

NBER WORKING PAPER SERIES

THE VALUE OF TIME IN THE UNITED STATES:
ESTIMATES FROM NATIONWIDE NATURAL FIELD EXPERIMENTS

Ariel Goldszmidt
John A. List
Robert D. Metcalfe
Ian Muir
V. Kerry Smith
Jenny Wang

Working Paper 28208
<http://www.nber.org/papers/w28208>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2020

We thank Vittorio Bassi, Antonio Bento, Alec Brandon, Jonathan D. Hall, Nathan Hendren, Justin Holz, Matthew Kahn, Shanjun Li, Paulina Oliva, Joseph Shapiro, Ken Small, and Cliff Winston for excellent comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Ariel Goldszmidt, John A. List, Robert D. Metcalfe, Ian Muir, V. Kerry Smith, and Jenny Wang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments
Ariel Goldszmidt, John A. List, Robert D. Metcalfe, Ian Muir, V. Kerry Smith, and Jenny Wang
NBER Working Paper No. 28208
December 2020
JEL No. D0,D1,R4

ABSTRACT

The value of time determines relative prices of goods and services, investments, productivity, economic growth, and measurements of income inequality. Economists in the 1960s began to focus on the value of non-work time, pioneering a deep literature exploring the optimal allocation and value of time. By leveraging key features of these classic time allocation theories, we use a novel approach to estimate the value of time (VOT) via two large-scale natural field experiments with the ridesharing company Lyft. We use random variation in both wait times and prices to estimate a consumer's VOT with a data set of more than 14 million observations across consumers in U.S. cities. We find that the VOT is roughly \$19 per hour (or 75% (100%) of the after-tax mean (median) wage rate) and varies predictably with choice circumstances correlated with the opportunity cost of wait time. Our VOT estimate is larger than what is currently used by the U.S. Government, suggesting that society is under-valuing time improvements and subsequently under-investing public resources in time-saving infrastructure projects and technologies.

Ariel Goldszmidt
Lyft
agoldszmidt@lyft.com

Ian Muir
Lyft
muir.ian.m@gmail.com

John A. List
Department of Economics
University of Chicago
1126 East 59th
Chicago, IL 60637
and NBER
jlist@uchicago.edu

V. Kerry Smith
Department of Economics
W.P. Carey School of Business
P.O. Box 879801
Arizona State University
Tempe, AZ 85287-9801
and NBER
kerry.smith@asu.edu

Robert D. Metcalfe Questrom
School of Business Boston
University
595 Commonwealth Avenue
Boston, MA 02215
and NBER
rdmet@bu.edu

Jenny Wang
Lyft
jwang5797@gmail.com

“Remember that time is money. He that can earn ten shillings a day by his labour, and goes abroad, or sits idle one half of that day, tho’ he spends but sixpence during his diversion or idleness, ought not to reckon that the only expence; he has really spent or rather thrown away five shillings besides.” Benjamin Franklin, 1748

1. Introduction

Perhaps taking the lead from Franklin’s advice to a young tradesman, the concept of opportunity cost was leveraged by early economists as far removed as [Mill \(1848\)](#) who laid the foundations for the notion; [Bastiat \(1848\)](#) who cleverly elucidated the broken window fallacy; Austrian economists like [von Wieser \(1876\)](#) who applied it to the phenomenon of cost; and all the way to [Green \(1894\)](#) who made it a central feature of his economic decision makers. Today it would be difficult to find an economist who would not place opportunity cost on a short list of key economic concepts that every citizen should understand. The central actor in Franklin’s opportunity cost advice, of course, is the value of time (VOT).

Since these early writers, economists for roughly the next century assumed that the VOT for each person could be regarded as a constant, with their wage rate as a reasonable approximation. For example, college students’ investments in human capital should include not only the real resource outlays for schooling but also the value of their forgone earnings from spending time in school rather than working. Similarly, the full cost of on-the-job training includes not only the cost of the training itself but also the value of forgone productivity associated with the employees’ time spent in that training. Crucially, this wage assumption makes labor market activity the primary basis for how we should judge the contribution of time to economic welfare.

While focusing on labor market activity yields rich insights that elucidate key economic trade-offs, there was a movement in the 1960s to also explore the allocation, efficiency, and welfare considerations of non-working time. One of the most influential of these contributions was [Becker \(1965\)](#), who proposed a model in which the individual combines time with a market good to produce a flow of services that he labeled “basic commodities.” In this manner, he described households as being small factories at their core. His framework unleashed the full economic toolkit to allow analysis of a wide array of issues within the household. Stimulated by Becker, Mincer, and their students, this new “home economics” began to apply economic concepts to a broad set of issues that included selection of a partner, spacing of children, division of labor among household members, divorce, and decision authority within the family (see [Greenwood et al. \(2017\)](#) for a recent overview).

The growth and maturity of this literature has served economics well, yet one important aspect of the Becker model has received less attention: the key features of his approach can also be leveraged to estimate the VOT. Indeed, as we describe more patiently below, there are two major

assumptions in the classic early time allocation studies (Becker, 1965; Johnson, 1966; DeSerpa, 1971) that have allowed the analyst to place a value on time: (1) the degree to which a consumer has the ability to flexibly allocate time to different tasks and (2) the degree of complementarity between time and purchased commodities as part of the consumption process.

The extant body of research on time use falls into two primary categories. The first, associated with microeconomic applications (especially in transportation and infrastructure), has largely relied on observing a person's decisions in the face of time and money trade-offs - either actual trip choices made to reduce travel time delays associated with congestion or hypothetical decisions made in the context of stated choice surveys or small-scale experiments (see, e.g., Deacon and Sonstelie (1985); Smith and Mansfield (1998); DellaVigna et al. (2012)). The second has instead focused on tracking the trends in time use for different demographic groups (Aguiar and Hurst, 2007b; Ramey and Francis, 2009). In these latter studies, the VOT is a latent variable implied by the observed differences in time allocations across these groups. Thus when a VOT is estimated, there are specific assumptions made about how money is traded for time savings. For instance, in Aguiar and Hurst (2007b), cohorts of people with low opportunity cost of time, such as retirees, are found to spend more time searching for lower prices than consumers with more limited time availability. By assuming search time can lead to monetary savings, it becomes possible to infer the VOT (see also Ghez et al. (1975); Juster and Stafford (1991); Robinson and Godbey (1999)).

We take this literature in a new direction by combining Becker's work with the early pioneering work in non-market valuation that explored weak complementarity. Formally, weak complementarity relies on the assumption that a person would not value a quality change in something she does not use (Mäler, 1971). This restriction implies that there is a price equivalent (in welfare terms) for a quality change in that resource (Smith and Banzhaf, 2007). Importantly, the VOT literature has yet to appreciate this theoretical connection, even though the original models sought to establish how the allocation of time contributes to the total value created through each use of time. In a setting where an agent has the ability to make choices that reflect consideration of both the time required and the price of a service, this relationship provides an estimate for the opportunity cost of time. Deeper inspection of Becker's proposed empirical framework for household production highlights another feature of time use that can be leveraged: his "technological coefficients" for the time required for each non-work activity provide the ability to consider how the context of a time allocation affects its value. In particular, one interesting aspect of Becker's model is the proposition that choice characteristics matter, or likewise the fundamental assertion that the properties of the situation and the population might matter a great deal in both the allocation and value of time.

With these necessary theoretical conditions in hand, we sought an appropriate testing ground to empirically study the VOT. Our search concluded with the realization that the assumption in the classic studies applies to any situation where one must wait for a service or good. An activity

that nearly half of all American adults have used satisfies exactly that assumption: rideshare. Waiting sessions on rideshare naturally provide the basis for estimating the trade-offs between waiting time and price that underlie economic measures for the VOT. To our best knowledge, our paper is the first to recognize the applicability of weak complementarity to situations where one must wait for a good or service. Furthermore, the richness of the rideshare environment permits a deeper exploration into measuring important context specificity (Becker’s “technological coefficients”) across relevant situations and individuals.

To request a ride via a rideshare service, a prospective passenger opens the app on their phone, can input an intended destination to receive a price quote and estimated wait time, and then decides whether to request (purchase) the ride. Such “take it or leave it” decisions exploit the role of weak complementarity in extensive margin choices. More specifically, when an individual opens the Lyft app we are able to observe how various combinations of wait times and prices affect purchase decisions. In 2015/16 and 2017, Lyft ran two natural field experiments that randomly varied these prices and wait times that were shown to customers in the app as they made their purchase decision. The assumption of weak complementarity ensures that the observed trade-offs then allow us to estimate price changes that are equivalent to a change in wait times. Thus, with weak complementarity in place, the observed trade-off reveals the marginal VOT.

The field experiments spanned 13 cities in the United States that had constituted most of Lyft’s largest markets at the time. In total, they cover 3.7 million customers and 14.8 million customer sessions (defined as an interaction with the app in which the customer receives a price and time quote). The exogenous prices and wait times in our field experiments are independent of the passengers’ outside options as well as potential unobserved shifters to aggregate demand (i.e., passengers) and supply (i.e., drivers). The exogenous variation combined with passengers’ decisions and weak complementarity allow us to estimate the marginal VOT.¹

An important feature of our work is that Lyft retains the ability to control not only whether people have to wait longer but also how much longer they have to wait. In the experiments, customers are randomly assigned to wait at the market wait time or at least an additional 60, 150, or 240 seconds over the market wait time (additional wait time irrespective of market conditions). With this variation, we can causally estimate how the waiting time elasticity and the VOT vary over the length of the wait time - with variation in length that is unrelated to other market factors (e.g. when and where customers open their app). Our approach allows us to recover an estimate of the VOT over wait time gradients, providing insight into the shape of the VOT function.

¹Our research is complemented by two contemporaneous studies that use very different identification strategies and data sets (Castillo, 2019; Buchholz et al., 2020). Both studies observe market wait times and prices but use econometric structure as opposed to experimental variation to solve the identification problem for Houston and Prague consumers. Our experimental structure allows for exogenous variation in waiting time and prices, and provides greater resolution in the nature and spatial variation in VOT estimates, and how market conditions experienced by rideshare users (e.g., weather conditions, location (airport, downtown, public transit distance), business trips etc.) affect VOT. In addition, weak complementarity assures that we avoid sensitivity of the results to the specification decisions that generally are associated with a structural approach.

Several insights can be drawn from these two large-scale natural field experiments. We divide them into three main areas. First, we find that consumers are responsive to both wait times and prices. Across both field experiments, we find the time elasticity of demand to be approximately -0.043 (standard error of 0.003) and the price elasticity of demand to be -0.59 (standard error of 0.02). Price elasticities are significantly larger than time elasticities within every city. The estimated price equivalent implied by these two elasticities yields an average VOT across the U.S. of \$19.38 per hour (standard error of 1.39; all prices are in 2015 dollars). To explore how our results map across relevant populations, we re-weight our estimates for the elasticities matching the rideshare population with the broader US population. We find no meaningful difference in VOT estimates based on the weighted sample, tentatively supporting their external validity with respect to the broader population of all travelers. We also find that our VOT estimate is stable across the duration of our eight-week experiments.

When evaluating projects, the US Government currently values people’s time between 33% and 50% of the wage rate for project appraisal, suggesting a value of at most \$14.20 per hour.² This one figure is meant to cover all types of travel (e.g., leisure, personal, and commuting to work), except where driving is part of the job (e.g., freight travel where the wage rate is used to impute the VOT) (USDOT, 2015). Our VOT estimate is approximately 35% higher than the currently used rule of thumb by U.S. federal guidelines. More specifically, our estimates imply that the VOTs for different metro areas are approximately 75% of the after-tax mean wage rate and about 100% of the median after-tax wage rate. In every metro region, we find a VOT estimate that is statistically larger than $\frac{1}{2}$ of the after-tax mean wage rate.

A second set of results that emerges from our field experiments relates to waiting time elasticities and the shape of VOT. We find that both are larger over longer periods of wait time, independent of market conditions and the reasons for consumption. This finding implies that the VOT is convex over time. Such convexity is important when considering the appropriateness of transferring or generalizing VOT estimates across space, time, or situations, as we find that the VOT depends critically on the baseline level of time use. Taken together, our first two results are directly relevant for the analysis of both private and public investments. They suggest that society is under-valuing projects that involve time saving infrastructure or technologies and furthermore, the degree of this under-investment increases in the amount of time saved.

Our third set of results leverages non-experimental variation to explore how properties of the situation affect the VOT estimate. Using a simple framework that directs our exploration of how various choice characteristics affect the VOT, we find substantial heterogeneity across contexts.

²The recreational demand model literature and environmental regulation use $\frac{1}{3}$ (following Cesario (1976)), and the transportation and infrastructure literature and regulation use $\frac{1}{2}$ (following Small et al. (2005); Small (2013)) of the wage rate. There does not seem to be consistency within the government on the values used for different policies. At best, one might argue that the U.S. Department of Transport is valuing time primarily as it relates to congestion–infrastructure to address congestion related delays whereas recreation is considered leisure travel, therefore the opportunity cost may be argued to be lower since travel may be part of the trip experience.

For example, across regions we find a Spearman correlation of 0.73 ($p = 0.025$) between the mean wage and the VOT estimate. Furthermore, we find that the VOT critically relates to the availability of substitutes for consumers, in that individuals considering a Lyft trip near alternative modes of transportation are more time sensitive than those who do not have readily available substitutes. In addition, signatures of the trip matter a great deal, as the VOT is strongly related to purpose of trip: during the morning and afternoon peak commuting times, the VOT is 50% higher than during off-peak times.³ Relatedly, weekdays have a 10% higher VOT than weekends, and we find that the VOT is 20% higher in the central business districts of cities than in the suburbs. Finally, we explore several other types of heterogeneity based on our economic framework to find time elasticities consistent with predictions from economic theory and the broader literature (Small, 2013).

We view our research as contributing to several areas of import. For policymakers, our estimates are a significant improvement over those in the existing literature with respect to both the identification and the design of the wait and travel time changes. Our experiment varied the total travel time of the journey, so we are assuming the increase in wait time for the ride is valued in the same way as an increase in the in-car time.⁴ They also provide a granular view of how choices vary with the context of time use, consistent with Becker’s early insights. Our research also has direct policy implications, as we recommend that policies: (i) account for the VOT heterogeneity with respect to cities, locations within cities, day of week, and time of day when estimating the benefit profile of public projects; and (ii) when this is not possible, adjust the rule-of-thumb VOT estimates up to 75% of the after-tax mean wage rate otherwise.

In this spirit, our estimated VOT varies predictably with economic aspects of the marketplace. This finding has important implications for how we value time in various economic sub-fields. For instance, in transportation, the VOT is usually the pivotal factor in benefit-cost decisions, as it has been estimated that excess urban road congestion led U.S. consumers to spend 5.5 billion hours sitting in traffic. Indeed, some studies (Schrank et al., 2012; Couture et al., 2018) estimate the annual deadweight loss due to congestion in the United States alone at \$30 billion. Understanding how best to apply our estimated elasticities to construct Pigouvian taxes designed to reduce the deadweight loss of congestion is an important next step in such research (Arnott et al., 1993; Duranton and Turner, 2011; Finkelstein, 2009; Small, 2013). Furthermore, the VOT in

³We acknowledge that the context of the wait in our field experiments is heterogeneous. Some consumers will be waiting for the car on a street corner, while others will be getting ready to leave their building, others working, etc. Given that some consumers can do other things while they wait for the car (e.g., peruse their emails, texts, etc), the wait might not be necessarily boring or painful, so our estimates might be viewed as a lower bound of the VOT that may be measured in less comfortable or productive situations (i.e., caught in gridlock traffic).

⁴Once you assume a linear time constraint (i.e., total hours = sum of allocations) in the model, we assume each part being allocated is a perfect substitute for another. Because wait time exhibits weak complementarity with the Lyft ride, we now have the VOT derived at the margin by the price equivalent change in the price of the ride. Linearity of the time constraint allows us to use this margin and apply it to other types of time at the margin, so it becomes a general estimate of the VOT. Empirically, to test this, we need experimental random variation in price and time for people who are randomly allocated to either wait time or in-car time (where the base price and total time are the same in the two scenarios). Such a study does not exist.

car transit is an important parameter when estimating the demand for public transit, evaluating any proposal for federal funding of infrastructure, and evaluating the impact of other more climate-friendly (green) transit options, such as a carbon tax (Parry and Small, 2009; Chen and Whalley, 2012; Anderson, 2014; Basso and Silva, 2014). More generally, given that the VOT usually constitutes the largest share of total benefits in infrastructure projects, our estimates open up the possibility of more efficient allocation of resources in the economy.⁵

In terms of linking to the broader literature on a host of policy decisions, many view time as the ultimate scarce resource. These are many policy- and market-relevant activities that require a period of waiting time before ascertaining utility from the commodity, such as time waiting for: a table at a restaurant; a delivery of a good; a store or government office to open (e.g., renewing a driving license or ID, or waiting to be seen at a Veterans Administration center); a dentist or a doctor; a voting booth to open; accessibility to buildings (e.g., how much longer does the handicapped-accessible ramp/elevator/parking/office take); or a car or public transit to get to a destination. For the economy as a whole, the VOT depends on how different people respond to the market and non-market signals in allocating their monetary resources and time (Juster and Stafford, 1991). These allocation decisions of time impact where people live (Wheaton, 1977; Van Ommeren and Fosgerau, 2009; Su, 2018; Kreindler and Miyauchi, 2019), how they supply their labor (Aguiar and Hurst, 2007b; Aguiar et al., 2013, 2017; Benhabib et al., 1991; Gelber and Mitchell, 2012; Goldin, 2014; Gronau, 1973; Mas and Pallais, 2017, 2019), how they commute (Small et al., 2005; Bento et al., 2017; Hall, 2020), how they invest in their health (Besley et al., 1999; Miller and Urdinola, 2010; Philipson et al., 2010), and what goods they buy (Nevo and Wong, 2015). VOT estimates have also become increasingly important in international debates about productivity and national accounting (Krueger et al., 2009; Nordhaus, 2009; Aguiar and Hurst, 2016), for the welfare estimation of business cycles (Aguiar et al., 2013), and the VOT is a central feature in governments and companies as a basis for investment decisions for the supply of intangible and service goods within economies.

