooses their departure time $t - q(t)$ (or equivalently, arrival time t) to minimize costs. Let t_q and $t_{q'}$ denote the start and end times, respectively, of the queue. The equilibrium queue length function $q(t)$ satisfies:

$$q'(t) = \begin{cases} \frac{\beta}{\alpha}, & \text{for } t_q \leq t < t_0 \\ -\frac{\gamma}{\alpha}, & \text{for } t_0 \leq t < t_{q'} \end{cases} \tag{12}$$

where $t_0 = t_1 + r(t_2 - t_1)$ is the time at which the queue reaches its maximum length, and $r = \frac{\gamma}{\beta+\gamma}$ represents the fraction of commuters who arrive early. The maximum queueing time, occurring at t_0, is $q(t_0) = \frac{N}{\alpha n_0}\frac{\beta\gamma}{\beta+\gamma}$. Further details are provided in Appendix C.

6.1.2 Optimal Congestion Toll

The socially optimal policy eliminates queueing entirely while maintaining the same pattern of arrivals as in the no-toll equilibrium. This is achieved through a time-varying toll $p(t)$ to pass the bottleneck (or equivalently, arrive at their destination) at time t that exactly equals the queueing costs, $\alpha q(t)$, that would have occurred without the toll. As shown in Appendix C, the toll schedule is given by:

$$p(t) = \begin{cases} \beta(t - t_q), & \text{for } t_q \leq t < t_0 \\ \frac{N}{n_0}\frac{\beta\gamma}{\beta+\gamma} + \gamma(t_0 - t), & \text{for } t_0 \leq t \leq t_{q'} \end{cases} \tag{13}$$

The toll starts at zero at time t_q, rises linearly at rate β to a maximum of $p(t_0) = \alpha q(t_0)$ at time t_0, and then falls linearly at rate γ to zero at time $t_{q'}$.

The total revenue collected from tolls is $\frac{1}{2}N\alpha q(t_0) = \frac{1}{2}\frac{N^2}{n_o}\frac{\beta\gamma}{\beta+\gamma}$, and the average toll paid per commuter is:

$$\frac{1}{2}\left(\frac{N}{n_o}\right)\left(\frac{\beta\gamma}{\beta+\gamma}\right) \tag{14}$$

While Equation 14 shows that early and late arrival costs are sufficient to quantify the congestion toll (holding fixed the congestion level N/n_0), the activity values v that we recover serve as underlying primitives that generate these costs. This insight provides a mapping between VOT, activity values v, and early/late arrival costs β and γ. As Equation 10 and Equation 11 show, simple estimates of VOT may be sufficient to estimate β and γ. However, this connection obscures the important nuance that VOT estimates are directional and route-specific. Expressing them as a difference of vś clarifies why: because VOT encodes a difference in destination value versus origin value. Further, directly estimating β and γ with VOT data imposes stronger data requirements, as VOT is estimate on a route-by-route basis whereas v is estimated location-by-location (where the set of routes is the square of the number of locations).

To apply this theoretical framework to multi-route urban settings, we need to map the single-bottleneck model to more realistic street networks. In practice, multiple routes typically connect any two regions in a city. We consider a multi-route equilibrium where congestion equalizes across all routes between two locations, with route demand proportional to each route's throughput capacity. This yields a common effective N/n_o ratio for all congested routes between an origin-destination pair, where N represents total demand and n_o represents the effective bottleneck capacity. The optimal toll, given by Equation 14, depends on this ratio and the underlying activity values.

We calibrate traffic demand on a per-road basis according to Prague government statistics, which estimate 70,000 commuters per day on the city's busiest road. We thus set a baseline "average" traffic volume at 75% of this estimate, and further assume 25% of daily demand occurs during the morning rush hour. We assume a route throughput equal to 5,000 cars per hour, approximately halfway between a two- and three-lane road. These calibrations yield $N = 6560$ and $n_o = 83$ cars per minute. This baseline represents a "high" congestion state. We reduce demand by 40% to define a "medium" congestion state. Note that the optimal toll depends only on the ratio N/n_o, making it invariant to proportional scaling of both parameters. In this sense, this ratio may be proportionally scaled to accommodate the entire traffic network.

Figure 7 illustrates the distribution of optimal congestion tolls across Prague under two congestion scenarios. In Scenario 1 (panels a and c), core routes experience mild congestion ($N/n_o = 31$) while peripheral routes remain uncongested. The resulting toll distribution is bimodal: approximately half of all routes require no toll, while congested core routes face tolls clustering around \$2, with none exceeding \$2.50. In Scenario 2 (panels b and d), core routes experience high congestion ($N/v = 79$) and peripheral routes now face mild congestion ($N/n_o = 31$). This produces a trimodal toll distribution with almost no zero-toll routes, and three distinct pricing tiers: mildly congested

