Older Workers and the Gig Economy

By Cody Cook, Rebecca Diamond, and Paul Oyer*

The growth of the "Gig Economy" and the aging of the labor force are two trends that have had large effects on the labor market and may have even larger effects over the next few decades. The popular press has connected these two trends with numerous articles suggesting that older workers can continue to be involved in the labor market through freelance work.[1]

Given the serious demographic challenges pending in most developed countries, keeping older people working longer seems likely (along with increasing female labor force participation, immigration, and accelerating automation) to be an important part of maintaining a healthy economy. The Gig Economy is a promising way to increase labor supply of older workers and allow them to ease into retirement where they can choose hours and intensity of work that fit their needs and capabilities.

However, there is a critical difference between the Gig Economy and the traditional labor market: older workers in W-2 employment relationships are often reaping the benefits of the latter end of an implicit contract with an increasing age/earnings profile (as in Lazear (1979)) while Gig Economy workers are, in equilibrium, paid their marginal product in a spot labor market.

Looking at all workers and then focusing on the transportation sector, we empirically verify that age/earnings profiles are quite different between one large Gig Economy platform and traditional employment. We use data from the March Current Population Survey (CPS) to show that, for the broad working population, average hourly earnings increase steadily for about twenty years from labor market entry and then flatten out for the rest of careers (as has been shown by Murphy and Welch (1990) and Murphy and Welch (1992)). We show that a very similar pattern holds for transportation workers and for taxi drivers. For all these groups of workers, hourly earnings climb steadily for workers as they age from 21 to their early forties.

We then use data from Uber, the largest rideshare platform in the world. Uber's driver-partners have total flexibility as to the hours that they work, which may be an attractive feature for many older workers. Uber driving is a narrowly defined and homogeneous job that does not change in any fundamental way as long as the person continues to do it. We find that driver hourly earnings have very little relationship to age for drivers in their twenties and thirties. However, driver earnings decrease steeply and steadily as a function of age for drivers about forty or older. Drivers who are 60, for example, earn almost 10% less per hour than drivers who are age 30.

Using detailed trip-level data for Chicago, we are able to explain almost all of the Uber age/earnings relationship. Most of the decline in earnings with age are due to the fact that older drivers drive in different places (less congested areas and more in suburbs than in city center) and at different times. Outlying areas have less constant demand, so drivers spend more idle time and benefit less from surge pricing.

Moving to the Gig Economy can be a valuable way for older workers to continue earning money and to capture the value of

* Cook: Stanford University Graduate School of Business, 655 Knight Way, Stanford, CA 94305, codycook@stanford.edu. Diamond: Stanford University Graduate School of Business and NBER, 655 Knight Way, Stanford, CA 94305, diamondr@stanford.edu. Oyer: Stanford University Graduate School of Business and NBER, 655 Knight Way, Stanford, CA 94305, pauloyer@stanford.edu. We thank Michael Amodeo, Jonathan Hall, Libby Mishkin for comments.
The views expressed here are those of the authors and do not necessarily reflect those of Uber Technologies, Inc. Cook is a former Uber employee and retains equity in the company.

[1]See, for example, Madden (2018) and Shrikant (2018).

highly flexible work (Chen et al. (2019)). But some of the benefits of Uber driving (and likely Gig work more generally) are offset by loss of the value of human capital developed previously and by an age-related productivity disadvantage.

I. Data Sources

We use two primary sources of data. From the March CPS for 2017 and 2018, we gather information about each respondent's labor market outcomes for the calendar year 2016 or 2017. We follow the basic procedure in Murphy and Welch (1990), though we use looser limitations on hours and weeks worked to capture more workers that value the flexibility of contract work. We limit the analysis to non-student, non-military men who worked at least 20 weeks and averaged at least 10 hours per week when working in the previous year. To more closely mirror ridesharing, we depart from Murphy and Welch (1990) by keeping part-time workers, not imposing an earnings minimum (other than that earnings must be positive), and dropping all people under 21. We form "transportation" and "taxi" samples based on Census occupation codes and have samples of 77,680, 5,003, and 1,744 for our total, transportation, and taxi samples, respectively.[2] Our regressions use ASEC sampling weights.

