WORKING PAPER #583a
PRINCETON UNIVERSITY
INDUSTRIAL RELATIONS SECTION
October 2014
Version: February 24, 2015
<http://arks.princeton.edu/ark:/88435/dsp014b29b826g>

Why You Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers*

Henry S. Farber
Princeton University

Abstract

In a seminal paper, Camerer, Babcock, Loewenstein, and Thaler (1997) find that the wage elasticity of daily hours of work New York City (NYC) taxi drivers is negative and conclude that their labor supply behavior is consistent with target earning (having reference dependent preferences). I replicate and extend the CBLT analysis using data from all trips taken in all taxi cabs in NYC for the five years from 2009-2013. Using the model of expectations-based reference points of Koszegi and Rabin (2006), I distinguish between anticipated and unanticipated daily wage variation and present evidence that only a small fraction of wage variation (about 1/8) is unanticipated so that reference dependence (which is relevant only in response to unanticipated variation) can, at best, play a limited role in determining labor supply. The overall pattern in my data is clear: drivers tend to respond positively to unanticipated as well as anticipated increases in earnings opportunities. Additionally, using a discrete choice stopping model, the probability of a shift ending is strongly positively related to hours worked but at best weakly related to income earned. These results are consistent with the neoclassical optimizing model of labor supply and suggest that consideration of gain-loss utility and reference dependence is not an important factor in these labor supply decisions.

I explore heterogeneity across drivers in their labor supply elasticities and consider whether new drivers differ from more experienced drivers in their behavior. I find substantial heterogeneity across drivers in their elasticities, but the estimated elasticities are generally positive and only rarely substantially negative. I also find that new drivers with smaller elasticities are more likely to exit the industry while drivers who remain learn quickly to be better optimizers (have positive labor supply elasticities that grow with experience).

*Industrial Relations Section, Firestone Library, Princeton University, Princeton, NJ 08544. Phone: (609)258-4044. email: farber@princeton.edu. This paper is based on my Albert Rees Lecture at the annual meeting of the Society of Labor Economists, May 2, 2014, Arlington, VA. The author thanks participants in workshops at the Federal Reserve Bank of Atlanta, Princeton University, MIT, Harvard University, Brandeis University, and Yale University for their helpful comments.

1 Introduction

That it is difficult to find a taxi in the rain has been a standard complaint in Manhattan for as long as there have been taxis. If asked why this is the case, the answer from an economist 20 years ago would have been that rainy weather increases the demand for taxi rides and there is no or an insufficiently rapid supply response to meet this transitory demand increase. That answer may have changed in recent years. In their seminal work, Camerer, Babcock, Loewenstein, and Thaler (1997), referred to here as CBLT, present evidence, based on a regression of log daily hours on log average hourly earnings, suggesting that the daily labor supply function of taxi drivers is negatively sloped so that a transitory change in the wage results in a reduction in hours worked. On this basis, they characterize taxi drivers as having reference-dependent preferences, which can be summarized simply by saying that workers will set a daily income target and generally work until that target is met. Others have also found that labor supply curves for taxi drivers appear to slope downward, and the consensus of much of this work is that taxi drivers are, in fact, target earners.¹ This suggests an alternative answer to the question of why it is difficult to find a taxi in the rain: To the extent that drivers have a daily income target and a rain-induced increase in demand increases earnings, drivers will reach their targets sooner and quit driving for the day. The new view then is that at least part of the reason you can't find a taxi in the rain is because drivers reach their daily earnings targets quickly and go home so that the demand increase is exacerbated by the resulting decline in supply.²

The question of whether reference dependence plays a substantial role in labor supply decisions is important, both intellectually and in designing public tax and transfer policies. There is a natural tension between the standard neoclassical optimizing model of labor supply and the model based on reference dependent preferences. Setting income targets is an inefficient way to earn money because it implies working less on high wage days and working more on low wage days. The neoclassical model implies the opposite. Clearly, over a period of days, the neoclassical optimizer works fewer hours than the target earner to earn

¹ See, for example, Chou (2002) and Agarwal, Diao, Pan, and Sing (2013) for estimates of labor supply models for taxi drivers that find a negative slope. Using different approaches, Crawford and Meng (2011) and Doran (2014) also find support for reference-dependent preferences in analyses of labor supply of NYC taxi drivers. Koszegi and Rabin (2006) and Ordóñez, Schweitzer, Galinsky, and Bazerman (2009) are examples of how this result has been accepted generally.

² See Ordóñez, Schweitzer, Galinsky, and Bazerman (2009) for an explicit statements that reference-dependent preferences are part of the explanation for difficulty in finding a taxi in the rain.

the same income.³ In this study I use new data, consisting of the complete records of all taxi drivers in New York City over the 5-year period from 2009-2013 to estimate models of labor supply in order to determine whether reference dependent preferences play an important role in explaining labor supply in this setting or whether the standard neoclassical model can account for most of the observed variation. The use of these new data addresses a weakness of much of the earlier work on taxi drivers, including my own, that the analyses are based on very small convenience samples.

A continuing focus of this “battle of models” is the daily hours decisions of taxi drivers, and, with some exceptions, the model used is a regression of log daily hours on the log daily wage (the log of daily average hourly earnings). This is why I focus on estimating the slope of the daily hours function in much of what I do here. An alternative empirical approach is to model the stopping decision of a driver, where the driver decides at the completion of each trip whether to continue driving or to end the shift. This discrete-choice stopping decision may be a function of accumulated hours and income on the shift as well as factors that reflect variation in supply and demand factors such as time of day, day of week, and month of year among other variables. Using this reduced form approach in earlier work (Farber, 2005), I found that taxi driver labor supply is best characterized by the neoclassical model and that there is little evidence in support of reference-dependent preferences.⁴ Following this approach, I supplement my regression analysis of the slope of the labor supply function with a reduced form discrete choice analysis of the stopping decisions of drivers that uses the new data.

Given the efficiency advantage of optimizing behavior, after investigating the basic mod-

³ This is the point of the taxi driver example used by Ordonez, Schweitzer, Galinsky, and Bazerman (2009), and the inefficiency of target earnings behavior is highlighted by Camerer (1997).

⁴ In a later paper (Farber, 2008) I estimated a structural version of the stopping model using the same data on NYC taxi drivers and again found little evidence for reference dependent preferences. Crawford and Meng (2011), using the same data, estimate structural stopping model that allows for reference points in both daily income and daily hours, and they conclude that the data are consistent with this dual reference point model. The empirical evidence from other settings using a variety of methods is mixed. Oettinger (1999) examines the labor supply of stadium vendors at baseball games and finds evidence that labor supply on the extensive margin (number vendors showing up for games) is consistent with the neoclassical model. Fehr and Goette (2007) run a field experiment varying the piece rate paid to bicycle messengers. Their evidence is generally consistent with the neoclassical model in that messengers work more in months with high piece rates. On the other hand, they interpret evidence that messengers work fewer hours per day on days in months with high piece rates as evidence of reference-dependent preferences. See also Nguyen and Leung (2009) for an analysis of labor supply in fisheries and Chang and Gross (2014) for an analysis of labor supply in fruit picking that find evidence consistent with reference-dependent preferences. Camerer (2000) presents a nice survey of evidence from the field on loss aversion and reference dependence in a variety of settings.

els of labor supply, I investigate whether drivers differ in their labor supply behavior (are some drivers target earners while others are optimizers?). I then investigate whether new drivers learn to be better optimizers and whether drivers who not strong optimizers disproportionately quit the industry.