Figure 7: Optimal Congestion Toll Comparison

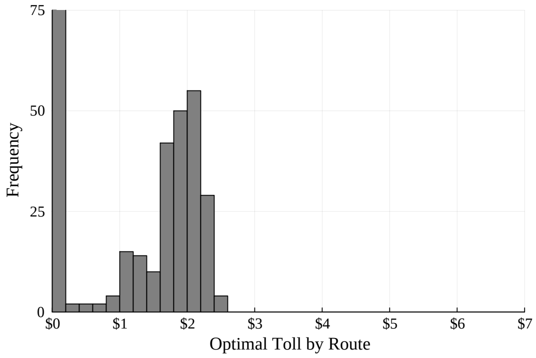

(a) Distribution of Tolls: Scenario 1

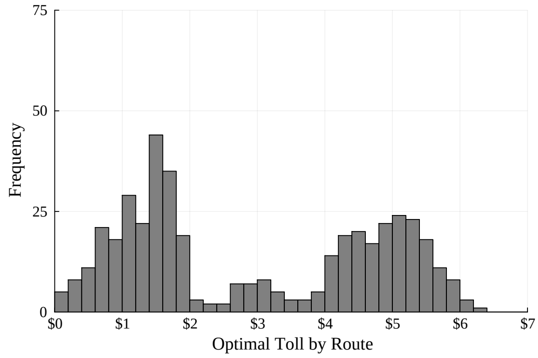

(b) Distribution of Tolls: Scenario 2

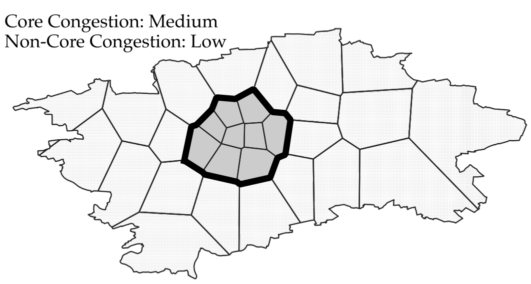

(c) Congestion Map: Scenario 1

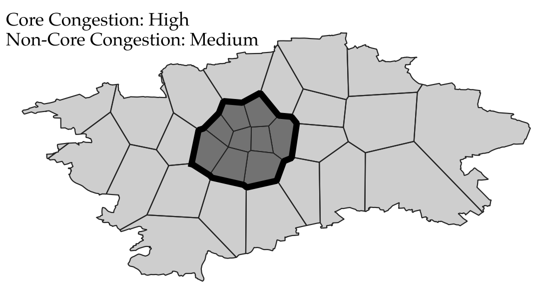

(d) Congestion Map: Scenario 2

NOTE: This figure shows the distribution of optimal congestion tolls in two scenarios. In scenario 1, routes between and within the inner core of Prague are mildly congested ($N/n_o = 31$) and outer regions are not congested. In scenario 2, core routes are highly congested ($N/n_o = 79$) and outer routes are mildly congested ($N/n_o = 31$).

peripheral routes around \$1.50, moderately congested transition zones around \$3.00, and highly congested core routes around \$5.00.

These patterns demonstrate how optimal congestion pricing varies substantially with both congestion levels and geographic heterogeneity in activity values. The doubling of core tolls between scenarios and the emergence of multiple pricing tiers highlight the importance of location-specific valuations in toll design. While our activity value estimates provide crucial inputs for such pricing schemes, several implementation challenges remain beyond our scope. Travelers may endogenously adjust their activities when value differentials are taxed away, potentially altering the equilibrium we analyze. Moreover, city engineers must determine precise N/n_o ratios for each route to effectively implement location-specific tolls. Finally, as the surplus gains from congestion pricing accrue entirely as government revenue, the optimal redistribution of these revenues remains an important open question for policy design.

6.2 Highway Procurement

Our second application considers the problem of using time-incentives in a highway procurement setting. By design, repairing and maintaining highways entails closures which slow down traffic and may induce substantial congestion. Some state and local authorities use time-incentives as part of procurement procedures, which enable contractors to earn higher payments for faster completion or require them to pay fines for delays. This problem was introduced to the economics literature by Lewis and Bajari (2011), where the authors examine CalTrans data to study the welfare gains associated with procurement auctions that incorporate bidding on both time to completion and price into a scoring auction framework. The authors find that such procurement mechanisms used by CalTrans raised social welfare by 30% of project costs by generating time savings despite higher project prices.

To quantify the social cost of traffic delays due to construction, the authors apply a uniform \$12 per hour value of time. This value is derived from engineering estimates and accounts for the possibility of multiple passengers in the same car.[4] Since our approach allows us to recover economic estimates of time valuation, we re-consider the problem of highway procurement and calculating the social costs of road work. Because road work impacts specific routes, we use our estimates to further explore how scoring rule incentives vary with underlying location-specific heterogeneity in activity values. By examining the impact of road closures in different areas of the city and,

[4] CalTrans does not specify the source of the underlying calculation. Updated estimates from the same source now put this value at \$13 per hour. See https://dot.ca.gov/-/media/dot-media/programs/maintenance/documents/office-of-concrete-pavement/life-cycle-cost-analysis/value-of-user-time-2013-a11y.pdf.

potentially, at only certain times of day, we show that the social cost of these closures can vary considerably. This result underscores the importance of incorporating value of time heterogeneity into optimal time incentives.

