

NBER WORKING PAPER SERIES

THE VALUE OF FLEXIBLE WORK:
EVIDENCE FROM UBER DRIVERS

M. Keith Chen
Judith A. Chevalier
Peter E. Rossi
Emily Oehlsen

Working Paper 23296
<http://www.nber.org/papers/w23296>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2017

The authors would like to thank Jonathan Hall, Elisa Long, and A. Ronald Gallant for valuable input. Peter Cohen and Dan Yavorsky provided invaluable technical and research assistance throughout this project. Through Emily Oehlsen, Uber Technologies provided the data used in this paper. Under the agreement granted the authors, Uber has the right to review the paper “solely to confirm that confidential information is being represented in a non-misleading fashion” but not to dispute or influence the findings or conclusions of the paper. Chevalier and Rossi have no material financial relationships with entities related to this research. Oehlsen is an employee of Uber Technologies. Chen is a former employee of Uber, and as a result, continues to hold stock options that may constitute a material financial position. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23296.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by M. Keith Chen, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen. 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 Flexible Work: Evidence from Uber Drivers
M. Keith Chen, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen
NBER Working Paper No. 23296
March 2017
JEL No. J22,L91

ABSTRACT

Participation in flexible contract work has increased dramatically over the last decade, often in settings where new technologies lower the transaction costs of providing labor flexibly. One prominent example of this is the ride-sharing company Uber, which allows drivers to provide (or not provide) rides anytime they are willing to accept prevailing prices for this service. An Uber-style arrangement offers workers flexibility in both setting a customized work schedule and also adjusting it throughout the day. Using high-frequency data of hourly earnings for Uber drivers, we document the ways in which drivers utilize this real-time flexibility and we estimate the driver surplus generated by this flexibility. We estimate how drivers' reservation wages vary in high frequency from hour to hour, which allows us to study the surplus and supply implications of both flexible and traditional work arrangements. Our results indicate that, while the Uber relationship may have other drawbacks, Uber drivers benefit significantly from real-time flexibility, earning more than twice the surplus they would in less flexible arrangements. If required to supply labor inflexibly at prevailing wages, they would also reduce the hours they supply by more than two-thirds. The implications of our findings for the future of flexible work are discussed.

M. Keith Chen
UCLA/Anderson School
110 Westwood Plaza
Cornell Hall, Suite D515
Los Angeles, CA 90095
keith.chen@anderson.ucla.edu

Peter E. Rossi
UCLA/Anderson School
110 Westwood Plaza B4.12
Los Angeles, CA 90095
peter.rossi@anderson.ucla.edu

Judith A. Chevalier
Yale School of Management
135 Prospect Street
New Haven, CT 06520
and NBER
judith.chevalier@yale.edu

Emily Oehlsen
1455 Market Street
San Francisco, CA 94102
eo@uber.com

1 Introduction

In recent years, a number of firms have launched business models that match demand for services to independent contractors providing those services. These businesses rely on independent contractors working intermittent or nonstandard hours. While these businesses typically do not offer many of the benefits of traditional employment relationships, they do provide an opportunity for service providers to earn compensation on a flexible schedule. Understanding the costs and benefits of such arrangements is of growing importance; recent survey evidence finds that 16% of US workers participate in flexible contract work as their primary job, a 56% increase over the last decade.¹ A much larger fraction, 30%, as estimated by Oyer [2016], participate in independent work as a primary or secondary activity.

The fastest growing part of this contract labor environment are digital platforms that instantaneously match buyers and sellers. In this paper, we use data from nearly two hundred thousand drivers on Uber (a popular ride-sharing platform), to examine the benefits to drivers from labor supply flexibility and the costs (if any) from nonstandard hours. To do this, we develop an approach in which we identify the taste for flexibility as being driven by (and equated with) time variation in a driver’s reservation wage. In our framework, the main benefit from flexible work is the ability to work only in those hours when reservation wages are lower than expected earnings. To our knowledge, while properties of reservation wages have been extensively investigated in labor economics and several aspects of non-standard work hours and flextime have been investigated, our paper is the first to frame the benefits of work flexibility as hour-to-hour variation in the reservation wage.

Uber is a platform on which drivers, once approved, can use their own cars to drive on Uber whenever they choose. There are no minimum hours requirements and only modest constraints on maximum hours. As ride requests arrive, the Uber platform allocates these requests to nearby drivers. When a trip is completed, riders pay a base fee plus a per-mile and minute fee. Setting aside taxes, fees, and promotions, drivers earn a fixed (typically 80) percent of this payment. Base prices are set at the city level and are dynamically updated in each “geofence” neighborhood at five minute intervals by a “surge” multiplier, which proportionately increases both rider fees and driver earnings. An algorithm is used to determine the level of surge pricing and both drivers and riders are informed of the surge multiplier before commencing a trip.² Thus, the driver’s compensation is a result of the driver’s labor supply and location decisions as well as the demand and supply of other riders and drivers.

There are several aspects of the contractual arrangement between Uber and drivers; we focus on the fact that drivers can choose their own hours in real time. In order to investigate the relative value of flexible work arrangements to Uber drivers, we construct and estimate a simple empirical model of individual drivers’ labor supply. In our model, the driver’s expected schedule is determined by the weekly pattern of expected payouts from driving and the weekly pattern of the driver’s reservation wage; deviations from that schedule are caused by either shocks to their reservation wage or shocks to their expected payouts.

¹ See Katz and Krueger [2016]. Katz and Krueger define alternative work arrangements as “temporary help agency workers, on-call workers, contract workers, and independent contractors or freelancers”. They find that the incidence of these types of work arrangements rose from 10.1 percent in February 2005 to 15.8 percent in late 2015.

² At regular intervals the Uber platform changes both how drivers and riders are matched and how both sides of the trip are priced. The description provided here however, accurately describes the Uber platform throughout our data period.

Our identification strategy, loosely speaking, is simple: if we see a driver supplying labor in an hour when the expected wage is \$15/hour and choosing not to supply labor in an hour when the expected wage is \$25/hour, controlling for a variety of other factors, we can infer that the driver's reservation wage is time-varying. Furthermore, in observing this, under various assumptions, we can make inferences about the driver's willingness to pay (if any) to avoid a counterfactual employment relationship that would require the driver to work during the driver's high reservation wage hours. We can also make inferences about driver distaste (if any) for nonstandard hours. Finally, we can analyze patterns of driver behavior that are prevalent in the data and provide preliminary evidence about the types of flexibility sought by these labor suppliers.

Our paper proceeds as follows. Section 2 briefly reviews the labor literature on labor supply, job flexibility and nonstandard hours. Section 3 describes our data sources and construction of the analysis dataset. Section 4 provides a first look at the habits of Uber drivers and suggestive evidence about the taste for flexibility. Section 5 introduces our labor supply model and outlines how we conduct inferences for that model. Expected labor surplus and expected labor supply is discussed in Section 6. In section 7 we examine sensitivity of our model to exogeneity assumptions and prior settings. Sections 8 provides a conclusion and summary of our findings.

2 Literature

For many jobs, work hours are fixed by the employer. This may be due to complementarities among employees, the shape of the hours/productivity function, and/or fixed costs in staffing and monitoring workers. If jobs are at least partially inflexible, this suggests that workers will often find both the total *quantity* of hours worked and the temporal *pattern* of hours worked mismatch their preferences. The hypothesis that the total quantity of hours is determined by the employer rather than as an individual negotiation between the employer and employee is supported by the findings in Altonji and Paxson [1988], Senesky [2005], and Altonji and Usui [2007]. In particular, there is evidence that many US workers would prefer to work fewer hours per week than the schedule set by their employer, if they could do so at their current hourly wages (see Rebitzer and Taylor [1995] and Reynolds [2004]). Consistent with this, The Council of Economic Advisors reports data from the National Study of Employers that suggests that 36% of firms with over 50 employees would allow "some employees" to transition for full-time to part-time work and back again while remaining in the same position or level, but that only 6% would allow it for "most or all" workers (see Council of Economic Advisors (2010)).

The literature on the scheduling of hours is sparser than the literature on the total quantity of hours. Kostiuk [1990] documents that workers receive compensating differentials for evening shift work while Hamermesh [1999] documents a secular decline in evening and weekend work from the early 1970s to the early 1990s. The pattern of observed changes is consistent with a model in which evening and weekend work is a disutility that higher productivity workers are willing to pay to avoid. This conclusion, that the reservation wage is on average higher in the evening and night due to disamenity effects, can be directly tested in our data.

More recently, a small literature has examined flexible workplace practices. For example, the Counsel of Economic Advisors (2010) reports that 81% of surveyed employers would allow some employees to periodically change starting and quitting times within some range of hours and 27% of employers

would do so for “most or all employees”.³ However, only 41% would allow “some employees” to change starting and quitting times on a daily basis and only 10% would do so for “most or all employees”. Thus, employers appear to typically have preferences for the particular hours of the day worked by employees. The reasons for this will vary across industries and jobs, but include at least monitoring costs, complementarities among coworkers, and the need for workers to interface with customers in real time. Interestingly, survey data from Bond and Galinsky [2011] suggests that lower wage employees have less flexibility than higher wage employees. The data on alternative work arrangements documented by Katz and Krueger [2016] similarly demonstrates that the recent growth in freelance and contract work largely excludes the low-wage sector. Indeed, lower wage workers in the retail sector frequently not only cannot choose their hours, but the hours chosen by their employers can frequently change from week to week, exacerbating work-life conflicts (see, for example, Henly and Lambert [2014]).

The paper that perhaps has the most overlap with our own is Mas and Pallais [2016]. Mas and Pallais conduct an experiment with job applicants to a call-center. Mas and Pallais find a low average willingness to pay for flexibility, although there was a substantial right tail of individuals whose willingness to pay was larger. They also find that job applicants have a high disutility for jobs with substantial employer discretion in scheduling; they attribute this largely to a large disutility for evening and weekend hours. They find that the average worker requires 14% more to work evenings and 19% more to work on the weekends.

The nature of traditional employment relationships pose challenges for researchers who might try to investigate worker willingness to pay for more flexibility using methods other than surveys. In particular, in examining labor data, one can infer that a worker’s total compensation exceeds the worker’s reservation value for the total hours worked. However, in conventional job settings, one cannot infer that the average hourly compensation exceeds the hourly reservation wage on an hour by hour basis. There may be hours for which the hourly wage paid the worker is less than the worker’s reservation wage, but the worker nonetheless works because the hour cannot be unbundled from other hours where the participation constraint is slack. Furthermore, there may be hours for which the worker would be willing to work at the worker’s “usual” wage, but for which working the hour is not an available option from the employer.

Because of these challenges, platforms such as Uber represent new opportunities both for individuals supplying labor and for researchers. Most importantly, neither the quantity nor pattern of hours worked are fixed. While contract and freelance work has been more prevalent in the economy, the evidence on flexible work arrangements in the low-wage sector cited above suggests that Uber (and its main competitor Lyft) are among a limited number of opportunities for fully flexible semi-skilled work. Hall and Krueger [2016] examine survey evidence and Uber administrative data. They document that drivers cite flexibility as a reason for working for Uber and that many Uber drivers report that Uber is a part-time activity secondary to more traditional employment. Consistent with this, they document that the hours supplied by drivers vary considerably from week to week. We examine drivers’ labor supply in more detail. Because of the flexibility of the platform, the driver can decide whether to supply labor minute by minute, which in turn allows us to infer time patterns of the driver’s reservation wage. If there are time periods for which there is on average a substantial disamenity value to driving, supply and demand should lead to an equilibrium of higher expected wages during the undesired hours. Both the typical weekly pattern and shocks to the driver’s reservation wage can in principle be

³ data from the National Study of Employers

extracted. We examine the driver’s labor supply decisions and estimate alternative scenarios which mimic the effects of traditional employment relationships.

While our focus is not on labor supply elasticities, this paper is also closely related to the literature that uses high frequency data on labor supply and wages to examine labor supply elasticities. In particular, Camerer et al. [1997] study the shift-ending decisions of New York City cabdrivers and find evidence for a negative labor supply elasticity. In contrast, Oettinger [1999] studies the decisions of individual stadium vendors to work or not work a particular game, and finds evidence of substantially positive labor supply elasticities. Farber [2005] and Farber [2015] reexamine New York City cabdriver data and finds that only a small fraction of drivers exhibit negative supply elasticities. Frechette et al. [2016] also use data from New York City taxis to estimate a dynamic general equilibrium model of a taxi market. Drivers make both a daily entry decision and a stopping decision (in contrast to our setting, where a driver could make multiple starting and stopping decisions). Stopping decisions are determined by comparing hourly earnings with the combination of a marginal cost of driving that is increasing in the length of a shift and a random terminal outside option.

As in our model, Frechette et al. treat individual drivers as competitive and, thus, only keep track of the aggregate state of the market (market-level hourly earnings) when making individual driving decisions. That is, the driver compares the opportunity cost and disutility of driving to the expected income, which is data. An important feature of their model, as opposed to ours, is the incorporation of the constraints imposed by the medallion system, in particular the scarcity of medallions and regulations that make driving multiple short shifts unattractive, and regulations that effectively cap shift length. Because of these features, Frechette et al. model the mean value of the outside option as a fixed value depending only on the eight hour shift and medallion type. The driver-specificity derives from the IID error in the opportunity cost of driving and the IID error in the utility of starting a shift. Thus the mean and variances of the outside option are identified by the decision to leave a medallion unused for an entire shift and the decision to stop driving before the shift ends. Unlike the analysis we conduct for this paper, driver heterogeneity in the value of the outside option is not a primary focus.⁴

A feature of our sample of contractors raises a point that has perhaps been underemphasized in the literature. In all of these papers, calculations of supply elasticities and the value of flexibility are not undertaken on a random sample of workers or potential workers and the estimates may not apply to other samples. For example, our study examines drivers on the Uber platform. While we have over a million drivers in our sample, they are all individuals who selected into providing labor in this flexible work environment. They likely have a higher taste for flexibility, for example, than individuals who answered the employment advertisement used to create the sample in Mas and Pallais [2016]. However, one reason the Uber sample is interesting is that technology platforms such as those used by Uber and its closest competitor, Lyft, create opportunities for relatively low-skilled flexible work on a scale that does not appear to have been previously possible.

3 Data Sources and Construction of Analysis Datasets

Our data are provided by agreement with Uber. We start with the universe of all Uber driver-hours in the United States from September 2015 to April 2016. We focus on the UberX platform, which is

⁴ There are other papers that use the New York City taxi data. However, both Buchholz [2015] and Lagos [2003] take the supply of taxis and drivers as exogenous and thus address economic questions distant from those that we consider.

Uber’s peer-to-peer service in the United States. We limit our study to UberX both because it is the service with the majority of Uber trips, and because other Uber services have characteristics which complicate studying drivers’ hours choices. For example, drivers on the more expensive UberBlack are licensed commercial drivers, some of who work for limo companies and may be paid a fixed salary for working hours not entirely under their control.

