Suppliers and Demanders of Flexibility: The Demographics of Gig Work*

M. Keith Chen

UCLA Anderson School of Management

Judith A. Chevalier

Yale School of Management

Peter E. Rossi

UCLA Anderson School of Management

Lindsey Currier

Uber Technologies

Abstract

Platform gig work such as rideshare driving involves workers supplying flexibility to the platform, for example, providing service when demand is high. It also can be attractive to workers who demand flexibility, for example, workers with irregular commitments in other jobs. Who benefits the most (and least) from flexible work arrangements? Workers who supply labor elastically provide flexibility to the platform and receive above the platform-average compensation. In contrast, workers who demand substantial flexibility from the platform may not achieve above-average remuneration but may nonetheless benefit from the availability of flexible work options. We characterize the demographics of Uber drivers; we demonstrate that the share of drivers that are women is highest for women from low-income Census tracks and among Black drivers. While within-demographic heterogeneity is large, we show that olders workers demand less flexibility from the system and women demand disproportionately more. We show that workers from low-income census tracts are important suppliers of flexibility to the system.

*Uber Technologies provided the data used in this paper. Under the agreement granted the authors, Uber has the right to review the paper for confidential information 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. Chen is a former employee of Uber, and as a result, continues to hold stock options that may constitute a material financial position. Currier is also a former Uber employee. This research was found to be IRB exempt (Yale approval case 2000024805).

1 Introduction

Disparities across worker demographics in the propensity to supply flexibility to their employers and to demand flexibility from their employers may be an important factor in job choice and earnings. Specifically, recent literature suggests that part of the wage gap between men and women stems from the differential propensity of men and women to provide after-hours work and on-demand work (Bertrand et al. (2010) and Chen and Chevalier (2008)). While women are disproportionately less likely to supply this type of flexibility *to* their employers, women may also require flexibility *from* their employers, as evidence suggests that childcare disruptions disproportionately impact womens' work. For example, (Houghton, 2020) documents that, for technology workers in Seattle, womens' hours worked drop dramatically on days with school snow cancellations while mens' hours do not. Recent evidence also demonstrates that childcare disruptions of the pandemic disproportionately impacted women at all wage levels, especially women of color (Hegeness (2020) and Frye (2020)). Relatedly, recent literature documents that the propensity of workers to undertake part-time contract work increases with age (Abraham et al. (2021b)), and suggests that workers would be inclined to continue working to an older age if more flexible work opportunities were available, particularly for less-educated older workers. (Ameriks et al. (2020)).

At-will labor relationships mediated by digital platforms have become a visible and growing part of the gig economy. Platform-based gig work provides a setting in which the remuneration received for supplying flexibility and the willingness to pay to obtain flexibility can potentially be extracted from workers' use of the platform. Platforms such as Uber, Lyft, TaskRabbit, Door Dash, Bird, Lime, and Instacart all rely on gig workers. Workers on these platforms perform services in response to demand but, critically, have flexibility in deciding whether or when to work. That is, for example, a worker who chooses to work recharging Bird or Lime scooters one night typically has no contractual obligation to charge them the next. However, if the worker does not respond to platform demand conditions, for example, by not collecting scooters after commuting time when uncharged scooters are plentiful, remuneration will be limited. While in other settings the rewards for supplying flexibility may be less clear, gig platforms explicitly reward workers who provide services in response to platform demand. However, many gig workers also demand flexibility. Many digital gig economy participants use gig work as a supplement to another economic activity such as a primary job, household production, entrepreneurial activities, or education. These workers may bypass the most remunerative hours in order to work when their opportunity cost of working is low. Some of these workers demand or consume flexibility in that they wouldn't work for the platform if the platform didn't allow for at-will working hours. While the demographics of gig workers have been addressed in the literature, it is less well-understood which types of workers supply work very flexibly, how that ultimately impacts their remuneration for the work, and which workers would not be willing to supply labor if these platforms were less flexible.

In this paper, we use data from nearly two hundred thousand drivers on Uber (a popular ride-sharing platform), to examine the heterogeneity across gig workers and the demographics of gig work flexibility. We will demonstrate that there is substantial heterogeneity across workers. Some of these workers disproportionately supply flexibility, providing labor that is very responsive to demand (and pay opportunities) on the site. Some gig workers, due perhaps to the importance of their other commitments, value the flexibility of gig work and we estimate that they would likely not commit to supply labor if the gig opportunity were less flexible.

To examine the supply of flexibility, we straightforwardly examine the propensity of drivers to supply labor at high payout times. High payout times will tend to be those in which there are a lot of demanders relative to drivers in the marketplace; driver utilization is high and per-mile fees may be elevated. Drivers can provide flexibility by consistently working high earning hours, for example, every Friday after midnight, or by responding to a random shock, such as a large sporting event. In our analyses, we separately characterize drivers with a high propensity to drive during "regular" high earnings periods and drivers who drive during "idiosyncratic" high earnings periods.

We examine the demand for flexibility by examining driver surplus from flexibility. Following the approach developed in Chen et. al. (2019), we identify the taste for flexibility as being driven by (and equated with) time variation in a worker's reservation wage. If a worker had a constant reservation wage in all hours, the worker would be indifferent between a job that prescribed which specific hours the worker worked and a job that let the worker choose his or her hours, holding all else constant. This time variation in a worker's reservation wage can result from stable differences in the mean reservation wage across time periods– for example, a preference to not work late nights. Time variation can also derive from transitory shocks to reservation wages. For example, a parent may have a very high reservation wage on a day that a child is home sick.

Following Chen et. al. (2019). We use data from drivers' decisions of whether and when to supply labor on the Uber platform to estimate each driver's pattern of mean reservation wages for different time blocks and also estimate the variance of each driver's reservation wage due to shocks. This allows us to estimate driver surplus from driving for Uber and to estimate changes to the driver's labor supply and total surplus that would result from requiring the driver to instead work specific patterns of hours. Using new data provided by Uber Technologies, anonymized driver data is matched to driver demographic characteristics. These data allow us to identify characteristics of drivers that particularly value flexibility and for whom participation in the platform is dependent on that flexibility. It also allows us to identify characteristics of drivers whose participation is particular price-elastic.

We are interested in both part-time Uber drivers and full-time Uber drivers. As documented by Campbell (2017), most rideshare drivers obtain a minority of their household income from driving, suggesting that driving is often a secondary economic activity. Many part-time drivers who demonstrate high reservation wages during some hours presumably do so due to the time demands and remunerativeness of the the driver's other economic activities. It is unsurprising that some drivers would not drive if it couldn't be worked around the primary economic activity.

Because of this, it may be ex ante difficult to speculate which demographic groups would be most likely to value flexibility and which groups would be most likely to withdraw their labor supply from Uber were it to adopt a less-flexible scheduling regime. For example, while a common intuition is that higher income or wealthier people and women value job flexibility more than men, that is not necessarily true in the Uber environment, where many drivers schedule their work around other, likely less flexible, jobs.

Section 2 briefly reviews the literatures on job flexibility, dual job-holding, and on gender and demographic issues in the gig economy. Section 3 describes our data sources and construction of the analysis dataset. Section 4 provides a first look at the habits of Uber drivers of different demographics. Section 5 describes our analysis of supply of flexibility. Section 6 briefly reviews the labor supply model introduced in Chen (2019) and outlines how we conduct inferences for that model. Expected labor surplus, labor supply, and expected labor supply are discussed in Section 7. Section 8 provides a conclusion and summary of our findings.