Our second data source is from Uber and draws from the set of all U.S. drivers for the years 2016 and 2017. To mimic the CPS data as closely as possible, we include only male drivers who work at least 20 weeks in a given year and average at least 10 hours per week on the platform. The 20 week criterion excludes a large share of the driver population given drivers exit the platform at a high rate (see Cook et al. (2018)). However, a high fraction of Uber rides are done by the highly attached drivers in our sample (54% of all Uber rides and 65% of rides with a male driver). Our sample includes 292,514 drivers and 368,358 driver-years. Using data on driver earnings and hours worked (which we define as hours during which they have the Uber app in operation), we calculate average hourly earnings for each driver-year. Driver net earnings are less than the gross earnings figures we use which include Uber's share of ride revenue, gas, and the depreciation and maintenance due to Uber mileage. However, the net/gross distinction should not materially affect the age/earnings relationship.[3]

Figure 1 shows the age distribution of the entire CPS sample, the CPS transportation sample, and the Uber sample. All groups get much smaller after about age sixty. The full CPS and Uber samples are remarkably similar in their age distribution while the entire transportation sample is somewhat older. The Uber age distribution makes it clear that Uber is not currently being used as a substantial source of retirement income. Few drivers are over sixty and there is not a more sizable share of drivers on Uber who have reached traditional retirement ages than the share of all workers of that age.

II. Age-Earnings Profiles

For both the CPS and Uber samples, we run regressions where the dependent variable is log of average hourly earnings for the year and the key explanatory variables are a quartic in age. In the CPS sample regressions, we interact the age variables with dummy variables for working in transportation and the taxi industries. We control for metropolitan area (or Uber "city") and year.

Figure 2 graphically captures the age-earnings profile from the CPS and Uber regressions. It shows how log hourly earnings change from a base of age 21. The pattern for all CPS groups is generally quite similar in that earnings rise steadily from age 21 to about age 40 and then are essentially flat

[2] The taxi sample is largely made up of independent contractors (which is also the status of the Uber driver sample) while we expect the vast majority of the other CPS samples to be "W-2" employees

[3] As we show below, older workers drive in less congested areas at higher speeds which may have minor effects on gasoline consumption and depreciation. Older drivers also operate more hours per week which means they have more incentive to invest in fuel-efficient vehicles.

from age 40 to age 70. Though the shapes of the age/earnings profiles are similar, the growth with age varies. The peak at age 40 is about 120% higher than the earnings at age 21 for the full CPS sample, 80% higher for transportation workers, and 65% higher for taxi workers. This suggests that work experience, while valuable for all groups, is slightly less valuable for transportation employees (and especially taxi drivers) than for the average worker.

The age/earnings profiles for drivers on Uber are dramatically different from the CPS samples. Uber earnings are increasing, though very slightly, in age for drivers in their 20's. Then hourly earnings drop steadily with age such that sixty-year-old drivers earn about 10% less than thirty-year-old drivers.

Table 1 shows the summary statistics for the Uber sample split into younger (under fifty) and older (fifty and over) drivers. The table shows that older drivers earn 24% more in weekly earnings than younger workers but the difference is due to older workers working 40% more hours per week. Younger drivers earn a premium of $1.50 (or a little over 8%) relative to older drivers on an hourly basis.

Figure 2 shows that workers that want to transition from traditional employment to Gig work at retirement ages face a challenge in that, at least for drivers, age is detrimental to earning power. In addition to losing whatever compensation benefits workers may have accumulated in their prior jobs, they will be starting from a lower base relative to younger drivers doing the same job. Overall, the figure shows that the earnings profiles in the traditional and Gig economies (at least in the case of Uber) make it challenging for retiring workers to replace a substantial share of their prior income doing Gig work.

III. Explaining the Age Earnings Relationship

Why are earnings higher for younger drivers than for those who are fifty and above? Identifying the mechanisms behind the age/earnings relationship can potentially provide insight into how productivity of workers more generally varies with age and, as a result, how we might expect semi-retirement Gig work to pay off for a broader population.

