Older Workers and the Gig Economy

By Cody Cook, Rebecca Diamond, and Paul Oyer*

Given the serious demographic challenges pending in most developed countries, keeping older people working longer seems likely to be an important part of maintaining a healthy economy (along with increasing female labor force participation, immigration, and accelerating automation). 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 traditional employment and one large Gig Economy platform. 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 (consistent with Murphy and Welch, 1990, 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 a driver gains experience on the platform. We find that driver hourly earnings have little relationship to age for drivers in their twenties and thirties but then decrease steeply and steadily as a function of age for drivers about forty or older. Drivers who are sixty, for example, earn almost 10% less per hour than drivers who are age thirty.

Using granular data for Chicago drivers, 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 at different times and in different places (less congested areas and more in outlying suburbs than in city center). These 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 in semi-retirement and to capture the value of 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, we gather infor-

* 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, Susan Houseman, and 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.

mation about labor market outcomes for the calendar year 2016 or 2017. We follow the basic procedure in Murphy and Welch (1990) and 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.[1]

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, though the majority of Uber rides are done by the highly attached drivers in our sample. Our sample includes 292,514 drivers and 368,358 driver-years. Using data on the earnings and hours worked (that is, hours with the Uber app in operation) of 292,514 drivers and 368,358 driver-years, we calculate average hourly earnings for each driver-year.[2]

The age distributions of the entire CPS, the transportation sample, and the Uber sample each include few individuals above the age of sixty. The full CPS and Uber samples are remarkably similar in their age distributions, while the entire transportation sample is somewhat older. 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 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 1 graphically captures the age-earnings profiles 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 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. Uber earnings are increasing, though very slightly, in age for drivers in their twenties and then drop steadily with age such that sixty-year-old drivers earn about 10% less than thirty-year-old drivers.

Figure 1 shows that workers transitioning from traditional employment to Gig work at retirement ages may 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 accrued 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 Gig Economy may make it challenging for retiring workers to replace a

[1] 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

[2] Driver net earnings are less than the gross earnings figures we use, which include Uber's commission rate, gas, and the depreciation and maintenance due to Uber mileage. However, the net/gross distinction should not materially affect the age/earnings relationship.

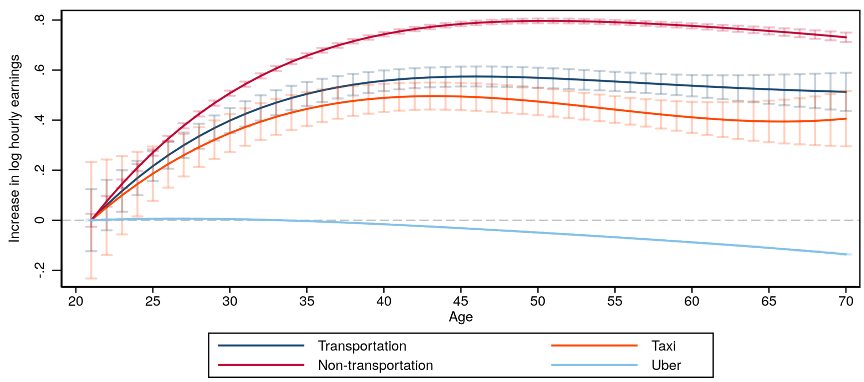

FIGURE 1. AGE-EARNINGS PROFILES

Note: Transportation, Taxi, and Non-transportation data are from CPS and cover the 2016 and 2017 calendar years. The Uber data have been sampled and aggregated to the driver-year level to mimic the CPS data. Regressions use ASEC weights and include controls for year, and metro area (CPS) or city (Uber).

substantial share of their prior income doing Gig work; however, the flexibility of Gig work may offer earnings when no other suitable job is available.

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 provide insight into how productivity of workers 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. As we describe in earlier work (Cook et al. (2018), Uber earnings are formulaic and driver earning variation reflects differences 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 drivers do not earn money). Further, earnings also increase in driving speed, as faster driving results in more trips per hour.

There are several reasons older workers could be less productive in this setting. 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, drive faster, and have more experience on the platform (which pays off through learning-by-doing).

We now concentrate on the Chicago area so we can use trip-level data to build a driver/hour dataset similar to the one used in Cook et al. (2018). 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 1, 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.[3]

[3]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.

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 ran a series of regressions of 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. When we control only for the week, drivers fifty and over earn about 8% less than those under fifty.[4] This earnings differences is even greater than the male/female difference. Introducing 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 reduces the older driver earnings coefficient by more than a third. A look at where drivers of different ages concentrate shows that youngest drivers are overrepresented closer to downtown where traffic is greatest, surge rates are higher on average, and wait times between rides are relatively short.

In a regression with a full set of indicator variables for all 168 hours in a week interacted with the calendar week and geographies worked, the coefficient on older drivers drops substantially. Older drivers are relatively likely to drive during daylight hours on weekdays and 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, older drivers make different choices than younger drivers about where and when to drive, choosing to operate disproportionately in outlying areas and avoiding high demand times. These decisions lead these drivers to have more idle (unpaid) time and lower surge rates.

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 2 shows how predicted Uber earnings vary with age based on regressions with coefficients for an age quartic. 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, driving intensity (hours worked 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 differential is due to some combination of our inability to fully capture all supply and demand variation 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, our results establish that younger workers have an earnings advantage in the largest independent worker platform. This advantage is substantial (8-10% per hour) at an absolute level. The differential becomes extremely large 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. A large share of the gap is driven by differences by where and when old/young drivers work.

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

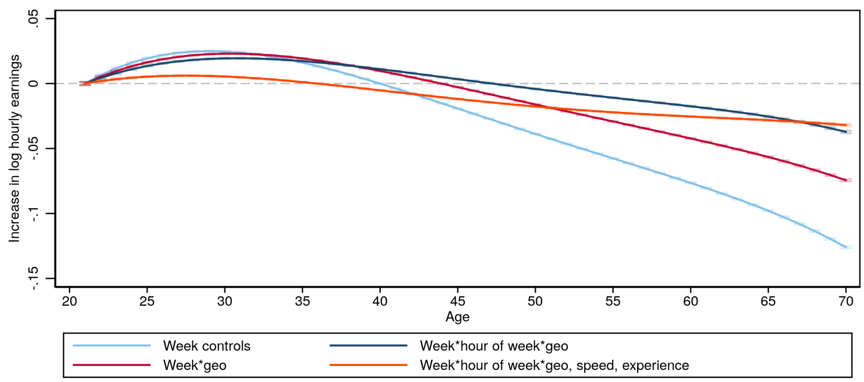

Figure 2. Age-Earnings Profile for Uber

Note: Data is at the driver-hour level and include all male Chicago UberX drivers from 01/2016-03/2017. Experience controls are bins for quartiles of 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. Standard errors are clustered at the driver-level.

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 is relatively new. It's possible some of the age/earnings relationship will change as the business matures. This does not affect the interpretation of our results, as people who use Uber to earn money after leaving a traditional job will be new to rideshare driving. Second, older people who drive for Uber are 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.

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 increase steeply with age, Uber earnings are essentially flat from age twenty to forty and steadily declining in age thereafter.

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. While rideshare makes up the majority of current Gig work, 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.

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.