2 Literature

Much of the literature on work flexibility centers on gender disparities in flexibility. (Bertrand et al. (2010) and Chen and Chevalier (2008)) suggest that part of the wage gap between men and women may stem from a willingness to provide flexible on-call work to employers and provide long hours of work.Goldin and Katz (2016) contrast the profession of pharmacist with the the professions described in (Bertrand et al. (2010). The analysis in Goldin and Katz (2016) suggests that the pharmacy profession exhibits narrow gender wage gaps while business professions contain wide wage gaps because the willingness to provide unexpected on-demand work and work long hours is not particularly valuable in pharmacy, in contrast to business and law.

A small but growing literature examines the issue or work flexibility and older workers. The share of workers who are self-employed rises with age. Survey data suggest that independent contractor opportunities are more common for the highly educated and that employment and self-employment would be larger for older workers if part-time and flexible work were readily available for less-educated workers (Abraham et al. (2021b)), (Ameriks et al. (2020)). However, as pointed out by Cook et al. (2019), older workers in traditional salaried employment may be benefiting from an implicit contract in which wages exceed the marginal product at the end of the year; gig workers are paid their marginal products making gig work potentially less attractive for older workers.

Uber data contain the minute by minute labor supply decisions of drivers. The relative remuneration of different times is directly and mechanically related to the demand for driver services at different times of day. This allows for precise measurement of the consequences for remuneration deriving from both the willingness to supply work flexibly and the unwillingness to supply work when the worker's time is demanded by other activities. While much of the literature focuses on differences between men and women in supply of and demand for flexibility, Uber data allows us to examine these differences across gender but also race, age, and a proxy for household income.

The exact magnitude of participation in the gig economy generally and online platforms specifically is poorly measured in core household surveys such as the Current Population Survey (CPS) (Abraham et al. (2021a)). Tax data suggest that, by 2016, about 2 millions Americans, or 1 percentage point of the workforce, had income from an online platform and that online platforms are an important driver of the increase in workers with self-employment income over the 2013-2015 period (Collins et al. (2019). However, Collins et al. (2019) show that workers in 2016 were no more likely to earn their livelihood through full-time gig work than they were a decade earlier.

As most rideshare drivers derive a minority of their overall earnings from driving, we consider driving to be closely linked to dual job holding. Substantial research suggests that multiple job holding has historically been limited to about 5 percent of the workforce, although it is more prevalent for workers in certain occupations (for example, Lale (2015) reports that multiple-jobholding rates for teachers are no less than 13 percent). While multiple-jobholding rates are low, a much larger number of workers transition in and out of multiple job holding over the lifecycle (see Paxson and Sicherman (1996), Renna and Oaxaca (2006), and Lale (2015)). Lale (2015) estimates that about 1 percent of full-time single jobholders and 2 percent of part-time single jobholders transition in to multiple jobholding each month. Uber drivers similarly have high churn. Cook et al. (2018) demonstrate that many workers drive for Uber only for a short time.The literature also suggests persistent geographic differences in dual job holding. Hirsch et. al. (2017) demonstrate that multiple job holding is weakly pro-cyclical, suggesting the importance of labor demand, and that it is negatively correlated with commuting times. This is at least suggestive that technologies that render secondary work more flexible may attract workers to secondary work.

Clearly, the flexibility of Uber is important for some drivers. 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 drivers report that Uber is a part-time activity secondary to more traditional employment. Their findings are consistent with the third party survey in Campbell (2018). Campbell (2018) finds that only about one-third of rideshare drivers report that the majority of their income derives from driving. Prior to the introduction of these platforms, there were clearly fewer opportunities to undertake secondary work that could be accommodated around the schedule of the primary work. Using data from individual bank and credit card accounts, Farrell and Greig (2016) present evidence that is strongly suggestive that workers supply more labor to online platforms such as Uber and Lyft when they receive negative shocks to their earnings in their other sources of employment.

Consistent with this, Hall and Kruger (2016) and Chen et. al. (2019) document that the hours supplied by drivers vary considerably from week to week. Chen et. al. (2019) examine drivers' labor supply in more detail. Because of the flexibility of the platform, a 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 in 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 extracted. Chen et. al. (2019) examine each driver's labor supply decisions and estimate driver response to alternative scenarios which mimic the effects of traditional employment relationships. They estimate substantial surplus from the flexible arrangement relative to a fixed hours arrangement for the group of workers who have opted in to driving.

There is a very small literature on driver demographics. Hall and Kruger (2016) document driver demographics from a survey. Cook et al. (2018) show that women Uber drivers earn somewhat lower wages per trip than do men. They demonstrate that this is largely due to women being in a lower position on the experience curve and due to women drivers driving more slowly. They also examine an issue closely related to our notion of supplying flexibility and show that, in four of the six cities they study, the average remuneration in hours of the week that women are overrepresented is not significantly lower than the hours of the week in which men are overrepresented. Caldwell and Oehlsen (2018) estimate Frisch labor supply elasticities for women vs. men and find that women's hours are more wage-elastic than are men's hours.Cook et al. (2019) demonstrate that older Uber drivers earn lower hourly compensation than younger ones, largely due to the times that they choose to drive. However, to our knowledge, there are not other papers that analyze the relationship between flexibility, demographics, and the decision to participate in gig work.

3 Data Sources and Construction of Analysis Datasets

Our data are provided by agreement with Uber. We use the same base data as in Chen et. al. (2019). For the period from September 2015 to April 2016, we start with the universe of all UberX driver-hours for drivers in the twenty cities in the United States with the most UberX trips. 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, where "active" is defined as having a passenger or being en route to a passenger. Note that a driver can meet our definition of active even in hours where a substantial amount of time is spent waiting for a passenger dispatch. We then measure the driver's discrete choice of being active in each of the 168 hour blocks. Thus for each driver, our measure of labor supply consists of a vector of 168 hours/week x 36 weeks of zeroes and ones. This is the variable we refer to when we consider the binary outcome of whether a driver "worked" in a given hour.

We also define the Uber "wage" faced by active drivers for each city for each hour. To calculate a prevailing wage for each city-hour, we calculate each driver's total earnings in that hour, divided by minutes worked, times sixty. This is averaged for all drivers in the UberX sample for the city-hour. Here, minutes worked is defined as the minutes a driver is available on the app, regardless of whether they are on a trip. This may be an overestimate, since drivers are free to do any non-work activities they want during idle time while the app is on. However, because we condition on drivers being active for at least 10 minutes each hour, we remove some cases of drivers ignoring the app, refusing to accept trips, or having the app on in remote locations. While individual driver's wages will vary from the city average, we treat the city average as what the driver can expect to make if she or he chooses to drive in a particular hour. It is important to remember that, on the Uber platform, drivers are expected to pay for both the capital costs of their vehicle and all costs of operating the vehicle. In our analysis, these costs are incorporated into the driver's reservation wages.

Because we will be evaluating patterns of activities over time, our analysis sample of drivers consists of drivers who are active in at least 1 hour for at least 16 of the 36 weeks that we have available in our data. We will refer to drivers who meet this criteria as "active drivers." This is an important filter. Our understanding is that many drivers try out driving for Uber but abandon the platform. Cook et al. (2018) find that 68 percent of Uber drivers who start driving for Uber have abandoned the platform after 6 months (though, because Uber drivers do not have to formally quit, it is possible that some are on an extended break). The platform was growing rapidly during the time of our data; our data requirements force us to oversample drivers who remain on the platform for a relatively long time. As mentioned above, we also restricted attention to the top 20 US cities by volume of labor supplied on the UberX platform. This gives us driver data for 196,198 drivers.