The remainder of our study is structured as follows. Section 2 provides key theoretical underpinnings of the classic literature and outlines how we leverage these features to estimate the VOT using two field experiments. Sections 3 and 4 report empirical results from our two natural field experiments and detail our framework for exploring heterogeneity. Section 5 addresses the issues in identification and external validity of our estimates, and section 6 provides a discussion linking our work to the current policy landscape. Section 7 concludes. The online appendix includes additional empirical analysis.

⁵Moreover, our VOT estimates relate to understanding the economics of online platforms (Goolsbee and Klenow, 2006; Chen et al., 2014; Allcott et al., 2019) and the amount of bureaucracy in government policymaking (Sunstein, 2018).

2. Theory and Market Context

In this section, we describe the theoretical landscape for valuing time and link it to our modeling approach. What results is a set of necessary experimental conditions that must hold to deliver theoretically-consistent estimates of the VOT. We then detail the market context for our two natural field experiments.

2.1 Theoretical Framework

The current modelling of time use and value in the literature stems from Becker’s (1965) insight that time is required for all consumption activities. Most discussions of his contribution equate it with the origin of home production, and focus on Becker’s argument that an individual “produces”—what Becker describes as “basic commodities”—which are then consumed.⁶ These basic commodities are the services derived when market goods are combined with time, and these services are what contribute to well-being and motivate choices for individuals. While the concept of home production is certainly important, Becker’s description of consumption also introduced two other features that are important to our research design. The first feature highlights the role of restrictions on how private goods and time enter preferences. This dimension of the classic framework is best illustrated by first considering a simple form of Becker’s model. Assume the household consumes two basic commodities, $Z_i, i = 1, 2$, which are service flows, and in our case, Z_1 is the service flow from rideshare travel, and Z_2 is the service flow from all other goods. The household production functions are Leontief as in equation (1):

$$T_1 = t_1 \cdot Z_1 \quad T_2 = t_2 \cdot Z_2 \quad x_1 = a_1 \cdot Z_1 \quad x_2 = a_2 \cdot Z_2 \quad (1)$$

Our specification assumes one private good x_i per basic commodity. In our case, x_1 is the rideshare trip, and x_2 are all other goods. Time is allocated exclusively to each activity and there is no multi-tasking. Both a_i and t_i are technological coefficients in Becker’s model. t_i is the time in discrete units for each unit of Z_i , and a_1 is the amount of rideshare service that is needed to produce Z_1 .⁷

Assume a time constraint with \bar{T} the total time available and T_w the amount of work time which is priced at w . The individual also faces a budget constraint with p_i , where $i = 1, 2$, the prices for private goods, the wage income, wT_w , and non-wage exogenous income (R). Becker’s

⁶There are a number of contributions using the household production logic to model consumption expenditures as well as in describing alternatives to the unitary model of individual behavior. A good access point is the review in Browning et al. (2014), where the work using household production in alternative models of individual choice is summarized.

⁷ a_1 is important because it allows us to interpret a local marginal condition that links the Becker model and weak complementarity (see below). Thus a_1 is a function of the waiting time, which allows us to illustrate if consumers are not producing" Z_1 (i.e., services from travel with ride share), they do not care about waiting time.

household production for consumption activities is equivalent to a restriction on preferences that treats time and private goods as perfect complements.

The second feature highlighted by Becker’s model arises from the fact that time use is described to be *specific to a particular consumption task*. That is, the framework allows time to be uniquely linked to each activity a person undertakes. As a result, a natural interpretation is that the analysis can consider how the context for using one’s time affects its value. For instance, time spent commuting on a weekday morning can be different from commuting in the evening or over weekends.

The importance of restrictions on how time and goods enter preferences, and the assumption an amount of time is uniquely linked to the use of each good (rather than allowing time to jointly produce two or more basic commodities), arises when we consider the two possible values for time. These different values are implied by the indirect utility function for Becker’s model. Substituting the time constraint and production functions into the budget constraint, we have the general form given by equation (2) (see appendix A for all equation derivations).

$$\bar{V} = V(wt_1 + p_1.a_1, wt_2 + p_2.a_2, w\bar{T} + R) \quad (2)$$

The first possible value of time is defined when we consider increasing or decreasing the time endowment \bar{T} . Such a shift in the time endowment implies the marginal value of time is equal to w . Such an approach is usually operationalized using stated preference surveys or by increasing the amount of time available for people to make choices (e.g., sleeping less during each day which is unlikely without a technology).⁸ The second possible value of time arises because we can define what might be termed the “supply” price or time cost of each activity, which depends on both the wage and the technology of home production (i.e. the t_i ’s). In this case, the supply price is equal to wt_1 for the first consumption activity and wt_2 for the second.

Since the developments in the literature in the 1960s, the contributions on the allocation and value of time have, for the most part, missed these key features of Becker’s model. They observed that the model had the same implications for VOT because the marginal value of adding to the time endowment remained the wage rate. This result is conditional on the optimal allocation of

⁸Small et al. (2005) impose a linearity assumption on the indirect utility function assumed to underlie an individual’s choice of whether or not to use an express lane for a trip. The express lane has a toll and an anticipated travel time while the conventional lane has no toll but a longer anticipated travel time. Linearity assures the ratio of the coefficients for the toll and the travel time in the choice model reveal a marginal value of time (see their equations (1) and (2)).

time among activities with different costs.⁹ Our conclusion is readily illustrated if we consider how the model describes adjustment in response to a change in the marginal value of time (w). As equation (3) illustrates, the supply of labor takes into account the reallocation of time among activities:

$$\frac{\bar{V}_w}{\bar{V}_R} = \bar{T} - t_1 Z_1^* - t_2 Z_2^*, \quad (3)$$

Equation 3 is derived from the partial derivative of equation 2 using Roy’s identity as amended for this model, where Z_1^* and Z_2^* are the utility maximizing choices for the two basic commodities (rideshare and all other commodities). Changes in the wage rate change the “prices” of each of the basic commodities as illustrated by the last two terms on the right side of equation (3). Adjustment in the amounts consumed determine the time requirements (with the assumed Leontief technology) and labor supply is the residual component of the time endowment.

To focus attention and simplify matters, Becker assumed the time requirements for each activity were fixed. The joint roles for preference restrictions between goods and time and the ability to take account of the context of when and how time is being used are the important elements in his framework for our research. Of course, demonstrating this point requires an ability to control both the time requirements for some set of activities and their prices. Such control allows the analysis to test whether time allocation is largely a matter of labor–leisure decisions. In Becker’s model, an agent’s choice of her allocation of non-work time matters according to equation (3). When the assumption of a Leontief technology for household production is relaxed, however, control over the prices and time requirements alone will not ensure recovery of the value of time for each use.¹⁰ Another preference restriction is needed.

Fortunately, for some activities a different form of complementarity, weak complementarity, provides sufficient information to value time. These are activities that require a period of waiting time before ascertaining utility from the commodity, such as waiting for: a table at a restaurant; a delivery of a good; a store to open; a dentist; a doctor; a voting booth to open; or a car or public transit to get to a destination. As Mäler (1971, 1974) showed, weak complementarity means that underlying changes in features of the good or service are only important to actual consumers of the good or service. He recognized that even without the assumption of perfect

⁹ Aguiar and Hurst (2007a) develop their approach for estimating a value of time by assuming that optimizing households exploit a shopping technology as another mechanism for substituting time for goods outside the labor market. They maintain that the price paid for private goods is a function of the time allocated to shopping. Greater search time yields lower prices. They also acknowledge that the price paid can depend on shopping needs or the number of items that might be involved in a household’s search activities. So the value of time is derived as the shadow price of time allocated between the shopping and household production technologies (see their equations (1) and (2)). The ability to freely substitute time between these two uses assures the marginal value of time is equalized between these activities “outside” the labor market. As a result they use estimates of this price function to estimate the value of time (see their figure 1).

¹⁰ When the household production technology is assumed to be more flexible, the VOT in each use depends on the marginal technical rate of substitution between time and goods at the optimal consumption levels for the basic commodities. To estimate this requires detailed information on the technologies involved as well as all the goods’ prices.

complementarity between a nonmarket good and a private good, weak complementarity makes it possible to estimate the demand price for the nonmarket good from the information contained in the demand function for the private good.¹¹

Our case provides a direct parallel to his example. If the demand for the rideshare service is zero, then the demand for a shorter wait time to obtain that service is also zero. Mäler was careful to spell out how the assumption of weak complementarity applies only at the individual level (or to an aggregate over homogeneous groups of consumers). He also noted that to use it in measuring the willingness to pay for a discrete change in a nonmarket good (like time), the specification of the Marshallian demand for the private good needs to ensure that the corresponding Hicksian demand function for that good has a finite choke price.

For our purposes, however, we are not attempting to estimate the demand for rideshare trips; we are simply interested in the marginal VOT. Our use of weak complementarity can be illustrated with a simple amendment to the Becker model. We acknowledge at the outset that our example abstracts from important details that could be used in a full structural model of these decisions.

We remain with the assumption that the first basic commodity is the travel services “produced” by using a rideshare company. We distinguish two types of time involved in ride share services. The first, designated T_1 , is a summary measure for the travel time associated with rideshare trips. The second type of time, T_1^a , is exogenous from the perspective of the individual. It represents a second summary measure for the waiting time for the rideshare trips while the individual waits for the drivers involved in these trips to arrive. The user knows these times on each trip occasion when selecting her trips. We assume waiting time is an indicator of the quality of the service. Longer waiting time implies a lower quality travel service can be produced. Weak complementarity implies that the individual does not value waiting time if she does not produce travel services. Thus, our modification to the Becker model must embed this assumption. For this example assume the technical coefficient linking Z_1 (our measure of the produced travel services) to x_1 (our measure of the rideshare services required) is a function of waiting time. So we replace

¹¹For certain environmental applications this condition can be controversial—a person may want to maintain high levels of air quality (and visibility) at the Grand Canyon but never plan to visit the site, or want to protect the Arctic Wildlife Preserve but not consider a wilderness adventure there. These omitted values are the existence values of such goods. As he noted:

It is, however, not necessary that the environmental service and the private good be perfect complements in order to carry out the steps involved in determining demand price. A much weaker condition of complementarity is the following: if the demand for a private good is zero, then the demand for some environmental service will also be zero. If, for example, the private good is swimming in the lake and the environmental service is the quality of that lake, then it is very reasonable to assume that if a person does not use this lake for recreation, he is indifferent to the quality of the water. ... It would therefore seem that this weak complementarity condition has very broad applications, although it cannot be applied in cases where option values are involved (Mäler 1974, 183).

this component of the Becker model with:

$$x_1 = a_1(T_1^a)Z_1 \quad (4)$$

We assume $a'_1 > 0$, implying increased waiting times reduces the amount of “constant quality” travel services produced by ride share (i.e., Z_1 gets smaller) This specification assures that if $Z_1 = 0$, then changes in T_1^a will not affect the individual’s well-being, consistent with weak complementarity between waiting time and the travel service basic commodity.¹²

We can now return to equation (2) and adjust the indirect utility implied by Becker’s model to reflect the assumed role of T_1^a in producing constant quality travel services. Equation (5) provides the modified indirect utility:

$$\bar{V} = V(wt_1 + p_1a_1(T_1^a), wt_2 + p_2a_2, w(\bar{T} - T_1^a) + R) \quad (5)$$

When we make this change and use duality to consider the marginal value of reducing waiting time, equation (6) results:

$$\frac{V_{T_1^a}}{V_R} = \frac{V_1}{V_R}p_1a'_1 - w \quad (6)$$

The terms on the right-hand side of the equation corresponds to the marginal “cost” to the user of increases in the waiting time. We can see this by noting that Roy’s identity implies $-V_1/V_R = Z_1$. Our argument suggests that the price equivalent value of waiting time is $(p_1Z_1a'_1)$, and that could be equal to, less than, or greater than the wage rate. This is an empirical question. Shorter (longer) waiting time reduces (increases) the incremental cost of producing travel services with the ride share, Z_1 . This intuition leads to Equation (7):

$$\pi = -\frac{V_{T_1^a}}{V_R} = p_1Z_1a'_1 + w \quad (7)$$

Of course, this case represents *one* simple example. Our general point is to show that the example “works” because of weak complementarity. From equation (5), the specification of the measure for ride share trips (x_1) implies that there is an equivalent change in p_1 that can be represented by a change in T_1^a .

Our main theoretical contribution to the Becker model is that a weaker preference restriction

¹²Several qualifications apply to this stylized example. If we were attempting to model both the decisions to use rideshare and the number of trips to take, our specification would need to reflect that T_1^a and a measure for the count of trips are related. Depending on how we added these details, connecting the model to measures of rideshare services, such as trips, the time used in traveling, and the time spent waiting, nonlinearity could be introduced into the budget constraint. Such a detailed formulation would rely on the specific functional assumptions made in identifying estimates for the value of time. Since our goal is simply to illustrate how weak complementarity allows us to measure the marginal value of time by implying a welfare equivalent link between price changes and wait time changes, we avoid spelling out these connections.

than was required in Becker’s model is possible for valuing time. In the context of rideshare, we have this form of weak complementarity as people wait for the ride. An additional important feature of the model is the specificity of the value of time can be linked to its particular use. This arises in rideshare because the timing and position of rideshare trips can be linked to a wide range of purposes, tasks, and contexts, allowing for an assessment of how the properties of the situation affect VOT measures. Taken together, we now have the theoretical machinery to value time through changes in prices and wait times of the good.

In this spirit, leveraging the Lyft rideshare platform, our model is specific to a trip session, in which a passenger opens up the app and receives a price and waiting time quote for a potential trip. Individual responses to the terms of a ride are indexed by passenger (i) and session (j). We assume that incomes and other prices faced by individuals opening the Lyft application are fixed across the experimental groups due to randomization, and focus our attention on the utility realized with and without requesting a ride. Equation (8) begins the process of formalizing the decision associated with requesting a ride:

$$V_{ij} = v(P_{ij}, T_{ij}^a) + \varepsilon_{ij} \quad (8)$$

Let V_{ij} be the utility associated with selecting a Lyft ride at a price of P_{ij} for individual i and session j . Let T_{ij} be the wait time indicated to individual i in session j , and ε_{ij} be a random error capturing unobserved (to the analyst) features of the circumstances of choice.

We assume that each individual compares the realized utility from a Lyft ride with that of a default condition that we do not observe. We assume the default option is specific to each individual and session and provides utility W_{ij} . A Lyft is requested if and only if $V_{ij} > W_{ij}$. By independently randomizing both P and T^a , we can recover estimates for the parameters used to describe the choice process in equation (8).

The price of a Lyft ride is the product of a base price, B_{ij} , that depends on the characteristics of the request (timing, route, and vehicle type) and a price multiplier (called Prime Time) $(1 + PT_{ij})$. This multiplier is set dynamically in response to local demand and supply conditions. B_{ij} and PT_{ij} are indexed by i and j because the records of sessions allow the timing of the session and the individual to be distinguished.