Travel Time Valuations Our framework distinguishes between the value of activities at origins and destinations, but does not directly address how time spent in transit is valued. In the context of road construction delays, we assume travel time value is valued at zero, implicitly assuming equal values to that of the normalized unproductive location (e.g., a peripheral location as an origin). However, for more granular results, the value of time in a car and any route-specific heterogeneity in it could be estimated by combining our methodology with tollway pricing studies (e.g., Hall (2018), Bento et al. (2024), Cook and Li (2023)) that directly value car time versus a given destination.

Computing the Social Cost of Delays As in Lewis and Bajari (2011), the social cost of a delay associated with a project j can be written as

$$\text{social cost}_j = \text{delay}_j \cdot \text{traffic volume}_j \cdot \text{time value}_j.$$

Repair projects cause delays due to reduced travel speeds. The delay depends on the nature of road repairs, for example the number of lanes that are closed and the length of the closure. Although road work often creates delays by lowering speed limits and reducing lanes, the realized traffic volumes and speeds should still be regarded as equilibrium quantities; some travelers will re-optimize plans and activities in the face of delays. Since we do not observe specific roadways or construction projects, we follow Lewis and Bajari (2011) in considering a hypothetical setting where the equilibrium impact of a closure results in three minutes of added delay, which is roughly equivalent to a decline in speed from 55mph to 35mph for five miles.

To map delays to time costs, we use our v measures by origin and destination. When delays are unanticipated, and travelers cannot plan ahead, this leads to a late arrival at the cost of time at the destination, or v. Conversely, anticipated delays lead to early departure, at the cost of time at the origin, i.e., $\delta \cdot v$. By examining overall averages we derive that anticipated delay costs to travelers on our platform at 9am are equal to \$0.15 per trip while unanticipated costs are equal to \$0.76 per trip. At 9pm, however, these costs drop to \$0.09 and \$0.54.

Underlying these averages, there are important spatial differences. On some routes, the average cost per trip is as high as \$2.50. How much construction delays are anticipated by drivers will vary according to several factors, such as how long the delays are in effect, how many drivers are repeated users of the road, and how many check traffic conditions before driving. We assume that half of drivers face expected congestion and half do not, motivated by the proportion of trips during an average weekday in Prague that occur during the commuting times.

Selection Adjustments To extrapolate from our setting to the population of Prague drivers, we need to address the selection of users in our platform. The population within Prague from which our estimates are drawn (i.e., taxi passengers) is likely to be selected on high-v individuals compared with the average Prague commuter. In order to extrapolate from our hourly congestion cost estimates, we first adjust these costs to reflect the opportunity cost of car commuters who are *not* on the Liftago platform.

We take advantage of the fact that the value of time that we measure at 9am is very close to the wage rate of those individuals as measured by a small survey conducted by the platform among passengers. The mean wage rate among survey participants is \$15.23, very close to our measurement of the v during work hours, which is \$15.44. This time period also most plausibly captures the value assigned to work related activities. We therefore assume that at 9:00am, a typical workday starting time, our v estimate is equal to users' average wage. The mean wage among Prague residents, however, is \$9.15 per-hour. We specify an adjustment factor $\kappa = 9.15/15.44 = 0.59$ and scale all v results by κ. Combining these assumptions allows us to infer, for example, that at 5pm when Liftago riders face an unanticipated congestion cost of \$1.50 per trip, the average car commuter would face a cost of \$0.89 per trip. Finally, we also scale by the average occupancy rate of cars in Prague, which is 1.3 as reported by the 2018 Prague Transportation Yearbook. The final scaling factor is therefore $1.3\kappa = 0.767$. The average scaled per-trip delay costs due to lost origin and destination v are \$0.098 and \$0.51, respectively.

Table 5 summarizes the heterogeneity in travel costs after scaling by 1.3κ. Panel A displays the baseline cost as the mean value of time for a three minute delay with fifty percent weights on both v and $\delta \cdot v$, or \$0.30 per trip. This value represents our best estimate of the average cost per-trip of any delay, equivalent to \$6 per hour, or around 60 percent of mean wages in Prague. Panel B shows the average costs, which we now allow to vary by time-of-day. Costs during the middle of the day are around 20 percent higher than in our uniform cost baseline, and costs in the evening are 20 percent lower. The second row of Panel B re-weights trip-costs by the overall share of Prague traffic volume by time-of-day, which we obtained from the Prague Transportation Yearbook. This row takes into account that traffic volume is very different during different hours, and thus if Prague were to apply time-of-day rules or incentives to construction, the