Uber internally, of course, maintains data on the identity of their drivers. While Uber did not share driver identities with the researchers, they provided anonymized information on driver demographics for our study. Using this information, we have information on driver age, driver gender, and estimates of driver race/ethnicity and driver neighborhood income. Driver age and gender come directly from Uber's internal records. Driver neighborhood income is estimated by matching the driver's address to the Census Bureau's geocode resource. This method allows matching to Census tracts for a total of 91% of drivers (mismatches may be due to street misspellings). This matching was done internally by Uber. The median household income for each driver Census tract was extracted using the 2013-2017 ACS summary file. Because of concerns from Uber that geographic sparsity could cause the exact median income figure to compromise driver privacy, the income estimates were further aggregated. Specifically, the Census tract median income across all drivers were grouped into buckets of 20 drivers and each driver was then assigned the average value of their bucket.

Uber does not collect data on driver race or ethnicity. In order to impute driver race and ethnicity, a procedure similar to that reported in Diamond et. al. (2019) was undertaken. Specifically, for each driver, the racial/ethnic composition of the driver's Census block is extracted from the 2010 Census. Driver full names were used to predict ethnicity/race using the python package ethnicolr. The package ethnicolr is built on a neural network applied to two character chunks in names that is trained on Florida Voting Registration data. It is described in detail in Sood and Laohaprapanon (2018). The driver race/ethnicity that we use in our summary statistics is the categorization of white, Black, hispanic, and asian. We assign the race/ethnicity with the highest Bayes' rule posterior given the driver's name and census block, where the name based predictions are treated as the prior. Due to the inability to match all driver addresses to Census files, plus missing gender data and age data for a small number of drivers, we are left with 178,401 drivers with complete data.

There are many ways in which this is an imperfect measure of race and ethnicity. For example, this methodology has no mechanism to capture multi-race individuals and treats Hispanic origin as mutually exclusive from Black or white. The category for Asian drivers is broad, including East Asian and South Asian and Middle-Eastern drivers. Finally, demographic inferences are an understandably sensitive issue (Andriotis and Ensign, 2015). Nonetheless, these techniques have been widely used to better understand ethnic inequalities by government agencies like the Department of Justice's Civil Rights Division, the Consumer Finance Protection Bureau, and the Office of Minority Health in the Department of Health and Human Services (CFPB (2014), Martino et al. (2013)). While demographic inference is imperfect, it is in some cases the best option, and can be used to increase understanding of important social issues. Finally, note that no method, including self-reporting, is perfect (Arday et al., 2000).

We briefly summarize the demographic characteristics of our driver sample.

Table 1 and Table 2 show some characteristics of our "dedicated driver" sample. Our overall female driver percentage of 15.2% is much lower than the 27% found by Cook et. al. (2019). However, their data runs into a later time period than does ours and they also show that women have a substantially higher attrition rate in the first 6 months of driving and our methodology requires driving at least 16 of 36 weeks. In contrast, the New York City Taxicab Fact Book (2014) reports that 1.1% of New York cab drivers are female. The share of drivers that are women are strikingly different across racial/ethnic

Category	Men	Women	% of row Women
Age<=60	135867	24998	15.6%
Age>60	15474	2062	11.8%
White	76808	12699	14.2%
Hispanic	30950	4920	13.7%
Black	24985	8340	25.0%
Asian	18598	1101	5.6%

Tab. 1: Number of drivers in our dataset of 178401 drivers with various characteristics. The percent female is the percent of the row that is female.

Category	Mean Income	% Women
Income Quintile 1	33730	18.3%
Income Quintile 2	49494	16.5%
Income Quintile 3	63305	15.2%
Income Quintile 4	80621	13.7%
Income Quintile 5	116616	12.1%

Tab. 2: Household income quintiles and share women by income quintile for our dataset of 178401 drivers.

groups and also differ markedly across the income quintiles. The share of drivers that are women is much higher among Black drivers. The share of drivers that are women decreases monotonically with income quintile.

Our sample of Uber drivers is about half non-Hispanic White. In comparison, the overall metropolitan population of the United States was approximately 58 percent non-Hispanic White in 2016 (see (Frey, 2017)). Relative to the overall US Metropolitan populations, Uber drivers are somewhat less likely to be non-Hispanic White, with the other racial and ethnic groups correspondingly over-sampled in the driver population. The Census tracts of the drivers are also surprisingly representative of the income distribution of the MSAs in which the drivers live. Coding each driver's Census tract's median income as a share of the overall MSA's median income shows that the median driver lives in a Census tract with a household income equal to 91 percent of the median household income in the overall MSA. Quintiles for this measure are shown in Table 3.

5%	25%	50%	75%	95%
0.44	0.68	0.91	1.20	1.76

Tab. 3: Quintiles of each driver's Census tract's median income as a share of the MSA's median income.

Finally, Figure 1 shows that Uber drivers are fairly similar to the US working population in age.

4 Model-free Evidence on Driver Demand for Flexibility for Demographic Groups

4.1 Uber Driver Labor Supply by Demographic Group

Chen et al. (2019) demonstrate the tremendous variation in Uber driver hours as well as the volatility across weeks in driving behavior for individual drivers. Here, we summarize differences across various demographic groups in driving behavior.

Table 4 examines weeks in which our sample drivers drive. We examine the distribution of average hours worked by various driving types.

Recall that we consider a driver active in any hour when she was active for at least 10 minutes, and we count how many of the 168 hour blocks in the week the driver was active. Figure 3 displays the

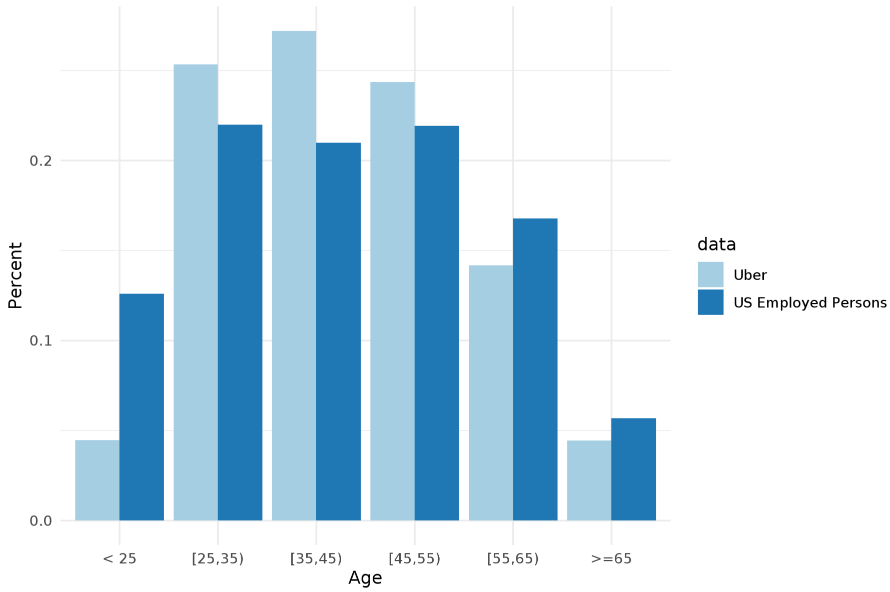

Fig. 1: Uber drivers' age compared to the age of employed persons from the 2015 CPS

	Share of drivers averaging N hours/week			
	Age<=60	Age<=60	Age<=60	Age>60
Total hours (N)	Female	Male	All	All
1-4	2%	1%	1%	1%
5-12	41%	28%	30%	21%
13-20	32%	28%	29%	27%
21-30	17%	22%	21%	54%
31-40	6%	12%	11%	15%
41+	2%	8%	7%	11%
Total drivers in category	24998	135867	160885	17536

Tab. 4: Distribution of average active hours using only weeks in which the driver works at least one hour

percentage of hours, averaged across those weeks that each driver was active. First, we see that, for younger drivers, women tend to work considerably fewer hours than men. Second, we see that drivers over the age of 60 work more hours in weeks that they work than drivers under the age of 60. This would be surprising in an environment in which most drivers were using Uber as a primary source of income, but unsurprising given the extent to which Uber is used as a secondary economic activity, perhaps especially among prime working age drivers.