I turn now to my replication and extension, based on the new data, of the earlier work estimating the slope of the labor supply function. The next section outlines the competing theories of labor supply, where I rely heavily on the work of Koszegi and Rabin (2006), who provide a clear theoretical analysis of the determination of reference points and a guide to the kind of wage variation that drives reference dependence. In section 3, I present some background information on the taxi industry in New York City and a broad description of the data. Section 4 contains my analysis of why taxis are hard to find in the rain. In section 5, I describe the construction of the analysis sample and present estimates of the labor supply function, and, in section 6, I investigate how the slope of the labor supply function is related to the deviation of the daily wage from its expected value. I then present an analysis of the discrete choice stopping model in section 7. Section 8 contains an analysis of the extent to which the elasticity of labor supply and the implied underlying model differs across drivers. In section 9, I investigate whether drivers learn to be better optimizers (exhibit larger labor supply elasticities as experience is gained), and, in section 10, I ask whether drivers who appear to be weaker optimizers (have smaller labor supply elasticities) at the start of their time as a driver are more likely to exit the industry. Section 11 contains final remarks and conclusions.

2 Setting the Stage: Competing Theories of Labor Supply

In the standard neoclassical inter-temporal model of labor supply, individuals work in period t until the shadow value of time (a function of lifetime wealth/income and increasing in hours in a given period) equals the period t wage rate.⁵ The model implies that there is intertemporal substitution in labor supply across periods so that a transitory increase in the wage rate in period t implies an increase in period t labor supply because the shadow value of time conditional on hours is unaffected. However, a permanent change in the wage will have an offsetting income effect as lifetime wealth increases, increasing the shadow value of

⁵ See MaCurdy (1981) for an early empirical analysis of inter-temporal substitution in labor supply. Blundell and MaCurdy (1999) present a relatively recent survey of the literature on labor supply.

time conditional on hours.⁶

The reference dependent model of choice has its roots in the literature on loss-aversion (Kahneman and Tversky, 1979; Tversky and Kahneman, 1991). In the context of the daily labor supply decisions of taxi drivers, the basic idea of the reference dependent preference model is that a driver has in mind a particular reference level of daily income, and utility as a function of income is evaluated relative to this reference level. The loss in utility from failing to reach the reference income level by some amount exceeds the gain in utility from exceeding the reference income level by the same amount. In other words, the individual is loss averse. There is a kink in the utility function at the reference income level, with higher marginal utility below the kink and lower marginal utility above the kink.

Consider the following simple model of labor supply with reference-dependent preferences. Individuals facing a wage rate W receive utility from income ($Y = Wh$) and disutility from hours of work (h). Individuals have a kink in their utility function at some reference level of income (T):

$$U(Y, h) = (1 + \alpha)(Y - T) - \frac{\theta}{1 + \nu} h^{1+\nu} \quad Y < T \quad (1)$$

$$U(Y, h) = (1 - \alpha)(Y - T) - \frac{\theta}{1 + \nu} h^{1+\nu} \quad Y \geq T \quad (2)$$

where the parameter $\alpha > 0$ controls the change in marginal utility at the reference point, θ indexes the disutility of hours, and ν is a parameter related to the elasticity of labor supply. This specification follows the model of Koszegi and Rabin (2006), where utility with reference dependent preferences is additive in the usual “consumption utility” and in a “gain-loss” utility around the reference point. The model is then based on a neoclassical utility function $((Y - T) - \frac{\theta}{1+\nu} h^{1+\nu})$ augmented with a gain-loss component $(\pm\alpha(Y - T))$.

The purely neoclassical model is the special case where there is no kink in the utility function (no gain-loss utility, $\alpha = 0$). The labor supply function in this case is $h = (\frac{W}{\theta})^{1/\nu}$, and the elasticity of labor supply is $\frac{1}{\nu}$. The extent to which labor supply behavior is neoclassical or influenced by reference dependence (gain-loss utility) depends on both α and the wage as follows: Maximizing the utility function in equations 1 and 2 with respect to hours of work yields three distinct labor supply functions depending on the wage.

1. For sufficiently low wages ($W < W^*$), the reference point is not relevant because the hours required to reach the reference point at such a low wage yield sufficient disutility

⁶ See Ashenfelter, Doran, and Schaller (2010) for an analysis of the effect of two fare increases (1996 and 2004) on the labor supply of NYC taxi drivers. They find an elasticity of -0.2 in response to these permanent fare increases.

of hours that it is optimal to stop on the high marginal utility section of the utility function (short of the reference point). In this region, the labor supply function is neoclassical:

$$h = \left(\frac{(1 + \alpha)W}{\theta} \right)^{1/\nu} \quad (3)$$

with elasticity of labor supply $1/\nu > 0$.

2. For intermediate wage levels ($W^* < W < W^{**}$), it is optimal to stop working upon reaching the reference income level. This is because the wage is in a range that is high enough to reward working when marginal utility is high ($Y < T$) but too low to reward working when marginal utility is low ($Y \geq T$). In this region, the individual is a target earner with labor supply function

$$h = \frac{T}{W} \quad (4)$$

and elasticity of labor supply is -1.

3. For sufficiently high wages ($W > W^{**}$), the reference point is not relevant because the wage is sufficiently high that it is optimal to operate on the low marginal utility section of the utility function (beyond the reference point). In this region, the labor supply function is neoclassical:

$$h = \left(\frac{(1 - \alpha)W}{\theta} \right)^{1/\nu} \quad (5)$$

with elasticity of labor supply $1/\nu > 0$.

The bounds of the range where reference-dependent preferences are relevant (where gain-loss utility plays a role) are derived from the optimizing behaviour of the individuals. Consider first the lower bound, W^* . The value of W^* is defined as the wage at which an individual with the “steep” utility function defined in equation 1 would choose hours so as to earn T . Based on the labor supply function in equation 3, this is

$$W^* = \left(\frac{\theta}{1 + \alpha} \right)^{\frac{1}{1+\nu}} T^{\frac{\nu}{1+\nu}}. \quad (6)$$

Consider next the upper bound, W^{**} . The value of W^{**} is defined as the wage at which an individual with the “flat” utility function defined in equation 2 would choose hours so as to earn T . Based on the labor supply function in equation 5, this is

$$W^{**} = \left(\frac{\theta}{1 - \alpha} \right)^{\frac{1}{1+\nu}} T^{\frac{\nu}{1+\nu}}. \quad (7)$$

The ratio of the upper bound to the lower bound is

$$\frac{W^{**}}{W^*} = \left(\frac{1 + \alpha}{1 - \alpha} \right)^{\left(\frac{1}{1+\nu} \right)}, \quad (8)$$

and this is directly related the degree of gain-loss utility (indexed by α) and the labor supply elasticity (measured inversely by ν). If α is close to zero, indicating little gain-loss utility, the range of wages where the individual is a target earner is very small, and reference dependence is not important. But as α grows, reference dependence and target earnings behavior have more relevance.