In our case, we observe whether an individual selected a Lyft ride but cannot completely characterize the features of the alternative set when the ride is not chosen. This is where our assumption of weak complementarity provides the “traction” needed to recover a VOT. As noted earlier, when waiting time is a weak complement to the rideshare service associated with each trip, a change in wait time is equivalent, from a welfare perspective, to a change in the price of the trip (Smith and Banzhaf, 2007). The important implication is that we do not need to know anything about a passenger’s labor supply decisions to recover an estimate for the opportunity cost for their time; the price of the Lyft ride serves this role. Thus, with appropriate exogenous variation in

both trip prices and wait times, consumers’ actual choices allow us to identify the key threshold trade-offs between time and money that serve to recover the VOT.¹³

The decision process in our model begins with the assumption that each individual is considering a “local” trip, that is within his or her metropolitan area. We observe everyone who opens the Lyft application during our experimental period. Yet, we do not know their complete set of outside options, and because of this limitation we consider a variety of approaches to organize sessions to account for our hypothesized differences in how these outside alternatives influence W_{ij} for each individual. There are several implications of this constraint on what can be observed. The first of these is the selection effect directly associated with knowing only those who open the Lyft app; we discuss the implications of this external validity issue in Appendix Section H. Moreover, because we cannot characterize W_{ij} , we do not have a reference or baseline condition that we would expect in defining a reference utility level to measure a VOT. This limitation influences our interpretation of the estimates. Without specifying the outside alternative, we estimate choice probabilities relative to a normalizing alternative.¹⁴

2.2 Empirical Model

Our primary model specification for the choice process is contained in equation (8), which compares V_{ij} with W_{ij} using the log transformation for the price and wait time as determinants of the observed trip request, $R_{ij}(0, 1)$. We assume that passenger i in their j^{th} session receives greater utility from requesting a ride, V_{ij} , than the alternative W_{ij} , and thus the request, R_{ij} , is made according to:

$$R_{ij} = \beta_1 \ln P_{ij} + \beta_2 \ln T_{ij}^a + \varepsilon_{ij} = \beta_0 \ln B_{ij} + \beta_1 \ln(1 + PT_{ij}) + \beta_2 \ln T_{ij}^a + \varepsilon_{ij} \quad (9)$$

Here P_{ij} is the offered price, T_{ij}^a the offered wait time for the ride, B_{ij} is the base price, and ε_{ij} is the unobserved error term. The price of a ride is the product of a base price B_{ij} and a price multiplier, $(1 + PT_{ij})$, which is set dynamically in response to local supply and demand conditions. By using the log transformation for the price and wait time in equation (9), the choices reveal the preference parameters needed to recover our measure for the price equivalent of waiting time without knowing the base price for each request.¹⁵ This formulation allows a separation of the base price and the price multiplier as determinants of the request, and identification of the price

¹³We do not need to specify a global model of time use with every commodity and price in it. Weak complementarity with random prices and wait time at the point of purchase is enough.

¹⁴In general, we do not know if a person actually considered the alternatives that were available, and only know that they were feasible when a decision to select a mode was made. As McFadden (1974) demonstrated, given the specification of the factors influencing an individual’s choices and an assumed choice set, the inability to know a specific default alternative or all the possibilities does not prevent one from assessing the relative importance of each determinant using a random sample of the hypothesized alternatives, together with assumptions that characterize the choice process. For many applications, this constraint on the information available is not important to the results.

¹⁵We considered other functional forms as robustness checks; see Table C.23 in Appendix C. We also discuss this in section 3.4.1.

effect through the experimental variation in the price multiplier. Since we observe the base price only for a self-selected sub-sample of sessions (as described below in the first field experiment), we include this effect as one of the components of the model’s error.¹⁶

Our *price equivalent of a unit change in the wait time* is defined by the marginal rate at which the passenger is indifferent in trading off units of waiting time with units of monetary cost for the trip. For our basic model with time and price in logs, as in (9), this rate is (minus) the marginal rate of substitution of T_{ij}^a and P_{ij} :

$$-\frac{dP_{ij}}{dT_{ij}} = \frac{\partial U_{ij}/\partial T_{ij}^a}{\partial U_{ij}/\partial P_{ij}} = \frac{\beta_1}{\beta_2} \frac{P_{ij}}{T_{ij}^a} \quad (10)$$

The value for this price equivalent depends on the values of P_{ij} and T_{ij}^a as well as the values of β_1 and β_2 .

Given that ε_{ij} in equation (9) includes the default alternative for each person opening the Lyft app, as well as many unobserved factors that will simultaneously affect both the utility of rides and that of the default alternatives, it is reasonable to expect that it will not be exogenous to wait time and prices. We address this limitation in two ways. First, we use experimental variation to construct instruments for the price multiplier and waiting time terms and use two-stage least squares to estimate our model. As a result, consistent estimation of β_1, β_2 can be realized with instruments that provide exogenous variation in $1 + PT_{ij}$ and T_{ij}^a that are independent of W_{ij} , B_{ij} , and ε_{ij} . The power of using a field experiment with randomization of both prices and wait time is that we do not need to make any further assumptions about customer behavior. Second, our large sample allows the definition of a variety of sub-samples that identify different circumstances in which the Lyft app is opened. Some of these outside factors can also be expected to affect the (baseline) ride prices and waiting times.

Our model is therefore estimated with two-stage least squares (2SLS), with the first-stage equations as:

$$\ln(\text{ETA}_{ij}) = \gamma_0 + \gamma_1 \mathbf{T1}_{ij} + \gamma_2 \mathbf{T2}_{ij} + \gamma_3 \mathbf{T3}_{ij} + \gamma_4 \mathbf{T4}_{ij} + \gamma_5 \mathbf{T5}_{ij} + (\text{controls}) + \eta_{ij} \quad (11)$$

¹⁶In early discussions of the random utility model, the framework was used to describe a choice among a discrete set of alternatives. As a rule, when there were more than two possibilities, a logit framework was often adopted. More recently, these estimators have been generalized to allow for unobserved heterogeneity in preferences by specifying some coefficients as random variables and relying on mixed logit estimators. Mixed logit avoids the restrictive assumption of independence of irrelevant alternatives (IIA) with a simple logit approach. When used for measuring willingness to pay the distributions for the coefficient for the price of alternatives need to be restricted to assure consistent welfare measures. When the alternatives are limited to two possibilities (taking a Lyft trip or not), and one is not attempting to account for a set of observed characteristics distinguishing the alternatives, then linear regression methods are often used. Ordinary least squares (OLS) and two-stage least squares (2SLS) provide robust strategies for estimating the parameters needed to recover the marginal VOT and evaluate the sensitivity of the results to the circumstances characterizing the context of the choice. We report in Appendix C the sensitivity of our conclusions to alternative estimators in Table C.24

$$\ln(1 + \text{PT}_{ij}) = \delta_0 + \delta_1 \text{T1}_{it} + \delta_2 \text{T2}_{it} + \delta_3 \text{T3}_{ij} + \delta_4 \text{T4}_{ij} + \delta_5 \text{T5}_{ij} + (\text{controls}) + \zeta_{ij} \quad (12)$$

and the second-stage equation is:

$$\text{Request}_{ij} = \beta_0 + \beta_1 \ln(1 + \text{PT}_{ij}) + \beta_2 \ln(\text{ETA}_{ij}) + (\text{controls}) + \varepsilon_{ij}. \quad (13)$$

Here T1 through T5 are dummy indicators of the experimental treatment assignments (described below), and *controls* include a vector of fixed effects for passenger and session types (e.g., a user’s number of past rides with Lyft) as well as time and location, controlling for the unobserved variation in W_{ij} and $\ln B_{ij}$.

The estimated β_1 and β_2 , as well as assumed values for P_{ij} and ETA_{ij} , allow recovery of the price equivalent of a unit of waiting time. We evaluate our price equivalent at the control average waiting time and control average price on completed rides during our experimental period. The incremental price equivalent is estimated by:

$$\frac{\hat{\beta}_1 \bar{P}}{\hat{\beta}_2 \bar{T}^a},$$

where \bar{P} and \bar{T}^a are the average price and waiting time for control units.¹⁷

In Section 3 below, we fully describe our two field experiments to identify how agents trade off units of waiting time with units of monetary cost. Our first field experiment randomly assigns Lyft users to a control or one of five treatments (T1 through T5 in our 2SLS model) varying (increase, decrease, or market) price and/or wait time (increase or market). Each individual remains with the assignment they are given at the time of their first opening of the Lyft application during our experimental time period (eight weeks). The realized values for both wait time and price depend upon the circumstances of the session; the Prime Time multiplier and waiting time are thus endogenous variables to the request decision. We use the features defining our treatments to define instruments for both variables. We also add additional controls that identify the features of the context in which a person opens the app. We use the 2SLS approach above to estimate our choice equations, but also report alternative estimates in Appendix Section C.

A second, complementary, field experiment—in which only ETA increases are randomized (while prices remain at market values)—explores more fully the circumstances of choice and how they allow us to measure heterogeneity. In particular, extra wait time is varied across location-time blocks (as opposed to across users) which provides (i) a robustness check on the results from our first field experiment and (ii) additional insights into how users’ time elasticities may vary

¹⁷The decision to evaluate the VOT at the control average price and waiting time is somewhat arbitrary, and due to the nonlinearity of the expression for the VOT in price and waiting time, the VOT expression evaluated at the average price and waiting time may differ from the average VOT. We address this concern in Appendix L by constructing a VOT estimate for each observation using observation-specific price and waiting time predictions (and semi-elasticities); this process produces a full distribution of VOTs across sessions rather than a single estimate.

over a different set of changes in wait time. Both field experiments are discussed more patiently below. Overall, our approach of randomizing price and wait time across consumers, together with knowing whether they decided to take the ride (based on both price and time), provides a set of VOT estimates unique to the literature in terms of approach, diversity of situation, and scale.

2.3 Context: Background on Lyft and Data

Lyft is a ridesharing platform that matches consumers (passengers) searching for motor vehicle transportation with independent contractors (drivers) providing the service.¹⁸ Passengers access Lyft through a smartphone app (Figure 1 shows the passenger user interface at the time of the experiment). The app shows the nearest driver’s estimated time to arrival (ETA) and the Prime Time (PT) price multiplier active for the potential passenger’s current location. ETA is an estimate of the time, in minutes, that it would take the nearest driver to reach the passenger’s location from the moment their request is accepted by the driver. Prime Time (PT) is a dynamic mechanism that increases prices from a base level to balance the local amount of Lyft ride requests (demand) and the local available pool of Lyft drivers (supply). A PT multiplier of +25% means that a ride will cost 25% more than the usual base fare (which is a deterministic function of distance and time of travel for each metro area).

In addition to showing an ETA estimate and the current PT level, the app also allows the passenger to input a destination. At the time of the first field experiment, users could see the ETA estimate and PT multiplier and request a ride without entering a destination. If a user did enter a destination, they were shown an estimated range for the trip’s cost, not including the effects of the PT multiplier on the price. These estimated ranges were only seen by the passenger in approximately one-third of sessions in our sample for the first experiment. Because the ranges were expressed in terms of the base price, they were not affected by the experimental treatments. Our primary model for describing the factors influencing Lyft requests estimates the price effect using the experimental variation in the PT multiplier and does not require information about the base price.¹⁹

Each opening of the Lyft app by a passenger starts a session, which is the primary unit of observation for our analysis. A session ends either when the passenger takes a trip or after 30 minutes of inactivity. So, for example, if a user closes the app without taking a ride and reopens it within 30 minutes, both these interactions count as a single session. For each session, Lyft records the following information: passenger’s unique ID code, whether the rider is registered as

¹⁸The following discussion of the Lyft ridesharing platform describes what existed in late 2015 through early 2016, the time period during which the first experiment was conducted. Lyft also now provides options for bikes, scooters, transit, and rental cars on its platform.

¹⁹Because price ranges are given only when a potential rider enters a destination, this information could affect the response to the experimental variation in the multiplier; we address this issue by considering the sample of respondents who enter a destination separately from those who did not as part of our robustness analysis in Tables C.26, C.27, and C.28 in Appendix Section C.

a business user, the local start time, the passenger’s current location (latitude and longitude), counts of how many requests the passenger makes and rides the passenger completes, as well as the ETA and PT shown to the passenger in the session. ETA and PT may vary over the course of one session due to real-time changes in local supply and demand. As such, our analysis focuses on the last shown ETA and PT in the session, as these are the ones faced by the passenger at their final decision node of whether to request a ride.²⁰ Thus, all of our discussion of the ETA and PT will refer to the last value presented in a session (unless otherwise specified).

From these data, we can define a number of variables to characterize the circumstances facing the potential passengers as they made their choices. For example, using a deidentified, unique ID for a passenger, we can determine how many Lyft rides that passenger has taken. Using session start times, we can categorize a session as taking place on a weekend evening, during the morning commute, or in any other time category. And, using location data, we can determine if a passenger is at an airport or at a downtown or suburban location for each metro area in our sample. Finally, note that Lyft offers various ride modes, including Classic (the standard mode), Lyft Line, now called Shared, (in which several passengers share a single car with multiple pickups and dropoffs), and Lyft XL (which offers larger vehicles). At the time of the experiment, the majority of rides (about three-quarters) were Classics. We include all ride types in our analysis, and consider exclusion of non-Classic sessions as a robustness check (see Appendix Section K).

3. Design and Results For Field Experiment 1

To identify how agents trade off time and money, we begin with a first natural field experiment (see Harrison and List (2004) for the various field experiment definitions) that randomly assigned consumers to one of several treatments that differ in both the realized wait time and price (i.e. Field Experiment 1). The second natural field experiment, which we denote as Field Experiment 2, is described in Section 4.

3.1 Design of Field Experiment 1

The process that determines both the price multiplier and the waiting time implies that both variables are endogenous. The PT algorithm is designed to raise prices during periods of relative high demand/low supply; similarly, the number of available drivers at the time each potential rider opens the app in relation to the others who do so at the same time in a location will determine the estimated wait time. Our first field experiment involved nine cities in the U.S. (San Francisco, Austin, Atlanta, Miami, Los Angeles, San Diego, Boston, Seattle, and New York City)

²⁰ETAs vary over the course of a session in 51.3% of sessions, while PT varies in 11.9% of sessions. The intra-session variation in ETAs is caused by drivers continuously moving during the course of a session, and is generally small. For example, in 76.4% of sessions, the difference between the maximum and minimum ETAs shown is one minute or less.

for eight weeks between December 2015 and January 2016, which involved 720,059 customers and 5,177,358 individual sessions.²¹

At the start of the experiment, 37% of all users in each city were randomly assigned to a control or one of five treatment groups. The experimental treatments ensure that there is exogenous variation in the components of each algorithm determining the PT multiplier and the wait time. The algorithm for wait time can increase the wait time above the normal arrival time, but cannot reduce the wait time. As a result, the wait time variations are limited to a high and normal (ETA). Each is matched with three possible treatments for the price algorithm: low, normal, and high. The sub-sample assigned to the normal ETA and normal price combination is treated as the control group. The remaining 63% of users associated with other sessions are excluded from our main analysis.²²

The proportion of users randomly assigned to each group is shown in Table 1.²³ Table C.1 compares the features of the control and the treatment groups in terms of the available variables for describing each session. Based on these variables, the randomization achieved balance on the observable covariates.

Throughout the eight weeks of the experiment, each individual remained in the same treatment. Thus, a user assigned to the low price, high ETA treatment group had all of his or her sessions during the eight weeks subjected to the same algorithm, which would lead to potentially lower price and higher waiting times than what would be the case if they were in the control group. Of course, in practice, specific values for the wait time and price multiplier vary depending on how local conditions affected the outcomes produced by each algorithm. We understand the potential selection effects that might occur from this long-term design (although we use a different design in the second field experiment), and we analyze such effects thoroughly in Appendix Section I.

Price variation was achieved by modifying each user’s PT multiplier, increasing it (for high price treatment groups) or decreasing it (for low price treatment groups) based on local market imbalances. This modification has the effect of raising or lowering a user’s effective price, but not necessarily in every session.²⁴ Waiting time variation was achieved by removing drivers from the nearest driver queue of each affected passenger: the nearest driver, all drivers whose ETA

²¹Sessions in one city, Nashville, were dropped due to implementation problems with the experimental treatments. Including Nashville data in our full-sample regressions does not significantly change our point estimates. See Figure B.3 for the number of sessions per day over the full duration of the experiment (eight weeks).

²²These sessions were not used as additional control observations because they may have been subject to other experiments conducted at Lyft concurrently with our field experiment.

²³Treatment group assignments are determined by applying a hash function to each user’s unique Lyft ID code. Since the assignment of each user to a treatment was random, we can assume that users in each treatment group are a representative sample of ride share users who open the Lyft app in the affected cities during the time of the first experiment.

²⁴More concretely, because PT takes values in a fixed, discrete set (0%, 25%, 50%, etc.), the change in the algorithm’s sensitivity to market conditions may not always result in a different PT level. For example, if the market has much more supply than demand, both the normal and the more sensitive, high price algorithm may find that the optimal PT level in the allowed set is 0%.

was within 30 seconds of that of the nearest driver, and one additional driver were removed from the queue. This removal has the effect of increasing ETA by at least 30 seconds for all passengers subject to the high ETA treatment, but potentially by considerably more than 30 seconds, especially when there are few drivers near a passenger. Because each passenger’s price and ETA treatments are independent, we can identify the coefficients for both price and time effects on the demand for Lyft rides.

Table 2 displays average ETA, PT, and completed ride price by treatment group. The randomization was successful: our high ETA treatment increases the ETAs by an average of approximately 1.6 minutes, an increase of about 52%. The high price treatment increases average PT levels from about 10.0% to 16.3%, while the low price treatment decreases the average PT levels to about 3.6%. These PT differences result in completed ride prices that are about 2.5% higher for passengers receiving high price treatments and 2.0% lower for passengers receiving low price treatments.²⁵

Table C.5 and Figures 2 and 3 show the distributions of PT and ETAs for each of the six treatment groups.²⁶ Appendix B provides graphs of the distributions for each variable, indicating that each of the treatments shifts the distributions in the intended directions. Tables C.3 and C.4 in the Appendix report p -values from Kolmogorov-Smirnov tests of the hypotheses that the distributions of average ETA and PT across users differs between treatment groups; the results suggest that the ETA distributions are approximately identical across PT treatments and vice versa, consistent with the independence of these components of the experimental treatments.²⁷

3.2 Results of Field Experiment 1

3.2.1 Summary Statistics

Lyft customers in the experiment had, on average, about seven sessions and about four sessions with a completed ride (see Table C.2). Customers in the high ETA and high price treatments had fewer sessions, ride requests, and completed rides than control passengers, while passengers in the normal ETA, low price treatment group had slightly more than the control. These differences are consistent with a treatment effect on passenger behavior, which we explore further below.²⁸

Before moving to the formal analysis, we consider the effects of the various treatments on passengers’ demand behavior. The lower panel in Table C.2 shows the average number of sessions

²⁵The effects of treatment on completed ride price are smaller than the effects on quoted PT because PT itself affects the probability that a session will result in a completed ride.