Data on Uber drivers is stored in two large tables: 1. a “trips” table which records both logins/logouts and trips made by the drivers, and 2. a “payouts” table which records the earnings and payments made to the drivers. Neither of these tables are in a form amenable to analysis with a labor-supply model. Our first decision was to convert this data to an hour-by-hour record of driver activity and payments. Specifically, our data consist of an anonymized driver identifier and an hour-by-hour record of time spent active on to the system, time driving, city, and payouts. For the purposes of standardizing analyses across cities we convert all data to the driver’s local time. This poses challenges in five Uber cities where the greater metropolitan area spans a time zone border (and therefore where drivers drive back and forth over the timezone border frequently); we omit these cities from our analysis.⁵

One issue in evaluating these data is that there is more than one way to define labor supply. For example, a driver is “working” if the driver has the Uber app on and has indicated that she is willing to accept passengers. However, if the driver has the app on, she may or may not be summoned to accept a passenger and the driver may or may not accept the passenger. A driver is “active” if he or she is en route to a passenger or carrying a passenger. We consider both the “working” state and the “active” states as alternative definitions of labor supply. The “working” state is likely an overestimate of labor supply. In particular, while some Uber drivers who have the app on are quite focused on securing a passenger and will even drive to a location that is likely to be productive, other Uber partners may have the app on while working another job (say, gardening) with the expectation that the location is not optimized to obtain riders. In contrast, the “active” state is probably an underestimate of labor supply because it includes minutes on trip en route to obtain a passenger but does not include minutes in which the driver is waiting to be matched with a passengers.

3.1 Data Construction and Definitions

For our analysis, we divide time into discrete hours as the unit of observation, 168 hours per week. We define a driver to be active in an hour if she is active for at least 10 minutes within that hour, and measure the driver’s discrete choice of being active in each of the 168 hour blocks.

We calculate the “wage” in an hour as a driver’s total earnings in that hour, divided by minutes worked, times sixty. Our use of the broader measure of work in our calculation likely understates the effective wage. However, because differences in overall utilization and time spent waiting to be called is an important difference in the profitability of different hours, we think it is important to use time working rather than time active in calculating wage metrics. However, by screening for a minimum level of 10 minutes active in the hour, we screen out drivers who have the app on but are not accepting trips, or who have the app on in remote locations where they may not be trying to find trips. There are a very small number of large payouts of more than \$250 (less than .001%) and we capped or winsorized these values at \$250.

⁵ Uber markets which cross time zones are Yuma AZ, NW Indiana, Louisville KY, Cincinnati OH, and South Bend IN.

On the Uber platform, drivers are expected to pay for both the capital costs of their vehicle and the cost of gas used while driving. In our analysis, these costs are incorporated into the driver's reservation wages. In part for this reason, our analysis of labor supply and surplus should be thought of as short-run; drivers can be thought of as making a longer-run vehicle choice, then choosing labor supply conditional on that capital stock. Some differences in the equilibrium wages across hours may well be driven by common cost differences in driving those hours.

3.2 Filters Applied to the Data

Our full data set consists of 1,047,176 drivers who are "active" or working in 183,608,194 hours. Because we will be evaluating patterns of activities over time, we create a sample of drivers who are active in at least 1 hour for at least 16 of the 36 weeks we have available in our data. We will refer to drivers who meet this criteria as "active drivers." There are 260,605 drivers who are classified as "active." These active drivers are responsible for 140,282,451 hours, or 76% of total hours. We removed holidays and holiday periods⁶ as these are unusual periods of Uber demand which occur at irregular intervals and we did not wish to expend parameters on accommodating this shift in demand. We also found irregularities with four hours (9 p.m. to midnight on 4/20/16) likely caused by database errors, and removed these from our data. After removing holidays and these four hours, we have 130,557,951 hours remaining, or roughly 70% of our original data.

In order to form estimates of expected wages that we use to estimate labor supply, we computed average wages by city, week, day, and hour. In order to insure enough observations to reliably compute these averages, we restricted attention to the top 20 US cities by volume of labor supplied on the UberX platform. This means that, in estimating our model, our final estimation data set has data on 197,517 drivers who supply 102,280,904 of the total hours observed in the original data pull (or 55%).

Our final analysis data set to estimate our model consists of information by driver, week, day, and hour of labor supplied and expected wage for each of the 197,517 drivers. This data set has 881,826,744 hourly observations.⁷ Expected wage is merged in from our table of expected wages on the basis of the modal city for the driver in that week. If there are periods of inactivity that last more than one week (i.e. a gap of a week or more), we impute an expected wage equal to the average of the wage for the first non-missing modal city before the week in question and the wage for the first non-missing modal city after the week in question. Before turning to the estimation of our model, we present some model-free evidence on labor supply flexibility using the full sample of 260,605 "active" drivers.

4 Model-free Evidence on Labor Supply Flexibility

4.1 Uber Driver Labor Supply Patterns

Our research is motivated by the unusual characteristics of this market, particularly the enormous flexibility allowed to Uber drivers. As discussed above, most workers in the economy must find

⁶ 19 days were excluded - Labor day, Halloween and the day after, Thanksgiving and day after, the week of Christmas, New Year's Eve and New Year's day, MLK day, Presidents Day, Columbus Day, and Veterans Day.

⁷ If we observe at least one active hour in the day, we fill in all of the non-active hours for that driver with a labor supply of 0.

Total hours	Share of driver-weeks
0	19%
1-4	11%
5-12	21%
13-20	17%
21-30	14%
31-40	9%
41+	9%

Tab. 1: Distribution of Active Hours for “Committed” Drivers Sample

t	$t + 1$	0	1-4	5-8	9-12	13-16	17-20	21-30	31-36	37-40	41-45	46-50	>50
0	0	0.49	0.17	0.11	0.07	0.05	0.03	0.04	0.01	0	0	0	0
1-4	1-4	0.28	0.26	0.18	0.11	0.07	0.04	0.04	0.01	0	0	0	0
5-8	5-8	0.18	0.18	0.21	0.16	0.10	0.06	0.07	0.02	0.01	0	0	0
9-12	9-12	0.13	0.12	0.17	0.18	0.14	0.09	0.11	0.03	0.01	0.01	0	0.01
13-16	13-16	0.10	0.09	0.12	0.16	0.16	0.13	0.17	0.04	0.02	0.01	0.01	0.01
17-20	17-20	0.08	0.06	0.09	0.12	0.14	0.14	0.23	0.06	0.02	0.02	0.01	0.01
21-30	21-30	0.06	0.04	0.06	0.07	0.10	0.12	0.30	0.12	0.05	0.04	0.02	0.03
31-36	31-36	0.05	0.03	0.03	0.04	0.05	0.07	0.27	0.18	0.09	0.08	0.05	0.06
37-40	37-40	0.05	0.02	0.03	0.03	0.04	0.05	0.21	0.18	0.12	0.11	0.08	0.09
41-45	41-45	0.05	0.02	0.03	0.02	0.03	0.04	0.16	0.15	0.12	0.14	0.11	0.14
46-50	46-50	0.05	0.02	0.02	0.02	0.02	0.03	0.12	0.12	0.10	0.14	0.13	0.23
>50	>50	0.04	0.01	0.02	0.02	0.01	0.02	0.07	0.06	0.06	0.09	0.11	0.47

Tab. 2: Transition Matrix of Hours Worked in Contiguous Weeks

employment choosing between employers who offer fixed wage-hour bundles. Unsurprisingly, the hours supplied by Uber drivers are not identical to the hours worked by workers in more conventional job settings and they vary from week to week. Table 1 examines the universe of drivers in our sample over the course of our 36 week sample. The distribution of total weekly hour blocks supplied by drivers over the sample is summarized in Table 1.

Recall that we consider a driver active in any hour block for which the driver was active for at least 10 minutes and count how many of the 168 hour blocks in the week the driver was active. Column 1 displays the share of the drivers who were active on the system for various time bins. We use our base sample of drivers who are active at least 16 weeks during our 36 week study, but eliminate drivers before their first week of activity. Our summary results are similar to Hall and Krueger [2016].

Table 2 is the transition matrix of hours worked in contiguous weeks for drivers who meet our active driver criterion. This illustrates how much a driver’s total activity varies from week to week, an issue also studied by Hall and Kruger (2016) in their earlier sample.

The first column shows bins of hours of driver activity in a week, and columns 1 through 6 show the share of the drivers who are active on the system for various time bins during the subsequent week. For example, of the drivers active from 21 to 30 hours in a week, 30% fall into the same time-supplied bin in the subsequent week.

Tables 1 and 2 reveal three interesting patterns. First, the overwhelming majority of Uber drivers are working part-time hours on Uber. Indeed, even among active drivers the majority of drivers are active

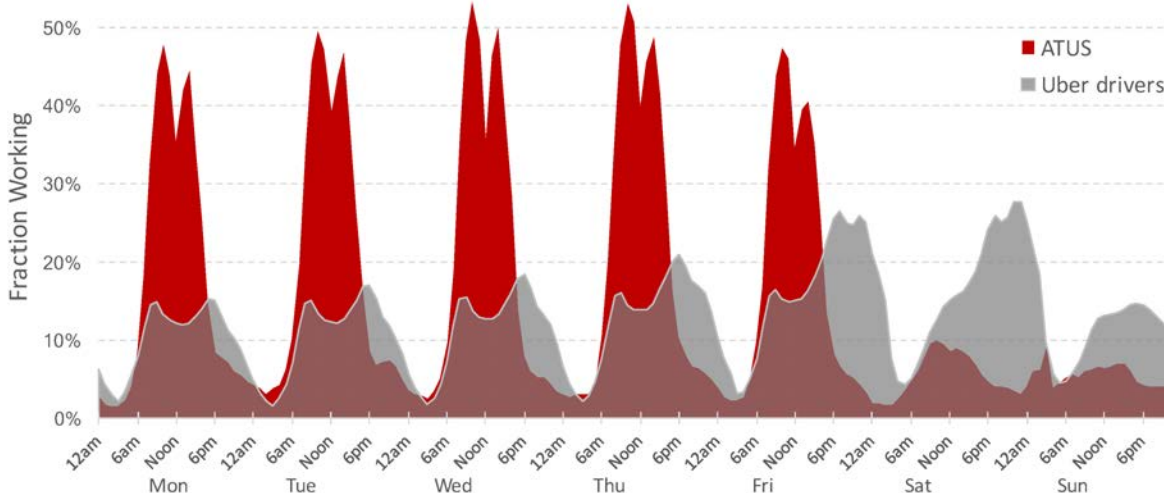


Fig. 1: Comparison of Uber Driver Activity to Workers in the American Time Use Survey

fewer than 12 hours a week. This is unsurprising given the survey evidence in Hall and Kruger (2016) which suggests that driving is complementary to other economic activities such as school attendance, caregiving, or employment. Second, a substantial fraction of drivers active in one week are simply not active the subsequent week; this is particularly true for drivers who had low activity the first week. Finally, there is some tendency for drivers to work a similar number of total hours from week to week, although there is certainly substantial weekly variation.

4.2 Comparisons with Standard Work Schedules

Given that Uber drivers are largely working part-time hours, it is not surprising that the pattern of hours worked does not necessarily resemble the hours worked by workers with conventional jobs. Figure 1 below compares the working habits of Uber drivers to the working habits of employed males over age 20 surveyed in the American Time Use Survey (ATUS) for 2014. On the x-axis, are the 168 hours of the week, beginning with Monday morning. On the y-axis, are the ATUS and Uber data. The ATUS data report the fraction of the surveyed employed individuals who report working in a given hour of the given surveyed day. Thus, while most employed individuals in the ATUS report working at some point in the week, the ATUS data also includes vacations, furloughs, and holidays. For this analysis, the universe of Uber drivers is our standard universe of committed Uber drivers who have logged at least one ten minute session in the week. The graph reports the share of such drivers working in each of the 168 hours of the week (averaged over all of the weeks).

While the overall levels are difficult to compare, it is clear that work in the ATUS largely takes place between 9 a.m. and 5 p.m., whereas Uber drivers are more likely to be working at 6 or 7 p.m. than they are at 2 or 3 p.m.. While male ATUS respondents are about half as likely to be working Saturday afternoon as in the afternoon on a weekday, Uber drivers are more likely to be working Saturday afternoon and evening than a weekday afternoon or evening. Of course, this pattern of driving is the outcome of both supply and demand factors, and so we will incorporate payout information in our

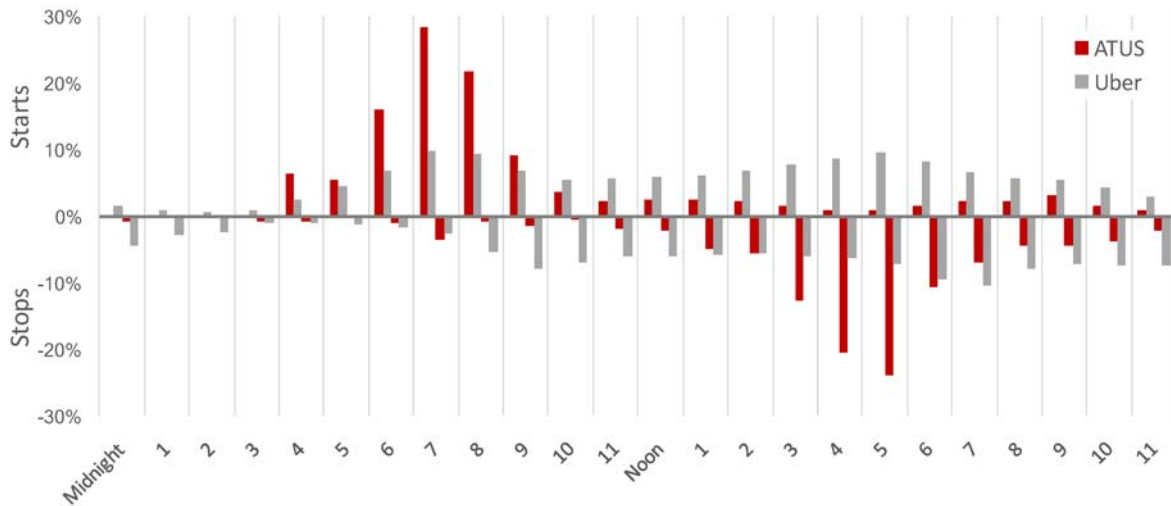


Fig. 2: Comparison of Wednesday Start and Stop Times on Uber vs ATUS

formal analysis in order to isolate labor supply factors.

We obtain further circumstantial evidence of the complementarity of Uber driving with more traditional work by examining work start and stop times in the ATUS versus Uber. Figure 2 uses data for employed individuals over 20 in the ATUS, and shows the fraction who start working and stop working at a particular hour, averaged over all Wednesdays in 2014.⁸ This is graphed against the hours of the day that Uber drivers begin and ends driving sessions averaged across all Wednesdays in our data sample.⁹ The figure expands on the Figure 1 and suggests that many Uber drivers begin driving during conventional work hours, but many also begin when conventional work hours end.

As we define starting and stopping, the majority of workers in the ATUS have only one start in a day; we find 114 starts in a day per 100 workers. However, Uber drivers are much more likely to drive in multiple sessions. We find 131 starts per 100 drivers who driver within a day. In part this stems from the short sessions driven by many drivers; indeed, 20% of our starts are also stops, meaning that the driver starts and ends a driving session within a given clock hour. Of course, these behaviors are a function of both supply of drivers and demand for rides, so our more formal analysis will attempt to separate out supply and demand contributions for Uber driving.

4.3 Types of Driving Schedules

Additional evidence of the value of flexible scheduling among Uber drivers comes from variation among drivers in the types of schedules they choose. The data suggest that there are several different types of schedules drivers naturally employ, both between Uber drivers and within a driver across weeks. To document this, Figure 3 shows the results of a k-means cluster analysis of the hours worked by

⁸ Stopping working is defined as ceasing reporting working in the ATUS without resuming working fewer than 2 hours later. Starting working is defined as working in an hour in when the worker did not work in either of the prior two hours.