Koszegi and Rabin (2006) suggest that the reference income level will be based on expected income.⁷ Expected income will be driven by the income level generated by the expected wage and the hours choice made by the individual based on the expected wage. Defining reference points as reflecting expected income has important implications, both for thinking about the potential importance of reference dependence in determining labor supply and in designing an empirical analysis that reflects appropriate variation.

The Koszegi-Rabin model of expectation-based reference income levels suggests importantly that labor supply, consistent with the neoclassical model, will be positively related to *anticipated* transitory wage changes. They argue that reference dependence (gain-loss utility in their terms) is related only to unanticipated variation in the wage. In periods where high wages are expected, individuals will have higher reference points implying higher labor supply as the neoclassical model predicts. The prediction of the reference dependent preferences model, that the elasticity of labor supply is -1, is relevant only with regard to *unanticipated* transitory wage changes that are close to the expected wage.

This limits how much of labor supply behavior can be accounted for by reference dependent preferences and suggests that much of the variability in labor supply is likely to be consistent with the neoclassical model.⁸ Later, I decompose variation in average hourly earnings into components that are plausibly interpreted as permanent, anticipated transitory, and unanticipated transitory.

I now use the Koszegi-Rabin formulation of expectation-based reference points to derive the bounds (W^* and W^{**} in equations 6 and 7 respectively) as a function of the expected

⁷ Abeler, Falk, Goette, and Huffman (2011) present experimental evidence that variation in work effort is consistent with reference points based on expectations. Crawford and Meng (2011) rely on expectation-based reference points in their analysis of taxi driver labor supply.

⁸ Crawford and Meng (2011), building on the Koszegi-Rabin work, also recognize that the response to anticipated transitory wage changes will be neoclassical.

wage. In order to simplify the exposition given the multiplicative functional forms, I work with the logarithms of the wage and of hours and the expectation of the logarithms. The expected log labor supply based on the consumption part of the utility function is

$$E(\ln(h)) = \frac{1}{\nu} E(\ln(W)) - \frac{1}{\nu} \theta, \quad (9)$$

and the implied log income reference point is

$$\begin{aligned} \ln(T) &= E(\ln(h)) + E(\ln(W)) \\ &= \frac{1+\nu}{\nu} E(\ln(W)) - \frac{1}{\nu} \theta. \end{aligned} \quad (10)$$

The bounds are related to the expected wage through the reference income level. Substituting the expression in equation 10 for the log reference income level into the logarithms of equations 6 and 7 yields particularly simple expressions for the logarithms of the bounds (W^* and W^{**}). These are

$$\ln(W^*) = E(\ln(W)) - \left(\frac{1}{1+\nu} \right) \ln(1+\alpha), \quad (11)$$

and

$$\ln(W^{**}) = E(\ln(W)) - \left(\frac{1}{1+\nu} \right) \ln(1-\alpha). \quad (12)$$

The likelihood that the realized wage is outside the bounds, yielding neoclassical behavior (positive labor supply elasticity) is a function of how variable the wage is around its expected value. If the wage has only small unanticipated variation (since anticipated variation is built into the reference income level through the expected wage), then behavior in response to unanticipated transitory wage variation will generally look like target earning. On the other hand, if the wage has substantial unanticipated variation, then behavior will be neoclassical.

This formulation of the labor supply model with reference dependent preferences has a direct empirical prediction: On days when the wage rate unexpectedly varies substantially from its expected value, labor supply will be more likely to vary directly with the wage rate. But on days where the wage rate is relatively close to expectation, hours worked will be more likely to vary inversely with the relatively small unanticipated wage variation.

3 Background and Data on Taxi Drivers

There are 13,238 taxi medallions in New York City. This number is set by regulation. Roughly speaking, there are two types of medallions.

1. Fleet medallions, of which there are 7,664, are attached (literally) to taxis generally operated through a fleet garage and leased on a daily shift basis to individual drivers with hack licenses.⁹ Owners of fleet medallions must own at least two medallions, and each fleet medallion must operate for two shifts of at least 9 hours per day. It is not clear if or how the latter requirement is enforced.
2. Independent medallions, of which there are 5,574, are owned by individuals who may own no more than one medallion. A subset of these medallions are “owner-driver” medallions, which have a requirement that a substantial number of shifts in taxis with such medallions be driven by the owner. Other independent medallions have no such restriction, and these taxis may or may not be driven by the owner. In either case, taxis with independent medallions may be leased for shifts to drivers with hack licenses who are not the owner. Again, it is not clear if or how these requirements are enforced.

The standard employment arrangement of New York City cab drivers who do not own their own cabs/medallions is that a driver leases a cab for a fixed period, usually a 12-hour shift. The driver pays a fixed fee for the cab plus fuel, and he keeps 100 percent of the fare income plus tips. The driver is free to work as few or as many hours as he wishes within a 12-hour shift. Thus, the driver internalizes the costs and benefits of working in a way that is largely consistent with an economist’s first-best solution to the agency problem with risk-neutral agents. In a manner of speaking, the employer has “sold the firm to the worker.” Because these drivers are free to set their hours once they have leased a taxi for a shift, analysis of their labor supply is fertile ground for learning about behavioral models.

Taxi drivers earn income only when there is a passenger in the cab. My data cover the 2009-2013 period. Prior to September 4, 2012, income was earned at the rate of \$2.50 for the first one-fifth mile (the “meter drop”) plus \$.40 per additional one-fifth of a mile when travelling at 12 miles per hour or more plus \$.40 for each minute when travelling at less than 12 miles per hour (waiting time). From September 4, 2012 through 2013, the rate for additional fifths of a mile and waiting minutes was increased from \$.40 income to \$.50. Throughout the period, there was also a night surcharge of \$.50 per trip between 8PM and 6AM and peak-hour weekday surcharge of \$1.00 Monday-Friday between 4PM and 8PM.¹⁰

⁹ Less commonly, taxis with fleet medallions are leased on a weekly basis.

¹⁰ There is a flat fare of \$52.00 (\$45.00 prior to September 4, 2012) plus tolls between Manhattan and JFK International Airport in either direction and a surcharge of \$17.50 (\$15.00 prior to September 4, 2012) on trips to Newark Liberty International Airport.

Clearly, a central factor in earnings is the ease/speed with which new fares are located.

The earlier studies of taxi driver labor supply were based on analysis of relatively small numbers of hand-written “trip sheets” (one per shift) that drivers were required to fill out with information on the fare and trip start and end times and locations.¹¹ These sheets were difficult to read accurately, and the limited number of sheets available severely constrained analysis.