It is reasonable to interpret the earnings differentials by age as reflecting productivity or marginal product of labor, given that drivers are largely paid a flat share of the revenues that they generate on the platform. There are several reasons older workers could be less productive in this setting. As we describe in earlier work (Cook et al. (2018), Uber earnings are formulaic and driver earning variation reflects differences across time and place in the parameters that comprise the earnings formula. For example, earnings vary with a "surge multiplier" that responds to supply and demand conditions in a given location at a given time. Even at times with no surge, earnings vary with supply and demand because this leads to variation in idle time (during which driver do not earn money). Also, driving faster generates more trips per hour, which increases earnings.

In our earlier work, we showed that female drivers make about 7% less per hour than male drivers and that this can be entirely explained by the facts that, on average, men drive in more lucrative areas, they drive faster, and they have more experience on the platform (which pays off through learning-by-doing). The earnings differences between drivers around thirty years old and those around sixty years old are even greater than the male/female difference. We now consider what explains the age/earnings relationship for drivers on Uber and what it tells us about earnings for older workers in the Gig Economy more broadly.

We follow the logic used by Cook et al. (2018) to compare drivers by gender to determine what factors explain why older drivers earn less than younger drivers. To look carefully at this in a way that allows us to control for location and other localized factors, we concentrate on the Chicago area. We use trip-level data on drivers to build a driver/hour dataset similar to the one used in Cook et al. (2018). See that paper for further details. The only differences in the data we use here are that we look only at men and we do not use 2015 data here. Unlike for the dataset used in Figure 2, where we wanted to compare Uber drivers to CPS respondents, we do not restrict by the hours or weeks worked in a year.

As detailed in Cook et al. (2018), the hourly earnings of a driver on Uber can be described by six underlying parameters – wait time, distance to pick up passengers, distance on trips, speed, surge multiplier, and "incentive" payments earned.[4] In Table 2, we show the average of each of these parameters (as well as the total earnings per trip) for drivers under fifty and those fifty and over.

The table indicates that younger drivers dominate (that is, the difference is in favor of them earning more) four out these six factors. They wait almost a full minute (13%) less for each ride, are closer to their passenger when they accept the ride, have a higher average surge multiplier, and earn higher incentive pay.

Older drivers go at a higher average speed. Holding other things constant, that leads to higher earnings for drivers. However, the reason older drivers go faster on average is that they tend to drive in less crowded (and, therefore, often less lucrative) areas. They also have longer trips, on average, reflecting the fact that they are more likely to drive in outlying areas than in central Chicago.

We regress log hourly earnings on an indicator variable for being fifty or older, adding controls to determine which factors lead to the baseline differences in earnings for older and younger workers. The results of these regressions are in Table 3. Column 1 of the table shows that, when we control only for the week, drivers fifty and over earn about 8% less than those under fifty.[5]

In Column 2, we introduce a set of fifty indicator variables for "geohashes" (each approximately three miles by three miles) that comprise about 90% of pickup locations for Chicago-area Uber rides. The coefficient shows that more than a third of the 8% differential between ages can be explained by where drivers work. Figure 4 graphically shows this age/geography relationship. Greener areas have a higher fraction of older drivers while red indicates younger drivers. The youngest areas are those in downtown where traffic is greatest. This is also an area with high surge rates and short wait times between rides.

Column 3 adds a full set of indicator variables for all 168 hours in a week interacted with the calendar week and geographies worked. These controls also drop the coefficient on fifty-plus substantially, indicating older workers drive at less lucrative times. Figure 5 shows more detail on how time of day and week driving choices differ with age. Older drivers are relatively likely to drive during daylight hours on weekdays. They are much less likely than younger drivers to drive in the evening and especially on Friday and Saturday nights. As a result, they miss out on some high demand hours. Overall, Figures 4 and 5 and Columns 2 and 3 of Table 3 show that older drivers make different choices than younger drivers about where and when to drive, choosing to operate disproportionately in the outlying parts of Chicago and in the suburbs and avoiding high demand times. As a result, these drivers have more idle (unpaid) time and lower surge rates.

Column 4 shows that controlling for driving speed and for experience driving on the Uber platform (a series of dummy variables for accumulated trips) has little effect on the age coefficient. This stands in sharp contrast to gender earnings differentials as Cook et al. (2018) showed that experience and driving speed explain about 80% of the gender earnings gap for a similar group of drivers.