⁹ A driving session begins if the driver wasn't driving in the prior two hours but begins driving in an hour, and ends if a driver drives in an hour, but not in the subsequent two.

Type of Driver	% of Active Drivers	$t + 1$					
		t	Eve	Morn	Late	Wknd	Infr
Evening Driver	16.1	Eve	38.3	11.0	9.3	21.0	20.4
Morning Driver	5.9	Morn	12.1	13.4	8.3	18.1	48.1
Late-Night Driver	6.1	Late	17.2	16.8	19.4	17.3	29.2
Weekend & Evening Driver	19.3	Wknd	8.7	5.4	3.2	70.6	12.1
Infrequent Driver	52.6	Infr	7.1	11.1	4.9	9.8	67.0

Tab. 3: Transition Matrix Across Types of Weekly Schedules

Uber drivers. Specifically, for each driver in our sample, we construct a 168-degree vector, with each component representing what fraction of those hour of the week the driver was active across the entire data sample. We then perform a k-means clustering analysis across the set of all drivers, increasing the number of clusters until the addition of an additional cluster produces two clusters that differ only in their intensity of driving, rather than a difference in shift pattern. This leads to a set of 5 clusters, which are displayed in Figure 3 below. For each driver cluster, the y-axis displays the mean fraction of drivers working that hour of the week. Also noted is the hour block for each day in which the greatest number of that cluster’s drivers are driving.

These patterns are instructive of the numerous ways that Uber drivers choose to drive. “Morning” and “Evening” drivers drive slightly less than half the days of the week, but drive sizable shifts when they do drive, and focus on driving during the morning or the evening rush hours, respectively. “Late-night” drivers drive shorter shifts peaking around 10 p.m. most nights, shifting that schedule back to midnight on Thursday, Friday, and Saturday nights. “Weekend and evening” drivers drive shorter shifts than these other types, and tend to drive at times that can fit around a full-time job; shorter shifts around 7 p.m. on weeknights, and longer, late evenings shifts on weekends. Finally, “Infrequent” drivers tend to drive opportunistic schedules that are relatively flat across waking hours.

4.4 Within-Driver Variation in Schedules

In Figure 3 we collapsed each driver’s complete driving history in our data to a single average week, in order to identify major driver “types”. In contrast, for Table 3, we take each driver-week and maps it to the closest of these “types”. Table 3 displays both how many drivers are in each type, and the week-to-week transition matrix of the closest “type” for each driver-week.

Table 3 shows that while there are a variety of common schedules, Uber drivers regularly switch which type of schedule they work from week to week. Indeed, for both “Morning” and “Late-Night” drivers, their current type of schedule is not the most common type the following week. Even when considering only active Uber drivers, over half of driver-weeks are of the “Infrequent” type, where a driver puts in a few hours relatively evenly distributed across waking hours. In fact, for every type of driver-week, an “Infrequent” next week is either the most or second-most common pattern. All together, across all types of driver-weeks, the next driver-week switches type 43% of the time.

We can also express the extent to which drivers vary the particular hours worked from week to week as an average probability of repeatedly working an hour block across weeks. In our model section below, we will be contemplating the idea that drivers potentially face a hierarchy of shocks to the reservation wage. Later in our formal model, we model a week shock (a shock to reservation wages that impact

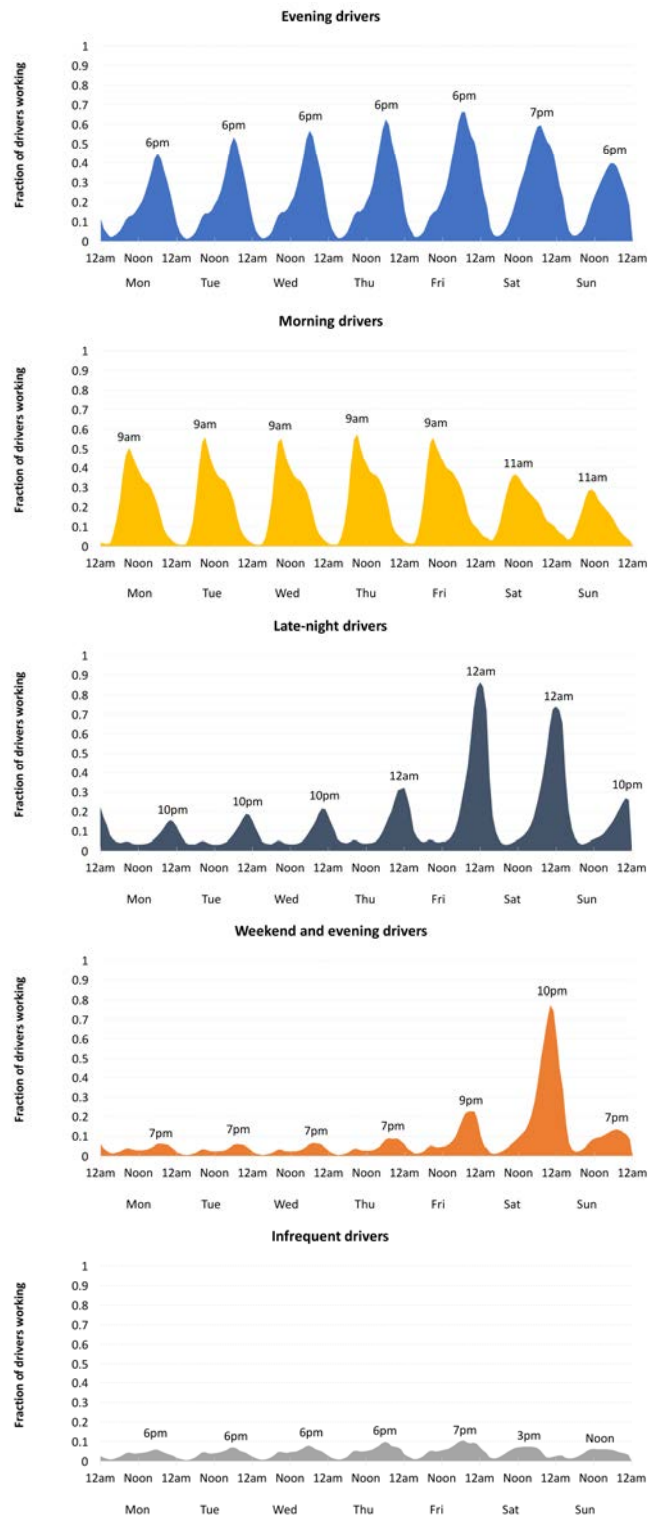


Fig. 3: Driver Types from K-Means Analysis

		Percent who worked that block in week $t+1$		
Did a driver work a block in week t ?	Yes	47.3	51.3	66.0
	No	9.3	11.9	22.0
Conditional on Working:		unconditional	in week t	that day

Tab. 4: Conditional Probabilities of Working a Block in Consecutive Weeks

the whole week), a day shock that impacts a whole day, and an hour shock that idiosyncratically impacts a single hour. Here, we give an intuition of driver flexibility over time.

To do this, we divide the 168 hours of the week into 56 three-hour blocks ordered sequentially from the beginning of the week. We then ask: if a driver drives in a block in week t , what is the probability that the driver drives in that same block in week $t + 1$? Then, to provide insight into the ways that a driver can alter her schedule, we ask the same question, but condition on the driver working at some point in week $t + 1$. The idea is to separate out how much of the not working from one week to the next is due to sitting out the entire week. Next, we again trace working in the same block across weeks, but condition on driving sometime in the day. The results are shown in Table 4.

Table 4 shows that a driver who works in a particular block has a roughly 47% chance of working in that same block on the following week. This probability increases very little when excluding drivers who take the entire next week off. However, conditional on working sometime that day in the next week, the probability that a driver works in the same three-hour block that he or she worked in the prior week rises to about two-thirds. This suggests that the particular hours driven by a given driver vary considerably, even conditioning on the driver working sometime in the day.

To examine this further, we divided the 168 hour week into coarser blocks, for a total of nine blocks of hours in the week. The blocks of hours are natural groupings like weekday mornings, weekday evenings, etc. and are the same ones that we will use below in the model section. We code a driver as driving in a block if he is working at any time during the block for the week. We then measure, between all potential pairs of weeks, the number of common hour blocks driven in the two weeks divided by the average number of hour blocks driven in the two weeks. Thus, for example, if a driver drives the exact same blocks in two weeks, this measure is 1. If the driver drives in four distinct blocks in week 1 and drives in one of those blocks again in week 2 plus one of the additional blocks, the measure would be $1/3$. This measure is graphed on the x axis of Figure 4 with the density measure on the y axis. Figure 4 separately graphs pairs of adjacent weeks and all weeks.

Clearly, the driver habits are more similar for two adjacent week than weeks farther apart in time. However, for either measure, it is clear that the fraction is one only a very small fraction of the time and also infrequently zero. Despite the relatively coarse hour blocks, there is substantially more mass to the left of 0.75 than to to the right of it, even for adjacent weeks. Drivers do frequently alter the particular pattern of hours worked from week to week.

As mentioned above, the pattern of driver hours is driven by both labor supply factors and labor demand factors. Of course, we will evaluate this more formally below. However, we provide a summary graph which is suggestive of the importance of time variation in the driver reservation wage in leading to driving outcomes. Figure 5 graphs the share of drivers working in each of the 168 hours of the week (on the left y-axis, which is the same as in the figure above) against a measure of the payout of driving in the hour (on the right y-axis). The measure we use for realized payouts is the total payout per minute worked for drivers working in the hour block demeaned by the overall city mean (across

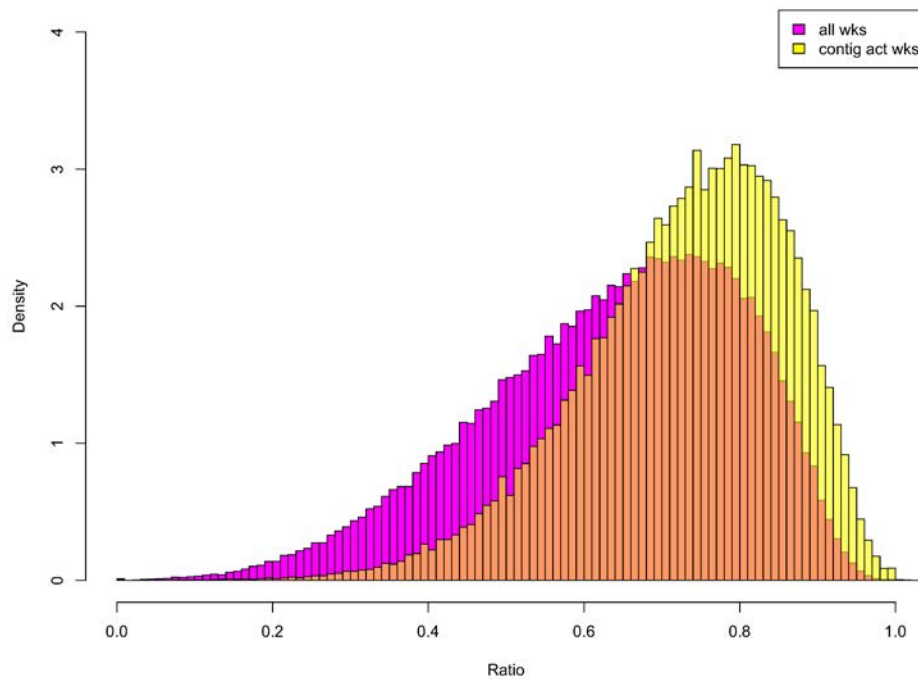


Fig. 4: Percent by Which Driver's Schedules Overlap Week to Week

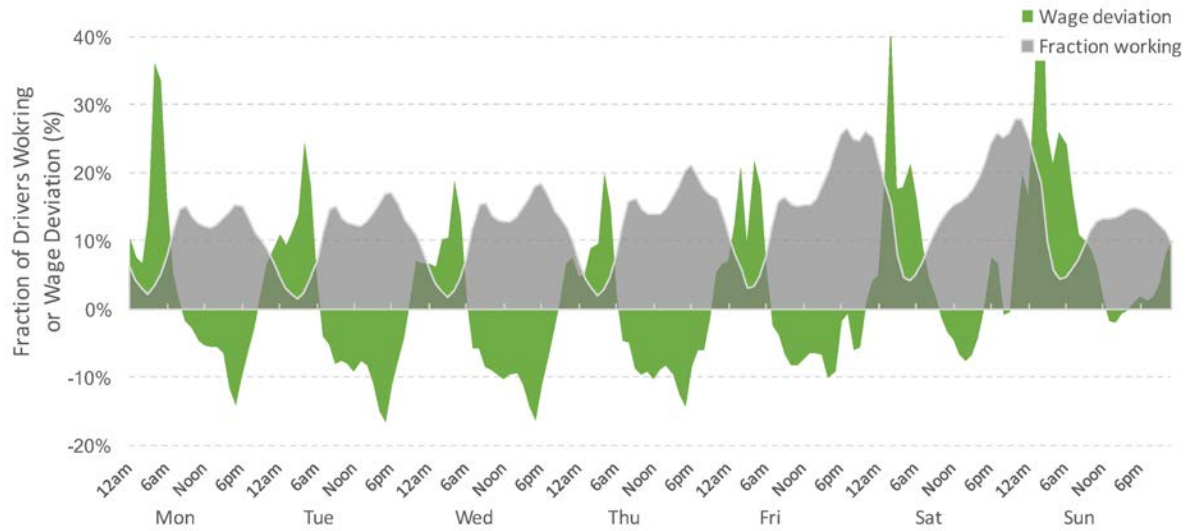


Fig. 5: Correlation Between When Drivers Work and When Earnings are High

all drivers and time periods), divided by the city mean. Thus, an hour with a value of 0 is an hour where a driver working would expect to earn the weekly mean payout of his or her city.

The results are quite striking. Absent time variation in the reservation wage, one would expect either a zero or positive correlation between share working and the payout per minute worked. As demand and payouts increased, drivers would be expected to work more (perhaps to the point where the payout per minute equilibrates across hours). However, in contrast, there appears to be a negative correlation between payout per minute and the share of drivers driving. This seems suggestive of the idea that payouts are high, in part, in periods where drivers have higher reservation wages and choose not to drive.

Our data also allow us to examine a longstanding issue in labor economics, the extent to which there is a disamenity premium for shift work. Prior literature has documented a wage premium for shift work. In the closest obtainable to a causal estimate, Kostiuk [1990] estimates a 5 percent premium for working the overnight hours. Mas and Pallais (2016) estimate that applicants to a call center position demand premia of 14 to 19 percent for evenings and weekends. It is not clear whether such a premium would emerge as an equilibrium outcome on Uber. Shift work in more conventional job settings (as in these studies) implies a commitment to working or being available in nontraditional hours. On Uber, such hours are not obviously a disamenity given that the driver pool has selected into this arrangement and these hours are ones where other employment obligations are at their minimum. Nonetheless, the figure above suggests that an hour of labor supplied in late evening/early morning hour, especially on the weekends is more remunerative than an hour of labor supplied during the day, consistent with the findings of Kostiuk et. al and the findings of Mas and Pallais (2016).