Table 5: Estimated Per-Trip Closure Costs by Time of Day

	Time-of-Day							
	3:00am	**6:00am**	**9:00am**	**12:00pm**	**3:00pm**	**6:00pm**	**9:00pm**	**12:00am**
A. Uniform Cost Baseline								
Uniform Price	$0.30	$0.30	$0.30	$0.30	$0.30	$0.30	$0.30	$0.30
B. All Routes with Time Variation								
All Routes	$0.31	$0.29	$0.36	$0.36	$0.37	$0.34	$0.27	$0.24
% change	0.02	-0.05	0.17	0.19	0.21	0.12	-0.1	-0.2
All, Volume Weighted	$0.05	$0.06	$0.51	$0.52	$0.54	$0.56	$0.33	$0.12
% change	-0.83	-0.8	0.68	0.71	0.77	0.85	0.08	-0.6
C. Routes by Destination and Time								
Highest-VOT Destination	$0.26	$0.31	$0.42	$0.35	$0.35	$0.39	$0.36	$0.26
% change	-0.13	0.01	0.38	0.16	0.15	0.28	0.19	-0.14
Median-VOT Destination	$0.20	$0.20	$0.26	$0.27	$0.30	$0.32	$0.27	$0.24
% change	-0.35	-0.35	-0.13	-0.12	-0.0	0.07	-0.1	-0.2
Lowest-VOT Destination	$0.07	$0.02	$0.11	$0.08	$0.13	$0.11	$0.13	$0.12
% change	-0.78	-0.93	-0.65	-0.73	-0.58	-0.62	-0.58	-0.59

NOTE: This table summarizes estimates of per-trip costs of three-minute delays by time of day. Each column aggregates results into three-hour bins so that, for example, *3:00am* refers to the period between 12:01am-3:00am.

hours from 9pm–6am would only cost 10–20% of the short-run welfare lost to a similar delay in mid-day. Panel C further decomposes costs into different routes. It reports the mean cost of trips ending in the highest-, median-, and lowest-VOT destinations. The highest-VOT destination trips are 2–4 times as expensive as the lowest-VOT destination trips, which would lead the uniform baseline to over-estimate by nearly double (93%) in some times-of-day and locations and under-estimate by up to 40 percent in other times-of-day and places. Panels B and C thus illustrate that there could be significant saving from timing road closures optimally, and from tailoring fines and rewards to projects based on their locations and the hours of closure.

While Table 5 summarizes the importance of the underlying heterogeneity and of time in particular, costs are borne on a route-by-route basis. Since we estimate valuations over 30 origins and destinations, Figure 8 shows the mean per-trip costs of delays across each of 900 routes. In this figure we report weighted averages over time of day and person-specific variation. This is the appropriate metric for measuring delay costs when closures cannot be conducted during selected times of day. Even projects that require closures throughout the day, such as road re-paving, exhibit substantial heterogeneity in the social cost imposed on road users. To maintain the right incentives and avoid lengthy closures that would result in high delay costs, procurers might need to impose fines that vary according to the impacted routes.

Using information on the average daily traffic flows across a few key routes in the city, we can derive some illustrative examples. We pair reported traffic flows with our

Figure 8: Estimated Per-Trip Closure Costs by Route

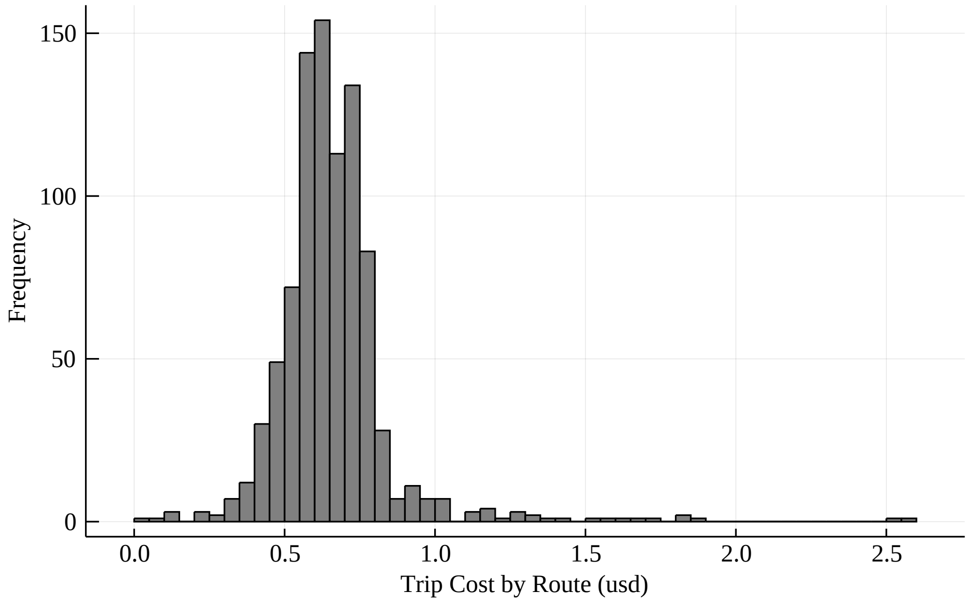

NOTE: This histogram shows route-specific estimates of mean per-trip costs of a three-minute delay across each of 900 distinct routes.