²⁶Two sessions had recorded ETAs of 0 minutes. These were dropped from the data, so that the log transformation could be applied to ETA.

²⁷Figures B.1 and B.2 in the Appendix indicate that these treatments are in effect consistently throughout the course of the experiment.

²⁸Figures B.7, B.8, B.9, and B.10 highlight the heterogeneity in our sample. While the majority of our sessions come from San Francisco and Los Angeles, we have a large number of observations from the other six cities in the experiment. The distributions of sessions over days of the week and hours of the day are relatively balanced, though weekend and late afternoon/early evening times are the best represented time periods in our sample.

which had ride requests for each treatment group. Figures 4 and B.11 display the demand rate, defined in two ways. The first, Figure 4, uses the total requests for service compared to those opening the Lyft app. The second definition uses completed rides in place of requests. Some requests are not completed because a passenger is not matched to a driver, or the passenger or driver cancels a ride before it is finalized. These situations amount to 0.6% and 10.0% of the total requests during experiment 1, respectively.²⁹ As expected, the high ETA and price treatments decrease request rates relative to the control, while the low price treatment increases request rates. The small confidence intervals around the means suggest that these differences are statistically significant at the 95% level, and the magnitude of the differences of request rate between treatment groups—which is as large as four percentage points between the normal ETA, low price treatment and the high ETA, high price treatment suggest the outcomes reflect economically consistent responses to the differences in the circumstances of choice. For example, the high ETA, low price treatment had a slightly higher request rate than control, and the high ETA, normal price treatment had a slightly higher request rate than the normal ETA, high price treatment.³⁰

3.2.2 Empirical Estimation

Our dependent variable is a discrete indicator for a request for the service (1 for request, 0 otherwise) with $\ln(\text{ETA})$ and $\ln(1 + \text{PT})$ as the independent variables of direct interest. We also have a set of controls that include fixed effects for the location, local hour of week and week of year, user experience with Lyft (decile of pre-experiment lifetime rides), and user type (whether the user has a business profile). Since the data generating process implies ETA and PT will be endogenous, 2SLS is our preferred estimator. As part of a robustness analysis, we estimated a probit model (instrumenting $\ln(\text{ETA})$ and $\ln(1 + \text{PT})$) of our main specification and report these results in Table C.24 in Appendix C. Standard errors are estimated clustering within passengers

²⁹Figure B.4 in the Appendix considers how the average number of rides a passenger in each treatment group takes evolves over the course of the experiment, relative to the average number of rides a control passenger takes. The five series are approximately comparable at the outset of the experiment, appear to spread out with the duration of the experiment, and then appear to stabilize about 10 days after a user’s first session in the experiment.

Figure B.5 in the Appendix is a similar plot comparing session rates (that is, percentage of passengers opening the app) between the treatment groups over time. Expected effects of the treatments on rates of use of the platform are observed: passengers facing the higher prices and waiting times become less likely to return to the platform over time. Price and waiting time thus have both intensive- and extensive-margin effects on demand, impacting not only the passenger’s probability of requesting after opening the app, but also the probability that the passenger opens the app again in the future. We explore how this extensive margin behavior impacts the VOT estimation in section 3.4.

Finally, our observations are distributed evenly over passenger/experience levels. That is, we observe a near equal number of sessions for passengers with 0 rides and over 50 rides before the start of the experiment. The presence of this heterogeneity across regions, time, and passengers allows us to investigate how the value of time varies across circumstances.

³⁰In addition to these during-experiment demand effects, we also find some evidence of treatment effects persisting beyond the end of the experiment; see Appendix E.

and are assumed independent across passengers.³¹

Our estimate for the VOT uses the average price and ETA for the control treatments associated with each sample definition. For most of our models, this is defined as:

$$VOT = \frac{\beta_1 \overline{Price}}{\beta_2 \overline{ETA}} \quad (14)$$

β_1 and β_2 are semi-elasticities of demand (request rate) with respect to waiting time and price, respectively. To recover estimates for the relevant elasticities, we divide these semi-elasticities by the average request rate $\overline{Request}$. As previously noted, both $\ln(ETA_{ij})$ and $\ln(1+PT_{ij})$ are endogenous in (13), so we estimate the β s via two-stage least squares (2SLS), with first-stage equations (11) and (12).^{32, 33} The F statistics from the first stage regressions, both with and without other controlling covariates, confirm the strength and relevance of the instruments. Standard errors for our estimates for the VOT are derived using the delta method.³⁴

Our approach to exploring the sensitivity of our VOT estimates follows from the basic structure

³¹This covariance structure would arise if, for example, the true data generating process had β_1 and β_2 as individual-level random effects. Some studies, including Small et al. (2005), have estimated such random (or mixed) effects models directly. Endogeneity of our explanatory variables and the individual-level randomization of our experimental treatments render this approach difficult in our context. Instead, we simply cluster our standard errors to account for the possibility that β_1 and β_2 vary between passengers. Note also that clustering at the individual user-level is consistent with our experimental treatments, which are randomized at the individual.

³²First stage results by region can be found in the Appendix Section C in Table C.29 and C.30.

³³The necessary conditions for 2SLS to give consistent estimates of β s in (13) are that our instruments be orthogonal to the error term ε_{ij} and correlated with the endogenous variables $\ln(ETA_{ij})$ and $\ln(1+PT_{ij})$. Exogeneity follows from the fact treatment randomization, but may fail if exposure to the treatment has a cumulative effect on passengers, which affects their future behavior outside of the effect on their received ETA and Prime Time within each session. Appendix Figure B.5 suggests that such cumulative effects on demand may be present. As part of our robustness analysis, we estimate the model using only each passenger's first session in the experiment and also only observations in the first week of the experiment. We find little variation in the estimates for the VOT over the number of sessions or weeks of the experiment, suggesting no clear selection effect due to consumers who decide not to use the app. As a result, the appearance of small cumulative effects do not appear to affect the time/price tradeoffs we estimate; see Tables C.19, C.20, and C.21 in the Appendix.

³⁴See, e.g., Davidson et al. (2004). Treating price and ETA as fixed, define the nonlinear function $g : \mathbb{R}^3 \rightarrow \mathbb{R}$ by:

$$g(\beta_0, \beta_1, \beta_2) = \frac{\beta_1 \overline{Price}}{\beta_2 \overline{ETA}} \quad (15)$$

so that $g(\hat{\beta})$ is our estimator of the VOT. The Jacobian of g is

$$Dg(\beta_0, \beta_1, \beta_2) = \left(0, \frac{1}{\beta_2} \frac{\overline{Price}}{\overline{ETA}}, -\frac{\beta_1}{\beta_2^2} \frac{\overline{Price}}{\overline{ETA}} \right) \quad (16)$$

which exists (provided $\beta_2 \neq 0$) and is always nonzero. Assume that Σ is the asymptotic variance-covariance matrix of $\hat{\beta}$, that is, $\sqrt{n}(\hat{\beta} - \beta) \xrightarrow{d} N(0, \Sigma)$. By the delta method,

$$\sqrt{n}(g(\hat{\beta}) - g(\beta)) \xrightarrow{d} N(0, [Dg(\beta)]\Sigma[Dg(\beta)]^T) \quad (17)$$

Then if $\hat{\Sigma}/n$ is any estimator of $\text{Var}[\hat{\beta}]$ with $\hat{\Sigma} \xrightarrow{p} \Sigma$, a consistent estimator of $\text{Var}[g(\hat{\beta})]$ is:

$$\frac{1}{n} [Dg(\hat{\beta})] \hat{\Sigma} [Dg(\hat{\beta})]^T \quad (18)$$

of the Becker model. The first of these is the selection of functional form describing how wait time and price influence requests. By transforming the price and time using logs we can estimate the parameter describing how people respond to price knowing only the PT multiplier. Second, to recover a measure of the VOT we need to select a point for evaluating the implied time/price trade-off. As noted earlier, we use the average values for these variables from the control treatments to estimate the semi-elasticities. Finally, to develop insights into how the circumstances of each person’s choice affect the VOT, we estimate the VOT using a set of sub-samples motivated by an extension of the Becker theory, which is presented below.

3.2.3 VOT Estimates

Tables 3 and 4 provide the first- and second-stage results for the main model respectively, estimated for the full experiment 1 sample.³⁵ Table 3 demonstrates that our experiment worked as expected in changing wait times and prices in the correct directions as specified by our experimental groups. In addition, we have strong instruments in the first-stage regressions. In the second-stage results in Table 4, coefficients on $\ln(\text{ETA})$ and $\ln(1 + \text{PT})$ are estimated to be -0.026 and -0.330 when controls are included; without controls, the coefficient on $\ln(\text{ETA})$ is smaller in absolute magnitude, and the coefficient on $\ln(1 + \text{PT})$ slightly larger in absolute magnitude.

Using the full sample where the overall average request rate is 64.2%, we find estimates implying that wait time and price elasticities of demand are -0.0427 and -0.5942 , respectively.³⁶ Our results provide strong support for the conclusion that Lyft requests are influenced by differences in both wait time and the price multiplier associated with the trip. Taken together, our estimated coefficients on $\ln(\text{ETA})$ and $\ln(1 + \text{PT})$ imply a VOT of \$19.38 per hour (s.e. = \$1.39 per hour³⁷) at baseline (i.e. control) waiting time of 3.08 minutes and price of \$13.83 (the average actual fare paid by control riders in the sample).

3.3 How do the Properties of the Situation Affect the VOT?

In this section we combine the logic of the Becker model with the temporal and spatial delineation in our sample to consider how situational features affect our estimates of the elasticities and the VOT. Recall from Section 2.1 that in Becker’s model with Leontief household production, time and

In practice, we use the standard cluster-robust 2SLS variance-covariance estimator for $\hat{\beta}$, with clustering at the passenger level; see [Baum et al. \(2003\)](#).

This standard error ignores the uncertainty introduced by estimating $\overline{\text{ETA}}$ and $\overline{\text{Price}}$ from the data. We also estimated the standard error for our main specification using a pairs-cluster bootstrap ([Cameron and Miller, 2015](#)) on our full estimation procedure, and the results were similar to those returned by the delta method.

³⁵Standard diagnostic tests of endogeneity and overidentifying restrictions are provided in Tables C.17 and C.18 in the Appendix. The results of OLS estimation of equation (13) are in Table C.6.

³⁶Our price elasticity estimate is consistent with [Cohen et al. \(2016\)](#), whose mean point estimate for the price elasticity of demand using Uber records for 2015 is -0.57 .

³⁷A pairs-cluster bootstrap ([Cameron and Miller, 2015](#)) with $B = 999$ replications yielded a standard error of \$1.40 per hour and a bootstrap- t confidence interval ([Efron and Tibshirani, 1994](#)) of (\$16.71, \$22.19).

private goods are perfect complements. When this restriction is relaxed, the wait time and price of a rideshare trip enter the indirect utility function without restriction, and we may consider how each of these influences our measure of the value of time.

Consider the willingness to pay WTP for a price increase ($1 + PT_0$ to $1 + PT_1$) and waiting time decrease (ETA_0 to ETA_1). W is defined by equation (19).

$$V(1 + PT_1, ETA_1, m - WTP) = V(1 + PT_0, ETA_0, m) \quad (19)$$

Denoting the marginal value of time (V_{ETA}/V_m) by π and the demand for rideshare trips by T , we have the following second-order expansion for WTP :

$$\begin{aligned} WTP \approx & \pi(ETA_1 - ETA_0) - T(PT_1 - PT_0) + \frac{1}{2}(\pi_{ETA} - \pi\pi_m)(ETA_1 - ETA_0)^2 \\ & - \frac{1}{2}(T_{PT} + TT_m)(PT_1 - PT_0)^2 + (\pi_{PT} + T\pi_m)(ETA_1 - ETA_0)(PT_1 - PT_0). \end{aligned} \quad (20)$$

When the changes in price and waiting time exactly offset each other, we have $WTP = 0$. If we also assume that $T = 1$ and that the adjustment of trips to price (T_{PT}) and income (T_m) are negligible, we can solve for the ratio of the price change to the equivalent waiting time change:

$$\frac{PT_1 - PT_0}{ETA_1 - ETA_0} \approx \pi + \frac{1}{2}(\pi_{ETA} - \pi\pi_m)(ETA_1 - ETA_0) + (\pi_{PT} + \pi_m)(PT_1 - PT_0). \quad (21)$$

Equation (21) shows that, to first order, the ratio of the price change to the equivalent waiting time change equals the value of time π . The equation further shows that this ratio is also influenced by the size of the waiting time change ($ETA_1 - ETA_0$), the sensitivity of the VOT to the waiting time (π_{ETA}), and the sensitivity of the VOT to income (π_m).

This conventional model fails to capture another of Becker's insights: the context of choice. The simple Leontief version of his model captured this effect by allowing time to be linked in different ways to different private goods. We observe only one private service: rideshare trips. Nonetheless, we can observe how the price equivalent changes in different situations, indexed by location, time of day, and weather condition. This rich detail in the circumstances of people's choices allows the data to add resolution to what a largely unrestricted model may imply for the VOT. By selecting sub-samples distinguished by the location, timing, and other features of passenger sessions, we can evaluate how each affects potential users' decisions with randomly varied prices and wait times, yielding heterogeneity tests across three bins of the data:

1. **Individual trip features** that are correlated with the opportunity cost of time, which include purpose of trip, time and day of week of trip, and specific location of trip. Such monetary values are driven by what activities are crowded out by greater wait times.
2. **Other trip signatures** such as weather conditions, the baseline ETA, and who bears the

marginal cost of the trip. In this case, when the cost burden is shared by an employer the rider should be more sensitive to temporal changes.

3. **Market variables** that are correlated with the opportunity cost of time, which include local wages and the available substitutes, such as access to alternative public transportation. Such variables affect the nature and shape of the demand curve for time.

Taken together, data generated across both natural field experiments combine to shed important insights on factors within each bin. In each bin, a mixture of experimental and non-experimental variation is utilized, and we take caution in interpreting results based on the naturally-occurring variation and selection below. The features which define each bin may be correlated with other features affecting preferences and the circumstance of choice; in Appendix J we attempt to address this concern in a secondary experiment by reweighting observations within each bin to balance the distributions of other relevant observable covariates.³⁸

3.3.1 Individual Trip Features

Table 5 summarizes our findings concerning trip and market features.³⁹ Consider first the different ways of describing the timing for choices: by day of the week and by an hour of the day. In the latter case, we also distinguish weekends (Saturday and Sunday) from weekdays. Our reasoning above yields predictions on how the various situations should affect VOT estimates. More specifically, in the case of distinguishing days of the week we would expect that the opportunity cost of time is higher during weekdays than on weekends. Our estimates in Table 5 support this conclusion, as weekdays experience ETA elasticities larger than we observe on weekends. And, VOTs average nearly \$20 during the week versus around \$18 on the weekend. While these differences are meaningful, overall they are not consistently different at conventional significance levels.

Yet, when we consider the time of the day, distinguishing weekdays from weekends, we find pronounced and statistically significant differences in the VOT estimates at the $p < 0.01$ level. These differences mainly arise in the morning and evening commuting rush hour periods for weekdays (6 to 10 AM and 4 to 7 PM) where we observe larger ETA elasticities (especially in the mornings) and higher VOTs—averaging above \$24—compared to non-commuting time blocks. For weekends, we observe less variation throughout the day, as we might expect.

These estimates suggest that when the wait time crowds out work time the ETA elasticity and the VOT estimates are higher than when it crowds out non-work time. In this manner, the

³⁸We also indirectly examine the impact of reliability by checking the robustness of our results when we include controls for late arrivals in passengers’ previous interactions with the Lyft app in Appendix Section D. Similarly, we attempt to address the potential bias that may arise from not accounting for the effect of in-vehicle time on passenger demand in Appendix Section F. Finally, we check the sensitivity of our results to the presence of always- or never-requesters in our sample in Appendix Section G, and to user selection in Appendix Section I.

³⁹The corresponding full regression tables are C.7 and C.8 in the Appendix.

pattern of results observed in Table 5 are consonant with the notion that individual trip features predictably affect the relevant primitives in the model.⁴⁰

3.3.2 Other Trip Characteristics

Table 6 considers other features of the trip that potentially affect the marginal sensitivities to price and waiting time.⁴¹ Our data from this first field experiment allows us to distinguish choices along the dimension of wait time unpleasantness as well as who bears the burden of the trip. Field experiment 2 provides the necessary variation to explore how baseline ETA variation impacts VOT estimates and therefore provides a sense of the curvature of the VOT. We discuss those results in Section 4.

The top panel of Table 6 provides our full sample estimates, and as a comparison the middle panel provides relevant estimates when there is rain, snow, and no precipitation.⁴² We find that the trips with adverse weather conditions suggest a clear and discernible impact on the VOT estimates, suggesting that precipitation makes longer waiting times considerably less desirable. For instance, in the presence of snow, the ETA elasticity is more than two times higher than when there is no precipitation. This enhanced time sensitivity leads the VOT estimate to be nearly 40% higher than no precipitation trips (\$26.56 versus \$19.04). The effect of rain, while muted compared to snow, is also significantly larger than when there is no precipitation.