Of course, the equilibrium premium for working overnight at Uber may or may not be driven by the same factors that lead to higher compensation for other overnight jobs. Indeed, the peak payouts are for the 2 a.m. hour on Saturday a.m. and Sunday a.m.; this corresponds to statutory or practical bar closing hours in many of our cities. There may be disamenities of interacting with these passengers that deters many drivers from working then. Smaller but still substantial spikes over the mean wage

occur each morning from about 4 a.m. to 6 a.m., especially on Monday morning. Looking across cities, the pattern appears consistent with a large spike in airport trips in this time period (especially on Monday). There, a disamenity of the passengers specifically rather than the time seem less likely. Premia over the mean wage for consistently driving a four hour shift including the airport peak averages around 13 percent.

5 A Model of Labor Supply and Inference Procedures

A simple model of labor supply is that drivers will supply labor if their reservation wage is less than the prevailing expected wage. That is, for a given period of time (which we take as one hour), we observe the labor supply decision, Y_{it} , as well as the expected prevailing wage, w_{it} , where $Y_{it} = 1$ if driver i is observed to work in hour t and 0 if not. We define “working” as having their driver side app on and ready to receive ride requests from Uber customers. Expected wages are computed assuming that drivers are rational and have access to the distribution of wages in a particular city and time. We estimate expected wages by computing the average wage over all Uber drivers in that city and time (see 3 above for details).

It should be noted that our measure of prevailing wages is not net of the variable costs of operating a vehicle. Therefore, our reservation wages should be interpreted as a gross quantity as well. Note that, if a given driver has a car that is cheaper or more expensive to operate than the mean driver, this difference in expenses would be reflected in the driver’s mean reservation wage. Of course, the labor supply decision is based on the difference between prevailing and reservation wages which does not depend on assumptions regarding the incorporation of operating costs.

5.1 A Model of Labor Supply

The specification of the reservation wage process is crucial to determining to what extent drivers can exercise flexibility in labor supply. As we have documented in section 4, Uber drivers have both predictable and unpredictable patterns of labor supply. There is predictability by day of week and time of day; for example, we have shown that some drivers prefer to work on weekends or in the evenings and clearly use Uber to supplement other jobs or responsibilities that occupy the standard 9-5 weekday period. Equally important, drivers change their work schedules from week to week, from day to day, and even from hour to hour. That is, there appears to be a good deal of evidence that, ex post, drivers behave as though they respond to events that are not known prior to a particular period of time. For these reasons, we postulate a model of reservation wages with both a predictable mean component as well as a random component that is unobserved by the econometrician but revealed to the drivers at particular points of time such as the beginning of the week, day, or hour.

$$w_{it}^* = \mu_i(t) + \varepsilon_{it} \quad (1)$$

Here w_{it}^* is the reservation wage of driver i in time t , $\mu_i(t)$ is the mean reservation wage at time t , and ε_{it} is a random shock to the reservation wage that will be resolved, for Uber drivers, before time t . That is, we assume that at the beginning of each time period (hour) each Uber driver has realized the shock and therefore simply compares their reservation wage to the expected wage to make a labor

supply decision. Note that the expected wage in a given period can incorporate common knowledge by drivers about predictable events (such as concerts, conventions, and sporting events) which create peaks in demand for Uber services.

Mean Function

The mean portion of the reservation wage process drives the predictable portion of labor supply. For example, if a driver has a regular weekday job, the model can accommodate this with high reservation wages during the 9-5 hours of each weekday. Since these patterns of labor supply vary widely across drivers, we must provide mean function parameters that vary at the driver level. Even though we have a relatively large number of driver-hour observations, the censoring mechanism applied to the reservation process means that the information content of even thousands of observations is limited. We use a parsimonious specification by 1) grouping hours into blocks associated with a common shift in the mean reservation wage and 2) we assume driver preferences are stable and do not allow for trends or other time shifts. This implies that our mean function is only a function of the day and hour corresponding to time interval t , $\mu_i(t) = \mu_i(d, h)$.

Our mean specification allows for 9 parameters corresponding to the following blocks of hours.

1. MF_am: Monday-Friday, 7 a.m. - 12 noon
2. MF_afternoon: Monday-Friday, 1 - 4 p.m.
3. MF_rush_hour: Monday-Friday, 5 - 8 p.m.
4. MTh_evening: Monday-Thursday, 9 p.m. - 12 a.m.
5. MTh_late_night: Monday-Thursday, 12 - 3 a.m.
6. FS_evening: Friday-Saturday, 9 p.m. - 12 a.m.
7. FS_late_night: Friday-Saturday, 12 - 3 a.m.
8. MSu_don¹⁰: Monday-Sunday, 4 a.m. - 6 a.m.
9. Base: all remaining hours in the week¹¹

Error Components

We have observed that labor supply behavior of Uber drivers has an unpredictable component at the weekly, daily and hourly frequencies. To accommodate these patterns of behavior, we employ a three-part variance components model for the shock to reservation wages.

$$\varepsilon_{it} = v_w + v_d + v_h \quad (2)$$

In this model, each of the error components is *iid* Normal over their respective frequencies with standard deviations, $\sigma_w, \sigma_d, \sigma_h$ respectively. "w" denotes weekly, "d" denotes daily, and "h" denotes

¹⁰ Dead-of-night

¹¹ Note that each hour block extends from the first minute of the first hour in the block to the last minute of the second hour in the block specification; for example, the MF_am block extends from 7:00 a.m. until 12:59 p.m.

hourly. Since each day within each week share the common week shock as does each hour within each day, this creates the well known variance components covariance structure which can exhibit very high correlation between periods within each broader timeframe. For example, hours within the same day have a correlation of $\rho_d = \frac{\sigma_w^2 + \sigma_d^2}{\sigma_w^2 + \sigma_d^2 + \sigma_h^2}$. These correlations are driven by the relative magnitudes of the error components. The error covariance matrix of the reservation wage shock in (2) is block diagonal across weeks with hours within a week having a co-variance structure given by

$$\Omega_w = (I_{nd} \otimes \Sigma_{day}) + \sigma_w^2 \mathbf{1}\mathbf{1}^t \quad (3)$$

$$\Sigma_{day} = \sigma_h^2 I_{24} + \sigma_d^2 \mathbf{1}\mathbf{1}^t \quad (4)$$

nd is the number of days in a week and allows for weeks with less than seven complete days. While Ω_w is high dimensional, the patterns of covariance are generated by only three variance component parameters. Variance component models have been criticized on the grounds of inflexibility (all covariances are positive and the same within error component block (e.g. within a day)). In our case, the variance components are interpreted as shocks to reservation wages that come at various frequencies in the lives of Uber drivers. Each component has meaning due to its association with a dimension of labor supply predictability.

Flexibility is conceptualized as the ability to respond to different kinds of shocks. Clearly, benefits of flexibility will be related to the relative magnitudes of these shocks. While Uber drivers can respond to each kind of shock, this is not true for many other labor supply arrangements. For example, a “standard” 9-5 factory shift job does not offer flexibility from hour to hour or from day to day. Indeed, if there are pre-commitments to work for more than one week in advance, the factory job does not offer any of the three dimensions of flexibility we have built into the model of reservation wages. On the other hand, some “taxi” style jobs allow for day to day flexibility. If a taxi driver enters into daily rental agreements with the taxi owner, and the driver receives a large positive daily shock to reservation wages, then the driver simply does not work in that day. However, to amortize the fixed fee of renting the taxi, the driver would typically work almost all of the hours of a particular shift. Thus, the taxi driver can respond to weekly and daily shocks but is constrained in his or her ability to respond to hourly shocks. If Uber drivers experience very small hourly shocks but large daily shocks, the Uber system will not afford drivers much value in terms of captured surplus relative to traditional taxi style arrangements.

5.2 Likelihood

Our model is a latent normal¹² and correlated reservation process coupled with a censoring function that indicates whether or not the observed wage rate exceeds or is less than the reservation wage.

¹² Mean reservation wages are apt to be large relative to the variance of the error components. We seldom find any negative reservation wages, even though this is theoretically possible with a normal distribution.

$$y_{it} = \begin{cases} 1 & w_{it}^* < w_{it} \\ 0 & w_{it}^* \geq w_{it} \end{cases} \quad (5)$$

$$w_i^* = \mu_i + \varepsilon_i \quad (6)$$

$$\text{Var}(\varepsilon_i) = I_{N_i} \otimes \Omega_w \quad (7)$$

Here Ω_w is given by (3) and N_i is the number of weeks we observe driver i .¹³ Given a specification of the mean function, the likelihood function for this model can be written down by first observing that, conditional on the observed wage vector, each week is independent of each other week.

$$\ell(\mu, \sigma_w, \sigma_d, \sigma_h | y, w) = \prod_{wk=1}^{N_i} \ell_{wk}(\mu, \sigma_w, \sigma_d, \sigma_h | y_w, w_{wk})$$

The computational problems of evaluating this likelihood are associated with the likelihood for each week. We observe a vector of 168 indicator variables corresponding to each hour of the week. These are censored version of the vector of latent normal variables where the censoring is dependent on the observed wages for each hour. The variance component structure results in a potentially highly correlated latent normal vector. This means that probability of the observed vector of labor supply decisions must be computed as the integral over specific region of a 168-variate normal distribution.

$$\ell_{wk} = \int_{A(y_1, w_1)} \int_{A(y_2, w_2)} \cdots \int_{A(y_{168}, w_{168})} \phi(x | \mu, \Omega_w) dx_1 dx_2 \dots dx_{168} \quad (8)$$

Here we have suppressed the notation for the i th driver and the regions are defined by

$$A(y, w) = \begin{cases} \{w^* : w^* \leq w\} & y = 1 \\ \{w^* : w^* > w\} & y = 0 \end{cases} \quad (9)$$

There are no reliable methods for calculating (with a reasonable degree of accuracy) such high dimensional integrals of a multivariate normal over a cone. Instead, we will employ a data augmentation strategy in an hybrid MCMC method as outlined below.

5.3 Identification

If the censoring point, w_{it} , is fixed at 0, our model is a multivariate Probit model with a variance component structure. As is well known, the multivariate Probit model does not allow for even parametric identification of variances, only correlations. However, the key difference between our setting and the multivariate Probit is that the censoring point is the observed wage and this is varying from observation to observation. Variation in the truncation point allows for parametric identification of all model parameters. The extent to which the model parameters are well-identified in our sample depends on extent of variation in expected wages.

¹³ We impose the constraint that we observe drivers for complete days only. We “fill-out” the days with zeros for those hours we do not observe labor supply. Typically, Uber drivers enter and exit our sample of 38 weeks on days of the week that make for incomplete (less than seven days) weeks. For these weeks, we use the same variance component model and Ω_w is modified to include “blocks” corresponding only to the number of days in the incomplete week.

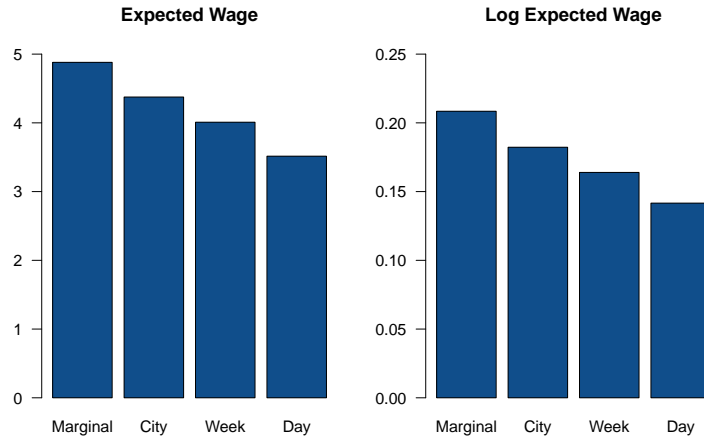


Fig. 6: Variation in Expected Wages

Expected wages for Uber drivers are centered around \$20/hr with a great deal of dispersion. This dispersion is essential to our ability to estimate our model at the driver level. To explore the extent of variation in expected wages, we conduct an analysis of variance by factor, including city, week, and hour. That is, we regress expected wages on the appropriate set of dummies and report the standard deviation of the residuals. Figure 6 shows bar-plots of these residuals variances for both expected wage and log expected wage. We see that even within city, week and hour, there is a great deal of variation with a standard deviation of over \$3 per hour which corresponds to variation in wages of at least 10 percent.

While our model is parametrically well-identified, the question of non-parametric identification remains. We observe about 35 weeks of hourly data on most drivers which means that we have something on the order of 5000 observations per driver. We also have a very large cross-section of drivers (about 197,000). Given that we are making inferences on the driver level, the cross-sectional sample size does not help with identification unless we impose some sort of further structure such as a particular random coefficient distribution. Non-parametric identification can only be achieved as the number of observations per driver tends to infinity. If we fix a particular hour (this holds the mean reservation wage constant) and if we have a continuous distribution of truncation points (observed wages), then, as the number of observations per driver increases, we can estimate the distribution of reservation wages non-parametrically for any location-scale family of distributions.

5.4 MCMC Method

Given our specification of the reservation wage process, our goal is to make driver-level inferences regarding the parameters of the mean reservation wage function and the standard deviation of the three error components. These parameters will allow us to compute the amount of surplus each driver can expect to receive from the Uber labor supply arrangement. To fix notation, we write the

reservation wage process as

$$\begin{aligned} w_i^* &= X_i \beta_i + \varepsilon_i \\ \varepsilon_i &\sim N(0, \Omega(\sigma_i)) \end{aligned}$$

Here $\sigma_i^t = (\sigma_{w,i}, \sigma_{d,i}, \sigma_{h,i})$. The X matrix contains dummies that allow the mean reservation wage to vary for each of our nine defined hour blocks. The variance-covariance matrix for shocks to reservation wages is populated according to our error component model (3) with the standard deviations for each shock. Thus, our goal is to provide a procedure which will deliver inferences for each of 12 parameters for each of our 197,000 drivers. We adopt a Bayesian approach which is the only feasible approach for models which have likelihoods involving high dimensional integrals.

Given that the likelihood function cannot be directly evaluated, we utilize data augmentation. For each driver, we augment the 12 mean and variance parameters with the vector of unobserved reservation wages as has become standard in the treatment of either the Multinomial Probit or Multivariate Probit models.¹⁴ The key insight is that given the reservation wages, inference for the β and σ parameters is straightforward, involving either standard Bayesian treatment of regression for the mean function parameters or a random walk in only three dimensions for the σ parameters.

Our MCMC algorithm is a hybrid ‘‘Gibbs-style’’ method which cycles through three sets of simulations – two of which are standard Gibbs sampler draws comprised of one-for-one draws from conditional posteriors and one of which is a random-walk Metropolis draw. Combining all three sets of draws creates a continuous state-space Markov process which has the posterior distribution for driver i as its equilibrium distribution.

$$w_i^* | X_i, w_i, y_i, \beta_i, \sigma_i \tag{10}$$

$$\beta_i | w_i^*, X_i, \sigma_i \tag{11}$$

$$\sigma_i | w_i^*, X_i, \beta_i \tag{12}$$

In the Bayesian treatment of linear models with error components, it is common to introduce each of the random effects as parameters.¹⁵ Here we have random effects corresponding to each of the weeks and days that we observe driver i . In our setting, it would be appropriate to impose a zero mean restriction on the random effects so that we can interpret them as a shock. The alternative to this approach is to integrate out the random effects, by using a correlated error term. Our experimentation with the approach which augments with the random effects is that the associated Markov chain is very highly autocorrelated. Since we have no direct use for estimates of the driver specific random week and day effects, we prefer the approach that integrates them out. This requires a random walk step instead of a pure Gibbs sampler (conditional on the draw of the reservation wage vector). However, our random walk step is only in three dimensions and performs remarkably well.