Figure 3 shows how predicted Uber earnings vary with age based on regressions

[4]Incentive payments are primarily derived from Uber promising drivers they will earn a certain amount if they do some specific number of rides over a period of a few days. The goals are set based on drivers' past driving intensity, so are roughly equally attainable for all drivers.

[5]Throughout our discussion of our results when looking at Uber data, we do not mention standard errors as all our estimates are extremely precise.

analogous to those in Table 3 except that coefficients for age and age-squared replace the coefficient for fifty-plus. The figure shows that the decline in earnings, both with and without controls, is slow and steady from age thirty to age seventy.

We experimented with other specifications that interact some of the variables, that include driver and passenger cancellations, driver work intensity (number of hours per week), and other variables we consider in Cook et al. (2018). However, none had an economically meaningful effect on the results and the older driver coefficient remained at about -2%. Our conjecture is that the remaining earnings differential is due to some combination of our inability to fully capture all supply and demand variation that affects idle time and the fact that older drivers are likely to be somewhat less adept at using the app and getting passengers in and out of the car quickly.

Overall, the regression results show that older drivers earn less than younger drivers and that a large part of this earnings differential can be explained by differences in where drivers choose to work. Older drivers display a preference for avoiding denser areas which could reflect a relatively high cost of traffic and/or that they live in less central areas.[6] While it is unclear if these factors would lead to older workers being less productive in other Gig economy situations, our results establish that younger workers have an earnings advantage in the largest independent worker platform. This disadvantage is substantial (8-10% per hour) at an absolute level. The differential becomes extremely large (on the order of at least 50%) when comparing the earnings differentials of, for example, a thirty-year-old to a sixty-five-year-old driving for Uber compared to people of these ages doing other jobs in the economy.

We should add two important caveats. First, Uber will, at least at this point in its history, naturally have a different age/earnings profile than other jobs because the job of rideshare driver has only existed for a few years. It's possible some of the age/earnings relationship will change as the business matures. This does not really affect the interpretation of our results, however, because people who use Uber as a means of earning money after leaving the traditional workforce will, as a result, be new to rideshare driving. Second, older people who drive for Uber are obviously not a random sample. Perhaps relatively low productivity people are more likely to become Uber drivers in retirement. Though we have no reason to believe that is the case, it is a further reason to pause before applying our results to other jobs.

[6] Cook et al. (2018) show that drivers tend to work close to home.

IV. Conclusion

Using data from Uber, we have shown that semi-retirement to the Gig Economy will put older workers in a new labor market where they are at a disadvantage. Whereas earnings for people in traditional jobs (broadly, in the transportation sector, or specifically focusing on taxi drivers) increase steeply with age from twenty to forty and hold steady thereafter, Uber earnings are essentially flat from age twenty to forty and steadily declining in age thereafter.

Though some portion of our findings may be specific to the nature of Uber, our results suggest that the Gig Economy's compensation-based-on-productivity nature can pose a challenge for older workers — especially those who benefited from increasing age/earnings profiles due to implicit contracts in traditional jobs. Rideshare is a substantial portion of the Gig Economy so our results are important regardless of their external validity to other Gig settings. More research is needed to understand how broadly our results apply. Other segments of the Gig Economy might have less stark earnings decreases with age if, for example, age and experience are more valuable in higher-skill freelancing that is done through sites such as Upwork.

REFERENCES

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

2019. “The Value of Flexible Work: Evidence from Uber Drivers.” *Journal of Political Economy*. forthcoming.

Cook, Cody, Rebecca Diamond, Jonathan Hall, John A. List, and Paul Oyer. 2018. “The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers.” NBER Working Paper, Number 24,732.

Lazear, Edward P. 1979. “Why is There Mandatory Retirement?” *Journal of Political Economy*, 87.

Madden, Janice Fanning. 2018. “Gig economy increasingly concentrated among older workers.” *San Francisco Chronicle*. June 7.

Murphy, Kevin M., and Finis Welch. 1990. “Empirical Age-Earnings Profiles.” *Journal of Labor Economics*, 2.

Murphy, Kevin M., and Finis Welch. 1992. “The Structure of Wages.” *Quarterly Journal of Economics*, 97.

Shrikant, Aditi. 2018. “The gig economy isnt going anywhere. Four experts explain why.” *Vox*. October 1.