This situation has changed dramatically. The New York City Taxi and Limousine Commission (TLC), the agency charged with regulating the industry, now requires all taxis to be equipped with electronic devices that record all trip information including fares, times, and locations. The (currently two) companies that supply these devices report all this information to the TLC on a regular basis, and I have obtained full information for all trips taken in NYC taxi cabs for the five years from 2009-2013. These data, called TPEP by the TLC, identify drivers by encrypted hack license number and medallions (cabs) by encrypted medallion number.

There are approximately 180 million trips taken during approximately 8 million shifts each year in taxi cabs in NYC. About 62,000 drivers had at least one fare in a cab over the five-year period, with about 40,000 drivers in any single year. About 25,000 drivers worked in all five years. In fact, these are more data than I can use efficiently, and much of my analysis is based on a random sub-sample of 2/15 of the drivers.

One limitation of the data is that they do not identify fleet and independent medallions or isolate the owner-driver medallions among the independent medallions. It may be that owner-drivers face different incentives regarding labor supply than do drivers who lease their cabs, whether from a fleet owner or an individual. For this reason, it would be interesting to analyze their labor supply separately.

Another limitation of the data is that complete information is not available on tip receipts. Data on tips are available only for fares with tips paid by credit card. While the percentage of fares paid by credit card has been increasing over time, about 46 percent of fares in the most recent year (2013) were paid in cash and so there is no tip information. Tips on credit card transactions averaged about 20 percent over the 2009-2013 period. I proceed ignoring tips due to the missing data for cash transactions, this will not cause a problem for my

¹¹ An exception is Agarwal, et al (2013), who have administrative data on taxi drivers in Singapore, but these data lack information on fares.

analysis unless variation the rate of tipping is correlated with hourly fare income.¹²

4 Why Can't You Find a Taxi in the Rain?

I begin with a direct analysis of the question of whether or not target earnings behavior can account for some of the perceived increased difficulty in finding a cab in the rain. The underlying idea is that the demand for taxis increases when it is raining, which, by itself, makes it more difficult to find a cab and that this increase in demand results in higher hourly earnings for taxi drivers. The result is that the drivers reach their daily earnings targets sooner and quit for the day, exacerbating the increased difficulty in finding a cab.

Rain may have a number of effects on the market for taxi rides. First, it likely increases demand. This will make it easier for drivers to find passengers and increase hourly income. However, rain may decrease speed due to congestion and a deterioration in general driving conditions, which could result in lower earnings. Rain may also make driving less pleasant, implying a reduction in the number of cabs on the road having nothing to do with target earnings behavior.

For this analysis, I use a random sub-sample of the drivers in the data described above. These data include all trips taken by 2/15 of the drivers I see between 2009 and 2013. This sub-sample includes 116,177,329 trips for 8,802 drivers.

In order to examine the market, I use the trip-level data to calculate the average of income, time with a passenger in the cab, miles travelled, and number of taxis on the street for each of the 48,824 clock hours in the 5 year period from 2009-2013. I then merge these data with data on hourly rainfall in Central Park.¹³ I calculate the hourly wage for a given driver in a given hour using the trip-level fare and time information. The hourly wage was computed by dividing each shift into minutes and assigning a “minute wage” to each minute. For minutes during trips, the minute wage is computed as the fare divided by the number of minutes for that trip. For minutes of waiting time (between fares), the minute wage is set to zero. The hourly wage for each clock hour is computed as the sum of the minute wages during that hour. I calculate time with a passenger in the cab and miles travelled during a given clock hour in an analogous fashion. The count of taxis on the street is a count of taxis

¹² The lack of observability of tips does introduce measurement error in income and the hourly wage, and I address this problem in my analysis. See Haggag and Paci (2014) for an analysis of tipping behavior in NYC taxi cabs in 2009.

¹³ Source: National Weather Service (NWS) of the National Oceanic and Atmospheric Administration (NOA) and the Network for Environment and Weather Applications, Cornell University.

who had a passenger in the cab for at least one minute during the clock hour.

A necessary condition for target earnings behavior to contribute to difficulty in finding a taxi in the rain is that hourly earnings be higher when it is raining. The first row of table 1 contains the coefficient of an hourly indicator for precipitation in Central Park from OLS regression analyses of average log hourly earnings.¹⁴ The estimates in column (A) include no other variables, while the estimates in column (B) include indicators for hour of day by day of week (167), month of year (11), year (4), the period subsequent to the September 4, 2012 fare increase (1), and major holiday (1).¹⁵ The hourly wage is not significantly correlated with whether it rained in Central Park.¹⁶ There is considerable systematic variation in average hourly earnings over the course of the day, week, month, and year as measured by the R-squared of the regression in column B, but whether or not it was raining does not significantly improve the fit of the model.

The finding of no relationship between earnings and rainfall is sufficient to reject the hypothesis that target earnings behavior contributes to the difficulty of finding a taxi in the rain, but it leaves a puzzle: If demand for taxis is higher in the rain, why are earnings not higher? I investigate this by examining the relationship of other measures of taxi activity with rainfall.

The second row of table 1 contains OLS estimates of the coefficient of the precipitation indicator from a regression of average log minutes per hour spent with a passenger in the cab. This is a measure of how busy the cabs on the street are (and how easy it is for drivers to find passengers). Taxi occupancy rates are 4.8 percent higher when it is raining, accounting for systematic variation over time (column B). Since drivers make money only when passengers are in the cab, it is clear that demand is higher relative to supply when it is raining. However, this does not translate into higher earnings for drivers.

One possibility is that traffic and other driving conditions are worse when it is raining so that the taxis drive more slowly. Since income is roughly proportional to miles traveled (with some payment for idle time in traffic), this would imply lower income than would ordinarily

¹⁴ The indicator for precipitation equals one if there is any precipitation in Central Park recorded during the hour. The results I present are qualitatively unaffected by using a graduated measure of the quantity of precipitation. Not surprisingly, on those rare days when there is a major storm (e.g., Hurricane Sandy), dramatically fewer taxis are on the street.

¹⁵ Major holidays are defined as New Years Day, Easter Sunday, Memorial Day, Fourth of July, Labor Day, Thanksgiving, and Christmas Day.

¹⁶ Consideration of tips does not change this conclusion. Examining the information on tips for credit card transactions, the tip rate (which averages about 20 percent of the fare inclusive of surcharges) is about 0.2 percentage points higher, on average, in hours with rain.