While both rain and snow cause the ETA elasticities to be higher, and in turn increasing the VOT estimates, tempering these increases are the price elasticity increases that we observe. While such increases are indeed small, we expected that as weather conditions worsen consumers would become less sensitive to price changes. These estimates may reflect that baseline PT is higher in such cases (15.2% in rain and 20.0% in snow).

The bottom panel of Table 6 provides insights into how the burden of cost affects the relevant elasticities. We use information on whether the trip was taken by a business or a non-business user, and make the assumption that the cost burden is lower for the former either because their firm pays the expense or if that is not the case at least the expense can be used a tax write-off. In both cases, the consumer bears less than 100% of the cost of trip burden.

We find that business users' ETA elasticities and VOTs are different from those of non-business users in expected ways: non-business users have a price elasticity that is roughly 50% larger (-0.605 versus -0.399), and a value of time about 25% higher. Inspecting the relevant elasticities, we note that this difference is driven by business users showing much less sensitivity to price, presumably reflecting the fact that business users may not themselves bear the full monetary

⁴⁰In Appendix I, we find that these heterogeneity results are robust to correcting for user selection into different trip contexts.

⁴¹The corresponding full regression tables are C.9 and C.10 in the Appendix.

⁴²We obtain historical realized weather data from Dark Sky at the geohash4-hour level, and match this information to sessions based on the passenger's geohash4 and the starting hour of the session.

cost of a trip.⁴³

3.3.3 Market Factors

Results on the effects of market factors on the VOT are summarized in Table 7.⁴⁴ We distinguish choices by location in three ways: by metro area; by prominent locations within each metro area (airports and downtown locations); and by distance to the nearest public transport station/stop.⁴⁵ Point estimates of the VOT vary between metro areas, but we fall short of rejecting the null of homogeneity across regions ($p = 0.154$).⁴⁶ While the estimates for Atlanta, Austin, and Seattle are distinctive, with the VOT for the first two metro areas smaller than our overall sample estimate, only the results for Atlanta would be judged as having a marginally significant difference from the estimate for the rest of the sample. We discuss more of the city-level estimates with respect to wages in Section 5.1.

Distinguishing downtown locations, the estimates indicate a greater VOT for downtown locations compared to non-downtown.⁴⁷ Trips initiating at airports have numerically larger VOT estimates, but the small sample size for these cases prevent judgement of statistical significance. The grouping of trips based on distance to nearest public transit stop yields VOT estimates that decrease with distance, consistent with available substitute reasoning. However, this category is also likely to serve as a proxy for non-downtown locations, which have lower VOT estimates.⁴⁸ In the second experiment below (section 4.2.1), we have random variation in wait time that is completely orthogonal to the base ETA.

4. Design and Results for Field Experiment 2

4.1 Design of Field Experiment 2

Our analysis of the second natural field experiment has three primary goals: (1) replicate the main insights gained from the first field experiment on ETA elasticities and various heterogeneity tests; (2) assess the robustness of the observed heterogeneity to sample reweighting adjustments; and (3) explore how the base level of wait time affects ETA elasticities and how exogenous changes

⁴³We are unable to assess the robustness of precipitation and business/non-business results to user selection into contexts because our weather data does not extend far past the beginning of the first experiment and business status does not vary within user.

⁴⁴The corresponding full regression tables are C.11, C.12, C.13, and C.14 in the Appendix.

⁴⁵Public transit stop/station locations are taken from the National Transit Map, and distance is calculated as the great circle distance between a passenger’s location and a stop.

⁴⁶When considering the marginal demand effects of price and waiting time separately, we do reject the nulls that time ($p = 0.046$) and price ($p < 0.0001$) semi-elasticities are homogeneous across regions.

⁴⁷Maps of the areas tagged as “downtown” can be found in Figure B.6.

⁴⁸Results in Appendix I suggest that the distance to transit results are not primarily driven by user selection, while the downtown/non-downtown and airport/non-airport may be. We are unable to effectively control for selection into regions due to limited within-user variation in region (only 19% of users have a pre-experiment session in more than one region. Full tables of regression results are available in Appendix C.

in the length of wait time affect the VOT. That is, by precisely varying wait time at the pickup location and hour (as opposed to across users in Field Experiment 1), we develop insights into how the timing of different wait times along with their location affects time elasticities. To do these chores, we randomized location-hour blocks into one of the following four groups: (1) Control (10% of the city sample); (2) ETA plus at least 60 seconds (5%); (3) ETA plus at least 150 seconds (3%); and (4) ETA plus at least 240 seconds (2%). To increase ETAs, we removed the closest drivers from the dispatch queue until the ETA was increased by at least the treatment amount.

The experiment took place for eight weeks from April 2017 to June 2017 in Los Angeles, San Francisco, New York City, Chicago, New Jersey, Boston, Philadelphia, Washington, D.C., Miami, and Atlanta, involving 3.3 million passengers and 9.7 million sessions (5.9 million of which had requests for service). The randomization used in this second experiment differs from that of our first experiment in that treatments were assigned at the pickup location-hour level across the whole eight-week period. Twenty percent of all location-hour blocks in these cities in this eight week period were in the experiment.⁴⁹ Price multipliers were not randomized in this experiment; consequently, our focus is on waiting time elasticities rather than VOT estimates. Figures 5 and 6 and Table C.32 show that the experimental conditions and the treatment groups have the expected variation and overall effects on demand.⁵⁰

4.2 Results of Field Experiment 2

To summarize the insights gained from our second field experiment, we split our results into two groups. First, we summarize the overall time elasticity estimates from field experiment 2, and explore how the properties of the situation affect the estimated ETA elasticity. In doing so, we explore how well results from field experiment 2 map into those from field experiment 1. Second, we delve into the shape of the estimated time elasticities. Section 5 focuses on the policy implications of our results from the combination of both field experiments, digging deeper into the relationship between VOT and local wages as well as exploring how our preferred estimates relate to preferred VOT values used by policymakers.

⁴⁹The location unit used for this randomization is a geohash7, which has a width and height of at most 153 meters. Counts of geohash7s per region can be found in Table C.31. In all the analysis of data from this experiment, standard errors are clustered at the geohash7-hour level, to match the clustering of the randomization scheme.

⁵⁰Figures B.12 and B.13 show the timeline of the ETAs and PT respectively over the course of the experiment. It is clear that both the ETAs and PT remain pretty stable over time apart from the holiday season, especially New Years Eve where both prices and ETA increase dramatically.

4.2.1 Time elasticities and how they vary: a replication

The first column of Table 8 provides estimates of the time elasticity from field experiment 2. As shown in the top row of Table 8, the estimated ETA elasticity for the full sample is -0.043 .⁵¹ This aggregate estimate is quite close to the ETA elasticity estimate from field experiment 1; indeed, we fail to reject the homogeneity null at the $p < 0.05$ level for these two estimates. When we focus on data drawn from the roughly 200,000 customers who are involved in both field experiments, the estimates are also statistically indistinguishable: we fail to reject the homogeneity null at the 5% level.⁵²

The remaining estimates in column 1 of Table 8 explore heterogeneity in ETA elasticities along days and various times of the week. Concerning such trip features that affect the opportunity cost of time, we find qualitatively similar patterns in the ETA elasticities as we found in the first experiment: weekdays have larger time elasticities than weekends, and weekday elasticities are largest during peak commuting times. Similar to experiment 1, these results show the increased importance of time when it relates to opportunity cost during commute hours versus other hours.

Column 2 in Table 8 shows how the ETA elasticities vary with other trip characteristics such as inclement weather and business users. In this case, we again find that time elasticities are higher in rainy conditions: -0.047 versus -0.042 when there is no precipitation.⁵³ In this case, we find that business users exhibit a larger time elasticity than non-business users (-0.049 vs. -0.043) at the $p < 0.01$ level, despite facing lower ETAs on average (see Table C.40).

In column 3 of Table 8, we show how ETA elasticities vary spatially. Consistent with our findings in experiment 1, we find differences across cities in our sample, with New York and Washington, D.C. having the largest elasticities (-0.086 and -0.054) and Miami and San Francisco having the lowest elasticities (-0.021 and -0.031). Furthermore, non-downtown sessions have a larger time elasticity than downtown sessions (-0.048 vs. -0.039) at the $p < 0.01$ level, though the average ETA is higher in the non-downtown sessions than the downtown sessions (Table C.42).

Finally, we again find that distance to public transit matters: the largest ETA elasticities are observed for passengers more than 800 meters from their nearest public transit stop. This result runs counter to economic intuition: passengers with better outside transportation options should respond more elastically to higher waiting times. Distance to transit, however, is correlated with other relevant characteristics of the trip, such as base ETA (Table C.43). In Appendix J, we

⁵¹In these regressions, we include $\ln(1 + PT)$ as a control variable for consistency with the model used for the first experiment, though it is endogenous to the decision to request. Our estimates of time semi-elasticities replace ETA with an instrument constructed as predictions from a first stage regression including the treatment indicators and fixed effect controls for region, geohash5, local hour of week, local week of year, business user, and decile of user lifetime rides (see Table C.33). These controls are also included in the second stage equation for Lyft requests. Consistency of the estimates for our time semi-elasticity relies on the exogenous treatment effects and the common controls included in both the first and second stage equations. The endogenous price multiplier will not be correlated with our experimental variation in the ETAs each user faces, and as a result will not affect the consistency of our ETA semi-elasticity of demand estimate. The ETA treatments do not affect $\ln(1 + PT)$, as we show in Table C.34.

⁵²These results can be found in table C.35.

⁵³The second experiment had no sessions for which the reported precipitation type was snow.

adjust for these differences in the distribution of other covariates through sample weighting, and find that the ETA elasticity is largest for passenger near public transit, as expected from the theoretical model (Table J.6).⁵⁴

4.2.2 The shape of time elasticities

An important feature of the second experiment is that we have the ability to observe multiple levels of experimental ETA increases. As a result, we can estimate the elasticities and VOTs implied by different increases in the wait times and can gauge the implied nonlinearity in the relationship, independent of the base wait time.

To operationalize our approach, we consider that our four levels of ETA in the field experiment provide the ability to contrast three levels of wait time, each of which may serve as an instrument to estimate the ETA elasticity of demand at different points of the demand curve. We consider three sub-samples of our data: all Control and Plus 60 sessions, all Control and Plus 150 sessions, and all Control and Plus 240 sessions. In each sub-sample, we estimate our main demand equation, using the indicator of treatment (respectively Plus 60, Plus 150, and Plus 240) as a single instrument for $\ln(\text{ETA})$. The resulting coefficient estimate is a weighted average of the expected derivative of demand with respect to $\ln(\text{ETA})$ over the ETAs between the two treatment levels in the sub-sample (Angrist et al., 2000). We divide this quantity by the average request rate in the sub-sample to obtain an estimate of the ETA elasticity of demand over the ETAs in each sub-sample. As ETAs increase from one sub-sample to the next, comparing these elasticities gives insight into how the ETA elasticity of demand responds to increases in ETA.

Table 9 and 10 provide empirical results of this analysis. Overall, across every base ETA level we find that increasing the wait time reduces demand. This finding represents a good rationality test. In addition, we find that time elasticities are increasing in base ETA: at a one minute base ETA, the elasticity ranges between -0.005 and -0.018 , whereas at a ten minute base ETA the elasticity ranges between -0.230 and -0.289 .

The interpretation of the elasticity of -0.026 in the first row and third column of 9, for example, is that toward the lower end of the demand curve with a 3 minute base ETA, a 1% waiting time increase translates to a 0.026% decrease in the quantity demanded. As we can see in the second row, a 1% increase in ETA up to the middle portion of the demand curve reduces demand by 0.044%, and the third row shows that a 1% increase in ETA up to the upper end of the demand curve reduces demand by 0.066%. This pattern suggests that time elasticities increase as we move further above the base waiting time.

We find that this result is consistent at all but one base ETA between one and ten minutes; time elasticities are larger for larger increases (see Table 9). In Table 11, we also find that this

⁵⁴In experiment 1, we find that correcting for user selection into different distance bins does *not* significantly affect our ETA elasticity estimates; see Table I.5.

finding holds across all cities in our sample. We also find that ETA elasticities are increasing in ETA: (i) in both downtown and non-downtown sessions (Table C.46); (ii) across all days of the week (Table C.47); and (iii) across all hours of the day (Table C.48 for weekdays and Table C.49 for weekends).

There could be many mechanisms underlying the observed shape of the estimated time elasticities, but we are unaware of any previous causal evidence that shows the marginal value of time varies dramatically over base wait time. Nevertheless, one concern in interpreting our analysis on the shape of the time elasticity is that consumers' responses may reflect a belief that a longer ETA implies greater congestion, and then erroneously expect that their travel time would be longer than what was stated on the app. Under this interpretation consumers might be responding to the changes in the expected ETD in addition to changes in ETA.

We believe that this interpretation is not supported by the structure of the experiment. First, if ETA increases by one minute on the app, then the ETD is also updated by one minute, to reflect the increase in pickup time. Second, it seems unreasonable to expect that consumers believe the quoted ETA but not the quoted ETD. Third, the consumer might believe that as time moves further away from the present, predictions become less accurate, and as a result the variance in ETD increases. We lack the necessary data to seriously investigate such an effect, but leave the question open for future research.⁵⁵

5 Issues in identification and external validity of our estimates

We next discuss potential threats to the internal and external validity of our natural field experiments. We focus on functional form assumptions (section 5.1), app competition (section 5.2), endogenous opening up of the app (section 5.3), experimental contamination (section 5.4), and external validity of our experimental results (section 5.5).

5.1 Functional form and specification

Our main econometric specification has wait time entering the demand equation in logarithmic form. This functional form assumption implies diminishing marginal cost of wait time. We assess the robustness of this assumption in Tables C.23 and C.25 by running our main regressions with ETA entering in levels, and find slightly larger VOTs (but not at conventional significance levels). The first experiment involves only a binary ETA treatment which does not precisely pin down an exact ETA increase (in sparse markets, for example, the treatment may induce larger absolute ETA increases than in dense markets); as a result, the experimental data provide no empirical basis for deciding between different functional forms. The second experiment randomizes sessions

⁵⁵Table F.1 suggests that, in the 2017 experiment, ETA elasticities do not vary with the estimated trip time (pickup to dropoff), which suggests that the elasticities and VOTs we have estimated capture passengers' responsiveness to pre-trip waiting time, and not some combination of pre-trip and in-vehicle time.

between multiple levels of ETA increases, which provides some rough insight into the local shape of the demand curve, but is still insufficient for pinning down the exact functional form.

Another decision we made is to use a simple linear probability model as opposed to a more sophisticated hierarchical model. A potential issue with LPM is that of aggregation. Suppose that there are two types of consumers, one elastic to price and inelastic to wait time and the other inelastic to price and elastic to wait time. The average VOT may be underestimated when taking the ratio of average time and price responsiveness. In Appendix L, we instead use an interacted model to estimate price and time coefficients for different combinations of session features, compute VOTs within each combination, and then average across the population. The resulting mean VOT is within two standard errors of our main estimate.⁵⁶ A more complex random-coefficients approach would require imposing further structural assumptions on the model that are difficult to verify.⁵⁷ One attractive feature of our approach is that we can recover the VOT with the Becker model and weak complementarity without structural assumptions.

5.2 Competition

Another potential concern is that Lyft had competition in the market during the experiment, so multi-homing behavior could potentially cause the outside option to be non-static. Importantly, our randomization is orthogonal to the outside option at that time in which the customer uses the app. It is unlikely that outside providers would have been able to track and respond to the random variation induced at the session-level in real time.

5.3 Endogenous opening of the app

Our VOT is estimated only when riders are considering using the service. That is a key feature of the weak complementarity assumption in the Becker model for estimating a VOT. However, we do not observe behavior outside of opening and using the app. If users only open Lyft when particularly in a hurry, then our VOT estimates may tend to overstate the overall average VOT.

We explored this potential issue with two distinct approaches. First, as a proxy for urgency of use, we compare VOTs across quartiles of user rides taken in the 28 days before the first experiment in Table C.16. The thought is that those who use the app frequently are using it for reasons beyond being in a hurry. We find a slight decreasing trend of VOT with respect to prior 28 day rides. However, the VOT of the most regular users remains only about 20% lower than our main estimate for the full sample; and this difference is not significant at the 5% level, suggesting that selective opening of the app does not greatly bias our results. One shortcoming with this first

⁵⁶These results may be sensitive to the choice of features used to define groups of sessions.

⁵⁷For example, Buchholz et al. (2020) utilizes a hierarchical model with individual-specific coefficients to estimate the entire distribution of VOT by assuming (i) a particular parametric form to the distribution of individual-specific heterogeneity and (ii) that heterogeneity in time and price responses not captured by additive effects for a few locations and time categories is individual-specific; the full model is then specified by a (relatively) small number of parameters, which can be estimated by Bayesian or maximum likelihood methods.

approach is that usage could be endogenous to other things that are related to VOT, and so one would need to randomize urgency, wait time, and price to provide exact insights.