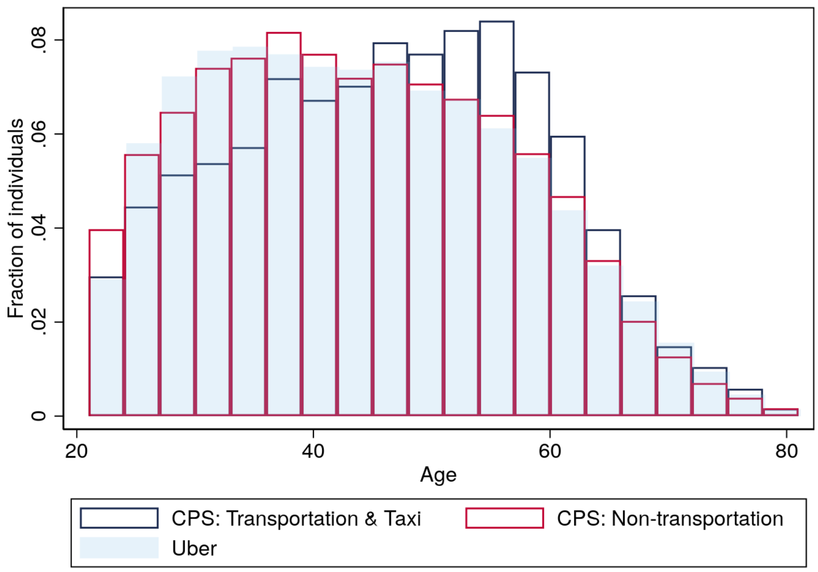

Figure 1. Age Distributions

Note: This figure documents the distribution of ages for the population of CPS workers and drivers on Uber. Transportation and taxi categories are determined according to industry and occupation codes.

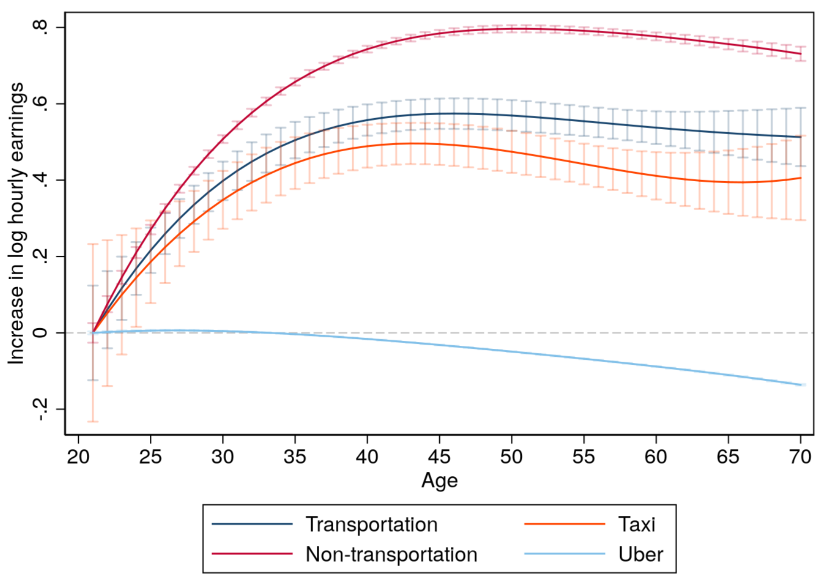

Figure 2. Age-Earnings Profile

Note: This figure shows the age-earnings profiles for drivers on Uber and workers in the CPS. CPS data covers the 2016 and 2017 calendar years and includes all non-military males over the age of 21 who worked at least 20 weeks of the year and averaged over 10 hours per week. Transportation and taxi categories are determined according to industry and occupation codes. The Uber data have been sampled and aggregated to the driver-year level to mimic the CPS data; the data include all male drivers over 21 who worked 20 weeks of the year and averaged over 10 hours per week, 2016-2017. Regressions include controls for year, and metro area (CPS) or city (Uber).