We also examine differences in driving habits for drivers by race and income.

	Share of drivers averaging N hours/week				
Total Hours	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
1-4	1%	1%	1%	1%	2%
5-12	29%	29%	31%	31%	31%
13-20	30%	29%	29%	29%	29%
21-30	22%	22%	21%	21%	21%
31-40	11%	11%	11%	11%	11%
41+	7%	7%	7%	7%	7%
Total drivers in category	33021	32854	32438	31745	30807

	Share of drivers averaging N hours/week			
Total Hours	white	black	asian	hispanic
1-4	1%	1%	1%	1%
5-12	32%	31%	23%	30%
13-20	29%	30%	25%	30%
21-30	21%	20%	23%	22%
31-40	10%	10%	15%	11%
41+	7%	7%	12%	6%
Total drivers in category	78113	30573	18384	33816

Tab. 5: Distribution of average active hours using only weeks in which the driver works at least one hour

Table 5 shows very similar patterns of driving across racial groups, with the exception that a higher percentage of Asian drivers appear to use Uber as a more full-time activity than any other group. We also examine driving by quintile of census tract income. The propensity to drive many versus few hours is virtually identical across the five income quintiles.

4.2 Within-Driver Variation in Schedules

We examine the extent to which drivers of various demographics vary their schedules from week to week. Intuitively, the variation in driver schedules (combined with available wages and the economic incentives to work particular hours) will motivate our analysis of driver behavior.

To summarize the data, we divide the 168 hours of the week into 56 three-hour blocks ordered sequentially from the beginning of the week. We examine the question: 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 identify the extent to which week-to-week variability is due to sitting out the entire week. Next, we trace working in the same block across weeks, but condition on driving sometime in the relevant day. The results for men vs. women are shown in Table 6.

Table 6 shows that a male driver who works in a particular block has a roughly 53 percent chance of working in that same block on the following week. If the driver did not work in a particular time block in week t, he has only a 10 percent chance of working in it the following week. The probability that a driver who worked in a block in week t will work in it again in week t+1 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. This is true for both men and women, with the probability that a woman works the same block in the adjacent week being lower. Of course, the lower consistency of women's hours could be motivated by demand (women drivers being more likely to have unpredictable demands on their time) or supply (women drivers being more elastic in chasing high-wage opportunities on Uber). Our analysis will allow us to view these possibilities separately.

Table 7 shows this same transition comparison for older versus younger drivers. Here, we can see real differences in the propensity of drivers to drive at the same time from week to week. For example, drivers over age 60 who are working in a particular day are less likely to work in a block they didn't work in prior week than are younger drivers, but substantially more likely to work in a block that they did work in the week past.

This evidence of the volatility in driver hours presented above fundamentally does not allow us to disentangle two sources of week-to-week variation in hours worked for specific drivers. Drivers who drive less predictably could be supplying flexibility to Uber– that is, driving when expected payouts are high. However, they could also be taking advantage of the flexibility of Uber– driving when it is convenient for them. We pursue a model to allow us to disentangle these factors. Drivers of different demographics may be differentially unpredictable either because they systematically differ in their propensity to chase high wage opportunities or because their reservation wages are more volatile.

5 Suppliers of Flexibility

Drivers contribute to the Uber platform by providing labor. An additional driver in a local area is particularly valuable to riders on the platform when wait times for rides are high and/or when prices

		Male			Female		
		% working that block in week t+1			% working that block in week t+1		
Did a driver work	Yes	52.9	55.3	67.9	45.1	48.0	64.1
a block in week t?	No	10.0	12.4	21.2	8.5	10.9	21.0
Conditional on working		unconditional	in week t+1	that day in wk t+1	unconditional	in week t+1	that day in wk t+1

Tab. 6: Conditional Probabilities of Working a Block in Consecutive Weeks

		Age > 60 % working that block in week t+1			Age <= 60 % working that block in week t+1		
Did a driver work a block in week t?	Yes	58.0	60.1	72.8	51.4	53.8	66.9
	No	10.2	12.0	19.7	9.8	12.2	21.4
Conditional on working		unconditional	in week t+1	that day in wk t+1	unconditional	in week t+1	that day in wk t+1

Tab. 7: Conditional Probabilities of Working a Block in Consecutive Weeks

on the platform are high (when there is surge). In contrast, if at a particular time in a particular city, many drivers are waiting for riders, an additional driver is not particularly valuable to consumers. The average payout earned by drivers in a particular city-hour is a proxy for the value of an incremental driver in that hour. This is because driver payouts are high when drivers have a high utilization (defined as active time divided by time a driver is on the app) and when there is surge. We have already seen that driving patterns across drivers of different demographics tend to differ. This raises the question of whether drivers differ systematically in the value they provide to the platform by driving during high-value versus low-value periods. A driver that tends to drive when payouts are high is disproportionately supplying valuable flexibility to the Uber system.

We create several measures to capture the supply of flexibility. Consider a driver i who drives in week l in city c. Let H_l be the 168 hours of week l, and $N_{i,l,c} \subset H_l$ be the set of hours that i drove in city c during week l. So:

$$|N_{i,l,c}| \leqq 168$$

Those hours might be ones in which the driver is particularly valuable to the system or they may be ones that are convenient for the driver but not particularly valuable for the system. As a first step we calculate the payout that the i would have earned if they had earned the city average for each of the hours they actually drove. That is, for each hour h in week l, let $w_{h,l,c}$ be the average observed earnings of all drivers driving that hour of that week in that city. Define the total expected wages for driver i in week l as:

$$TW_{i,l,c} = \sum_{\tilde{h} \in N_{i,l,c}} w_{\tilde{h},l,c} \tag{1}$$

That is, $TW_{i,l,c}$ is the total wages driver i would have made had they earned average city-hourly wages for each of the hours they drove in week l. By replacing actual earnings with city-average earnings for observed hours, $TW_{i,l,c}$ removes the effect of the driver being particularly lucky or unlucky (or skilled or unskilled) relative to other drivers driving at identical times.

Now we ask, suppose driver i had driven the same number of hours in week l, but had chosen those hours in which average earnings in city c were highest; what could they have earned? Define the potential wage $\widehat{PW}_{i,l,c}$ to be:

$$\widehat{PW}_{i,l,c} = max_{H \subset H_l, |H| = |N_{i,l,c}|} \sum_{\tilde{h} \in H} w_{\tilde{h},l,c} \tag{2}$$

That is, $\widehat{PW}_{i,l,c}$ is the total earnings of the top wage-paying $|N_{i,l,c}|$ hours in week l.

The more able a driver is to concentrate their driving in the hours that are the most lucrative, the closer their total wages will be to their potential wages. Our first measure of the propensity of the driver i to supply flexibility in city c is, for a driver who drove in all 36 weeks of our sample, the average share of potential wages $\widehat{PW}_{i,l,c}$ earned by the driver across all 36 weeks in our sample. That is:

$$AvgShare_{i,c} = (\frac{1}{36})\sum_{\tilde{l}=1}^{36}\frac{TW_{i,\tilde{l},c}}{\widehat{PW}_{i,\tilde{l},c}} \tag{3}$$

We can further decompose the share of potential wages earned by the driver into two components, the propensity of the driver to drive in the hours in which an additional driver is *typically* very valuable, and the propensity to drive in hours which are *idiosyncratically* valuable. For example, there are hours, such as 5 a.m. on weekdays, which are typically very lucrative due to airport trips. There are other hours—such as 3 p.m. on a Saturday, which are not typically lucrative, but could be in a given week due to a sports event or concert. Uber often informs drivers of upcoming potential busy times and drivers make varying investments themselves in learning about these opportunities. We hypothesize that drivers in different demographic groups may be systematically different in their willingness to pay attention to and respond to these opportunities. We are interested in, and measure here, both the propensity to drive during "regular" lucrative hours, versus idiosyncratically lucrative hours.