Table 1: Effect of Rain on Hourly Taxi Market Outcomes, OLS Estimates
Coefficient of *Precipitation* (=1 if Rain in Central Park During Hour)

Outcome	(A)	R-Squared	(B)	R-Squared
1) Log Hourly Earnings	-0.0047 (0.0054)	0.000	-0.0003 (0.0021)	0.879
2) Log Time w/Passenger	0.0508 (0.0077)	0.002	0.0475 (0.0028)	0.904
3) Log Miles w/Passenger	-0.0306 (0.0046)	0.002	-0.0241 (0.0026)	0.792
4) Log Number of Hacks	-0.0671 (0.0149)	0.001	-0.0711 (0.0089)	0.878
Includes Other Controls?	No		Yes	

Note: Estimated using data for each of 43,824 hours in the years 2009-2013 derived from trip-level data for the 2/15 sample of all drivers of NYC taxi cabs. This sample contains 8,802 drivers on 116,117,329 trips. *Precipitation* is an indicator for hours where there is positive precipitation in Central Park. Other controls include indicators for hour of day by day of week (167), week of year (51), year (4), the period subsequent to the September 4, 2012 fare increase (1), and major holiday (1). HAC standard errors are in parentheses.

result from higher occupancy. In order to investigate this, the third row of table 1 contains OLS estimates of the coefficient of the precipitation indicator from a regression of average log miles travelled per hour with a passenger in the cab. Miles travelled with a passenger are about 2.4 percent lower when it is raining accounting for systematic variation over time (column B). That miles traveled per hour with a passenger are lower despite the fact that the occupancy rate is higher is clear evidence that driving conditions are worse in the rain. This is the factor that offsets the increase in demand and results in no change in average hourly earnings when it is raining.

In order to investigate any supply response to rain, the fourth row of table 1 contains OLS estimates of the coefficient of the precipitation indicator from a regression of the log number of hacks that take at least one trip during the hour in question. The number of cabs on the street is about 7.1 percent lower when it is raining. This reduction in supply with no change in average earnings despite the increase in demand likely reflects added disutility of driving when it is raining. Some drivers stop, but this is not due to their reaching their income target. Some drivers stop simply because it is less pleasant to drive in the rain and there is no additional benefit in continuing to drive.

The increase in taxi utilization measured by time with a passenger of 4.8 percent is more than offset by the decline in supply of cabs of 7.1 percent. This means that the supply of

rides is lower in rainy hours, and any surge in demand is unmet. One logical response would be to have a rain surcharge in order to encourage an increase in supply.¹⁷

5 Estimating the Wage Elasticity of Labor Supply

5.1 How Much Wage Variation is Unanticipated?

I noted earlier that the role of reference dependence in determining labor supply is limited to the response of labor supply to unanticipated wage variation. Before proceeding with estimation of the labor supply model, I present some evidence on the magnitude of unanticipated wage variation based on the data on average earnings and number of hacks on the road by hour that I used for the analysis of taxis and rain in section 4.

Table 2 contains simple variance decompositions for the average log wage and log number of hacks by hour for the 43,824 hours from 2009-2013. This decomposition is carried out in two stages. In the first stage, I regress average log hourly earnings ($\ln W$) on a set of year indicators (4) and an indicator for period subsequent to the September 4, 2012 fare increase (1). These variables capture permanent wage variation, and the variance of the predicted values from this regression is my measure of the variance of permanent wage variation. The residuals from this regression include both anticipated and unanticipated transitory wage variation. In the second stage, I regress these residuals on a set of controls including hour of day by day of week (167), week of year (51), and holiday (1).¹⁸ These controls capture anticipated transitory variation in the wage, and the variance of the predicted values from this regression is my measure of the variance of anticipated transitory wage variation. The residuals from this regression capture unanticipated transitory wage variation, and the variance of these residuals is my measure of the variance of unanticipated transitory wage variation.¹⁹

The first row of table 2 contains the variance decomposition for the average log wage

¹⁷ An example of real-time adjustment of rates to meet fluctuations in demand is Uber’s “surge pricing.” For example, if the labor supply elasticity was 0.5, a 14.2 percent “rain surcharge” could get supply back to the dry-weather level. And one might want a larger surcharge in order to meet increased demand and offset slower driving in the rain. Of course, this depends on there being reasonable very short run elasticity of labor supply.

¹⁸ See footnote 15 for a list of major holidays used in defining the major-holiday indicator.

¹⁹ Of course, drivers do not perform regression analyses to calculate expected earnings opportunities. While there is likely to be some difference between drivers’ calculations of the expected wage and the model I use, the information used in these models (date, time, etc.) are just the sort of factors drivers are likely to use in forecasting earnings opportunities.

Table 2: Variance Decomposition of Hourly Averages (Across 43,824 Hours, 2009-2013)

Variable	Permanent		Transitory	Transitory
	Total	Anticipated	Anticipated	Unanticipated
$\ln(W)$	0.0521	0.0058	0.0400	0.0063
$\ln(Hacks)$	0.2716	0.0011	0.2369	0.0337

Note: The decomposition is based on two stages: 1) OLS regression of each variable on indicators for year (4) and the period subsequent to the September 4, 2012 fare increase (1), and 2) OLS Regression of the the residuals from the first regression on a set of controls including hour of day by day of week (167), week of year (51), and holiday (1). “Total” is the raw variance. “Permanent Predicted” are the variances of the predicted values from the first regression. “Transitory Anticipated” and “Transitory Unanticipated” are, respectively, the variances of the predicted values and residuals from the second regression.

by hour. Most variation (76.8 percent of total variation) is anticipated transitory variation. This is largely variation by hour-of-day by day-of-week. 11.1 percent of total variation is permanent variation, driven largely by the September 4, 2012 fare increase. The remaining 12.1 percent of total variation is unanticipated transitory variation. This decomposition clearly limits the scope of reference dependence to account for variation in labor supply. About 7/8 of variation in average hourly earnings is anticipated so that it does not involve gain-loss utility, and its effects on labor supply are consistent with the neo-classical model. Only the remaining 1/8 of total variation in average hourly earnings is due to unanticipated factors.

The second row of table 2 repeats this decomposition for variation in the log number of hacks on the road by hour. Interestingly, the decomposition yields similar results. An even larger share of total variation (87.2 percent) is anticipated transitory variation. Again, this is largely variation by hour-of-day by day-of-week. Almost no variation (0.4 percent of total variation) is due to permanent changes. The remaining 12.4 percent (about 1/8) of total variation is unanticipated transitory variation.

The clear implication of this simple decomposition is that reference dependence (gain-loss utility) is not relevant for explaining the broad patterns of variation in labor supply, which largely result from anticipated variation in demand by hour of day and day of week. Only about 1/8 of total variation in the wage and labor supply could potentially be influenced by reference dependence and gain-loss utility.

5.2 Shift Definition and Creation of the Analysis Sample

In order to create an analysis sample of shifts, it is necessary to define a shift based on the available data. This was straightforward for the data used in the earlier small-sample studies, because the data were transcribed from hand-written trip sheets, each of which represented a distinct shift. The new electronic data upon which I rely is simply a running list of all trips by a particular driver. The assignment of trips to particular shifts is an analytic decision.

The definition of a shift is necessarily subjective. Given the approach here, which is to model the shift-level labor supply decision of a driver (perhaps with some reference income level for the shift), it makes sense to define a shift as composed of all trips that are part of a sequence that the driver considers to be a single shift. Absent a clear guide, I define any gap between trips of more than six hours (more than 360 minutes) as marking the end of one shift and the beginning of the next.²⁰ Defining shifts in this way yields a sample of 5,047,343 shifts for 8,802 drivers over the period from the 2/15 random sample of all drivers from 2009-2013. I now present some data that support this choice as a reasonable way to identify shifts.