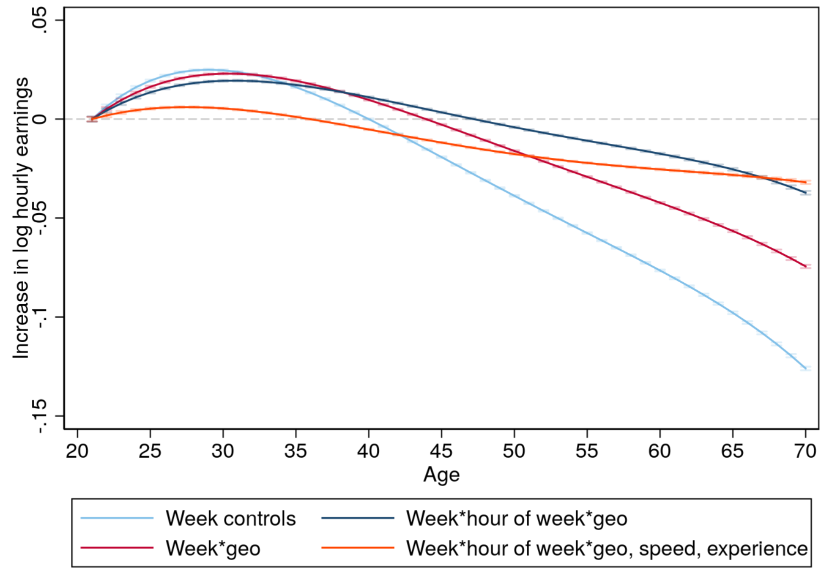

Figure 3. Age-Earnings Profile for Uber

Note: This figure documents the age-earnings profile for drivers, under various controls. Data is at the driver-hour level and includes all male Chicago UberX/UberPool drivers from January 2016 to March 2017. Experience controls are bins for quartiles of lifetime trips completed. Geo controls are dummies for the geohashes in which a driver had a trip that hour. Speed is the log of the average speed on-trip. For more information on the data, see Cook et al. (2018). Standard errors are clustered at the driver-level.

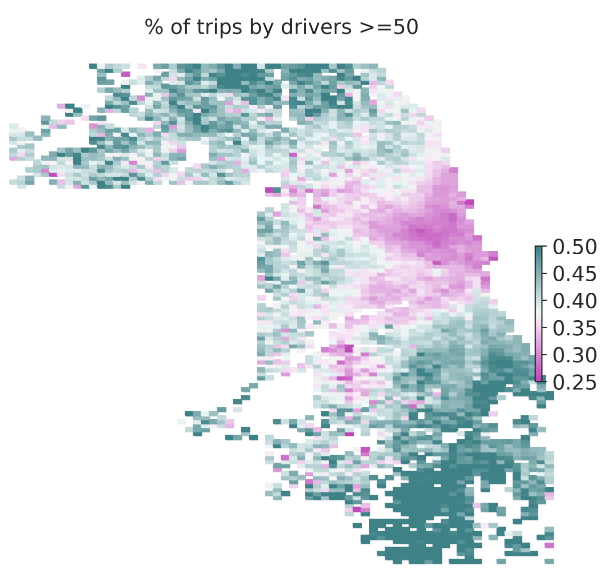

Figure 4. Geographies worked, Uber

Note: This figure maps the percent of trips in a given geohash that are completed by drivers over fifty years old. The geohashes used are more precise than those used in regressions, measuring about 0.75 miles on each side.