Draw of Reservation Wage

(10) could be accomplished by a direct draw from a truncated multivariate normal for each week of the driver data. There is no method which can efficiently draw from a truncated 168-dimensional normal

¹⁴ See Rossi et al. [2005], 4.2-4.3, for example.

¹⁵ See, for example, Gelfand et al. [1995].

distribution. Instead, we employ the Gibbs sampler of McCulloch and Rossi [1994] to “Gibbs-thru” each element of the reservation wage vector using univariate normal draws (see equation 4.2.5 of Rossi et al. [2005] for details). These univariate draws are either truncated from below or above depending on the value of the labor supply indicator variable (truncated from above by the wage if $y = 1$ and from below by the wage if $y = 0$). The mean and variances of the truncated normal distribution can be obtained from $X\beta$ and the inverse of the weekly covariance matrix, Ω_w . The weekly covariance matrix of the error components is easily computed from σ using the eigen-value, eigen-vector decomposition of Ω_w as shown below in equation (22,24,25).

The standard inverse CDF method cannot be used to draw the univariate truncated normal draws because of the possibility of draws in the extreme tails of the normal distribution and possible numerical errors and overflows. Truncated normal draws were made with a three part method that depends on where the truncation point (w_{it}) is relative to the center of the normal distribution. The algorithm breaks the draw problem into three regions (defined here in terms of truncation from above); 1: if the truncation point is more than 4 standard deviations above the mean, then normal based rejection sampling is done, 2: if the truncation point is within $+/4$ standard deviations from the mean, then an inverse CDF method is used, and finally 3: if the truncation point is more than 4 standard deviations below the mean then rejection sampling based on an exponential envelope is used.

Draw of β

Given the draw of the latent reservation wage vector and Ω_w , the draw of the β parameter which provides different means for different hour-blocks can be accomplished by a standard Bayesian analysis of a regression model with a known variance covariance matrix. We assume a standard normal prior for β .

$$\beta \sim N(0, A) \tag{13}$$

$$A = \text{diag}(0.01, .1\iota_8) \tag{14}$$

The notation, $\text{diag}(\text{vec})$, means to form a diagonal matrix with vec as the diagonal. The setting of A is a very diffuse prior on the element of the intercept (overall mean reservation wage) coupled with tighter priors on the mean reservation wages for each hour-block. This is a conservative choice of prior in that it puts low probability on situations in which very large excesses of wages over reservation wage (large surplus) are possible. Given σ , we are able to compute the variance-covariance matrix of the error terms on a week by week basis using (3). This matrix can be used to transform the regression equation into a equation with uncorrelated and unit variance error terms. The posterior for this transformed system is given by

$$\beta|X, w^*, y, \sigma \sim N\left(\tilde{\beta}, \left(\tilde{X}^t \tilde{X} + A\right)^{-1}\right) \tag{15}$$

with $\tilde{\beta} = (\tilde{X}^t \tilde{X} + A)^{-1} (\tilde{X}^t \tilde{w}^*)$, $\tilde{w}^* = Hw^*$, and $\tilde{X} = HX$. H is the root of the inverse of the covariance matrix of the error terms.

$$\Sigma = Var(\varepsilon) = \text{diag}(\Omega_{w_1}, \dots, \Omega_{w_{N_i}}) \quad (16)$$

$$\Sigma^{-1} = HH^t = \text{diag}(H_{w_1}H_{w_1}^t, \dots, H_{w_{N_i}}H_{w_{N_i}}^t) \quad (17)$$

Here Ω_{w_j} is the variance-covariance matrix of week j in driver i 's work history (N_i is the number of weeks we observe driver i). (3) provides the formula for each driver week's covariance matrix. Each of these matrices can easily be computed from knowledge of σ and the number of days in each week.

Draw of σ

The draw of each of the three variance components in (12) is accomplished by a random-walk Metropolis step. Since each of sigmas must be positive, we reparameterize as $\tau^t = \log(\sigma^t)$. We assume that each of the three taus have an independent and normal prior. This means that we are assuming that, a priori, each of the variance components is independent and log-normally distributed.

$$\sigma_j \sim \exp(N(\mu_{lnsigma}, \sigma_{lnsigma}^2)), j = w, d, h \quad (18)$$

$\mu_{lnsigma} = 1$ and $\sigma_{lnsigma}^2 = .5$. This a relatively tight prior that shrinks each of the variance components towards 0. Again, this is a conservative choice which reduces surplus as large values of the reservation wage variance yield greater surplus in a flexible work arrangement.

Given the log-normal prior on σ , we implement a standard random-walk Metropolis method for drawing $\tau = \log(\sigma)$. The likelihood function is simply the product of N_i multivariate normal distributions (one for each week). Each normal distribution has a mean extracted from $X_i\beta$ and variance-covariance matrix determined by the number of days in each week and the current value of σ . Candidate values of τ are drawn as follows:

$$\tau_c = \tau_{old} + s_{R\&R} R_{inc}^t v \quad (19)$$

$$v \sim N(0, I) \quad (20)$$

$s_{R\&R} = \frac{2.93}{\sqrt{3}}$ is the scaling constant of Roberts and Rosenthal [2001]. The RW increments variance covariance matrix is chosen to approximate the diagonal of the inverse of the negative Hessian of the log-likelihood. Experimentation with different values of the σ vector show that the Hessian is remarkably stable on a weekly basis and we take

$$R_{inc} = \frac{1}{\sqrt{N_i}} \text{diag}(1, 0.3, 0.05) \quad (21)$$

Efficient Computation of Covariance Structures In order to perform the draw of w^* (10) and the draw of β (11), we must have an efficient way of computing the inverse of variance-covariance matrix of the error term. As shown in (17), the error covariance matrix depends only on the number of

days in each week of the driver's work history and the three variance components in σ . We use the eigen-vector, eigen-value decomposition of the week covariance matrix to achieve this efficiently.¹⁶

$$\Omega_w = Q\Lambda Q^t \quad (22)$$

Here Q is an $(nd * 24) \times (nd * 24)$ orthogonal matrix of normalized eigenvectors and Λ is a diagonal matrix of of the $nd * 24$ eigenvalues. In order to perform the draws of required, we need the root of the inverse of this matrix.

$$\Omega_w^{-1} = Q\Lambda^{-1}Q^t = HH^t; H = Q\Lambda^{-\frac{1}{2}} \quad (23)$$

The key to efficient computation of H is that the matrix of eigenvectors does not depend on σ , only the eigenvalues in the depend on σ . Both the eigenvectors and eigenvalues depend the number of days in the week. Given that there are only 7 possibilities for the number of days in a week, we precompute the matrices of eigenvectors and store them in a container. As we draw new values of σ , we simply use the old matrices of eigenvectors and update the eigenvalues. The eigenvalues are given by the formulae

$$\begin{aligned} \lambda_1 &= (24 \times nd) \sigma_w^2 + nh\sigma_d^2 + \sigma_h^2 \\ \lambda_2 \dots \lambda_{nd} &= nh\sigma_d^2 + \sigma_h^2 \\ \lambda_{nd+1} \dots \lambda_{nd \times 24} &= \sigma_h^2 \end{aligned} \quad (24)$$

The eigenvectors (note: these are presented un-normalized to preserve the intuition of how they are derived) can be thought of grouped into one "overall" eigenvector, $nd - 1$ "across-day" eigenvectors and $nd \times (24 - 1)$ "within-day" eigenvectors

$$\begin{aligned} E &= [e_1 | E_A | E_W] \\ e_1 &= \iota_{nd \times 24} \\ E_A &= F_A \otimes \iota_{24} \\ E_W &= I_{nd} \otimes F_W \end{aligned} \quad (25)$$

with

$$F_A = \begin{bmatrix} 1 & 1 & & 1 \\ -1 & 1 & & 1 \\ 0 & 2 & \dots & \\ \vdots & 0 & & \vdots \\ & \vdots & & \\ 0 & 0 & & -(nd - 1) \end{bmatrix}$$

and

¹⁶ We thank A. Ronald Gallant for the suggestion to use the eigenvalue-vector decomposition.

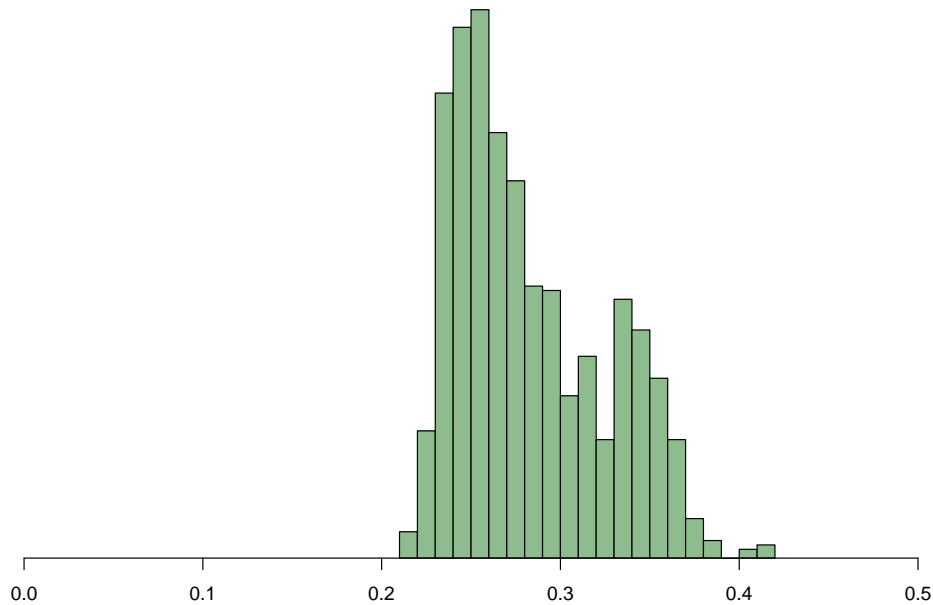


Fig. 7: RW Acceptance Rates

$$F_W = \begin{bmatrix} 1 & 1 & & 1 \\ -1 & 1 & & 1 \\ 0 & 2 & \cdots & \\ \vdots & 0 & & \vdots \\ \vdots & & & \\ 0 & 0 & & -(24 - 1) \end{bmatrix}.$$

5.5 Results of Model Fits

Our MCMC procedure shows excellent convergence and fit properties. A RW Metropolis chain with poorly tuned increment density will exhibit either slow navigation because of steps that are smaller than optimal or high autocorrelation because of high rejection rates. Figure 7 shows a histogram of acceptance rates over an initial run of 1000 randomly selected drivers. For a normal target, an acceptance rate of 30 percent is optimal. Our chain shows a very tight distribution of acceptance rates near 30 percent.

Model fit is also an important consideration. Figure 8 shows a plot of the our model's predicted or expected labor supply in terms of hours per week versus the observed labor supply. Our model certainly captures well the average labor supply. Figure 9 shows actual vs predicted labor supply by hour of the week. The black line connects the observed labor supply for each of the 168 hours of the week, averaged over the 1000 drivers in our initial estimation sample. The red line provides the expected labor supply by hour from our model fits. The vertical colored bands correspond to the hour

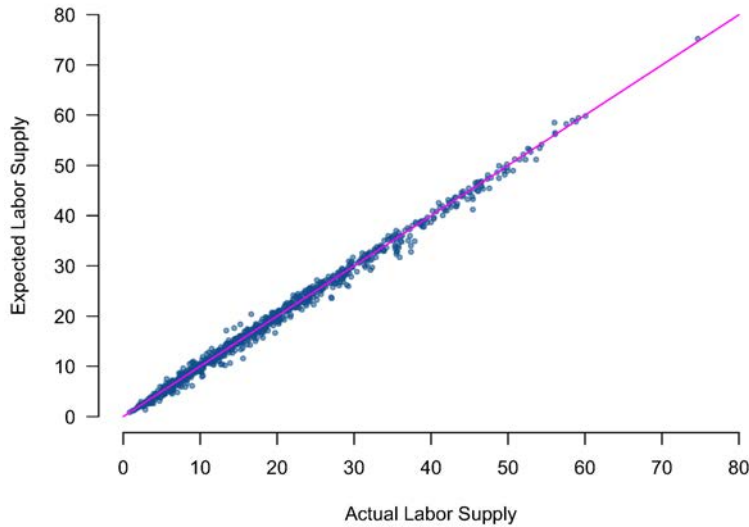


Fig. 8: Expected Vs. Actual Labor Supply (hrs/wk)

blocking used for estimation of the mean wage. In spite of our rather dramatic simplification in which we blocked 168 hours each week into only 9 groups, the model tracks the observed labor supply quite well.

Our MCMC procedure provides draws from each of the 1000 driver posteriors in our estimation sample. We summarize these draws by computing the mean draw which is a simulation-based estimate of the posterior mean. The posterior mean is often used as a Bayesian estimate. Figure 10 shows the estimates of each of the three variance components associated with week, day and hour shocks. We see that all shocks are large with estimates of standard deviations over \$10. The daily and hourly shocks are very large and larger than weekly shocks. This suggests that the drivers in our sample experience large shifts in reservation wage that are likely not predictable and, thus, may place a large value on a flexible work arrangement.

Not only does it appear that Uber drivers are subject to large unpredictable changes in reservation wages, but Uber drivers do not have homogeneous preferences for time of day and day of week. Figure 11 provides scatterplots of various mean preference parameter estimates. Each driver has a separate, and possibly, unique mean reservation wage for all of the nine hour-blocks. For example, the left panel of Figure 11 shows that preferences for the Monday-Friday Rush-Hour block are very heterogeneous (from large negative to large positive mean change in reservation wage relative to the base period). In addition, there is a positive correlation between preferences for the Monday-Friday Afternoon (horizontal axis) and Monday-Friday Rush-Hour block as might be expected for two contiguous hour blocks. Other the other hand, there is a clear negative correlation between preferences for the weekday afternoons versus late night on Friday and Saturday (right panel).