- Figure 1 shows the distribution of shift length in hours (truncated, e.g., 8.5 hours is shown as 8). The modal shift duration is in the 9th hour, and shift durations are concentrated between hours 5 and 11 (81.6 percent). 10.7 percent of shifts are less than 6 hours, and 7.7 percent of shifts are 12 hours or longer.
- Figure 2 shows the fraction of shifts in my sample that start in each clock hour. There are two daily spikes, with a surge of shifts (44.5 percent) starting in the six-hour segment from 4AM-9:59AM and another surge (42.4 percent) in the six-hour segment from 2PM-7:59PM. Only a small fraction of shifts (4.6 percent) start in the eight hours between 8PM and 3:59AM. About 8.6 percent of shifts start in the four hours between 10AM and 1:59PM.
- Figure 3 shows average shift duration in hours by hour of shift start. Interestingly,

²⁰ In my first analysis of these data, presented as the Albert Rees Lecture at the 2014 meeting of the Society of Labor Economists, I defined a shift mechanically as the collection of all trips in the 24 hour period between 5AM one day and 4:59AM the next. However, this definition clearly does not mirror drivers conceptions of a shift. For example, a driver who started two 10-hour shifts at 4AM on two consecutive days, will have the hour from 4AM-4:59AM on the second day as the end of the previous day's shift (starting at 5AM on the first day). This "shift" will look like it lasted 24 hours when, in fact, it lasted only 11 hours. While there are some differences between the results I presented in the Rees Lecture and those I present here, the general conclusions are not sensitive to this choice of shift definition.

Figure 1: Distribution of Shift Length in Hours

Figure 2: Distribution of Shift Start Time, by Clock Hour

Figure 3: Average Duration of Shift, by Clock Hour of Start

shifts that start very early in the morning last longer on average (more than 10 hours) than those starting later in the day (8 to 9 hours). Shifts starting later in the evening (from 7PM on), last progressively less time.

Daily leases typically run for twelve hours, with two such shifts per day. The “day shift” runs from very early morning to mid- to late afternoon. The “night shift” runs from mid- to late afternoon until early morning. The day-shift drivers generally do not start driving at the first available moment, but they must stop no later than when their shift ends so that the night shift driver can take over. Similarly, the night-shift drivers may take over when the cab is first available but they generally stop before they are required to do so. It is clear from figure 1 that drivers who lease their cabs daily do not use the entire 12 hours to which their lease entitles them.²¹

I assign shifts as “day” or “night” based on the information in figures 1-3. I define shifts that start between 4AM and 9:59AM as day shifts (44.5 percent of shifts). I define shifts that start between 2PM and 7:59PM as night shifts (42.4 percent of shifts). The remaining 13.1 percent of shifts, starting between 10AM and 1:59PM or between 8PM and 3:59AM, are unclassified (not classified by me as day or night). Drivers who start in late morning or early afternoon may be day shift drivers who are getting a late start or night shift drivers who are getting an early start.

There may be an important difference in the margin on which labor supply can be adjusted on day shifts versus night shifts. Day shift drivers are more likely to be constrained at the end of their shift when the cab must be turned over to a night shift driver. This implies that day shift drivers may adjust hours by changing their start time and is consistent with the decrease in average shift duration, shown in figure 3, as the early morning progresses from 4AM through 10AM. In contrast, night shift drivers may not be able to start early but can adjust hours by changing their end time. As demand declines late at night, many of these drivers stop. The decline in demand late is reflected in the declining average duration of shifts apparent in figure 3 as starting time moves later in the evening (after 6PM). These patterns are a natural consequence of the 2-shift structure of the day.

The difference in the active margin of decision making between day and night shifts is important for analyzing labor supply. Day shift drivers often select hours before information

²¹ Owner-drivers and those drivers who lease by the week or month from medallion owners with owner-driver medallions may not be constrained to 12 hour shifts. Their constraints depend on whether there is a second driver who leases (or sub-leases) the cab. Drivers who lease taxis with fleet medallions daily are so constrained.

Table 3: Mean Hours, Income, and Average Hourly Earnings. By Shift

Shift	Hours	Income	Wage	# Shifts
Day	9.48 (3.97)	248.41 (106.03)	26.39 (5.75)	2247417
Night	8.78 (3.05)	262.03 (93.06)	30.13 (6.55)	2137499
Unassigned	8.30 (5.33)	228.12 (138.46)	28.26 (8.47)	662427

Note: Day shifts start between 4AM and 9:59AM. Night Shifts start between 2PM and 7:59PM. Shifts starting at other times are Unassigned. Standard deviations are in parentheses.

on unanticipated daily earnings opportunities is revealed. In contrast, night shift drivers can experience the evolution of earnings opportunities and decide when to quit for the day. If labor supply is affected by unanticipated transitory variation in earnings opportunities and if information about these earnings opportunities are learned by the driver only after he has started the shift, then night shift drivers will have more opportunity than day shift drivers to adjust their labor supply in response. Operationally, this suggests that estimated labor supply elasticities could be larger for night shift drivers.

Table 3 contains mean hours, income, and average hourly earnings by shift type.²² Day shifts are longer than night shifts by about 0.7 hours, but more money is earned on night shifts (about \$13 more). Average hourly earnings (the wage) is about \$3.73 higher on the night shift. All differences in means across shifts are statistically significant (p -value $< 10^{-100}$).

Labor supply and earnings have distinct patterns over the week by shift. Figure 4 contains plots of the average shift length by day of week for day and night shifts.²³ Day shift drivers work longest on Sundays, with average hours declining from about 10 hours on Sunday to 9 hours on Saturday. In contrast night shift drivers work their shortest days on Sunday, with average hours increasing sharply from about 8 hours per shift early in the week to about 10 hours per shift on Friday and Saturday. Day shift drivers work longer hours than night shift drivers on all days but Friday and Saturday. Total income per shift, shown in figure 5, generally follows hours: Day shift drivers earn the most on Sunday and the least on Saturday, and night shift drivers earn the least on Sunday and the most on Friday and

²² Income is defined as the sum of fare income and surcharges. Tip income is not included.

²³ Some shifts span days of the week. I assign each shift to a particular day of the week based on the date of the first trip in the shift.

Figure 4: Distribution of Shift Length in Hours, by Day of Week and Shift

Figure 5: Distribution of Shift Income, by Day of Week and Shift

Saturday. Day shift drivers earn more than night shift drivers early in the week, while night shift drivers earn relatively more Thursday through Saturday. Average hourly earnings, are shown in figure 6. Night shift drivers earn more per hour throughout the week, with the gap being larger late in the week.

The difference in earnings patterns between day and night shift drivers over the course of the week are substantial and says something more broadly about variation in demand over the course of week. The fact that average hourly earnings are higher on night shifts throughout the week is reflected in higher caps set by the TLC on daily taxi lease rates for night shifts than for day shifts. Additionally, the cap on lease rates for night shifts are higher

Figure 6: Average Hourly Earnings, by Day of Week and Shift

later in the week than earlier in the week.²⁴

Given the sharp differences in labor supply and earnings patterns for day shift and night shift drivers and the potential differences in available information regarding earnings opportunities when making labor supply decisions, I analyze the labor supply of day shift and night shift drivers separately in what follows.