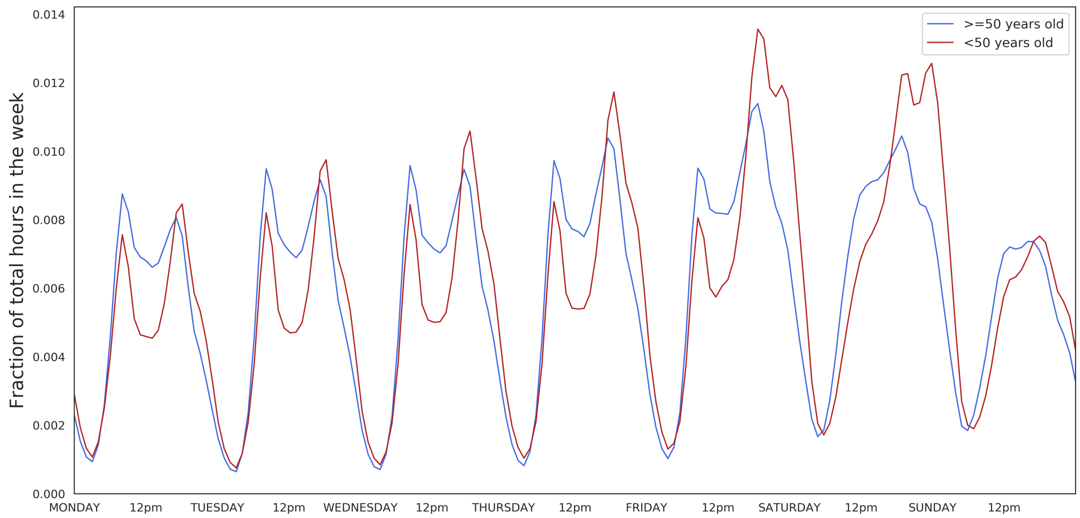

FIGURE 5. TIME OF WEEK WORKED, UBER

Note: This figure shows which hours of the week older and younger drivers work; each point represents the fraction of their total hours in the week that older (or younger) drivers spend working in that specific hour of the week. Data are limited to male Chicago UberX/UberPOOL drivers in Chicago, January 2016-March 2017.

TABLE 1—BASIC SUMMARY STATISTICS, UBER

	All	$\geq$50 years old	$<$ 50 years old
Weekly earnings	237.89	279.92	226.46
Hourly earnings	19.041	17.887	19.355
Hourly earnings (no incentive pay)	18.228	17.176	18.514
Hours per week	11.772	14.858	10.937
Weeks worked in year	10.425	13.449	9.602
Trips per week	20.362	24.176	19.324
Number of driver-years	2,674,977	580,231	2,094,746

Note: This table documents summary statistics for the year-level data from Uber, split by whether the driver was over fifty years old (as measured by age during first week driven that year). Data include all male drivers over the age of 21 who drove in 2016-2017. Values are weighted such that each year of data is equally represented. Earnings are gross of any expenses, such as the Uber commission rate and gasoline.

Table 2—Average for certain trip characteristics, Uber

	≥50 years old	< 50 years old	Difference
Wait time (min)	6.373	5.516	0.857
	(0.0020)	(0.0020)	
Accepts-to-pickup distance (mi)	0.608	0.546	0.062
	(0.0006)	(0.0006)	
Trip distance (mi)	6.401	6.036	0.365
	(0.0013)	(0.0013)	
Speed (mph)	18.359	18.197	0.162
	(0.0019)	(0.0019)	
Surge multiplier	1.065	1.089	-0.024
	(0.0001)	(0.0001)	
Per-trip incentive pay	1.169	1.368	-0.199
	(0.0003)	(0.0003)	
Per-trip total pay	12.397	12.396	0.001
	(0.0018)	(0.0018)	

Note: This table documents averages for many parameters that affect earnings for older and younger drivers. Data include UberX and UberPOOL trips in Chicago for May 2016 through December 2017 (accurate data for some of these values are not available before May 2016) by male drivers over 21 years old. To avoid issues with possible changes in the composition of driver ages over time, averages are weighted such that each week of data contributes equally. Wait time is based on time between either coming online or completing previous trip and picking up passenger for new trip. Trip distance is based on actual route taken; however, accepts-to-pickup distance is the Haversine distance between corresponding coordinates. Standard errors in parentheses.

Table 3—Hourly Earnings, Chicago

	(1)	(2)	(3)	(4)
≥50 years old	-0.0779	-0.0449	-0.0272	-0.0215
	(0.002)	(0.002)	(0.003)	(0.002)
Week	✓			
Week*geo		✓		
Week*hour of week*geo			✓	✓
Speed				✓
Experience				✓
R^2	0.0502	0.1672	0.3909	0.5428
N	13,514,221	13,514,221	13,514,221	13,514,221

Note: This table documents the gap between earnings for older and younger drivers under various controls. Data is at the driver-hour level and includes all male Chicago UberX/UberPool drivers from January 2016 to March 2017. Experience controls are bins for quartiles of lifetime trips completed. Geo controls are dummies for the geohashes in which a driver had a trip that hour. Speed is the log of the average speed on-trip. For more information on the data, see Cook et al. (2018). Standard errors (in parentheses) are clustered at the driver-level.