Our second approach to the endogenous opening of the app is to examine VOT across individual usage rates. The idea being that those customers who use the Lyft app regularly (e.g., for daily commute) likely have a stronger preference for Lyft as a means of transport and so do not simply use it when they are particularly in a rush. Alternatively, those consumers who use it only a few times per month are primarily using it when they are in a rush. Importantly, we do not find significant variation in the VOT by lifetime prior rides (see Table C.15 in the Appendix).

5.4 Contamination

When a user decides whether or not to request a ride, their decision affects total supply in the market, which may in turn influence the request decisions of other users (this is analogous to the displacement effect found in Crépon et al. (2013)). This channel introduces interaction effects in the experiment: the outcome of one session can be dependent on the treatment assignments of other sessions, since those treatment assignments affect the overall available supply pool. While these interaction effects technically violate the SUTVA assumption, we argue that the violation is minor and relatively unimportant, because (i) these interaction effects will on average uniformly affect all sessions in the market, across all treatments and in and out the experiment; (ii) the closest data available to shed light on the size of these interaction effects (described in Appendix J) suggests that their magnitude is small (on the order of a few seconds); and (iii) as we argue in Appendix M, despite this violation of SUTVA, we can still interpret our estimates as weighted average causal derivatives of demand.

5.5 External validity

Many of the above considerations relate to generalizability, or external validity, of our VOT estimates. As List (2020) notes, “all results are externally valid to some setting, and no result will be externally valid to all settings.” Given that we view our results as speaking to not only tests of theory, but also to any application that involves time, the populations of people and situations to which our estimates apply merits serious consideration. In terms of sampled population, one key feature of our sample is that nearly 60% of individuals in the US have used rideshare, so consumption of the good itself is quite common. And, importantly, as shown in Appendix H, when we re-weight our user sample to more closely match the characteristics of the US population, our VOT estimates are not significantly changed. This suggests that the VOT of those customers who use ridesharing services is not significantly different from those individuals who do not use ridesharing services.

A next important consideration is whether there is sufficient similarity in relevant situational conditions to generalize our results (List, 2020). A first key feature is whether the experiment

places the agent on an artificial or natural margin when making key decisions. In this manner, it would be difficult to find a better setting than the one studied in this paper to estimate the VOT within the theoretical construct of the classic studies. Indeed, our field experiment leverages key features of common purchase decisions in any familiar setting that involves price and quality trade-offs: several alternative choice sets are provided naturally and purchase decisions are made in a familiar form. In this spirit, such choices would be made whether or not our natural field experiment was being conducted.⁵⁸ We view this aspect of our approach as quite attractive.

A next important consideration is whether our two field experiments present relevant conditions to explore temporal trade-offs that are sufficiently similar to other settings. The richness of our data provides several variants that speak to very different situations. We have estimated wait time trade-offs at various opportunity cost levels: when wait times crowd out work time, leisure time, airport time, and nearly any other moment. Likewise, we have considered how liberal changes in the substitute set affect VOT estimates and how weather and aspects associated with who pays for lower wait time affects VOT. Where our approach has less coverage revolves around the circumstances of the wait (standing curbside rather than in gridlock traffic) and for time-tradeoffs that involve considering the opportunity cost of hours, rather than minutes, of wait time. We trust that future work in different settings and over different choices can complement our VOT estimates.

6. Implications for public policy

In this section, we use our VOT estimates to consider their implications for public policy. First, we offer some comparisons with the city's wage rate to examine whether local wage rates provide an important correlate with our VOT estimates. Second, we provide examples of how our VOT estimates impact certain public policy analyses that rely on estimates of the opportunity cost of time to assess the associated public actions. Third, our estimated price and wait time elasticities are nicely complemented by contemporaneous work that leverages structural models to estimate

⁵⁸Importantly, our study does not require selection into surveys or the experiment, unlike several others in the empirical VOT literature. Our study also does not rely on strong structural assumptions (see above) for identification.

the VOT.⁵⁹

6.1 The relationship between VOT and wages

Our region-level estimates of the VOT allow a test of how well conventional rules-of-thumb for approximating VOT hold across different metropolitan areas in the United States. Many of these rules-of-thumb for the VOT are formulated as percentages of the wage rate. For example, the [USDOT \(2015\)](#) recommends that the VOT be approximated as 50% of hourly earnings for personal travel (the same as [Small \(2013\)](#)); others, particularly in the recreation demand literature stemming from early research by [Cesario \(1976\)](#), use 33% of hourly earnings as an approximation (see [Phaneuf and Smith \(2005\)](#)).

To assess the relevance of the various approximations, we use the estimates from our first field experiment across regions in an array of 16 possible rule-of-thumb approximations. The 16 rule-of-thumb approximations are given for all possible combinations of using the median or mean wage; using pre- or post-tax earnings; and using 1/3 or less, less than 1/2, greater than 1/2, or 100% or more of the resulting hourly earnings. We use one-sided tests to explore whether the VOT is below 1/3, below 1/2, above 1/2, and above 100% of earnings, for a total of 16 tests per region. The tests are all one-sided asymptotic z -tests with level $\alpha = 0.05$, based on the VOT estimates, standard errors, and median and mean wages reported for each region in Table [C.11](#).⁶⁰ For a robustness check, we explore all 16 approximations with both a log and a linear specification to estimate the VOT.⁶¹ For simplicity, we assume a constant 25% tax rate across all regions for both median and mean wages.

Table [12](#) summarizes our empirical results for the log and linear specifications. Each cell

⁵⁹[Buchholz et al. \(2020\)](#) identifies 30 locations within Prague and considers ride choices for different days of the week, time periods during the day, weather conditions, and the location (inside or outside of structure) where choices are made. The location choice for consumers is based on a random utility specification that assumes price and time effects have random coefficients with means varying based on the features used to distinguish trips. The random error in both the price and time coefficients is assumed to reflect individual heterogeneity. A VOT in a location is established in a second estimation step. The model maintains that the choice to move from one location to another identifies the difference in the values of time for the origin of the trip compared to other locations. Their estimates for the VOT's for each location exploit the panel nature of their sample and use the random coefficient component of their model to construct the data for a linear program that is used to estimate the values of time for classes of individuals in each location. They rely on a control function strategy to help take account of price endogeneity. Their VOT estimates display sensitivity to the timing of the time that is similar to our findings, with higher values for peak commuting times on weekdays. While the absolute magnitudes of their price and time elasticities are about 10 times larger than our estimates, the relative size of the relevant elasticities are comparable. Moreover, during work times, the value of time implied by their model is approximately equal to the average wage rate. The second study, due to [Castillo \(2019\)](#), develops a structural model for the demand, supply, and market matching of Uber rideshare users and suppliers in Houston. The demand component of the model is estimated with a linear choice model that relies on the location specific price multipliers together with the rounding of surge adjustments to provide sufficient exogenous variation in prices and avoid bias in estimating price effects on rideshare choices. This study finds a much greater value of time, about two dollars a minute. Yet, the qualitative distinctions isolated for weekend and weekdays conform well to our findings, and the absolute magnitude of the price elasticity is comparable to our estimates.

⁶⁰We take the median and mean wages as constants for the purposes of the hypothesis testing.

⁶¹To recall, the linear specification differs from the log specification because we replace $\log(\text{ETA})$ in the first and second stages (equations [\(11\)](#) and [\(13\)](#)) with the level of ETA in hours. This specification keeps the price (multiplier) in log, as base prices (which would be necessary for calculating the level of price) are generally unobserved in our data.

records the number of regions (of eight) for which the null hypothesis that the VOT in that region equals the rule-of-thumb approximation of that cell is rejected at the 5% level. Overall, we find evidence in favor of rejecting the 1/3 or less earnings rule for most regions, and have the least evidence in favor of rejecting the 1/2 or more pre-tax mean wage and 100% or more post-tax median wage earnings rules. We do not reject the hypothesis that the VOT is above 1/2 earnings for any region, but do reject the hypothesis that the VOT is above 100% earnings for some regions when using the linear specification. We repeat the same exercise using the city VOTs estimated using a linear functional form for the demand equation, of the same form considered in Section 3.3.2 (see Table C.25 for the city-level VOTs). The test results are broadly similar between specifications. Table 13 shows a similar pattern for the pre-tax median and mean wage values, with slightly more cities with a VOT over 100% of the pre-tax mean wage being rejected.

We provide an ocular depiction of these results in Figure 7, which plots our point estimates of the VOTs and their 95% confidence intervals alongside the post-tax mean and median wages (and the policy-relevant 1/2 and 1/3 of these quantities) for each region. Our region-level estimates also provide insights into the cross-sectional relationship of wages and the estimates for the opportunity cost of time. Using our log-functional-form estimates, we compute a Spearman (rank) correlation of 0.73 ($p = 0.025$) between the mean wage and the value of time estimate across regions. In Table 14, we apply the log transformation to the mean wage and our VOT estimates and run OLS and GLS regressions to estimate the cross-sectional elasticity of the VOT with respect to the wage level.⁶² Our results imply wage elasticities of the VOT of 1.7 and 1.3 with the OLS and GLS specifications, respectively.

While these results are particularly striking for the first experiment, results from the second experiment suggest that these comparisons might underestimate the estimated connection because there is convexity in the VOT. By combining the multiple levels of ETA increases in the second experiment with the price variation of the first experiment, we can quantify how the VOT varies by magnitude of the increase in the waiting time. Specifically, for the five regions which overlap in the two experiments, we use the first experiment to estimate the average price and price (semi-) elasticity of demand, and the second experiment to estimate the average ETAs and ETA (semi-) elasticities of demand between each of the four treatment levels. We then combine these results to estimate the VOT between each of the four experimental ETA levels of the second experiment. To ensure that our price and time elasticity estimates are comparable yet independent, we subset both experiments to observations from the five overlapping regions, remove airport sessions from the first data, and remove from the first experiment observations from any users who appear in the second experiment.

Table C.50 provides the results of this analysis. Here, low, medium, and high refer to estimates

⁶²The GLS specification uses as the weighting matrix the inverse of the estimated covariance matrix of the vector of log VOT estimates across regions, which is derived from the estimated covariance matrix of the vector of price and time semi-elasticities across regions using the delta method.

derived from contrasting Control and Plus 60 treatments, Control and Plus 150 treatments, and Control and Plus 240 treatments, respectively. We note that for the full sample and each individual region, the estimated VOT is highest when using the largest ETA treatment level as an instrument, supporting the conclusion that the value of a marginal minute is increasing as more minutes are added to the base wait time. Figure 8 provides a visual for these results, and plots them against the mean after-tax wages of the regions. The figure shows that the shape of the VOT elasticity over the changes in wait time impact whether the VOT is less than or equal to the after-tax wage rate in each city.

6.2 The impact of our VOT on public policies

As aforementioned, the VOT is an important input into many policy decisions. For example, transportation projects often have benefits that largely derive from the time savings associated with new or improved infrastructure, and similarly, the design of optimal road pricing depends crucially on a measure of travelers' VOTs. Policymakers often use cost-benefit analyses (CBAs) as a part of the information considered in prioritizing projects.⁶³ Many such analyses have embedded assumptions underlying the benefit (or the cost) estimates that rely on the types of proxy methods we reviewed earlier

Most cost-benefit analyses use the recommendations put forth by the Department of Transportation (DoT) in 1997, which was later revised in 2014. Both DoT reports tether the VOT to the median (or mean) wage in an area. While the recommendations allude to the complex set of factors that likely affect values of travel time across various contexts (e.g., comfort of transportation, reliability, magnitude of time change, etc.), there did not exist sufficient reliable data to incorporate these considerations more concretely into the estimated VOT. The reports recommend the VOT be set according to the trip purpose and typical wage rate: the median (or mean) wage should be used for business travel (i.e., on the clock) and half the median (or mean) wage should be used for personal travel.

Some of our analysis is consistent with the rule-of-thumb approach in the DoT reports: we find that the mean values of time are indeed correlated with the mean/median wage rates of various geographical regions, with an elasticity near 1. We also find that values of time are closer to the full mean/median wage rates during typical work commuting hours than at other times, consistent with the idea that travel related to work should be valued at a rate closer to the full hourly wage rate.

Our results, however, have the distinct advantage of being estimated with data that are exceptionally rich in source time variation, granularity, and breadth. As such, we are able to more directly address heterogeneity in VOT that has been posited but not empirically verified. In par-

⁶³Executive Orders beginning with EO12291 first issued in 1981 require cost-benefit analyses for major new rules. In addition agencies now routinely use CBA as part of the resource allocation and decision-making process.

ticular, we find that the VOT is generally higher than 50% of the median/mean wage across most cities, and there is strong evidence of substantial heterogeneity by context (e.g. time of day, day of week, within-city downtown vs non-downtown areas). Our specific recommendations to policy-makers are thus two-fold:

1. When feasible, account for the great deal of VOT heterogeneity with respect to cities, days of week, and times of day.
2. If data are unavailable or unreliable, adjust the mean VOT estimates up to 75% of the after-tax median/mean wage rate.

To demonstrate the potential relevance of our suggestions, we revisited a set of cost-benefit analyses from 1998–2017 that used a VOT for cities that were included in our experiment. We calculated how total costs, total benefits, and cost-benefit ratios change as we use (1) our region-specific empirical point estimates and (2) 75% of median/mean wage adjustment for value of time. We report the differences in Table C.51.

In each case, we find that the estimates used in the reports are lower than our region-specific estimates, so the total benefit (and thus benefit-cost ratio) tends to increase, by values ranging from 2% to 52%. While the qualitative result does not change for the cost-benefit analyses we examined (benefit-cost ratios are above 1 for all cases, both pre- and post-adjustments), it is still apparent from this exercise that a more directly measured, empirically-based VOT can substantially affect the relative benefits and costs of important projects compared to DOT’s rule-of-thumb estimate. Indeed, our estimates are large enough to suggest that there was an important lot of projects that were not undertaken that would have provided substantial net benefits. Such “missed” projects are those tilted towards ones that have significant time savings, suggesting that current mis-measurement of temporal values yields under-investment in time-saving projects.

We also compute benefit-cost ratios following the simpler rule using 75% of mean/median wage (adjusted from the 50% of mean/median wage recommended by DOT). Among this set of cost-benefit analyses, this adjustment tends to yield values of time and measures of benefit that are usually close to but sometimes higher than the region-specific empirical point estimates. We view our estimates—based on exogenous price and time variation and rich, granular observations—have the potential to more efficiently allocate of resources and encourage technologies that induce more time savings for individuals.

7. Conclusion

Having gotten this far in our study you have surely invested a fair amount of time. We hope that such time was indeed an investment, and not ill-spent. This is because time is the ultimate scarce resource, and its value has deep implications for a range of economic phenomena and investment

decisions. Our starting point is a literature from the 1960s that had deep implications for our understanding of the family, the household, and time allocation more generally. We leverage insights from these classic time allocation theories to provide a theoretically-consistent but updated approach to estimate the VOT. The theory carefully directs two large-scale natural field experiments on the Lyft platform to estimate the causal effects of wait time and price on ride-share demand.

We report several interesting insights. First, we estimate a VOT that is roughly \$19 per hour (2015 prices). This estimate is 75-80% of the mean wage rate for the various regions in our experiment, which is quantitatively different from the findings of previous empirical studies on the VOT (Small et al., 2007) and is greater than the existing US policy guidelines on the VOT (USDOT, 2015). Second, we document that, consistent with standard microeconomic models (Becker, 1965; DeSerpa, 1971), the VOT is related to the opportunity cost of time, the available substitute set, and other key features of the trip that impact marginal benefits and marginal costs. Third, taken in aggregate, our research has key implications for policy. Specifically, we recommend that policymakers: (i) account for the great deal of VOT heterogeneity with respect to cities, locations within cities, day of week, and time of day; and (ii) adjust the rule-of-thumb VOT estimates up to 75% of the after-tax mean wage rate otherwise.

We view our VOT estimates as not only adding unique measures to a rich literature, but also providing a key link to the classic time allocation literature. Important areas where this research agenda goes from here can be found in caveats to our research. First, we do not examine the value of reliability in passenger travel or through using ridesharing companies, such as Lyft. We acknowledge that this value could be important and separate from the VOT, but is beyond the scope of this paper. Future studies should consider the value of reliability—the extent to which waiting times vary about their mean—simultaneously with the VOT. On a rideshare platform, one possible approach would be to run a multi-modal waiting time and price experiment, and to model the passenger’s choice of not only whether to take a ride, but also of which ride mode (e.g., Classic or Shared) to take. Shared ride modes typically have more variable waiting times (today quoted as a range on the Lyft app), so such a cross-mode comparison might shed light on the value of reliability. Second, we focus on passenger travel in our data and ignore the VOT for rail, air, and freight travel. Combining the various travel modes and exploring their interplay is an ambitious research agenda that promises to lend deep positive and normative insights.

References

Aguiar, M., M. Bils, K. K. Charles, and E. Hurst (2017). Leisure luxuries and the labor supply of young men. Technical report, National Bureau of Economic Research.