5.3 The CBLT Analysis

Because my analysis is meant as a replication and extension (with new data) of the analysis of CBLT (1997), I start with a short summary of their analysis and results. CBLT base their analysis of taxi driver labor supply on three samples:

1. TRIP – 70 shifts for 13 drivers (8 with multiple shifts) from April 24 - May 14, 1994.
2. TLC1 – 1044 shifts for 484 drivers (234 with multiple shifts) from October 29 - November 5, 1990.
3. TLC2 – 712 shifts for 712 drivers (none with multiple shifts) from November 1-3, 1988.

In the simplest terms, for each of their samples, CBLT regress log hours on a given shift (defined as the difference in time between the drop-off time of the last fare and the pick-up time of the first fare) on log average hourly earnings (defined as the ratio of total fare

²⁴ The lease rate cap for day shifts has been \$115 since October 2012. The lease rate cap for night shifts varies by day of week. The night shift cap since October 2012 ranges from \$128 Sunday-Tuesday to \$142 Thursday-Saturday.

income divided by hours worked). These regressions also control for other shift characteristics including weekday vs. weekend, weather, day vs. night shift, and, in some specifications, driver fixed effects. Presumably, day-to-day variation in hourly earnings controlling for the other shift characteristics is the result of unanticipated transitory factors affecting earnings opportunities. Using OLS, the estimated labor supply elasticities found by CBLT range from -0.618 to -0.186, depending on sample used and other variables included in the model.

As CBLT recognize, the OLS estimate of the elasticity may well be downward biased due to “division bias” if there is any specification or measurement error. This is because the key explanatory variable, average hourly earnings, is calculated as the ratio of daily income to daily hours and daily hours is the dependent variable. There are several reasons to suspect measurement error. First, the data are imperfectly transcribed from potentially erroneous hand recordings of trip sheets. Second, the trip sheets do not record tip income, which likely varies trip-to-trip around some mean tipping rate. Third, there may be breaks taken in the middle of shifts that are not recorded as such. It is unclear how to handle these periods of time in the data. And there are likely to be other sources of error.

CBLT address this problem in a sensible way by instrumenting for average hourly earnings of a given driver with measures of the distribution of hourly earnings of other drivers on the same calendar date. An obvious choice for an instrument for average hourly earnings is the average log hourly earnings of other drivers on the same day, but this was not done. The measures CBLT use as instruments are the 25th, 50th, and 75th percentiles of the other-driver daily earnings distributions. One worry, particularly using the TRIP data, is that there cannot be very many other drivers on any given day on which to base computation of the instruments. Using their IV approach, CBLT find elasticities that are small and insignificantly different from zero for the very small TRIP data but range from -1.313 to -0.926 for the larger TCL1 and TCL2 samples. Interestingly, these IV elasticities are more negative than those found using OLS.

5.4 New Estimates of the Labor Supply Elasticity

I begin by presenting OLS estimates of the labor supply elasticity based on the much larger sample of over 5 million shifts from 2009-2013. These are regressions of log shift duration on log average hourly earnings during the shift and other variables as noted below. I use four samples: 1) all shifts, 2) day shifts, 3) night shifts, and 4) other (unclassified) shifts. These samples are based on the 2/15 random sample of drivers described above, which includes 5,047,343 shifts for 8802 drivers over the 2009-2013 period. For each sample, I estimate

Table 4: Wage Elasticity, OLS Regression of Average Log Daily Hours, by Shift

Model	Controls	Driver F.E's	Elasticity All Shifts	Elasticity Day Shifts	Elasticity Night Shifts	Elasticity Other Shifts
(1)	No	No	0.0159 (0.0154)	0.0485 (0.0177)	-0.0017 (0.0169)	0.0738 (0.0220)
(2)	Yes	No	-0.0034 (0.0177)	0.0505 (0.0203)	-0.0784 (0.0210)	0.0606 (0.0240)
(3)	Yes	Yes	-0.1002 (0.0089)	-0.0615 (0.0109)	-0.1487 (0.0077)	-0.0501 (0.0138)

Note: Each estimated elasticity is from a separate OLS regression. “Elasticity” is the estimated coefficient of log average hourly earnings from a regression of log shift duration. “Controls” include indicators for day of week (6), calendar week (51), year (4), the period subsequent to the September 4, 2012 fare increase (1), and major holiday (1). Estimated using sample of 5,047,343 shifts for 8,802 drivers from 2009-2013. Sample sizes are listed in table 3. Robust standard errors clustered by driver are in parentheses.

three specifications: 1) no controls, 2) a set of controls including indicators for day of week, calendar week, year, the period subsequent to the September 4, 2012 fare increase, and major holiday, and 3) additionally including driver fixed effects.²⁵ These controls account for anticipated wage variation and leave the average hourly earnings measure to account for unanticipated transitory variation in earnings opportunities. I present the estimated coefficient of log average hourly earnings, which is interpreted as the wage elasticity of labor supply.

Table 4 contains these OLS estimates. The estimates for the entire sample, contained in the first column of the table, show elasticities that are small and insignificantly different from zero in the first two specifications. When driver fixed effects are included, the estimated elasticity is negative and statistically significant but relatively small at -0.1. For day shifts, the estimated elasticities are small and positive but statistically significant in the first two specifications. When driver fixed effects are included, the estimated elasticity is again negative and statistically significant though small. The pattern for night shifts is that the elasticities are significantly negative when the controls are added. The estimated elasticity when driver fixed effects are included is more negative than for the day shift. The estimates for the unclassified (other) shifts are very close to those for day shifts.

While I do find some negative elasticities, none approach minus one as suggested by a target earnings model. My elasticities are much smaller in magnitude than those found using

²⁵ See footnote 15 for a list of major holidays used in defining the major-holiday indicator.

OLS by CBLT or Farber (2005), which may reflect a lower level of measurement error in my administrative data compared with the transcribed trip sheet data used in the earlier work.

Although the administrative data may have less measurement error than data derived from the paper trip sheets, it is not error free. Simple consistency checks of the data show more than a few instances of trips ending before they start and new trips starting before the previous trip ends.²⁶ Additionally, as I mentioned earlier, my income data do not include tips, which surely vary across trips as a proportion of fares (Haggag and Paci, 2014). On this basis, it makes sense to estimate the model using an instrument for average hourly earnings.

In the spirit of CBLT, I use as an instrument the average across other drivers of log average hourly earnings on the day each shift started. To avoid problems using an instrument derived from the dependent variable in the estimation sample, I use a non-overlapping randomly selected 2/15 subset of the drivers to generate the instruments.²⁷ The average of log average hourly earnings of shifts starting on date t in the non-overlapping sample ($\overline{\ell n W}_t$) serves as the instrument for the log average hourly earnings of driver i in my estimation sample for shifts that start on date t ($\ell n W_{it}$).