estimated delay cost in locations connected by each route and compute the weighted average v over the associated trips. For example, a hypothetical three-minute delay along the Zlichovsky tunnel, which sees an average of 84,000 cars per day, would impose social costs from $\$31,600$ to $\$35,500$ per day depending on the composition of the direction of travel. A delay along the nearby Brusnicky tunnel, which sees 77,000 cars per day, would impose social costs from $\$29,600$ to $\$31,800$ per day. In contrast, by assuming a uniform average value of time of \$0.30 per car trip, the costs of these two hypothetical delays are $\$25,200$ and $\$23,100$ respectively. The differences could lead the city to under-price time-savings related to construction on these specific routes by up to 40%, or about \$300,000 for a 30-day span.

These results underscore the importance of using economically relevant VOT measurements when designing procurement incentives. The ability to draw insights from market behavior to quantify the costs of delays can be valuable for local policy makers and regulators. In order to determine what incentives to offer, our method could be used to measure the costs of delays due to each additional day of construction, as long as anticipated delay time estimates for the particular project are available. Our cost estimates can also be used in determining the value of new infrastructure, such as the benefits from adding a new lane to a road, upgrading an intersection, or adding public transportation services. In each of these cases we would need to have information about the specific project to explicitly characterize its impact on traffic speeds and any spillover effects to nearby roads, but an equilibrium analysis quantifying the costs and benefits would require the same type of time valuations that we have obtained with our approach.

7 Conclusions

Measuring the value of time remains a persistent challenge in economics, even in settings where consumers directly trade off time against money. While markets like transportation can reveal consumers' willingness to pay for time savings, these choices alone do not tell us how consumers value time spent in different locations and contexts. This paper develops a framework for decomposing observed willingness-to-pay for waiting time reductions into location-specific values of time and demonstrates how these values reflect consumers' underlying trade-offs between activities and locations.

Our decomposition approach is broadly applicable in the various settings that obtain estimates of the value of waiting time. We demonstrate its application using detailed data from a ride-hailing platform, showing how choices over faster, more expensive options versus slower, cheaper alternatives reveal information about the relative value consumers place on time spent at different locations. Our framework provides both a theoretical foundation for interpreting such choices and an empirical approach for measuring location-specific values of time that can inform transportation policy and infrastructure decisions. Our rich environment allows for further decomposition by time of day and across a distribution of consumer types. We find heterogeneity by location, time-of-day, and especially individuals: 79% of the overall heterogeneity is driven by individual differences. In contrast with the literature on residential sorting, our results suggest that people with a higher value of time do not sort into locations with a higher value of time. This suggests that the valuation of places in the short-run may be more egalitarian than that implied by long-run decisions in housing markets.

We show how our approach can be used to help quantify optimal policies in transportation markets. Specifically, we use the fine-grained nature of our estimates to study congestion surcharges and optimal procurement incentives. In the first case, we show how optimal congestion prices vary across different urban densities and in different congestion scenarios. In the second case, by combining our estimates with information from Prague's 2018 Transportation Yearbook, we show that using a uniform, city-wide measure of v to infer the costs associated with road closures would mis-price delays by up to ninety percent along some routes.

Bibliography

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

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

_ , _ , and Loukas Karabarbounis, "Recent developments in the economics of time use," *Annu. Rev. Econ.*, 2012, *4* (1), 373–397.

Allen, Treb and Costas Arkolakis, "The welfare effects of transportation infrastructure improvements," *The Review of Economic Studies*, 2022, *89* (6), 2911–2957.

Almagro, Milena and Tomás Domınguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," 2019.

_ , Felipe Barbieri, Juan Camilo Castillo, Nathaniel G Hickok, and Tobias Salz, "Optimal Urban Transportation Policy: Evidence from Chicago," Technical Report, National Bureau of Economic Research 2024.

Arnott, Richard, André De Palma, and Robin Lindsey, "Schedule Delay and Departure Time Decisions with Heterogeneous Commuters," *Transportation Research Record*, 1988, *1197*, 56–67.

_ , _ , and _ , "Economics of a bottleneck," *Journal of urban economics*, 1990, *27* (1), 111–130.

_ , Andre De Palma, and Robin Lindsey, "A structural model of peak-period congestion: A traffic bottleneck with elastic demand," *American Economic Review*, 1993, pp. 161–179.

Athey, Susan, Billy Ferguson, Matthew Gentzkow, and Tobias Schmidt, "Experienced Segregation," Technical Report 2019.

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

Bento, Antonio, Kevin Roth, and Andrew Waxman, "The Value of Urgency: Evidence from Real-Time Congestion Pricing," *Journal of Political Economy Microeconomics*, 2024, *2* (4), 786–851.