- Aguiar, M. and E. Hurst (2007a). Life-cycle prices and production. *American Economic Review* 97(5), 1533–1559.
- Aguiar, M. and E. Hurst (2007b). Measuring trends in leisure: The allocation of time over five decades. *Quarterly Journal of Economics* 122(3), 969–1006.
- Aguiar, M. and E. Hurst (2016). The macroeconomics of time allocation. In *Handbook of Macroeconomics*, Volume 2, pp. 203–253. Elsevier.
- Aguiar, M., E. Hurst, and L. Karabarbounis (2013). Time use during the great recession. *American Economic Review* 103(5), 1664–96.
- Allcott, H., L. Braghieri, S. Eichmeyer, and M. Gentzkow (2019). The welfare effects of social media. Technical report, National Bureau of Economic Research.
- Anderson, M. L. (2014). Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion. *American Economic Review* 104(9), 2763–96.
- Angrist, J. D., K. Graddy, and G. W. Imbens (2000). The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish. *The Review of Economic Studies* 67(3), 499–527.
- Arnott, R., A. De Palma, and R. Lindsey (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *American Economic Review*, 161–179.
- Basso, L. J. and H. E. Silva (2014). Efficiency and substitutability of transit subsidies and other urban transport policies. *American Economic Journal: Economic Policy* 6(4), 1–33.
- Bastiat, F. (1848). *Propriété et loi, Justice et fraternité*. Paris.
- Baum, C. F., M. E. Schaffer, and S. Stillman (2003). Instrumental variables and GMM: Estimation and testing. *The Stata Journal* 3(1), 1–31.
- Becker, G. S. (1965). A theory of the allocation of time. *Economic Journal*, 493–517.
- Benhabib, J., R. Rogerson, and R. Wright (1991). Homework in macroeconomics: Household production and aggregate fluctuations. *Journal of Political Economy* 99(6), 1166–1187.
- Bento, A., K. Roth, and A. Waxman (2017). Avoiding traffic congestion externalities? the value of urgency. Technical report, Technical Report, Working Paper.
- Besley, T., J. Hall, and I. Preston (1999). The demand for private health insurance: do waiting lists matter? *Journal of Public Economics* 72(2), 155–181.
- Browning, M., T. Crossley, and J. K. Winter (2014, April). The measurement of household consumption expenditures. IFS Working Papers W14/07, Institute for Fiscal Studies.
- Buchholz, N., L. Doval, J. Kastl, F. Matějka, and T. Salz (2020). The value of time: Evidence from auctioned cab rides. Technical report, National Bureau of Economic Research.
- Cameron, C. A. and D. L. Miller (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources* 50(2), 317–372.
- Castillo, J. C. (2019). Who benefits from surge pricing? Available at SSRN 3245533.
- Cesario, F. J. (1976). Value of time in recreation benefit studies. *Land Economics* 52(1), 32–41.

- Chen, Y., G. Y. Jeon, and Y.-M. Kim (2014). A day without a search engine: an experimental study of online and offline searches. *Experimental Economics* 17(4), 512–536.
- Chen, Y. and A. Whalley (2012). Green infrastructure: The effects of urban rail transit on air quality. *American Economic Journal: Economic Policy* 4(1), 58–97.
- Cohen, P., R. Hahn, J. Hall, S. Levitt, and R. Metcalfe (2016). Using big data to estimate consumer surplus: The case of Uber. Technical report, National Bureau of Economic Research.
- Couture, V., G. Duranton, and M. A. Turner (2018). Speed. *Review of Economics and Statistics* 100(4), 725–739.
- Crépon, B., E. Duflo, M. Gurgand, R. Rathelot, and P. Zamora (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *Quarterly Journal of Economics* 128(2), 531–580.
- Davidson, R., J. G. MacKinnon, et al. (2004). *Econometric theory and methods*, Volume 5. Oxford University Press New York.
- Deacon, R. T. and J. Sonstelie (1985). Rationing by waiting and the value of time: Results from a natural experiment. *Journal of Political Economy* 93(4), 627–647.
- DellaVigna, S., J. A. List, and U. Malmendier (2012). Testing for Altruism and Social Pressure in Charitable Giving. *Quarterly Journal of Economics* 127(1), 1–56.
- DeSerpa, A. C. (1971). A theory of the economics of time. *Economic Journal* 81(324), 828–846.
- Duranton, G. and M. A. Turner (2011). The fundamental law of road congestion: Evidence from u.s. cities. *American Economic Review* 101(6), 2616–52.
- Efron, B. and R. J. Tibshirani (1994). *An introduction to the bootstrap*. CRC press.
- Finkelstein, A. (2009). E-z tax: Tax salience and tax rates. *Quarterly Journal of Economics* 124(3), 969–1010.
- Gelber, A. M. and J. W. Mitchell (2012). Taxes and time allocation: Evidence from single women and men. *Review of Economic Studies* 79(3), 863–897.
- Ghez, G., G. S. Becker, et al. (1975). The allocation of time and goods over the life cycle. *NBER Books*.
- Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review* 104(4), 1091–1119.
- Goolsbee, A. and P. J. Klenow (2006). Valuing consumer products by the time spent using them: An application to the internet. *American Economic Review* 96(2), 108–113.
- Green, D. I. (1894). Pain-cost and opportunity-cost. *Quarterly Journal of Economics* 8(2), 218–229.
- Greenwood, J., N. Guner, and G. Vandenbroucke (2017). Family economics writ large. *Journal of Economic Literature* 55(4), 1346–1434.
- Gronau, R. (1973). The intrafamily allocation of time: The value of the housewives' time. *American Economic Review* 63(4), 634–651.

- Hall, J. D. (2020). Can tolling help everyone? estimating the aggregate and distributional consequences of congestion pricing. *Journal of the European Economic Association*.
- Harrison, G. W. and J. A. List (2004). Field experiments. *Journal of Economic Literature* 42(4), 1009–1055.
- Johnson, M. B. (1966). Travel time and the price of leisure. *Economic Inquiry* 4(2), 135.
- Juster, F. T. and F. P. Stafford (1991). The allocation of time: Empirical findings, behavioral models, and problems of measurement. *Journal of Economic Literature* 29(2), 471–522.
- Kleibergen, F. and R. Paap (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics* 133(1), 97 – 126.
- Kreindler, G. E. and Y. Miyauchi (2019). Measuring commuting and economic activity inside cities with cell phone records.
- Krueger, A. B., D. Kahneman, D. Schkade, N. Schwarz, and A. A. Stone (2009). National time accounting: The currency of life. In *Measuring the subjective well-being of nations: National accounts of time use and well-being*, pp. 9–86. University of Chicago Press.
- List, J. A. (2020). Non est disputandum de generalizability? a glimpse into the external validity trial. Technical report, Working paper.
- Mäler, K.-G. (1971). *A Method of Estimating Social Benefits from Pollution Control*, pp. 106–118. London: Palgrave Macmillan UK.
- Mäler, K.-G. (1974). *Environmental Economics: A Theoretical Inquiry*. Johns Hopkins Press.
- Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review* 107(12), 3722–59.
- Mas, A. and A. Pallais (2019). Labor supply and the value of non-work time: Experimental estimates from the field. *American Economic Review: Insights* 1(1), 111–26.
- McFadden, D. (1974). The measurement of urban travel demand. *Journal of Public Economics* 3(4), 303–328.
- Mill, J. S. (1848). *Principles of Political Economy*. John W. Parker.
- Miller, G. and B. P. Urdinola (2010). Cyclicalities, mortality, and the value of time: The case of coffee price fluctuations and child survival in colombia. *Journal of Political Economy* 118(1), 113–155.
- Nevo, A. and A. Wong (2015). The elasticity of substitution between time and market goods: Evidence from the great recession. Technical report, National Bureau of Economic Research.
- Nordhaus, W. (2009). Measuring real income with leisure and household production. In *Measuring the subjective well-being of nations: National accounts of time use and well-being*, pp. 125–144. University of Chicago Press.
- Parry, I. W. and K. A. Small (2009). Should urban transit subsidies be reduced? *American Economic Review* 99(3), 700–724.

- Phaneuf, D. J. and V. K. Smith (2005). Recreation demand models. *Handbook of Environmental Economics* 2, 671–761.
- Philipson, T. J., G. Becker, D. Goldman, and K. M. Murphy (2010). Terminal care and the value of life near its end. Technical report, National Bureau of Economic Research.
- Ramey, V. A. and N. Francis (2009). A century of work and leisure. *American Economic Journal: Macroeconomics* 1(2), 189–224.
- Robinson, J. P. and G. Godbey (1999). Time for life: The surprising ways americans use their time. 2nd.
- Schrank, D., B. Eisele, and T. Lomax (2012). Tti 2012 urban mobility report. *Texas A&M Transportation Institute. The Texas A&M University System* 4.
- Small, K. (2013). *Urban Transportation Economics*. Taylor & Francis.
- Small, K. A., E. T. Verhoef, and R. Lindsey (2007). *The economics of urban transportation*. Routledge.
- Small, K. A., C. Winston, and J. Yan (2005). Uncovering the distribution of motorists' preferences for travel time and reliability. *Econometrica* 73(4), 1367–1382.
- Smith, V. and S. Banzhaf (2007). Quality adjusted price indexes and the willig condition. *Economics Letters* 94(1), 43–48.
- Smith, V. K. and C. Mansfield (1998). Buying time: Real and hypothetical offers. *Journal of Environmental Economics and Management* 36(3), 209–224.
- Su, Y. (2018). The rising value of time and the origin of urban gentrification.
- Sunstein, C. R. (2018). Sludge and ordeals. *Duke Law Journal* 68, 1843.
- USDOT (2015). Revised departmental guidance on valuation of travel time in economic analysis. *US Department of Transportation. Washington, DC*.
- Van Ommeren, J. and M. Fosgerau (2009). Workers' marginal costs of commuting. *Journal of Urban Economics* 65(1), 38–47.
- von Wieser, F. (1876). Über das verhältnis der kosten zum wert (on the relation of cost to value). Technical report, *Gesammelte Abhandlungen*, pp. 377-404.
- Wheaton, W. C. (1977). Income and urban residence: An analysis of consumer demand for location. *American Economic Review* 67(4), 620–631.

Figures

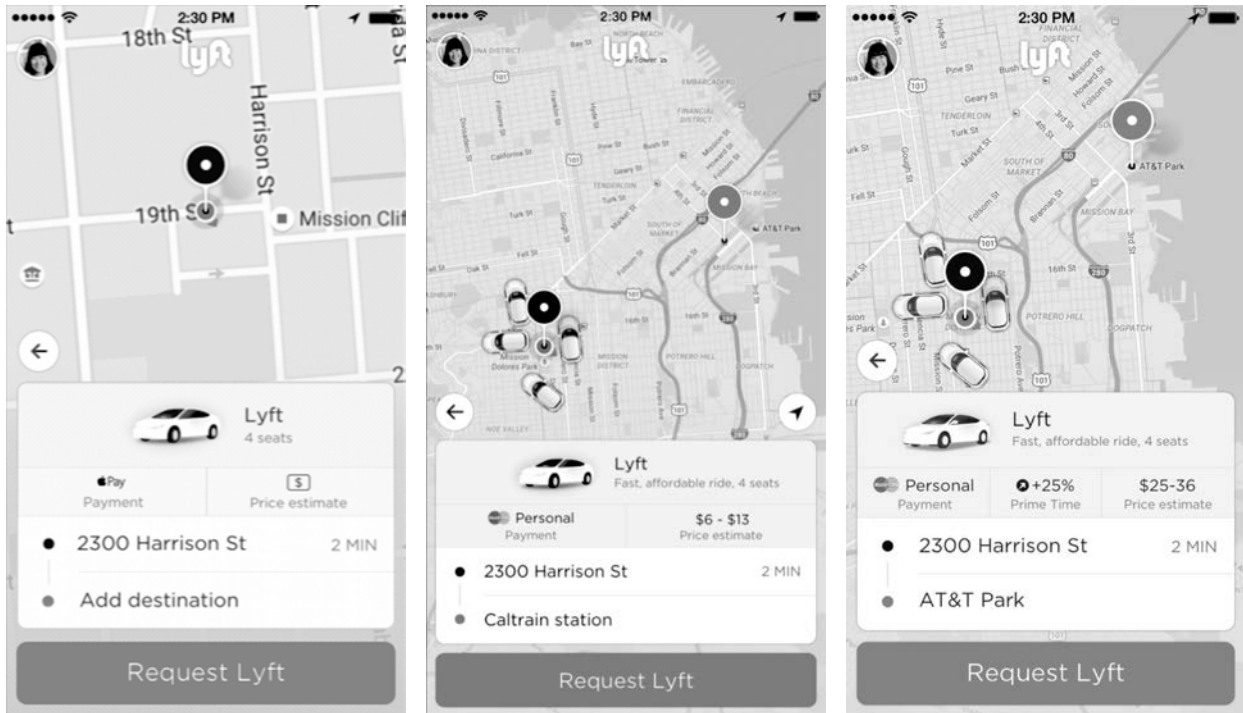


Figure 1
Screenshots of the passenger Lyft app's user interface, at the time of the experiment. Left: no destination entered and no Prime Time active. Middle: destination entered and no Prime Time active. Right: destination entered and +25% Prime Time active.




Figure 2
Distribution of ETAs across sessions in each exp group.

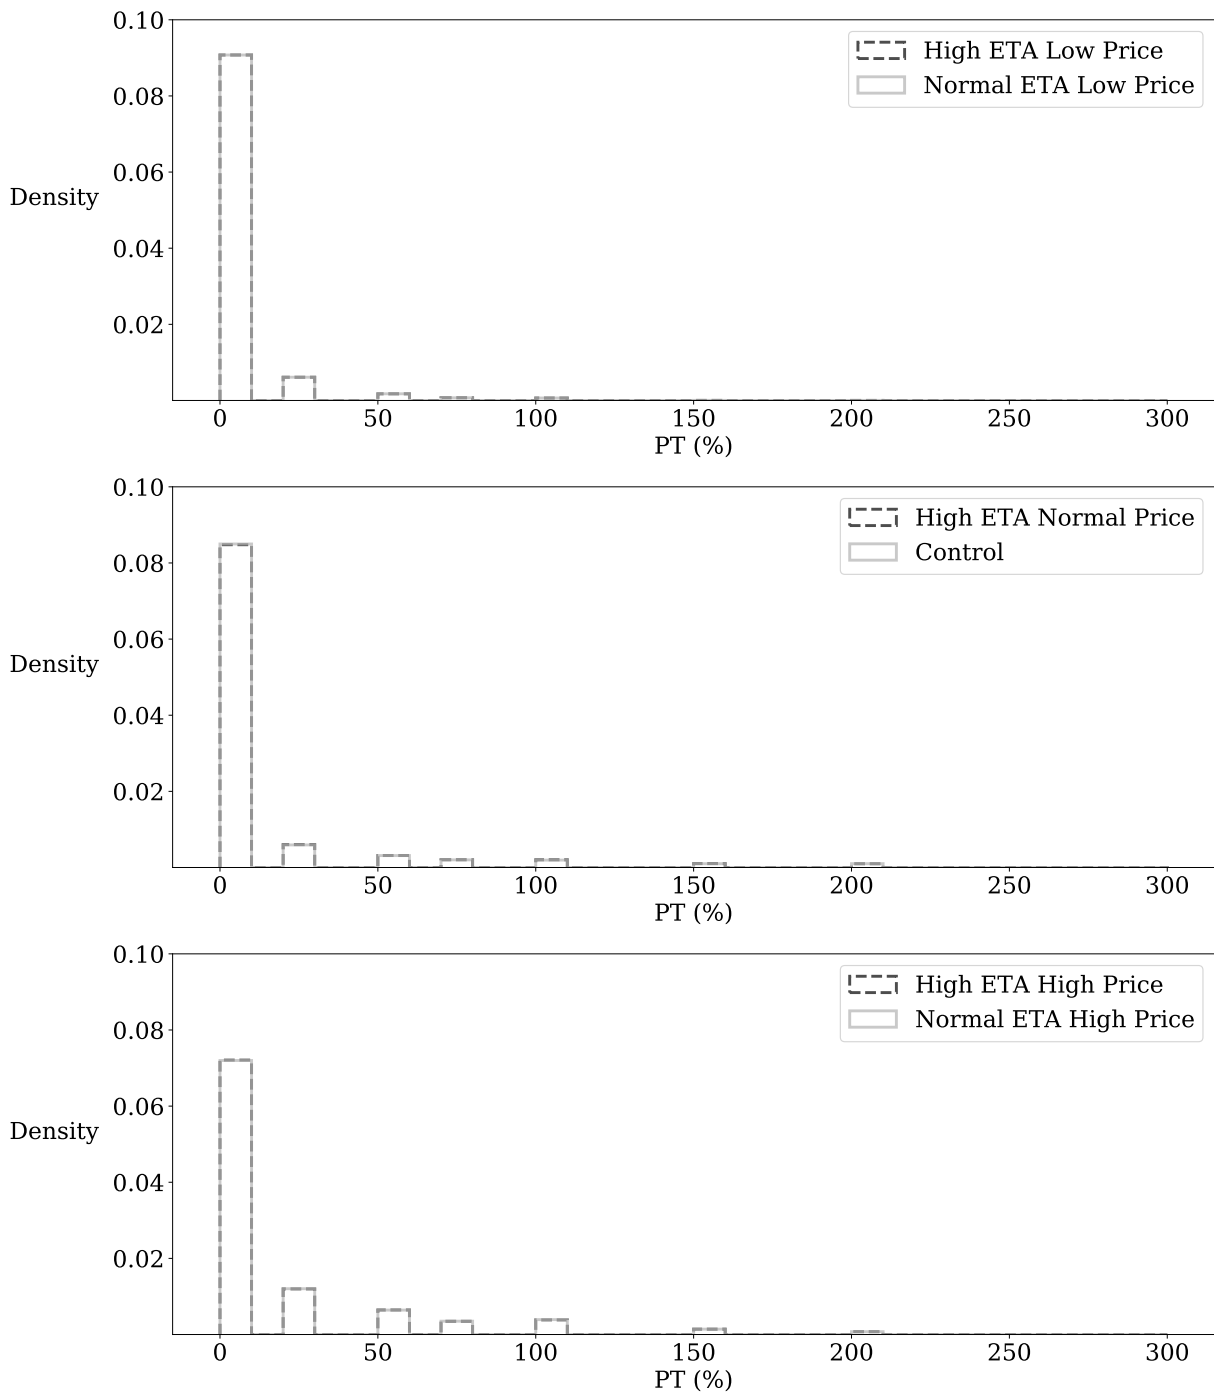


Figure 3
Distribution of PT across sessions in each exp group.