To examine this, we decompose $AvgShare_{i,c}$ into two components: $AvgShareExp_{i,c}$ in which $w_{h,l,c}$ is replaced with the mean wage of each particular hour of week over our 36 week sample in that city, and $AvgShareIdio_{i,c}$ in which $w_{h,l,c}$ is replaced with its deviation from that overall mean. Thus, for example, for the 5pm Wednesday hour in Chicago for week 10 in our sample, there is a "regular" wage which is the average payout for the 5pm Wednesday hour in Chicago for all 36 weeks of our sample and there is an "idiosyncratic" component which is, for week 10, week 10's deviation from the overall sample. We decompose both total wages earned by the driver, $TW_{i,l,c}$ and the potential wages $\widehat{PW}_{i,l,c}$ into their regular and idiosyncratic components. That is, let:

$$\overline{TW}_{i,l,c} = \sum_{\tilde{h}\in N_{i,l,c}}\overline{w}_{\tilde{h},c} \tag{4}$$

and:

$$\widehat{\overline{PW}}_{i,l,c} = max_{H\subset H_l,|H|=|N_{i,l,c}|}\sum_{\tilde{h}\in H}\overline{w}_{\tilde{h},c} \tag{5}$$

and:

$$\widehat{PW_{i,l,c} - \overline{PW}}_{i,l,c} = max_{H\subset H_l,|H|=|N_{i,l,c}|}\sum_{\tilde{h}\in H}(w_{\tilde{h},c} - \overline{w}_{\tilde{h},c}) \tag{6}$$

and similarly:

$$\underbrace{PW_{i,l,c} - \overline{PW}}{}_{i,l,c} = min_{H\subset H_l,|H|=|N_{i,l,c}|}\sum_{\tilde{h}\in H}(w_{\tilde{h},c} - \overline{w}_{\tilde{h},c}) \tag{7}$$

This allows us to define:

$$AvgShareExp_{i,c} = (\frac{1}{36})\sum_{\tilde{l}=1}^{36}\frac{\overline{TW}_{i,\tilde{l},c}}{\widehat{\overline{PW}}_{i,\tilde{l},c}} \tag{8}$$

and also:

$$AvgShareIdio_{i,c} = (\frac{1}{36})\sum_{\tilde{l}=1}^{36} \frac{TW_{i,\tilde{l},c} - \overline{TW}_{i,\tilde{l},c} + (PW_{i,\tilde{l},c} - \overline{PW}_{i,\tilde{l},c})}{PW_{i,\tilde{l},c} - \overline{PW}_{i,\tilde{l},c} + (PW_{i,\tilde{l},c} - \overline{PW}_{i,\tilde{l},c})} \quad (9)$$

where:

$$\overline{w}_{\tilde{h},c} = \frac{1}{36}\sum_{\check{l}=1}^{36} w_{\tilde{h},\check{l},c}$$

For the $AvgShareIdio_{i,c}$ measure, we scale the measure to be between zero and one. To do this, as can be seen in Equation 9, we add the minimum idiosyncratic wage achievable in the week to both the numerator and denominator of the idiosyncratic share calculation. The measure is then naturally interpreted for each week as the share of the idiosyncratic wage achievable in the week that is captured by the driver.

We summarize these measures for various driver demographics.

The results show interesting demographic distinctions. For example, women earn a lower share of the possible wages than do men, both because they work less-lucrative regular hours but also because they don't work during more lucrative idiosyncratic hours. Older workers have a lower share of possible wages and this is driven entirely by their propensity to drive less lucrative regular hours. The differences across income quintiles may not appear large, but, we will see later, are quite statistically robust when controlling for other factors. While there is not a clear pattern across the income quintiles in the propensity to capture a large share of the potential wages, there is a monotonically increasing relationship moving through the income quintiles in the propensity to obtain high "idiosyncratic" wages. We interpret this as suggesting that workers from lower income deciles are particularly likely to respond to opportunities to earn above average wages such as concerts or sporting events. In this sense, workers from lower income deciles are particularly important in supplying flexibility to the system.

We use a regression framework to decompose these differences across driver demographics systematically. We simply regress each of the potential wage measures discussed above on age (entered as a continuous variable), an indicator for female, an indicator for each of the four lower income quintiles, with the highest income quintile as the excluded category, and an indicator for each of the three racial groups, with white as the excluded category. The results are presented in Table 10. The second column of each pair excludes the race variables, some of which are highly correlated with the income variables. These results suggest that lower income and non-white workers are particular suppliers of

Group	% Possible AvgShare	% Expected Possible AvgShareExp	%Idiosyncratic Possible AvgShareIdio
All	73.3%	77.8%	42.1%
Age<=60	73.4%	78.0%	42.1%
Age>60	72.2%	76.7%	42.3%
<=60, female	71.0%	75.8%	41.2%
<=60, male	73.9%	78.4%	42.2%

Tab. 8: Fraction of theoretical max wage by group

Group	% Possible AvgShare	% Expected Possible AvgShareExp	%Idiosyncratic Possible AvgShareIdio
M,<=60,Inc Quint 1	74.0%	78.3%	42.4%
M,<=60,Inc Quint 2	74.0%	78.4%	42.3%
M,<=60,Inc Quint 3	73.8%	78.3%	42.2%
M,<=60,Inc Quint 4	73.7%	78.4%	42.2%
M,<=60,Inc Quint 5	73.7%	78.6%	42.0%

Tab. 9: Fraction of theoretical max wage by group

flexibility to the Uber system. However, the decompositions reveal very different patterns for "regular" wages and "idiosyncratic" wage innovations. The monotonic negative relationship between income and the propensity to supply hours when Uber opportunities are idiosyncratically lucrative persists in the regression specifications although these same workers do not work particularly lucrative "regular" hours. We hypothesize that workers from lower income census tracts might be disproportionately likely to have jobs or other obligations that constrain them from working in the predictable high earnings period but that these workers from lower income census tracts are disproportionately likely to take advantage of high-earning opportunities that arise episodically on the platform.

6 Demand for flexibility: A Review of the Model of Inference of Demand for flexibility

Here, we briefly review the model in (Chen et al., 2019), from which our flexibility demand results are derived. A simple model of labor supply specifies that drivers will supply labor if their reservation wages are 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" in a given hour \as having at least 10 minutes of "active" time engaged in picking up a rider or on a trip. Expected wages are computed assuming 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 Section 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 driving decision is based on the difference between prevailing and reservation wages which does not depend on assumptions regarding the incorporation of operating costs.