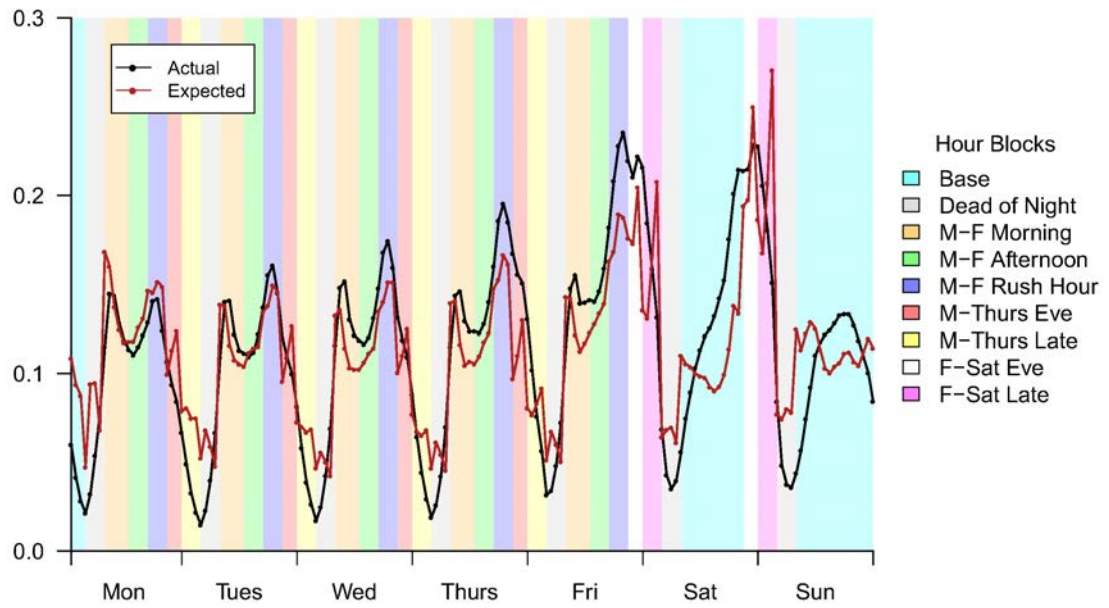


Fig. 9: Expected Vs. Actual Labor Supply by Hour of Week

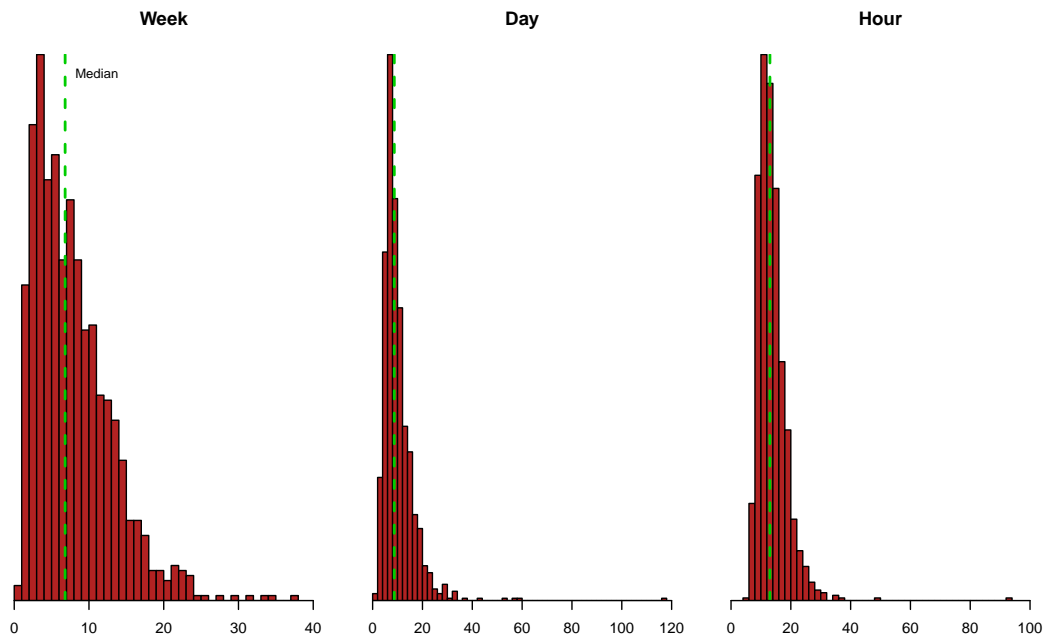


Fig. 10: Variance Component Estimates

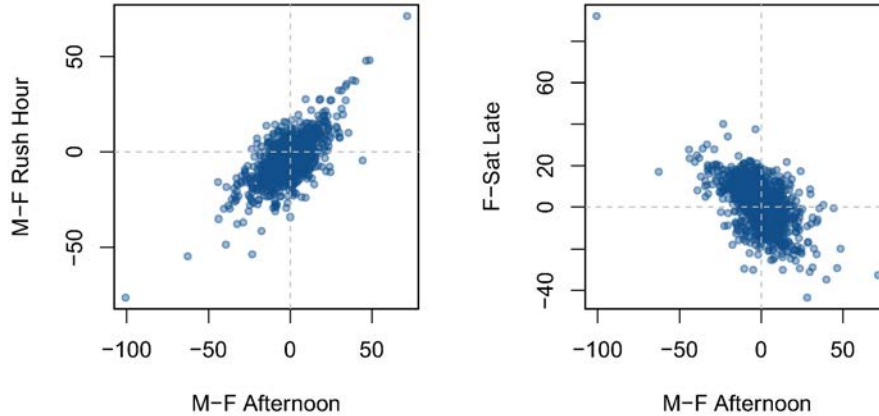


Fig. 11: Scatterplots of Mean Reservation Wage Parameters

6 Driver Surplus

We have postulated a flexible but parsimonious model of labor supply which allows for heterogeneity across drivers and features both predictable and unpredictable aspects of labor supply. It is also clear from our initial exploratory analysis that Uber drivers exercise ample flexibility in their labor supply both in the level or average number of hours per week as well as the pattern of which hours and days on which labor is more frequently supplied. In addition, our fitted model of labor supply shows very large error components, implying that flexibility to adjust to random shocks could be an important component of value to Uber drivers. In this section, we compute measure of expected surplus from the Uber labor arrangement as well as alternative arrangements that afford less flexibility.

6.1 Surplus Measure

Our goal is to compute the expected surplus for each driver. In our model, drivers will only work if their surplus (excess of wage over reservation wage) is positive. We will compute the expected surplus which is the probability that the surplus is positive (i.e. the driver decides to work) times the expected surplus conditional on working. Consider hour t with driver facing wage w_t in that hour, expected surplus can be written as

$$ES_{i,t} = [w_t - E[w_{i,t}^* | w_{i,t}^* < w_t]] \times Pr[w_{i,t}^* < w_t] \quad (26)$$

To produce a welfare measure for each driver, we sum the expected surplus to the driver-week level and compute the average of this measure over all weeks for which we observe the driver in our data. This averages the measure over the distribution of prevailing wages faced by each driver. In the end, we will have one expected surplus value for each driver and we can gauge the impact of various restrictions on the flexibility of labor supply on both the total Uber driver surplus as well as the distribution of this surplus across drivers.

6.2 Constraints on Flexibility

We start with the base case, in which the Uber system imposes no constraints on labor supply flexibility. At the granularity of hour blocks, Uber drivers can choose to work at any time, and because they can make moment-to-moment decisions about labor supply, it is natural to assume that they make these decisions with full knowledge of both the mean reservation wage as well as the realization of shocks to their reservation wages. That is, if there are weeks, days or hours where the cost of supplying labor is very high due to other time commitments, Uber drivers are free to choose not to work either for the whole week, specific days in the weeks, or even specific hours in the week. This flexibility means that a driver can make labor supply decisions based on the idiosyncrasies of the driver's pattern of mean reservation wages as well as the shocks.

For example, if a Uber driver holds down a traditional 9 to 5 job, then we would expect that driver's mean reservation wages for work at Uber are very high during the 9-5 weekday hours. In addition, the Uber system affords drivers flexibility with respect to unpredictable changes in time commitments. For example, while a driver might "normally" work a particular time-block within the week, if a relative or child falls ill, the Uber driver may decide not work for whatever time is required to care for this person. The ability to respond to unpredictable shocks could be a significant source of value for the Uber-style flexible work system and we might expect that individuals with high variances in their hourly reservation wages to find the Uber platform very attractive.

Thus, our approach to welfare calculations is to compare the Uber system "base case" to alternative arrangements in which constraints are imposed on the driver's labor supply flexibility. Some of these alternative arrangements are designed to mimic common labor market institutions. The Uber system allows each driver to make hour by hour labor supply decisions without any constraints. As such, the Uber system represents a base case with the highest degree of flexibility. We will compare the expected surplus under the most flexible "Uber-style" system with the a host of other labor supply arrangements that differ in the nature and severity of constraints on labor supply flexibility.

(Base) Drivers can adapt to week, daily, and hourly shocks with full knowledge of the prevailing wages for that city, week, day and hour and full knowledge of the realization of all of the shocks.

We will consider three basic types of constraints:

Lessened Adaptation to Driver Shocks

In the Base case, drivers make labor supply decisions with full knowledge of the realized value of all weekly, daily and hourly shocks. We consider two other scenarios of decreasing flexibility.

- (A) *Cannot adapt to hourly shocks.* In this scenario, we do not allow the driver to adapt to hourly shocks. One view of this is that the driver must make a decision at the beginning of each day as to what hours he/she will work in that day with knowledge of the distribution of hourly shocks to the reservation wage, but without knowledge of the realization of the shocks for each hour in that day. This case affords flexibility to adapt to weekly and daily shocks but not to hourly.
- (B) *Cannot adapt to daily and hourly shocks.* Here we do not allow the driver to alter labor supply in response to either daily or hourly shocks but only to the weekly shock. Here the driver can adapt to changes in shocks from week to week but not within the week.

It should be emphasized that these scenarios are only restrictions on the driver’s ability to adapt to shocks. We still allow the driver to respond to changes in the prevailing wage, and we assume that the drivers have perfect foresight as to the prevailing wage. We also allow the driver to have a driver-specific profile of mean reservation wages that can vary by day of week and hour of day.

Commitment over Longer Time Horizons

In scenarios (A) and (B) above, the driver can vary labor supply from week to week in response to weekly shocks as well as predictable changes in prevailing wages. Our scenario (Month) restricts this ability.

(Month) *Month-long Labor Supply Commitment.* At the beginning of each month the driver must make a commitment to work the same schedule each week and cannot respond to week, daily, or hourly shocks. In addition, the driver cannot respond to changes in prevailing wages from week to week in the month. We assume that the driver must make decisions based on average weekly profile of prevailing wages where this is averaged over the month ahead.

Institutional Constraints

In many labor markets, workers face institutional constraints that require a form of pre-commitment to specified blocks of time.

(Taxi) *“Taxi” Constraints.* In the “taxi” industry a driver must choose between one of each of three shifts on a daily basis. If taxi worker decides to work a shift, the driver is effectively obligated to work the entire shift by virtue of high lump sum rental prices for the taxi cab. Thus, the taxi labor system has flexibility from day to day and week to week but imposes block constraints on the hours within each day.

We model these “taxi” constraints in two ways. In both, drivers know their week and day shocks but cannot adapt to hourly shocks and must decide which if any of the three shifts to work based on the expected surplus for the entire shift.¹⁷ In most taxi settings, the high upfront fee charged to taxi drivers effectively necessitates working the whole shift, and we examine that simplified model. We also examine a different “Uber-shift” scenario, where a lower fee is pegged at 20% of the expected total earnings in a shift. Drivers who pay this lower fee purchase the right to pick up fares during that shift without any further fee, and while maintaining the ability to flexibly choose which hours to work (or not work) in that shift.

6.3 Calculating Expected Surplus

Base Case

In the base case, the driver sees the realization of week, day and hour shocks and is able to make a labor supply decision conditional on these shocks, mean reservation wages as well as prevailing wages

¹⁷ Indeed, our model of the “taxi” constraint still provides the driver more flexibility than he or she might have in a true taxi environment in which decisions to rent a cab for the shift are often bundled across days.

for that city, week, day and hour. Thus, expected surplus for each hour can be computed from the expectation of the reservation wage conditional on working which is the expectation of a truncated normal random variable. We sum these for each hour of the week and then average them over all complete weeks in the data (note: typically, there are incomplete weeks at the beginning and ending of the drivers data records as drivers do not begin and end their affiliation with Uber on the first hour of each week, which we take to be midnight to one a.m. on Sunday).

Consider week, l

$$ES_{i,l} = \sum_{d=1}^7 \sum_{h=1}^{24} [w_{ldh} - E[w_{i,ldh}^* | w_{i,ldh}^* \leq w_{ldh}]] \times Pr[w_{i,ldh}^* \leq w_{ldh}] \quad (27)$$

and the surplus for driver i is

$$ES_i = \frac{1}{N_i} \sum_{l=1}^{N_i} ES_{i,l} \quad (28)$$

where N_i is the number of complete weeks observed for driver i .

The conditional expectation of reservation wage given the decision to work can easily be computed as it is the mean of truncated normal random variable:

$$\begin{aligned} E[w_{i,ldh}^* | w_{i,ldh}^* \leq w_{ldh}] &= \mu_i(d, h) + E[\xi | \xi \leq w_{ldh} - \mu_i(d, h)] \\ &= \mu_i(d, h) - \sigma_\xi \frac{\phi\left(\frac{TP_{i,ldh}}{\sigma_\xi}\right)}{\Phi\left(\frac{TP_{i,ldh}}{\sigma_\xi}\right)} \end{aligned} \quad (29)$$

where $TP_{i,ldh} = w_{ldh} - \mu_i(d, h)$ is the truncation point and $\xi \sim N(0, \sigma_w^2 + \sigma_d^2 + \sigma_h^2)$ is the sum of the shocks. The probability of working is simply given by the normal CDF evaluated at the truncation point.

$$Pr[w_{i,ldh}^* \leq w_{ldh}] = \Phi\left(\frac{TP_{i,ldh}}{\sigma_\xi}\right) \quad (30)$$

Cases A, B: Restricted Ability to Adapt to Shocks

In cases A and B, the drivers cannot respond to all of the shocks. In case A, they cannot adapt labor supply to the hourly shock while in case B they cannot adapt to both the hourly and daily shocks. In these cases, the expected surplus can be computed in exactly the same manner except we must refine the random variable ξ . In case A, $\xi = v_w + v_d$ while in case B, $\xi = v_w$. That is, drivers can only use the expected value (0) of the shocks as they do not have access to the realization.

“Taxi” Constraints

Here the driver can observe weekly and daily shocks but not hourly shocks and must choose one of three eight hour shifts or not to work at all. This means that the driver must calculate the expected surplus for each shift. The driver chooses the shift with largest positive surplus. If no shifts have

expected surplus that is positive, then the driver doesn't work at all. That is, at the beginning of each day, the driver knows $v_w + v_d$, mean reservation wages for each hour of that day, $\mu(d, h)$, and prevailing wages for each hour of that day. To make the labor supply decision and compute expected surplus, we simply sum over all of the hours in each shift.

$$w_{S_j}^* = \sum_{t \in S_j} \mu_t + N_{S_j} (v_w + v_d) = \mu_{S_j} + \xi_{S_j} \quad (31)$$

The driver will find the shift with the maximum value of $(w_{S_j} - \mu_{S_j})$ where $w_{S_j} = \sum_{t \in S_j} w_t$. He will work that shift if

$$\begin{aligned} \max(w_{S_j} - \mu_{S_j}) - \xi_{S_j} &> 0 \\ \xi_{S_j} &\sim N\left(0, N_{S_j}^2 (\sigma_w^2 + \sigma_d^2)\right) \end{aligned} \quad (32)$$

From the inequality above, we can compute both the probability of working on that day and the expected surplus from the shift worked (note: we assume that under the "taxi" arrangement drivers only work one shift). Expected surplus for each day is given by

$$ES = \Phi\left(\frac{\max(w_{S_j} - \mu_{S_j})}{\sigma_\xi}\right) \times \left[\max(w_{S_j} - \mu_{S_j}) - \sigma_\xi \frac{\phi\left(\frac{\max(w_{S_j} - \mu_{S_j})}{\sigma_\xi}\right)}{\Phi\left(\frac{\max(w_{S_j} - \mu_{S_j})}{\sigma_\xi}\right)} \right] \quad (33)$$

Note that σ_ξ is determined by the number of hours that are in shift j . In our "taxi" scenario, all shifts are eight hours long. To compute total expected surplus for the "taxi" arrangement, we simply sum over all days (note that the wages will vary across days) and express this on a weekly basis.