Figure 4
Average request rates across the six experimental groups. Error bars are 95% confidence intervals based on user-level clustered standard errors.




Figure 5
Distribution of ETAs across treatments in experiment two.




Figure 6
Average demand across the four treatments in experiment two.





Figure 7
Estimated VOTs vs. mean and median wages for each region.




Figure 8
Values of time estimated by three different treatment level contrasts.

Table 1
Distribution of users across the six experimental groups.

	Low Price	Normal Price	High Price
Normal ETA	5%	15% (control)	5%
High ETA	5%	5%	2%

Notes: The remaining 63% of users were excluded from the experiment.

Table 2
Average session ETA and Prime Time by treatment group.

	Avg. ETA (minutes)	Avg. PT (%)	Avg. Completed Ride Price (\$)	Number of sessions
Control	3.082 (0.004)	10.014 (0.032)	13.835 (0.022)	2,123,671
Normal ETA High Price	3.103 (0.008)	16.362 (0.072)	14.204 (0.041)	689,325
Normal ETA Low Price	3.085 (0.008)	3.646 (0.024)	13.544 (0.037)	722,367
High ETA Low Price	4.640 (0.011)	3.647 (0.025)	13.529 (0.038)	693,141
High ETA Normal Price	4.649 (0.011)	10.136 (0.055)	13.788 (0.039)	683,159
High ETA High Price	4.661 (0.017)	16.271 (0.112)	14.062 (0.064)	265,695

Notes: Standard errors in parentheses, clustered at the user level.

Table 3
First-stage regression results for full data sample.

	ln(ETA)	ln(1 + PT)	ln(ETA)	ln(1 + PT)
High ETA High Price	0.474*** (0.004)	0.049*** (0.001)	0.480*** (0.001)	0.049*** (0.001)
High ETA Normal Price	0.473*** (0.003)	0.001** (0.000)	0.479*** (0.001)	0.001** (0.000)
High ETA Low Price	0.472*** (0.003)	−0.041*** (0.000)	0.479*** (0.001)	−0.042*** (0.000)
Normal ETA High Price	0.004 (0.003)	0.050*** (0.001)	0.001 (0.001)	0.050*** (0.000)
Normal ETA Low Price	0.001 (0.003)	−0.041*** (0.000)	0.001 (0.001)	−0.041*** (0.000)
Controls			x	x
<i>N</i>	5177358	5177358	5177358	5177358
<i>R</i> ²	0.101	0.026	0.534	0.212
<i>F</i>	13646***	12784***	116249***	20455***
Kleibergen and Paap (2006) rk-statistic		40175***		48791***

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Clustered standard errors in parentheses. Independent variables are indicators for each treatment group. Controls include dummy variables for week of year, hour of week, user geohash5, business, airport, and decile of pre-experiment lifetime rides. *F* statistics are scaled cluster-robust Wald statistics of the null that the coefficients on the treatment indicators are identically 0. **Kleibergen and Paap (2006)** statistics test underidentification of both endogenous regressors, and follow a $\chi^2(df = (5 - 1)(2 - 1))$ distribution under the null.

Table 4
Second-stage regression results for full data sample.

	(1)	(2)
$\ln(\text{ETA})$	-0.0231^{***} (0.0023)	-0.0264^{***} (0.0018)
$\ln(1 + \text{PT})$	-0.3601^{***} (0.0165)	-0.3665^{***} (0.0132)
ETA Elasticity	-0.0374^{***} (0.0025)	-0.0427^{***} (0.0029)
PT Elasticity	-0.5838^{***} (0.0213)	-0.5942^{***} (0.0213)
VOT	17.27^{***} (1.77)	19.38^{***} (1.39)
Control Avg. ETA	3.08	3.08
Control Avg. Price	13.83	13.83
Control Request Rate	0.620	0.620
Controls		x
N	5177358	5177358
R^2	-0.008	0.072

Notes: $***p < 0.01$, $**p < 0.05$, $*p < 0.1$. Clustered standard errors in parentheses. $\ln(\text{ETA})$ and $\ln(1 + \text{PT})$ instrumented by experimental group indicators. Value of time expressed in 2015 US dollars per hour. Controls include dummy variables for session geo-hash5, local hour of week and week of year, decile of user lifetime rides, whether the session is at and airport, and whether the user is a business user.

Table 5
Summary of Elasticity and VOT Results: Features Affecting Opportunity Cost

Sub-sample	ETA Elasticity	PT Elasticity	VOT
Full Sample	−0.043*** (0.003)	−0.594*** (0.021)	19.38*** (1.39)
Day of Week			
Monday	−0.048*** (0.005)	−0.627*** (0.046)	20.08*** (2.47)
Tuesday	−0.046*** (0.005)	−0.676*** (0.049)	17.95*** (2.22)
Wednesday	−0.046*** (0.005)	−0.653*** (0.048)	19.34*** (2.38)
Thursday	−0.041*** (0.005)	−0.585*** (0.035)	19.44*** (2.41)
Friday	−0.039*** (0.005)	−0.473*** (0.023)	21.26*** (2.45)
Saturday	−0.039*** (0.004)	−0.596*** (0.031)	18.74*** (2.04)
Sunday	−0.040*** (0.005)	−0.640*** (0.030)	17.12*** (2.01)
Time of Day (Weekends)			
6–10 AM	−0.056*** (0.011)	−0.687*** (0.081)	17.40*** (3.80)
10 AM–4 PM	−0.053*** (0.007)	−0.727*** (0.047)	18.92*** (2.49)
4 PM–7 PM	−0.048*** (0.007)	−0.616*** (0.039)	20.16*** (2.92)
7 PM–11 PM	−0.033*** (0.006)	−0.669*** (0.062)	15.04*** (2.87)
11 PM–6 AM	−0.023*** (0.005)	−0.500*** (0.028)	14.71*** (3.12)
Time of Day (Weekdays)			
6–10 AM	−0.068*** (0.006)	−0.594*** (0.034)	26.71*** (2.50)
10 AM–4 PM	−0.047*** (0.006)	−0.704*** (0.061)	18.76*** (2.64)
4 PM–7 PM	−0.046*** (0.005)	−0.522*** (0.034)	22.12*** (2.70)
7 PM–11 PM	−0.024*** (0.005)	−0.540*** (0.041)	13.62*** (2.59)
11 PM–6 AM	−0.040*** (0.006)	−0.539*** (0.031)	19.03*** (2.66)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Clustered standard errors in parentheses. Controls include local week of year, local hour of week, user geohash5, business user, airport, and decile of pre-experiment lifetime rides.

Table 6
Summary of Elasticity and VOT Results: Other Trip Characteristics

Sub-sample	ETA Elasticity	PT Elasticity	VOT
Full Sample	−0.043*** (0.003)	−0.594*** (0.021)	19.38*** (1.39)
Precipitation			
None	−0.042*** (0.003)	−0.597*** (0.024)	19.04*** (1.50)
Rain	−0.051*** (0.005)	−0.560*** (0.026)	23.10*** (2.26)
Snow	−0.103*** (0.029)	−0.711*** (0.119)	26.56*** (7.83)
User Type			
Business	−0.033** (0.014)	−0.399*** (0.073)	24.42** (10.55)
Non-Business	−0.043*** (0.003)	−0.605*** (0.022)	19.22*** (1.39)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Clustered standard errors in parentheses. Controls include local week of year, local hour of week, user geohash5, business user, and decile of pre-experiment lifetime rides.

Table 7
Summary of Elasticity and VOT Results: Market Factors

Sub-sample	ETA Elasticity	PT Elasticity	VOT
Full Sample	−0.043*** (0.003)	−0.594*** (0.021)	19.38*** (1.39)
Region			
Austin	−0.028** (0.012)	−0.709*** (0.161)	11.37** (4.76)
San Francisco	−0.026*** (0.005)	−0.413*** (0.026)	21.08*** (3.79)
Los Angeles	−0.040*** (0.006)	−0.522*** (0.071)	19.52*** (3.91)
Miami	−0.052*** (0.010)	−0.595*** (0.103)	16.36*** (3.95)
Boston	−0.074*** (0.011)	−0.670*** (0.072)	23.37*** (3.79)
Seattle	−0.053*** (0.011)	−0.516*** (0.077)	27.08*** (6.38)
San Diego	−0.044*** (0.011)	−0.723*** (0.152)	14.80*** (4.34)
Atlanta	−0.056*** (0.017)	−0.825*** (0.256)	11.24** (4.40)
New York City	−0.056*** (0.009)	−0.956*** (0.046)	23.85*** (3.81)
Location Type			
Airport	−0.024 (0.015)	−0.553*** (0.153)	27.55 (18.02)
Non-Airport	−0.043*** (0.003)	−0.595*** (0.021)	18.84*** (1.35)
Downtown	−0.034*** (0.003)	−0.524*** (0.019)	21.64*** (1.99)
Non-Downtown	−0.060*** (0.005)	−0.762*** (0.058)	17.88*** (1.97)
Distance to Nearest Public Transit			
Under 50 Meters	−0.039*** (0.004)	−0.495*** (0.026)	23.81*** (2.56)
50 to 200 Meters	−0.040*** (0.003)	−0.528*** (0.024)	20.81*** (1.91)
200 to 800 Meters	−0.051*** (0.006)	−0.802*** (0.049)	16.90*** (2.03)
Over 800 Meters	−0.058*** (0.014)	−1.162*** (0.161)	10.88*** (2.82)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Clustered standard errors in parentheses. Controls include local week of year, local hour of week, user geohash5, business user, and decile of pre-experiment lifetime rides.

Table 8
Summary of Elasticity Results: Heterogeneity (2017 Experiment)

Time of Week	(1)	Trip Characteristics	(2)	Location	(3)
Full Sample	−0.043*** (0.001)	Full Sample	−0.043*** (0.001)	Full Sample	−0.043*** (0.001)
Day of Week		Precipitation		Region	
Monday	−0.044*** (0.002)	None	−0.042*** (0.001)	San Francisco	−0.031*** (0.001)
Tuesday	−0.047*** (0.002)	Rain	−0.047*** (0.002)	New York City	−0.086*** (0.002)
Wednesday	−0.044*** (0.002)	User Type	Business	Chicago	−0.039*** (0.002)
Thursday	−0.047*** (0.002)			Washington, D.C.	−0.054*** (0.002)
Friday	−0.042*** (0.002)			Miami	−0.021*** (0.002)
Saturday	−0.039*** (0.002)			New Jersey	−0.053*** (0.004)
Sunday	−0.040*** (0.002)			Boston	−0.055*** (0.003)
Time of Day (Weekends)		Non-Business	−0.043*** (0.001)	Philadelphia	−0.040*** (0.003)
6–10 AM	−0.060*** (0.004)			Atlanta	−0.037*** (0.004)
10 AM–4 PM	−0.048*** (0.003)			Los Angeles	−0.033*** (0.002)
4 PM–7 PM	−0.044*** (0.003)			Location Type	
7 PM–11 PM	−0.033*** (0.003)			Downtown	−0.039*** (0.001)
11 PM–6 AM	−0.026*** (0.002)			Non-Downtown	−0.048*** (0.001)
Time of Day (Weekdays)				Distance to Public Transit	
6–10 AM	−0.061*** (0.002)			Under 50 Meters	−0.042*** (0.001)
10 AM–4 PM	−0.041*** (0.002)			50 to 200 Meters	−0.038*** (0.001)
4 PM–7 PM	−0.052*** (0.002)			200 to 800 Meters	−0.049*** (0.002)
7 PM–11 PM	−0.035*** (0.002)			Over 800 Meters	−0.084*** (0.004)
11 PM–6 AM	−0.037*** (0.002)				

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Clustered standard errors in parentheses. Controls include local week of year, local hour of week, user geohash5, business user, and decile of pre-experiment lifetime rides.

Table 9

Time elasticities estimated by three different levels of treatment contrasts, across different base ETA levels (minutes) (2017 experiment).

	1	2	3	4	5	6	7	8	9	10
Control — Plus 60	-0.005*** (0.001)	-0.014*** (0.002)	-0.026*** (0.002)	-0.058*** (0.003)	-0.087*** (0.003)	-0.124*** (0.004)	-0.142*** (0.005)	-0.188*** (0.005)	-0.208*** (0.006)	-0.230*** (0.007)
Control — Plus 150	-0.008*** (0.001)	-0.025*** (0.001)	-0.044*** (0.002)	-0.063*** (0.002)	-0.103*** (0.002)	-0.118*** (0.003)	-0.166*** (0.003)	-0.182*** (0.004)	-0.203*** (0.005)	-0.289*** (0.004)
Control — Plus 240	-0.018*** (0.001)	-0.042*** (0.001)	-0.066*** (0.001)	-0.075*** (0.002)	-0.108*** (0.002)	-0.137*** (0.002)	-0.170*** (0.003)	-0.217*** (0.003)	-0.232*** (0.004)	-0.289*** (0.004)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Clustered standard errors in parentheses. $\ln(\text{ETA})$ instrumented by experimental group indicators. Controls include local week of year, local hour of week, user geohash5, business user, and decile of pre-experiment lifetime rides.

Table 10

Time elasticities estimated by three different treatments.

	(1)
Control — Plus 60	-0.027*** (0.001)
Control — Plus 150	-0.037*** (0.001)
Control — Plus 240	-0.053*** (0.001)

Notes: Clustered standard errors in parentheses. $\ln(\text{ETA})$ instrumented by experimental group indicators. Controls include local week of year, local hour of week, user geohash5, business user, and decile of pre-experiment lifetime rides.

Table 11
Convexity by city.

	San Francisco	New York City	Chicago	D.C.	Miami	New Jersey	Boston	Philadelphia	Atlanta	Los Angeles
Control — Plus 60	-0.015*** (0.002)	-0.055*** (0.002)	-0.017*** (0.003)	-0.026*** (0.003)	-0.015*** (0.004)	-0.041*** (0.004)	-0.034*** (0.003)	-0.033*** (0.004)	-0.023*** (0.006)	-0.026*** (0.002)
Control — Plus 150	-0.025*** (0.002)	-0.081*** (0.002)	-0.035*** (0.002)	-0.048*** (0.002)	-0.019*** (0.003)	-0.043*** (0.004)	-0.044*** (0.002)	-0.031*** (0.004)	-0.031*** (0.005)	-0.026*** (0.002)
Control — Plus 240	-0.039*** (0.001)	-0.101*** (0.001)	-0.048*** (0.002)	-0.071*** (0.002)	-0.025*** (0.002)	-0.061*** (0.003)	-0.067*** (0.002)	-0.050*** (0.003)	-0.047*** (0.004)	-0.040*** (0.002)

Notes: Clustered standard errors in parentheses. $\ln(\text{ETA})$ instrumented by experimental group indicators. Controls include local week of year, local hour of week, user geohash5, business user, and decile of pre-experiment lifetime rides.

Table 12

Number of regions for which each hypothesis test rejects the null that the VOT equals the given rule-of-thumb approximation at the 5% level, using post-tax wages and assuming a constant 25% tax rate.

1/3- Rule			1/2- Rule			1/2+ Rule		
	Median	Mean		Median	Mean		Median	Mean
Log	7	7	Log	7	6	Log	0	0
Linear	7	7	Linear	6	6	Linear	0	0
			100%+ Rule					
	Median	Mean		Median	Mean		Median	Mean
Log			Log	0	0			
Linear			Linear	0	4			

Table 13

Number of regions for which each hypothesis test rejects the null that the VOT equals the given rule-of-thumb approximation at the 5% level, using pre-tax wages.

1/3- Rule			1/2- Rule			1/2+ Rule		
	Median	Mean		Median	Mean		Median	Mean
Log	7	6	Log	6	3	Log	0	0
Linear	7	6	Linear	6	2	Linear	0	0
			100%+ Rule					
	Median	Mean		Median	Mean		Median	Mean
Log			Log	0	7			
Linear			Linear	4	8			

Table 14

Regressions of log estimated VOT on log mean wage at the region level.

	OLS	GLS
$\ln(\text{MeanWage})$	1.73** (0.82)	1.28** (0.54)
N	9	9
R^2	0.586	0.534
Adj. R^2	0.527	0.467

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Heteroskedasticity-robust (HC3) standard errors in parentheses. Dependent variable is $\ln(\widehat{VOT})$.