6.1 A Model of Reservation wages and driving decisions

We start by providing a simple intuition of our identification strategy (taken from Chen et. al. (2019)). Consider a weekly one-hour period, say, Tuesday, 2 to 3 p.m. For concreteness, assume that the mean prevailing wage for that hour is \$20 in a particular city, and consider a driver who works that hour

VARIABLES	(1) % potential wage	(2) % potential wage	(3) % potential wage_regular	(4) % potential wage_regular	(5) % potential wage_idio	(6) % potential wage_idio
age	-0.0001***	-0.0002***	-0.0002***	-0.0002***	0.0002***	0.0002***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
female	-0.0272***	-0.0276***	-0.0255***	-0.0251***	-0.0089***	-0.0111***
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0003)	(0.0003)
incomeq1	0.0010**	0.0032***	-0.0039***	-0.0017***	0.0058***	0.0035***
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0004)
incomeq2	0.0020***	0.0025***	-0.0017***	-0.0015***	0.0045***	0.0033***
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0004)
incomeq3	0.0001	0.0000	-0.0024***	-0.0027***	0.0032***	0.0022***
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0004)
incomeq4	0.0002	-0.0001	-0.0017***	-0.0022***	0.0021***	0.0017***
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0004)
black	0.0175***		0.0214***		-0.0079***	
	(0.0005)		(0.0005)		(0.0004)	
hispanic	0.0022***		-0.0045***		0.0045***	
	(0.0005)		(0.0005)		(0.0004)	
asian	0.0332***		0.0311***		0.0143***	
	(0.0006)		(0.0006)		(0.0004)	
Constant	0.7338***	0.7429***	0.7855***	0.7930***	0.4098***	0.4135***
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0005)	(0.0005)
Observations	178,401	178,401	178,401	178,401	178,401	178,401
R-squared	0.0414	0.0226	0.0448	0.0201	0.0207	0.0091

Tab. 10: Regressions of potential wage measures on demographic variables. Standard errors in parentheses, *** p¡0.01, ** p¡0.05, * p¡0.1

most weeks. Our estimation would infer that the driver has a mean reservation wage for that hour that is less than \$20. Now, suppose that there are some slow weeks where the prevailing wage is around \$15 for that hour. If the driver drives most of those weeks too, that suggests that the driver has a mean reservation wage for the hour that is less than \$15, and thus, on the more typical \$20 weeks, she is getting at least \$5 in surplus. In contrast, if the driver does drive the \$20 weeks usually, but *doesn't* drive in the \$15 weeks, then our estimate of the mean reservation wage of the driver for that hour will be bounded between \$15 and \$20. This illustrates how the variation in the wage across weeks helps us to pinpoint the driver's reservation wage. With a lot of data, we'd be able to see the wage at which the driver "drops out" from working in the hour. For the driver who usually drives the Tuesday 2 to 3 p.m. hour when the prevailing wage is \$20, if the driver doesn't drive that hour in some of the \$20 weeks, given her other behavior, her not driving will have to be ascribed to some kind of shock. The extent to which it is attributed to a shock to her hour or day or week will largely be a function of whether the rest of her day/week are also outliers relative to her other behavior. The variance of the shocks experienced by the driver will be determined in part, loosely, by whether we sometimes observe the driver to not drive in that hour when it is *more* lucrative than a typical \$20 hour.

We now turn to a more specific description of our methods. The specification of the reservation wage process is crucial to determining the extent to which drivers are able to exercise flexibility in labor supply. As we have documented in section 4, Uber drivers have both predictable and unpredictable patterns of labor supply. 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.

$$w_{it}^{*} = \mu_i(t) + \varepsilon_{it} \tag{10}$$

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 by at least the beginning of each time period (hour) each Uber driver has realized the shock and therefore simply compares his or her reservation wage for the hour to the expected wage to make a labor supply decision.

While the reservation wage w_{it}^{*} is unobservable to the econometrician, both driver labor supply, y_{it}, and the expected wage, w_{it}, are observed. Driver labor supply, y_{it}, takes the value of one in any hour in which the driver works and zero in any hour in which the driver does not work. In an hour when the driver works, we can infer that the reservation wage is exceeded by the expected wage. 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) that 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) assuming driver preferences are stable and not allowing for trends or other time shifts. This implies that our mean function is a function only 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_don[1]: Monday-Sunday, 4 a.m. - 6 a.m.
9. Base: all remaining hours in the week[2]

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.

$$\varepsilon_t = v_w + v_d + v_h \tag{11}$$

In this model, each of the error components is *iid* normal[3] over its respective frequency with standard deviations, $\sigma_w, \sigma_d, \sigma_h$ respectively. "w" denotes weekly, "d" denotes daily, and "h" denotes hourly. Thus, each time period (an hour) sees a new realization of the hour shock, v_h, each day a new day shock, and each week a new week shock.

[1] Dead-of-night.

[2] 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.

[3] Normal error components imply that the reservation wage process is multivariate normal over the 168 hours that comprise one week. The assumption of normality allows us to specify a model in which the mean of reservation wages can be determined independently of the size or variability of the shocks or unpredictable component of reservation wages. One possible drawback to this assumption is that there is some probability that reservation wage realizations will be negative (this may be very small). Some might suggest modeling the log-reservation wages. While this certainly removes the possibility of negative reservation wages, this assumption creates other undesirable problems. If we assume log-normal reservation wages, then high mean reservation wages are also associated with high variances. This means that we cannot independently vary the degree to which drivers have unpredictable (large shock) patterns versus when they work on average. To take the example of someone with a high reservation wage during the day (due to another work opportunity), the log-normal model would also require that they be more unpredictable during the day then on weekends and evenings. We do not want to impose this sort of restriction on driver behavior.

Since each day within a week shares the common week shock and each hour within a day shares a common day shock, this creates the well-known variance components covariance structure that can exhibit very high correlation between periods within each broader timeframe. The error covariance matrix of the reservation wage shock in (11) is block diagonal across weeks.

Our focus will be on using our model to measure driver surplus both in the actual Uber labor arrangement as well as in various alternative scenarios. Uber driver surplus can derive from a variety of factors. First, some drivers will have low reservation wages overall and will derive surplus from the difference between those reservation wages and the prevailing hourly wage. For an extreme case, consider the lonely driver who enjoys driving and talking to customers. This driver is clearly not the marginal driver who sets the wage, and this inframarginal driver clearly earns surplus. Second, some drivers will have reservation wages that are systematically heterogeneous across the hour blocks, and the Uber structure allows the driver to drive only in the lower reservation wage hours. For example, a driver who always works a valuable noon to 8 p.m. job can systematically not work in those hours. This driver earns surplus by avoiding work in those hours but working in other hours when the primary job is unavailable. Third, some drivers will have significant variance in their reservation wages that differ from week to week and the Uber arrangement allows the driver to shift driving hours. For example, an actor can choose not to drive whenever he is called for an audition. Similarly,a retail worker can work when a shift has been cancelled.

Details of our estimation methodology are provided in (Chen et al., 2019) (2019).

Parameter results

Our results suggest that Uber drivers do not have homogeneous preferences for time of day and day of week. Figure 2 provides scatterplots of normalized mean reservation wage estimates. Recall that each driver has a separate, and possibly, unique mean reservation wage for all of the nine hour-blocks. For example, the four graphs of Figure 2 shows reservations wages for two time blocks for male vs female drivers and drivers over versus under 60 years old. The y-axis in each shows the mean reservation wage for the Friday-Saturday late night hour block relative to the base period (truncated to remove outliers). The x-axis shows the same measure for the Monday-Friday afternoon hour block. For bot hour blocks in all graphs, the mean reservation wages range from a large positive to large negative deviation from the base period estimates, suggesting that reservation wages for these hour blocks, even within driver demographic type, are very heterogeneous. In addition, for all groups, there is a negative correlation between preferences for the Monday-Friday Afternoon (horizontal axis) and the late night block. Drivers who like to drive one tend not to drive the other. Differences among the demographic groups are also apparent. The over-60 drivers tend to have lower reservation wages for the Monday-Friday afternoon time (graphs look shifted to the left relative to the younger group). While there are many women who have below-baseline reservation wages for the late night time periods, they are a smaller fraction of all women than of all men. Overall, 49 percent of prime age men have a preference for late night weekend driving relative to the baseline (have a negative coefficient for Fri-Sat late night) but only 22 percent of older women.