Buchholz, Nicholas, "Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry," *The Review of Economic Studies*, 2022, *89* (2), 556–591.

__ **, Laura Doval, Jakub Kastl, Filip Matějka, and Tobias Salz**, "Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides," 2024.

__ **, Matthew Shum, and Haiqing Xu**, "Rethinking reference dependence: Wage dynamics and optimal taxi labor supply," Technical Report, working paper, Princeton University Economics Department 2025.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor supply of New York City cabdrivers: One day at a time," *Quarterly Journal of Economics*, 1997, pp. 407–441.

Castillo, Juan Camilo, "Who benefits from surge pricing?," *Available at SSRN 3245533*, 2023.

Chen, M Keith, Peter E Rossi, Judith A Chevalier, and Emily Oehlsen, "The value of flexible work: Evidence from Uber drivers," *Journal of political economy*, 2019, *127* (6), 2735–2794.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter, "The opportunity atlas: Mapping the childhood roots of social mobility," Technical Report, National Bureau of Economic Research 2018.

Cook, Cody and Pearl Li, "Value Pricing or Lexus Lanes? The Distrributional Effects of Dynamic Tolling," Technical Report, Working Paper 2023.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst, "Income growth and the distributional effects of urban spatial sorting," *Review of Economic Studies*, 2024, *91* (2), 858–898.

Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales, "How Segregated Is Urban Consumption?," *Journal of Political Economy*, 2019, *127* (4), 1684–1738.

Dunning, Iain, Joey Huchette, and Miles Lubin, "JuMP: A modeling language for mathematical optimization," *SIAM Review*, 2017, *59* (2), 295–320.

Duranton, Gilles and Matthew A Turner, "The fundamental law of road congestion: Evidence from US cities," *American Economic Review*, 2011, *101* (6), 2616–52.

Fajgelbaum, Pablo D and Edouard Schaal, "Optimal transport networks in spatial equilibrium," *Econometrica*, 2020, *88* (4), 1411–1452.

Farber, Henry S., "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy*, 2005, *113* (1), 46–82.

Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz, "Frictions in a competitive, regulated market: Evidence from taxis," *American Economic Review*, 2019, *109* (8), 2954–92.

Goldszmidt, Ariel, John A List, Robert D Metcalfe, Ian Muir, V Kerry Smith, and Jenny Wang, "The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments," Technical Report, National Bureau of Economic Research 2020.

Hall, Jonathan D, "Pareto improvements from Lexus Lanes: The effects of pricing a portion of the lanes on congested highways," *Journal of Public Economics*, 2018, *158*, 113–125.

Heblich, Stephan, Stephen J Redding, and Daniel M Sturm, "The making of the modern metropolis: evidence from London," *The Quarterly Journal of Economics*, 2020, *135* (4), 2059–2133.

Kreindler, Gabriel, "Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications," *Econometrica*, 2024, *92* (4), 1233–1268.

Kreindler, Gabriel E and Yuhei Miyauchi, "Measuring commuting and economic activity inside cities with cell phone records," *Review of Economics and Statistics*, 2021, pp. 1–48.

Lewis, Gregory and Patrick Bajari, "Procurement contracting with time incentives: Theory and evidence," *Quarterly Journal of Economics*, 2011, *126* (3), 1173–1211.

Nevo, Aviv and Arlene Wong, "The elasticity of substitution between time and market goods: Evidence from the Great Recession," *International Economic Review*, 2019, *60* (1), 25–51.

Ostrovsky, Michael and Michael Schwarz, "Congestion Pricing, Carpooling, and Commuter Welfare," Technical Report, Working Paper 2023.

Rosaia, Nicola, "Competing platforms and transport equilibrium," 2025.

Small, Kenneth A, "The scheduling of consumer activities: work trips," *American Economic Review*, 1982, *72* (3), 467–479.

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

Vickrey, William S, "Congestion theory and transport investment," *American Economic Review*, 1969, *59* (2), 251–260.

Online Appendix

A Identification of Location Values

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

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

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

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

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

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

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

Theorem 1. (Sufficient Conditions for Identification)

1. *Under Assumption 1, if* $\exists\left(X^{o}, X^{d}, a\right) \subset X \times X \times \mathcal{A} : \Pr\left[V_{i,a}^{k} = \hat{v}_{k}|X^{k}, a\right] = 1$ *for* $k \in \{o, d\}$*, then* $F_{a,k}(v)$ *is identified* $\forall a \in \mathcal{A}, k \in \{o, d\}$*.*

2. *Under Assumptions 1 and 2 , when either* $\exists k \in \{o, d\} : v_{i,a,t}^{k} = v_{i',a,t}^{k} \forall (i, i', a, t)$ *or* $\exists k \in \{o, d\} : v_{i,a,t}^{k} = v_{i,a,t'}^{k} \forall (i, a, t, t')$*, then* $F_{a,k}(v)$ *is identified* $\forall a \in \mathcal{A}$ *and* $k \in \{o, d\}$ *whenever* $|\mathcal{A}| \geq 3$*.*

3. *Under Assumptions 2 and 3,* $\left(\mu_{a}^{k}, \sigma_{a}^{k}\right)$ *is identified* $\forall a \in \mathcal{A}, k \in \{o, d\}$ *whenever* $|\mathcal{A}| \geq 3$*.*

Proof:

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

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

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

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

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

□

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

B Data Details

B.1 Locations

Following Buchholz et al. (2024), we use the exact GPS points of trip origin and a *k-means* clustering procedure to partition our data into 30 locations. Figure 9 shows each location and its corresponding index value. Details about this procedure can be found in Buchholz et al. (2024).