Month-Long Time Commitment

In this case, the driver cannot respond to any shock and must commit to a work schedule for a month at time (note: for simplicity's sake, we define a month as a four week interval). In addition, the driver cannot respond to week to week changes in prevailing wages and must make decisions based on the average prevailing wages for that "month."

Let w_{mdh} be the average wage for each day of week and hour of day for 4 week period, m .

$$ES_{i,m} = 4 \times \sum_{d=1}^7 \sum_{h=1}^{24} [w_{mdh} - \mu_i(d, h)]^+ \quad (34)$$

Again, we simply average these measures over the number of four week periods observed for each driver to obtain, ES_i .

6.4 Results

For each of the 1000 drivers in our preliminary estimation sample, we compute Bayes estimates of the mean reservation wage parameter and Bayes estimates of each of the variance components necessary

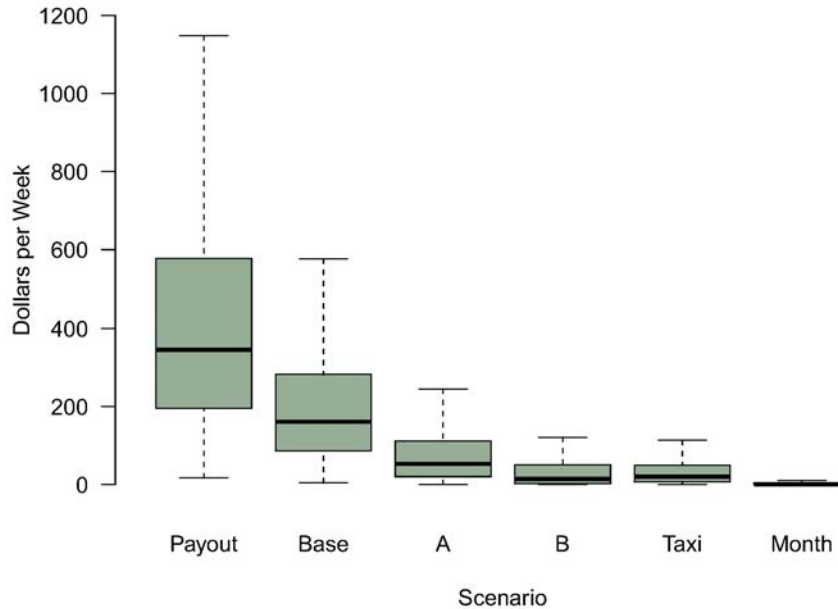


Fig. 12: Expected Labor Surplus

for the expected surplus computations. Figure 12 shows box plots of the distribution of surplus over the various labor supply flexibility arrangements or scenarios outlined in 6.2. The box plot labeled “Payout” is the expected total wages (in \$/wk) that drivers should have earned for each of the hours in which they were actually observed to supply labor.

Our abstraction of the Uber-style arrangement in which drivers can pick whatever hour, day, and week they choose to work generates a large surplus of about 40 percent of their total pay. Constraints on flexibility reduce surplus a great deal. For example, just inability to adapt hour by hour within the same day (contrast case “A” with the “Base”) dramatically reduces surplus. If drivers are further restricted to be unable to adapt to both hourly and daily shocks (scenario “B”), surplus is further reduced but by a smaller factor than the hourly case. In other words, adaptation to the daily shock is less valuable to drivers than adaptation to hourly shocks. The “taxi” case allows drivers to respond to daily shocks but constrains them to work a full 8 hour shift. This results in a large reduction in surplus over the case “A” where daily adaptation is possible but drivers do not have to work an entire shift.

Finally, pre-commitment to a “month” in advance reduces expected surplus to near zero for most Uber drivers. We should emphasize that commitment to a month in advance merely limits adaptation to weekly, daily and hourly shocks but still allows for a flexibility in the amount time worked and the allocation of that time worked over days of the week and hours of the day. It is not just a flexible work schedule, per say, that creates value to Uber drivers but it is the ability to adapt to events which are not predictable that is most valuable.

Figure 13 reports the distribution of expected labor supply for each of our proposed arrangements. It is clear that constraints on flexibility also reduce willingness to work. For example, the “taxi” constraint on adaptation to unpredictable shocks, reduces labor supply from about 15 hours per week to less

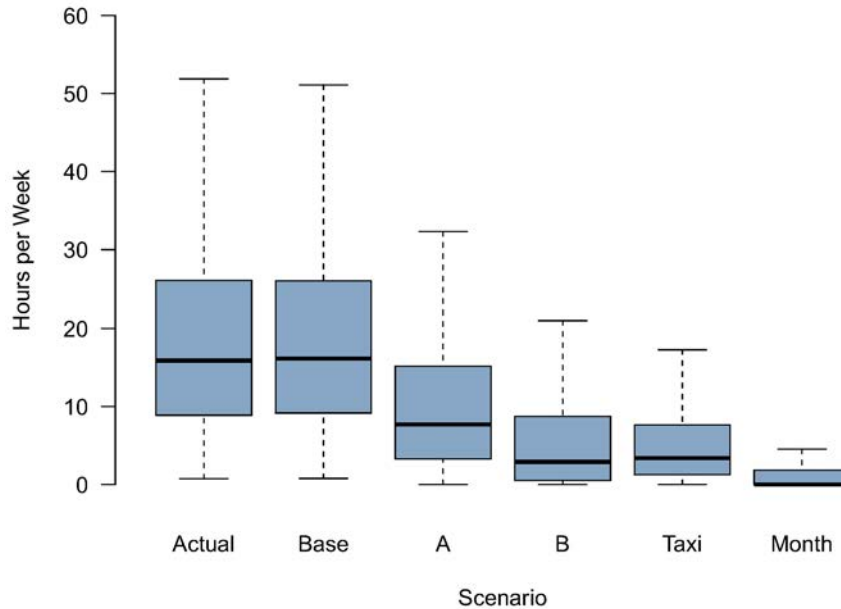


Fig. 13: Expected Labor Supply

than 5 hours per week. This is, perhaps, not that surprising as our sample consists of drivers who have selected the Uber arrangement by choice over a taxi arrangement. However, there are many reasons a driver might prefer to be an Uber driver that are not related to flexibility, including a superior dispatching driver application with suggested driving routes and many other features of potential value, direct deposit of payments to a checking account and the lower upfront cash requirements (no shift rental fee). What we have learned is that ability to adapt unpredictable shocks attracts Uber drivers and, without this adaptability, they will not participate much in the labor market.

Table 5 provides more detail on the distribution of expected labor supply (top panel) expected surplus (bottom panel). Of course, there is a distribution of surplus across drivers, some benefit much more than others from the Uber arrangement. At the median of the expected surplus distribution across drivers, the Uber arrangement (“base”) provides an expected surplus of \$160 per week while simply turning off the ability to adapt to hourly shocks reduces surplus to less than one third of that, about \$53 per week. Taxi arrangements are even worse, reducing surplus to one eighth of the Uber base case (\$20 per week).

We might expect that the value of the Uber arrangement should depend on the level of labor supply offered. If a driver is driving 50 or more hours per week, flexibility from day to day and hour to hour is limited almost by definition. Table 6 slices the expected surplus distribution based on decile of observed weekly labor supply. While drivers in the lower deciles of labor supply do benefit proportionately more (compare expected surplus in the “base” condition to payout) than drivers in the highest deciles, the proportion is relatively stable from 37 to 51 percent. It is remarkable that even drivers working more than 37 hours per week find the “taxi” arrangement to have an expected surplus of only one quarter of the Uber arrangement.

Our alternative scenario surplus calculations are designed to decompose the source of the labor surplus

Labor Supply		Expected Labor Supply				
Quantile	Actual	Base	A	B	Taxi	Month
0.99	55.1	53.6	45.9	46.1	30.8	46.0
0.95	45.2	44.7	35.2	30.6	21.2	27.6
0.90	37.1	36.3	26.9	20.6	14.4	13.3
0.75	26.1	26.0	15.1	8.7	7.6	1.8
0.50	15.9	16.1	7.7	2.9	3.4	0.0
0.25	8.9	9.2	3.3	0.5	1.2	0.0
0.10	5.6	5.5	1.4	0.1	0.4	0.0
0.05	4.1	4.1	0.7	0.0	0.2	0.0
0.01	1.9	2.1	0.2	0.0	0.0	0.0

Estimated Surplus		Expected Surplus				
Quantile	Payout	Base	A	B	Taxi	Month
0.99	1264.1	768.5	486.1	325.6	272.4	270.6
0.95	982.9	544.8	285.5	191.8	152.7	119.2
0.90	821.1	430.6	215.6	126.8	110.4	51.6
0.75	576.2	281.6	110.8	50.2	49.5	4.0
0.50	344.8	160.3	52.9	14.0	19.9	0.0
0.25	195.5	85.8	19.8	2.0	6.5	0.0
0.10	119.8	48.7	7.2	0.1	1.6	0.0
0.05	87.7	32.4	3.6	0.0	0.7	0.0
0.01	41.4	13.3	0.8	0.0	0.1	0.0

Tab. 5: Quantiles of the Expected Labor Supply and Surplus

Labor Supply			Expected Surplus				
Decile	Labor Supply Range	Payout	Base	A	B	Taxi	Month
10	(37.2 - 74.3)	987.0	500.3	239.0	144.2	125.7	84.3
9	(29.1 - 37.0)	715.3	345.1	146.3	60.6	61.3	4.4
8	(23.9 - 29.0)	569.3	278.6	98.9	43.9	44.7	2.5
7	(19.6 - 23.9)	455.7	214.0	71.7	25.2	32.4	0.6
6	(15.9 - 19.6)	382.1	185.1	71.2	23.3	33.6	0.0
5	(12.9 - 15.9)	308.0	135.5	43.5	7.1	16.7	0.0
4	(10.1 - 12.9)	253.1	117.7	36.8	6.5	13.7	0.0
3	(7.9 - 10.0)	192.7	87.7	22.4	2.4	8.4	0.0
2	(5.6 - 7.8)	150.1	64.0	15.3	1.9	4.5	0.0
1	(0.7 - 5.6)	86.8	32.7	5.6	0.4	1.5	0.0

Tab. 6: Expected Surplus by Labor Supply Decile

enjoyed by drivers under the Uber labor supply arrangement. That is, there are various dimensions of adaptation to shocks and lack of pre-commitment which contribute the Uber surplus value; and we use our model to quantify the value of each. We hasten to add that we are not simulating a new equilibrium in the labor market under each scenario considered. For example, if the Uber arrangement were outlawed (as it has been in some communities) and only taxis remain, there would be a new equilibrium in the labor market with a new distribution of labor supply and different equilibrium wages from those observed in our data. Many, if not most, of the Uber drivers would withdraw from the market which meets consumer transportation demand and work elsewhere or not at all. Clearly, much of the labor surplus would be lost but there would be surplus gains for taxi drivers as taxi utilization and wages might increase as well as surplus obtained by Uber drivers in alternate jobs. We are not undertaking such a calculation which would require a host of difficult-to-verify assumptions as well as detailed data on the labor supply decisions of taxi drivers.

Another limitation of our analysis is that our simplified alternative scenario analyses may not align with many employment relationships. Workers in conventional jobs typically have some flexibility to adjust for shocks. That is, while one would imagine that a shift worker would lose her job if she continuously adjusted her labor supply to small shocks to her reservation wage, the worker can, when faced with a large and unusual shock, for example, likely call in sick. Thus, some conventional jobs may well be more flexible than our counterfactual examples. However, the literature cited above suggests that flexibility for low-wage low-skilled workers may be particularly limited.

Finally, there are alternative ways to conceptualize a taxi-like scenario other than the one we have posed here. For example, we could consider a scenario in which an Uber driver purchases the right to work during a shift for a fixed fee, without having to subsequently pay Uber's commission. In this alternative "Uber-shift" scenario, the driver has the right, but is not obligated to work each individual hour in a purchased shift (as we have modeled it in our "taxi" scenario). This "Uber-shift" scenario amounts to recasting equations 33 through 31 to be conditioned on an individual hour being worked only if it generates a positive surplus at the higher effective wage. In this scenario, the driver compares the total surplus generated to the fixed fee to determine participation.

For an individual driver presented with this Uber-shift scenario, we can see intuitively how this alternative scenario would compare to those we have already presented. For a fixed fee that is very low, the driver behaves similarly to the base case with full flexibility (but supplies more labor through the wage increase that comes from not sharing with Uber). For sufficiently high levels of the fixed fee, this scenario should resemble our "taxi" simulation, as a driver will only supply labor if he or she expects to work the whole shift and earn enough surplus to cover the fixed fee. For fixed fees high enough, most drivers will simply refuse to participate and both labor supply and surplus will be low. In ongoing work, we explore additional flexibility scenarios, including one in which the driver pays an individualized fixed fee keyed to his or her individual mean number of hours driven.

7 Sensitivity Analyses

In this section, we consider two sensitivity analyses performed to assess the role of model assumptions and prior settings.

7.1 Exogeneity of Wages

We assume that our expected wage variable is exogeneous to the labor supply decisions made on an hour-to-hour basis by drivers.¹⁸ While there can be common demand shocks such as a concert or sporting event, this assumption rules out common supply shocks. A legitimate concern is that if there are common supply shocks of large magnitude, this can affect both the labor supply decisions of drivers as well as the prevailing expected wage. For example, if there was a large positive common supply shock which raised reservations prices simultaneously for all drivers in a city, then we would expect falling overall labor supply and increasing wages, particularly in light of Uber’s dynamic pricing policies which are designed, in part, to remedy supply deficiencies. Thus, we would expect that if there are common supply shocks (both positive and negative) that our model fitted under the exogeneity assumption would under-estimate the responsiveness of drivers’ to changes in wages. Of course, the extent of this “endogeneity” bias would depend on the relative magnitudes of common and idiosyncratic supply shocks.

To assess the importance of possible bias due to common supply shocks, we would ideally like some source of variation in wages which can be plausibly viewed as exogeneous. If wages were varied randomly, this would be the ideal source of wage variation which is indisputably exogenous. Uber has conducted some randomized changes in wages on a limited basis in several cities. In the range of our data, one such set of changes was conducted in Orange County, California during April of 2016. A random sample of drivers received an email with a 10 per cent increase in wages for a three week period.¹⁹ Another randomly selected group of drivers was selected for control purposes. The randomization was personally conducted by Keith Chen who was then an employee of Uber. Approximately three thousand drivers were assigned to both the “control” and “incentive” groups. In our analysis sample, we have 1,272 of the “incentive” group drivers and 1,240 of the “control group.”²⁰

To assess the extent of biases in our model estimates, we exploit this source of exogeneous variation to stress-test our model. We refit our model for each of the non-experimental weeks in the “incentive” group and use these fitted coefficients to predict the response of the drivers to the experimentally induced increase in incentives during the duration of the experiment. If our assumption of exogeneity of wages is incorrect and if there are large common shocks, we should expect to that our model would under-predict the actual labor supply response to the increase in effective wage rate.