The parameter results estimated here play a role both in our estimates of driver's supply of flexibility and the driver's demand for flexibility from the Uber platform. Drivers supply flexibility if they drive a times when wages are high. We will consider both hours when wages are typically high (such as late night) but also times when wages are idiosyncratically high (when there are concerts or other events).

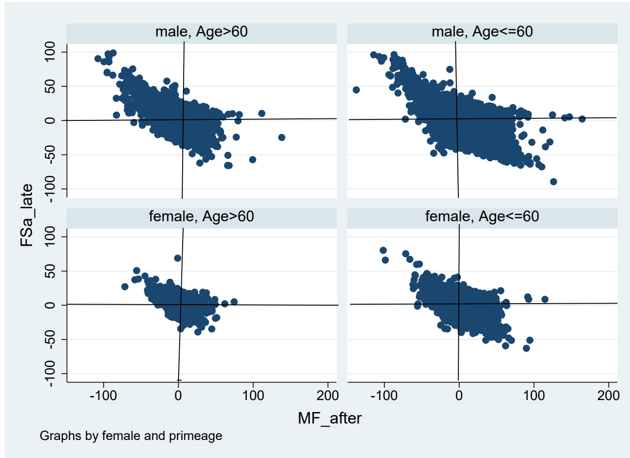

Fig. 2: Scatterplots of Mean Reservation Wage Parameters

Drivers who have a large distaste for particular time are not willing to drive that time even when it is lucrative.

7 Driver Surplus

The parameters of our model allow us to calculate driver surpus as the difference between the reservation wage and the expected wage in any hour in which the driver is driving. We also can also calculate this surplus for alternative arrangements that afford less flexibility and calculate whether or not a driver would be expected to drive a given hour, as the driver should drive only if the surplus from doing so is positive.

7.1 Surplus Measure

As in (Chen et al., 2019)), our goal is to compute the expected surplus for each driver. In our model, drivers will work only 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 in which a driver faces wage w_t, expected surplus can written as

$$ES_{i,t} = \left[w_t - E\left[w^*_{i,t}|w^*_{i,t} < w_t\right]\right] \times Pr\left[w^*_{i,t} < w_t\right] \tag{12}$$

To produce a surplus measure for each driver, we sum 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. We can gauge the impact of various flexibility restrictions on driver labor supply and the distribution of this surplus across drivers.

7.2 Constraints on Flexibility

We start with the base case, in which the Uber system imposes no constraints on labor supply flexibility. We will compare the expected surplus under this flexible system with the surplus under an alternative system in which the driver cannot adjust hours in response to the driver's hourly or daily shocks. Importantly, in these alternative scenarios, we are not examining a new equilibrium in which the system changes and the set of drivers on Uber changes, wages change, etc. Instead, our exercise is more modest. We consider what would happen to the surplus of each individual driver if the wages, etc. facing the driver remained constant, but the driver's ability to respond to daily or hourly shocks was eliminated. We can think of the difference between the unconstrained surplus and these hypothetical constrained surpluses as (some) of the value that the driver gets from the Uber system. This informs our understanding of the driver's demand for flexibility. Thus, we define our scenarios as:

(Base) Drivers can adapt to weekly, 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.

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 interpretation is the driver must make a decision about which hours she'll work at the beginning of each 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 shocks.

(B) *Cannot adapt to daily and hourly shocks.* Here, we do not allow the driver to adapt to daily or hourly shocks. 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 restrictions only 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 drivers have perfect foresight as to the prevailing wage. We will examine the driver's sensitivity to the prevailing wage. We also still allow the driver to have a driver-specific profile of mean reservation wages that can vary by day of week and hour of day. That is (A) and (B) are still much more flexible than most conventional work arrangements.

Details of methods for calculating the expected surplus under each of these scenarios is given in Chen et. al. (2019)

7.3 Expected Surplus and Labor Supply Computations

For each of the drivers, we compute Bayes estimates of the mean reservation wage parameter and Bayes estimates of each of the variance components necessary for the expected labor supply and expected surplus computations.

Surplus		Expected Surplus		
Group	Actual	Base	A	B
All	Surplus	202	84	41
	% of Base		41.6%	20.3%
Age<=60	Surplus	199	82	40
	% of Base		41.1%	20.0%
Age>60	Surplus	230	106	55
	% of Base		46.0%	23.9%
<=60, female	Surplus	147	53	20
	% of Base		36.3%	13.9%
<=60, male	Surplus	208	87	43
	% of Base		41.8%	20.8%

Tab. 11: Surplus under alternative scenarios by group

Surplus		Expected Surplus		
Group	Actual	Base	A	B
M,<=60,Black	Surplus	217	86	41
	% of Base		39.5%	18.8%
M, <=60,White	Surplus	202	84	41
	% of Base		41.8%	20.5%
M, <=60,Asian	Surplus	235	106	60
	% of Base		45.1%	25.6%
M,<=60,Hisp	Surplus	200	82	39
	% of Base		41.2%	19.6%

Tab. 12: Surplus under alternative scenarios by group

Surplus		Expected Surplus		
Group	Actual	Base	A	B
M,<=60,Inc Qunt 1	Surplus	210	84	41
	% Base		40%	19.5%
M,<=60,Inc Qunt 2	Surplus	212	88	43
	% Base		41.5%	20.3%
M,<=60,Inc Qunt 3	Surplus	207	86	43
	% Base		41.5%	20.7%
M,<=60,Inc Qunt 4	Surplus	208	88	44
	% Base		42.3%	21.2%
M,<=60, Inc Qunt 5	Surplus	206	88	44
	% Base		42.7%	21.4%

Tab. 13: Surplus under alternative scenarios by group

We start with some summary statistics on surplus for different demographic groups. The goal of this exercise is to estimate what demographic groups are those that take advantage of and value the flexibility of the platform.

The tables illustrate that, for all drivers, surplus declines precipitously in scenarios in which we disallow the driver to increase or decrease their labor supply in response to idiosyncratic shocks of the hourly or weekly frequency. The loss of surplus appears to be somewhat larger for women relative to men, for younger people relative to older people, and for lower income groups relative to higher ones.That is, these groups appear to disproportionately demand the type of flexibility available on the gig platform.

We can also examine the surplus from flexibility by demographic group using a regression framework. The observations are individual drivers and the right hand side variables are the same as in our potential wage specifications above. The left hand side variable in each specification is the share of surplus estimated to be retained by the driver in scenario A, where the driver is constrained from optimizing their driving in response to hourly shocks in scenario B, where the driver is constrained from optimizing their driving in response to hourly and daily shocks. That is, the left hand side variables take values between 0 and 1 and are computed from taking the A or B surplus divided by the base surplus for each driver. Thus, positive coefficients imply that the demographic type values flexibility *less*.

The results in Table 14 are similar to the univariate results. It is important to note that these results do not have a specific hard-wired relationship to the supply of flexibility results that we presented before. For example, we see that women lose a lot of surplus in the constrained scenarios and are thus demanders of flexibility. Our results above regarding potential wage suggest that women are also, on average, not suppliers of flexibility. In contrast, we showed above that the lowest wage quintile is, on average, an important supplier of flexibility. Members of that quintile capture a statistically larger fraction of potential wages, driven by their propensity to drive during idiosyncratically high-demand periods. However, this group is revealed in these specifications to also be statistically significant demanders of flexibility in that their surplus falls substantially in the more constrained alternative scenarios. Older workers demand less flexibility. This accords with our finding that they tend to work more predictable hours and perhaps with the intuition that older workers are less likely to be working Uber around childcare or other employment.