Figure 9: Locations in Prague

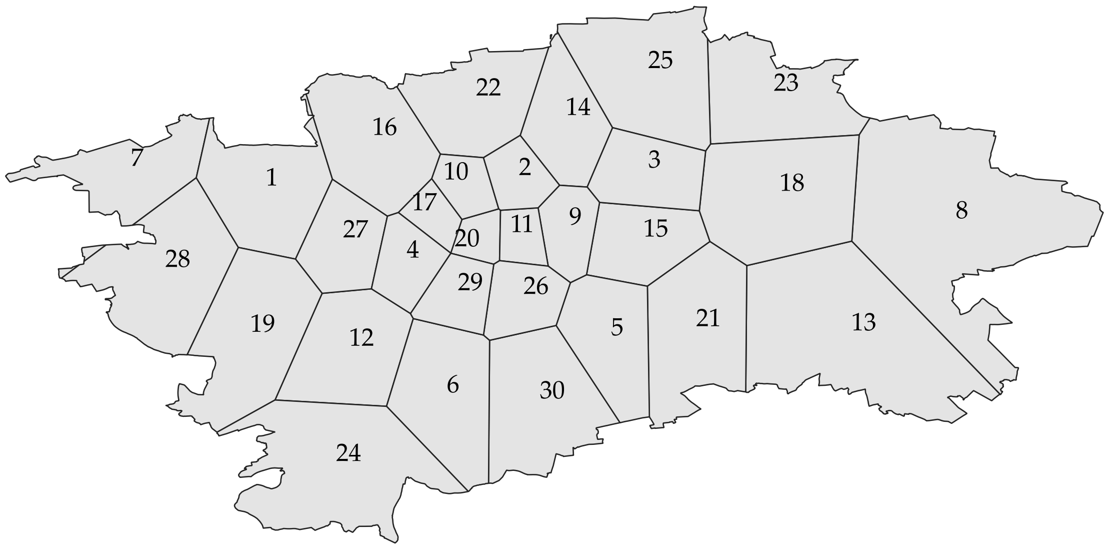

Note: This figure maps the boundaries of the city of Prague with locations defined by a k-means-clustering procedure on GPS-locations of trip origins and depicted as Voronoi cells that contain the clustered points. Displayed index values correspond to indices used in the paper.

B.2 Market Time Series Summary

Figure 10, panels (a) and (b) summarize the week-over-week trends in prices. Beginning in 2017, the ridership and drivers stop to grow and enter a relatively stationary period, although there are some large swings in the number of passengers towards the end of 2017. It shows the relationship between labor supply and price during the sample period. As expected, price decreases as supply increases during the winter holiday seasons and increases as supply falls during the summer. The second panel shows how average waiting time and average price evolve during the sample period.

Figure 10: Weekly Waiting Times and Number of Drivers Compared to Prices

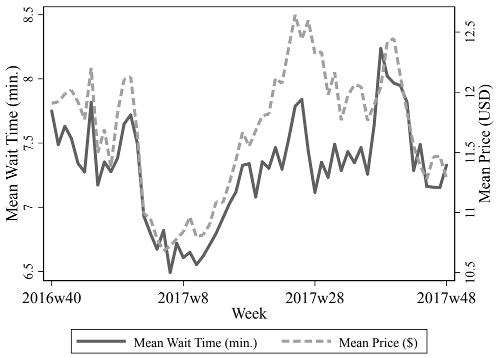

(a) Average Prices

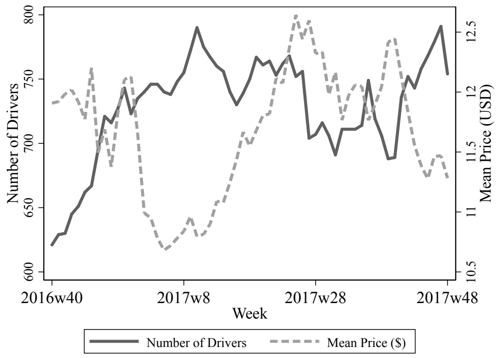

(b) Price and Driver Count

C Vickrey (1969) Model and Congestion Tolling

In the bottleneck congestion model of Vickrey (1969), each consumer chooses their arrival time t (equivalently, departure time $t - q(t)$) to minimize costs

$$C(t) = \alpha q(t) + \beta(t^* - t)\mathbb{1}[t < t^*] + \gamma(t - t^*)\mathbb{1}[t \geq t^*].$$