As in all field experiments, there are important implementation considerations. First, the incentive group could only qualify for the 10 per cent increase in earnings on trips which originate in Orange County. If, for example, an Uber driver picked up a fare from Irvine (in OC) to LAX airport, the driver would certainly try to pick up a return fare to maximize efficiency. The return trip would not qualify for the 10 percent incentive. This is a major issue for the Orange County market due to close proximity to Los Angeles and many other separate Uber markets. In addition to restrictions on the origin of trips, the Uber incentive offer contained additional qualifications. Drivers “must maintain a 90% acceptance and 25% completion rate over all hours on line to qualify for this offer.” It is not clear how binding these constraints are on Uber drivers. In sum, the restrictions on the incentive condition in the OC incentive experiment mean that drivers effectively faced an incentive that could be considerably lower than 10 per cent.

¹⁸ Frechette et al. [2016] make a similar assumption in the model of taxi driver labor supply used in their equilibrium simulations.

¹⁹ The same email offer was sent out once per week for a total of three weeks.

²⁰ Recall that we restrict attention to an “active” sample of drivers who work in at least 16 hours during our dataset.

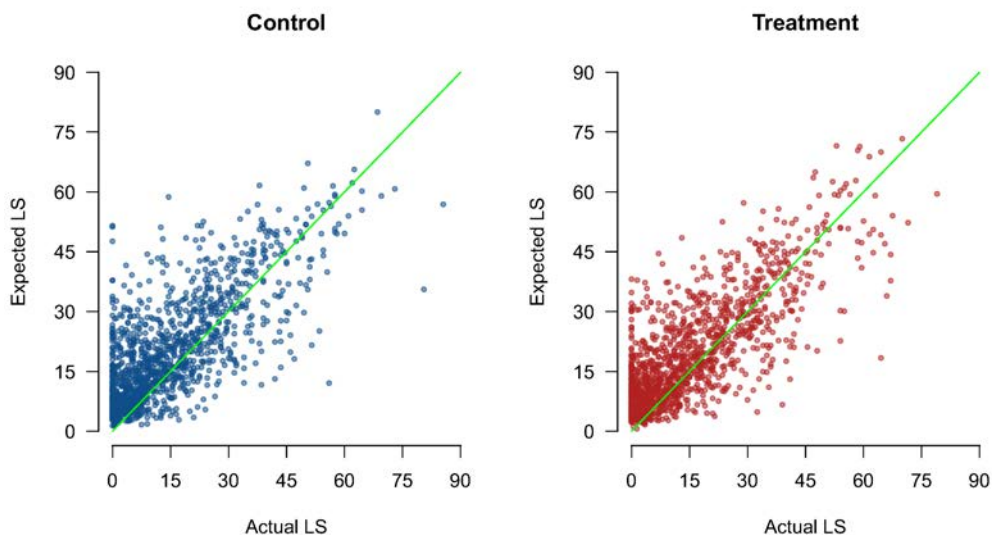


Fig. 14: Expected Vs. Actual Labor Supply: OC Experiment

To assess the effective incentive rate, we exploit the fact that the OC experiment data show that, on average, drivers do not appear to have responded to increased incentives in the first week of the experiment.²¹ In week 1, 1050 “control” drivers supplied 12,938 hours of labor or 12.32 hours per week on average, while 1063 “incentive” drivers supplied 13,060 hours or 12.28 hours per week.²² That is, there was no aggregate response to the incentive. We can then use the payouts made to drivers to estimate the effective incentive rate. We regressed $\log(\text{wages})$ on dummy variables for “incentive” eligibility and fixed effects for hour of day and week. We find that “incentive” eligible drivers earned a wage 2.3% higher than controls (std error of .003). We will use this wage differential to predict labor supply for the remaining two weeks of the experiment.

Figure 14 compares actual labor supplied (in hours/wk) with what is expected or predicted from our model fit to non-experimental weeks (and using the estimated 2.3 per cent increase in wages) for the two weeks of experimental data. We plot this separately for the control group of drivers versus treatment group. If there are no substantial biases in our model coefficient due to endogeneity, we should expect that controls and treated (incentive) group drivers would exhibit the same level of model fit. The two scatterplots are very similar.

However, it is possible that this model diagnostic based on true experimental variation is valid but of low discriminatory power due to the relatively small effective incentive of 2.3%. To assess the power of our diagnostic, we compare the predicted labor supply on experimental weeks under the assumption of a 2.3% increase in wages with the predicted labor supply with no increase in wages. Figure 15 provides a histogram of the difference in expected labor supply assuming a 2.3% increase in wages and a 0% increase. There is a discernible increase in labor supply predicted by our model from even this relatively small change in wage rates. This provides some validation that our diagnostic procedure

²¹ This is not uncommon with driver incentives at Uber, and can be explained by inattention or delayed attention to emails. Drivers appear to first become widely aware of incentives at the end of an incentive’s first week, when additional earnings appear on earnings statements.

²² This is not significant.

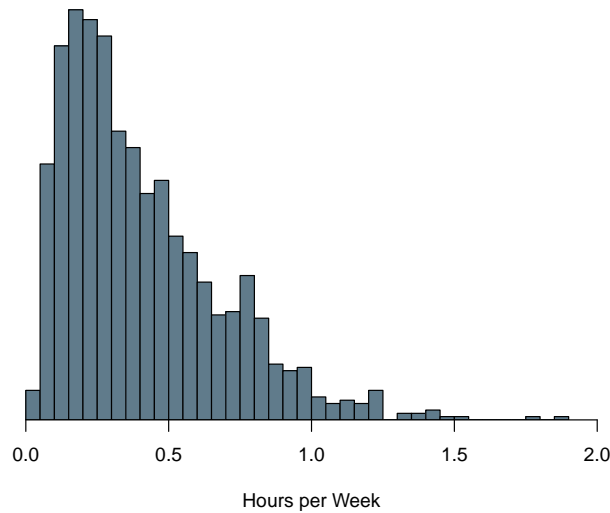


Fig. 15: Difference in Labor Supply for Observed Vs. No Incentive

has power in the relevant range of wage increases.

We conclude that possible biases due to common supply shocks are apt to be small relative to the predicted labor supply changes predicted by our model.

7.2 Prior Sensitivity

The prior we use in our analyses (5.4) is a relatively diffuse prior but still more informative than is typically used. In many Bayesian analyses, extremely diffuse “default” prior are used. The emphasis in typical Bayesian applications is inference regarding model parameters. Our emphasis is on estimated labor surplus which is a complicated function of the distribution of model parameters over drivers. We use a tighter prior than typical analyses in order to obtain conservative estimates of the extent of labor surplus. Large values of the reservation wage parameters would typically be associated with higher levels of surplus. For example, the prior on σ shown in (18) shrinks the estimates of the shock standard deviations toward zero. Smaller shock sizes will lower the value of flexibility. Similarly, the prior on the mean reservation wage parameters is designed to shrink away from higher values of surplus. We do not shrink the intercept in the mean reservation wage parameters but we do shrink the “slopes” or differences in mean reservation wages toward zero which reduces the labor surplus generated by periods with low mean reservation wage. Of course, all priors will be dominated by the data in situations where large amount of informative. Most drivers in our dataset are observed for 5000 or more hourly observations; however, there is a small subset (those drivers that worked only a short interval of time in our data) where there number of hourly observations is in the 1000 to 2000 range. For these drivers, our priors serve to reduce outlying parameter estimates which might be associated with erroneously high labor surpluses.

In this section, we explore more diffuse prior settings and gauge the impact of these “looser” prior settings on the distribution of estimated labor surplus. We employ four basic sets of prior settings.

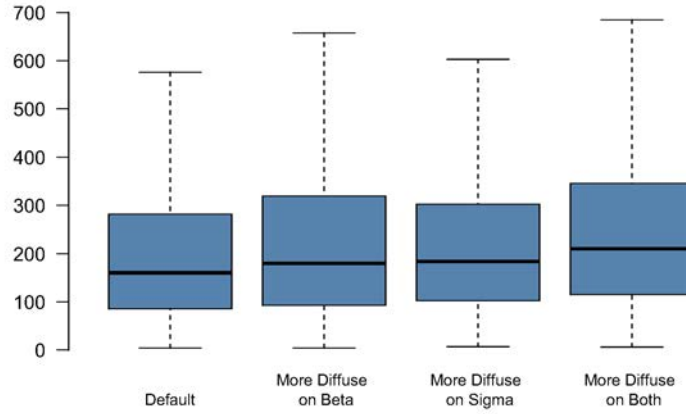


Fig. 16: Prior Sensitivity for Expected Labor Surplus

Recall that the priors are of the form:

$$\beta \sim N(0, A^{-1})$$

$$\sigma_j \sim iid exp(N(\mu_{lnsigma}, \sigma_{lnsigma}^2))$$

Default: $diag(A)^t = (.01, .1, \dots, .1)$, $\mu_{lnsigma} = 1$, and $\sigma_{lnsigma}^2 = .5$

Loosen_Beta: $diag(A)^t = (.01, .05, \dots, .05)$, $\mu_{lnsigma} = 1$, and $\sigma_{lnsigma}^2 = .5$

Loosen_Sigma: $diag(A)^t = (.01, .1, \dots, .1)$, $\mu_{lnsigma} = 3$, and $\sigma_{lnsigma}^2 = 2$

Loosen_Both: $diag(A)^t = (.01, .05, \dots, .05)$, $\mu_{lnsigma} = 3$, and $\sigma_{lnsigma}^2 = 2$

Figure 16 shows the distribution of expected labor surplus for the Base case with each of the four priors above. As expected the more diffuse priors shift the distribution of expected labor surplus to somewhat higher values. It does appear that, within a reasonable range, our results are not overly sensitive to the prior distributions used.

8 Conclusions

The Uber driver arrangement attracted more than a million drivers to offer labor supply during the 8 month period of our data, which is limited to only US UberX service. One of the attractions of Uber is the flexibility afforded drivers. Not only can drivers choose to supply relatively small number of hours per week but they can allocate these hours flexibly over the days and hours of the week. However, this is not the only or even most important source of flexibility provided to Uber drivers.

Another important source of flexibility is the ability of an Uber driver to adapt on an hour-by-hour basis to changes in demands on his/her time. While traditional workplaces do compete to provide flexibility to workers, the literature suggests that lower-wage lower-skill workers typically have limited ability to respond to everyday shocks. The goal of this paper is to estimate a model of labor supply which will allow for a quantification of the value of both flexibility and adaptability.

We postulate a model in which each driver has a reservation wage process with both predictable mean component as well as weekly, daily, and hourly shocks. This operationalizes the view that workers face unpredictable events which can change their labor supply decisions on an hourly basis. We assume that drivers form rational expectations regarding the expected wage and make labor supply decisions on an hour-by-hour basis by comparing their own reservation wage to the prevailing expected wage. This model is a multivariate Probit model with a time-varying censoring point which facilitates a greater degree of identification than the traditional Probit structure. Driver-level exact finite sample inference is possible using a hybrid MCMC approach.

Our analysis demonstrates very large shock variances, suggesting the potential for large driver surplus in Uber-like arrangements that allow drivers to decide, on an hour-by-hour basis, when to work. We compute driver labor surplus—accounting for 40% of total expected earnings, or \$150 per week on average—under the existing arrangement and alternative work arrangements, which limit drivers' ability to adapt to hourly and daily reservation wage shocks. Constraints on the ability to adapt to shocks have large effects on expected labor surplus; eliminating this ability reduces labor surplus by more than two-thirds. We also consider a “taxi” style arrangement in which drivers can decide on a daily basis whether or not to work and which of three shifts to work but must work an entire 8 hour shift. The “taxi” arrangement reduces expected labor surplus to one-eighth of the Uber arrangement.

In summary, we document an important source of value in flexible work arrangements—the ability to adapt work schedules to unpredictable shocks to reservation wages. Perhaps not surprisingly, this adaptability has high value to individuals who have selected into the Uber platform. Our expectation is that technology will enable the growth of more Uber-style work arrangements. While such arrangements may have important downsides relative to the traditional careers they supplant, we expect that flexibility will be an important source of value in such arrangements.

References

- Joseph G Altonji and Christina H Paxson. Labor supply preferences, hours constraints, and hours-wage trade-offs. *Journal of Labor Economics*, 6(2):254–276, 1988.
- Joseph G Altonji and Emiko Usui. Work hours, wages, and vacation leave. *ILR Review*, 60(3):408–428, 2007.
- James T Bond and Ellen Galinsky. Workplace flexibility and low-wage employees. *New York, NY: Families and Work Institute*, 2011.
- Nicholas Buchholz. Spatial equilibrium, search frictions and efficient regulation in the taxi industry. Technical report, University of Texas at Austin, 2015.
- Colin Camerer, Linda Babcock, George Loewenstein, and Richard Thaler. Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2):407–441, 1997.
- Henry S Farber. Is tomorrow another day? the labor supply of new york city cabdrivers. *Journal of political Economy*, 113(1):46–82, 2005.
- Henry S Farber. Why you cannot find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, 130(4):1975–2026, 2015.
- Guillaume R Frechette, Alessandro Lizzeri, and Tobias Salz. Frictions in a competitive, regulated market: Evidence from taxis. Technical report, NYU, 2016.
- Alan E. Gelfand, S. K. Sahu, and Brad P. Carlin. Efficient parameterizations for normal linear mixed models. *Biometrika*, 82(3):479–488, 1995.
- Jonathan V. Hall and Alan B. Krueger. An analysis of the labor market for uber driver-partners in the united states. Working Paper 22843, National Bureau of Economic Research, November 2016.
- Daniel S Hamermesh. The timing of work over time. *The Economic Journal*, 109(452):37–66, 1999.
- Julia R Henly and Susan J Lambert. Unpredictable work timing in retail jobs: Implications for employee work–life conflict. *ILR Review*, 67(3):986–1016, 2014.
- Lawrence F. Katz and Alan B. Krueger. The rise and nature of alternative work arrangements in the united states, 1995–2015. *NBER*, Working Paper(22667), September 2016.
- Peter F Kostiuk. Compensating differentials for shift work. *Journal of Political Economy*, 98(5, Part 1):1054–1075, 1990.
- Ricardo Lagos. An analysis of the market for taxicab rides in new york city. *International Economic Review*, 44(2):423–434, 2003.
- Alexandre Mas and Amanda Pallais. Valuing alternative work arrangements. Technical report, National Bureau of Economic Research, 2016.
- Robert E. McCulloch and Peter E. Rossi. An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics*, 64:207–240, 1994.

- Gerald S Oettinger. An empirical analysis of the daily labor supply of stadium vendors. *Journal of Political Economy*, 107(2):360–392, 1999.
- Council of Economic Advisors. Work-life balance and the economics of workplace flexibility. Working paper, Executive Office of the President, March 2010. URL <https://www.worldatwork.org/adimLink?id=50354>.
- Paul Oyer. The independent workforce in america: The economics of an increasingly flexible labor market. Technical report, Stanford University, 2016.
- James B Rebitzer and Lowell J Taylor. Do labor markets provide enough short-hour jobs? an analysis of work hours and work incentives. *Economic Inquiry*, 33(2):257–273, 1995.
- Jeremy Reynolds. When too much is not enough: Actual and preferred work hours in the united states and abroad. In *Sociological Forum*, volume 19, pages 89–120. Springer, 2004.
- Gareth Roberts and J. Rosenthal. Optimal scaling for various metropolis-hastings algorithms. *Statistical Science*, 16:351–367, 2001.
- Peter E. Rossi, Greg M. Allenby, and Robert E. McCulloch. *Bayesian Statistics and Marketing*. John Wiley and Sons, 2005.
- Sarah Senesky. Testing the intertemporal labor supply model: are jobs important? *Labour Economics*, 12(6):749–772, 2005.