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 the U.S. UberX service. One of the attractions of Uber is the flexibility afforded to drivers. However, a characteristic of Uber is that drivers who can respond to incentives on the system ("supply flexibility") will find driving more remunerative than those who cannot, and drivers who can respond to incentives on the system are also very valuable to riders. In this paper we examine the demographics of supplying flexibility as well as the demographics of demanding flexibility.

We see some patterns that are perhaps expected and intuitive. For example, older drivers appear not to value the ability to rearrange their schedules as much as do younger drivers. This points to the possibility that primary work and family obligations of younger drivers likely create a demand for flexibility. Similarly, we see that women appear to have a greater demand for flexibility than do men

VARIABLES	(1) A_frac_ES	(2) A_frac_ES	(3) B_frac_ES	(4) B_frac_ES
age	0.0016***	0.0015***	0.0013***	0.0013***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)
female	-0.0433***	-0.0495***	-0.0471***	-0.0540***
	(0.0012)	(0.0012)	(0.0010)	(0.0010)
incomeq1	-0.0040***	-0.0110***	0.0060***	-0.0033***
	(0.0015)	(0.0014)	(0.0013)	(0.0012)
incomeq2	0.0010	-0.0029**	0.0064***	0.0005
	(0.0014)	(0.0014)	(0.0012)	(0.0012)
incomeq3	-0.0002	-0.0033**	0.0047***	0.0001
	(0.0014)	(0.0014)	(0.0012)	(0.0012)
incomeq4	0.0007	-0.0007	0.0043***	0.0019*
	(0.0014)	(0.0014)	(0.0012)	(0.0012)
black	-0.0196***		-0.0182***	
	(0.0015)		(0.0013)	
hispanic	0.0106***		0.0023*	
	(0.0014)		(0.0012)	
asian	0.0484***		0.0635***	
	(0.0017)		(0.0015)	
Constant	0.2846***	0.2964***	0.0851***	0.0987***
	(0.0020)	(0.0019)	(0.0017)	(0.0016)
Observations	178,401	178,401	178,401	178,401
R-squared	0.0273	0.0205	0.0379	0.0246

Tab. 14: Regressions of alternative scenario surplus as a share of total surplus measures on demographic variables. Standard errors in parentheses, *** p¡0.01, ** p¡0.05, * p¡0.1

drivers. In this environment, the relationship between remuneration and the supply of flexibility is straightforward. The patterns that we document have been hypothesized to play an important role in labor markets more generally.

The patterns across income groups are particularly interesting. We see that those who live in the lowest-income census tracts are somewhat more likely to demand flexibility (they lose more surplus in the alternative less-flexible scenarios). However, these workers are important suppliers of flexibility to the system. That is, these workers are disproportionately represented in the provision of labor in response to "idiosyncratic" high-earnings opportunities on the system.

From tax data. survey data and the low mean hours supplied by Uber drivers, it is clear that most Uber drivers use the platform as a secondary economic activity. We interpret our results against the backdrop of the persistently low rate of true dual job-holding in the United States. The rapid uptake by drivers of the gig platform, even amongst drivers who reside in high-income Census tracts, points to a latent demand for secondary work that can be undertaken relatively flexibly. While we highlight the differences in usage of the platform across demographic groups, our results also demonstrate substantial heterogeneity within demographic groups, suggesting opportunities for further study of the circumstances that lead individuals to use gig platforms.

References

Katharine G. Abraham, John C. Haltiwanger, Claire Hou, Kristin Sandusky, and James R. Spletzer. Reconciling survey and administrative measures of self-employment. *Journal of Labor Economics*, 39(4):825–860, 2021a. doi: 10.1086/712187. URL https://doi.org/10.1086/712187.

Katharine G Abraham, Brad Hershbein, and Susan N Houseman. Contract work at older ages. *Journal of Pension Economics & Finance*, 20(3):426–447, 2021b.

John Ameriks, Joseph Briggs, Andrew Caplin, Minjoon Lee, Matthew D Shapiro, and Christopher Tonetti. Older americans would work longer if jobs were flexible. *American Economic Journal: Macroeconomics*, 12(1):174–209, 2020.

AnnaMaria Andriotis and Rachel Louise Ensign. U.s. government uses race test for $80 million in payments. *The Wall Street Journal*, 2015.

Susan Arday, David Arday, Stephanie Monroe, and Jianyi Zhang. Hcfa's racial and ethnic data: Current accuracy and recent improvements. *Health Care Financing Review*, 21(4):107–116, 2000.

Marianne Bertrand, Claudia Goldin, and Lawrence F Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American economic journal: applied economics*, 2(3):228–55, 2010.

Sydnee Caldwell and Emily Oehlsen. Monopsony and the gender wage gap: Experimental evidence from the gig economy. *Massachusetts Institute of Technology Working Paper*, 2018.

Harry Campbell. The rideshare guy 2018 reader survey. Technical report, Therideshareguy.com, 2018.

CFPB. Using publicly available information to proxy for unidentified race and ethnicity: A method and assessment. Technical report, Washington D.C., 2014.

M Keith Chen and Judith A Chevalier. The taste for leisure, career choice, and the returns to education. *Economics Letters*, 99(2):353–356, 2008.

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

Brett Collins, Andrew Garin, Emilie Jackson, Dmitri Koustas, and Mark Payne. Is gig work replacing traditional employment? evidence from two decades of tax returns. *Unpublished paper, IRS SOI Joint Statistical Research Program*, 2019.

Cody Cook, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. Technical report, Stanford University, 2018.

Cody Cook, Rebecca Diamond, and Paul Oyer. Older workers and the gig economy. In *AEA Papers and Proceedings*, volume 109, pages 372–76, 2019.

Diana Farrell and Fiona Greig. Paychecks, paydays, and the online platform economy: Big data on income volatility. *JP Morgan Chase Institute*, 2016.

William H. Frey. Brookings, the avenue: Census shows non-metropolitan america is whiter, getting older, and losing population, 2017. URL https://www.brookings.edu/blog/the-avenue/2017/06/27/census-shows-nonmetropolitan-america-is-whiter-getting-older-and-losing-population/.

Jocelyn Frye. On the frontlines at work and at home: The disproportionate economic effects of the coronavirus pandemic on women of color. *Center for American Progress. Available online: https://www. americanprogress. org/issues/women/reports/2020/04/23/483846/frontlines-work-home/(accessed on 15 May 2020)*, 2020.

Claudia Goldin and Lawrence F Katz. A most egalitarian profession: pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics*, 34(3):705–746, 2016.

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.

Misty L. Hegeness. Why is mommy so stressed? estimating the immediate impact of the covid-19 shock on parental attachment to the labor market and the double-bind of mothers. Technical report, Minneapolis, 2020.

Kendall Houghton. Childcare and the new part-time: Gender gaps in long-hour professions. 2020.

Etiennel Lale. Multiple jobholding over the past two decades. Technical report, Bureau of Labor Statistics, 2015.

Steven Martino, Robin Weinick, David Kanouse, Julie Brown, Amelia Haviland, Elizabeth Goldstein, John Adams, Katrin Hambarsoomian, David Klein, and Marc Elliott. Reporting cahps and hedis data by race/ethnicity for medicare beneficiaries. *Health Services Research*, 48(2):417–434, 2013.

Christina H Paxson and Nachum Sicherman. The dynamics of dual job holding and job mobility. *Journal of labor economics*, 14(3):357–393, 1996.

Francesco Renna and Ronald Oaxaca. The economics of dual job holding: A job portfolio model of labor supply. 2006.