The shape of $q(t)$ can be derived following the reasoning in Ostrovsky and Schwarz (2023). Consider a consumer who optimizes and arrives early in equilibrium. If they shift their arrival time forward by ε, they incur additional queueing costs of $\varepsilon\alpha q'(t) + o(\varepsilon)$ but reduce their early arrival cost by $\varepsilon\beta$. Since all consumers are optimizing in equilibrium, the first-order condition

$$\begin{aligned} \alpha q'(t) - \beta &= 0 \\ q'(t) &= \frac{\beta}{\alpha} \end{aligned} \tag{15}$$

holds for early arrivals. Similarly, consider a consumer who optimizes and arrives late in equilibrium. If they shift their arrival time forward by ε, they incur additional queueing costs of $\varepsilon\alpha q'(t) + o(\varepsilon)$ and increase their late arrival cost by $\varepsilon\gamma$. Since all consumers are optimizing in equilibrium, the first-order condition

$$\alpha q'(t) + \gamma = 0$$
$$q'(t) = -\frac{\gamma}{\alpha} \tag{16}$$

holds for late arrivals.

Let t_q and $t_{q'}$ denote the start and end times of the queue, respectively, and t_0 denote the time of the queue's peak. The queue function $q(t)$ is equal to zero at $t = t_q$, rises linearly with slope β/α until $t = t_0$, and falls linearly with slope γ/α until $t = t_{q'}$. The length of time that the queue persists is equal to the time required for all consumers to pass through the bottleneck: $t_{q'} = t_q + \frac{N}{n_0}$. Combining with Equation 15 and Equation 16 and the conditions $q(t_q) = q(t_{q'}) = 0$ yields

$$t_0 = t_q + \frac{\gamma}{\beta+\gamma}\frac{N}{n_0}, \tag{17}$$
$$q(t_0) = \frac{N}{\alpha n_0}\frac{\beta\gamma}{\beta+\gamma}. \tag{18}$$

A fraction $r = \frac{\gamma}{\beta+\gamma}$ of consumers queue while $q(t)$ is increasing and hence arrive early; the remaining fraction arrive late.

In the tolling equilibrium, early and late arrival costs remain (and are unavoidable because the no-toll equilibrium already achieves the maximum traffic flow through the bottleneck up to its capacity). However, traffic costs from waiting in the queue can be completely eliminated. This is achieved through a time-varying toll $p(t)$ to pass the bottleneck at time t that exactly equals the queueing costs $\alpha q(t)$ that would have occurred without the toll. The marginal toll $p'(t)$ for early and late arrivals is β and $-\gamma$, respectively, which exactly equals early and late arrival costs; therefore, consumers have no incentive to adjust their arrival times compared to the no-toll equilibrium.

From Equation 15 and Equation 16, the queue function in the no-toll equilibrium is

$$q(t) = \begin{cases} \frac{\beta}{\alpha}(t - q_t) & t_q \le t < t_0 \\ q(t_0) - \frac{\gamma}{\alpha}(t - t_0) & t_0 \le t \le t_{q'} \end{cases}$$

where t_0 and $q(t_0)$ are given in Equation 17 and Equation 18. Therefore, the optimal toll is equal to

$$p(t) = \alpha q(t) = \begin{cases} \beta(t - q_t) & t_q \leq t < t_0 \\ \frac{N}{n_0}\frac{\beta\gamma}{\beta+\gamma} - \gamma(t - t_0) & t_0 \leq t \leq t_{q'} \end{cases}$$

which does not depend on queueing costs α.

The departure rate of consumers is constant between t_q and t_0 and between t_0 and $t_{q'}$.[6] Therefore, the total value of tolls collected during the commute is

$$\frac{1}{2}Np(t_0) = \frac{1}{2}\left(\frac{N^2}{n_0}\right)\left(\frac{\beta\gamma}{\beta+\gamma}\right) \tag{19}$$

and the average toll per consumer is

$$\frac{1}{2}p(t_0) = \frac{1}{2}\left(\frac{N}{n_0}\right)\left(\frac{\beta\gamma}{\beta+\gamma}\right). \tag{20}$$

Note that Equation 19 is identical to the expression for "total travel costs" in Arnott et al. (1988) and Arnott et al. (1993), as the optimal toll simply collects as revenue each consumer's cost of queueing in the no-toll equilibrium.

[6]Waiting time in the queue satisfies $n_0 q(t) = \int_{t_q}^{t} d(u)du - n_0(t - t_q)$, where $n_0 q(t)$ is the number of consumers in the queue and $d(t)$ is the departure rate. For $t_q \leq t < t_0$, $q(t) = \frac{\beta}{\alpha}(t - q_t)$ is linear in t, so $d(t)$ is constant; a similar reasoning applies for $t_0 \leq t \leq t_{q'}$.