In order for $\overline{\ell n W}_t$ to be a valid instrument, two conditions need to be satisfied. First, $\overline{\ell n W}_t$ needs to be strongly correlated with $\ell n W_{it}$. While I don't present the first-stage estimates here, this is clearly satisfied here.²⁸ Second, the instrument must satisfy the exclusion restriction. In this case, $\overline{\ell n W}_t$ should not be correlated with log hours other than through $\ell n W_{it}$. The motivation to use an instrument here is the potential for division bias mainly due to measurement error. Since it is reasonable to assume that measurement errors are uncorrelated across drivers, the instrument satisfies the exclusion restriction on these grounds.

There is a more subtle argument for potential violation of the exclusion restriction. To the extent that $\overline{\ell n W}_t$ reflects unobserved demand conditions and information about these demand conditions is communicated across drivers, labor supply of individual drivers might be affected directly by this information. To the extent that this effect works through individual average hourly earnings (as it likely would to a substantial degree), this is not a violation of the exclusion restriction. However, to the extent that individual drivers use $\overline{\ell n W}_t$ to adjust their labor supply independently and individual average hourly earnings are

²⁶ I used some simple algorithms to adjust the data to eliminate these inconsistencies and serious outliers.

²⁷ This sample contains 115,733,041 trips on 5,012,244 shifts for 8,768 drivers.

²⁸ The first-stage t-statistic on the instrument is generally greater than 100, and the coefficient on the instrument in the first stage is generally close to one.

Table 5: Wage Elasticity, IV Regression of Average Log Daily Hours, by Shift

Model	Controls	Driver F.E's	Elasticity All Shifts	Elasticity Day Shifts	Elasticity Night Shifts	Elasticity Other Shifts
(1)	No	No	0.2288 (0.0101)	0.0202 (0.0134)	0.3484 (0.0117)	0.2913 (0.0306)
(2)	Yes	No	0.5709 (0.0100)	0.3683 (0.0119)	0.6182 (0.0132)	0.9383 (0.0329)
(3)	Yes	Yes	0.5890 (0.0099)	0.3672 (0.0112)	0.6344 (0.0124)	0.8751 (0.0281)

Note: Each estimated elasticity is from a separate IV regression. The instrument for average hourly earnings is the average of average hourly earnings for a non-overlapping sample of drivers on the same day. “Elasticity” is the estimated coefficient of log average hourly earnings from a regression of log shift duration. “Controls” include indicators for day of week (6), calendar week (51), year (4), the period subsequent to the September 4, 2012 fare increase (1), and major holiday (1). Estimated using sample of 5,047,343 shifts for 8,802 drivers from 2009-2013. Sample sizes are listed in table 3. Robust standard errors clustered by driver are in parentheses.

not perfectly correlated with the overall average, then there is scope for violation of the exclusion restriction. This stands as a potential caveat to the IV analysis.²⁹

The IV estimates of the labor supply elasticity are contained in Table 5. The results are striking in comparison with the OLS estimates in table 4. The estimated elasticities are substantially positive and strongly statistically significant. Adding the control variables raises the estimated elasticity for each sample, but controlling for driver fixed effects does not have much effect. The estimated elasticity on the day shift is about 0.36 while the elasticity on the night shift is about 0.62. The larger elasticity for the night shift is consistent with the observation that drivers on a night shift are more likely than drivers on a day shift to be able to adjust hours mid-shift in response to new information regarding earnings opportunities. Interestingly, the elasticity is even larger on unclassified shifts. It may be that these other shifts are less likely to be worked by lease drivers and more likely to be worked by owner-operators who have more flexibility in selecting hours.

Overall, the evidence presented so far is consistent with the neoclassical optimizing model. The positive estimated elasticities do not support the idea that reference dependent preferences or target earnings behavior are important factors in taxi driver labor supply decisions.

²⁹ In section 7, I present estimates of a discrete choice stopping model that does not depend on an instrument, and it yields results consistent with the IV estimates I present here.

6 Does the Labor Supply Elasticity Depend on How the Wage Compares With the Expected Wage?

The theoretical discussion in section 2 set bounds on the range within which one would expect to find behavior influenced by the gain-loss (reference dependent) component of utility. If the realized daily wage lies in the range defined by equations 11 and 12, then gain-loss utility is relevant and target earnings behavior may be observed. Otherwise, the labor supply elasticity with respect to unanticipated transitory wage variation will be positive and determined by consumption (neoclassical) component of utility. Intuitively, reference dependence is a local phenomenon. If the wage is far lower than what was expected, drivers will find it optimal to stop working before the reference income level is reached, and, if the wage is far higher than what was expected, drivers will find it optimal to continue working after the reference income level is reached.

In this section, I examine how the estimated labor supply elasticity varies with the level of unanticipated wage variation (the absolute deviation of the average daily log wage from its expected value). I calculate the expected log wage for each day using data on mean daily log average hourly earnings for drivers in the non-overlapping sample that I used to construct the instrument for the estimation of the labor supply model in section 5.4. The expected log wage is calculated as the predicted value of log average hourly earnings from an OLS regression of the daily average log average hourly earnings on indicators for day of week, week of year, year, the period after the fare increase of September 4, 2012 and major holiday. I then calculate the difference between the observed average daily log average hourly earnings in the non-overlapping sample and the predicted value. This difference is what I use as the deviation of the average daily log wage from its expectation.

Across the 1826 days in the sample from 2009-2013, the average deviation is zero by construction. Interestingly, the deviations appear relatively small. The average of the absolute deviation is 0.033, and the inter-quartile range of the absolute deviation runs from 0.011 to 0.043. The 90th percentile of the absolute deviation is 0.067 and the 95th percentile is 0.093. In other words, less than 5 percent of the days considered have an observed deviation from expected average hourly earnings of 10 percent or more.

The bounds on what is a sufficiently small deviation from the expected log wage are defined in equations 11 and 12. These bounds depend on the importance of gain-loss utility in the utility function (equations 1 and 2), which is controlled by the parameter α and by the neoclassical labor supply elasticity, which is controlled by the parameter ν . The parameter

α is directly related to the coefficient of loss aversion (λ) used in the behavioral economics literature. The coefficient of loss aversion is defined as the ratio of marginal utility below the reference point to marginal utility above the reference point. In the utility specification used here (equations 1 and 2), coefficient of loss aversion is $\lambda = \frac{(1+\alpha)}{(1-\alpha)}$, which implies that $\alpha = \frac{(\lambda-1)}{(\lambda+1)}$.

The bounds defined in equations 11 and 12 combined with definition of λ as a function of α in the previous paragraph imply that size of the range of deviations where reference dependence is relevant is

$$\ell n(W^{**}) - \ell n(W^*) = \frac{\ell n \lambda}{1 + \nu} \quad (13)$$

where ν is the inverse labor supply elasticity. Existing evidence, mostly from laboratory experiments, suggests that the coefficient of loss aversion (δ) is in the range of 1.5 to 2.5.³⁰ Assuming an elasticity of labor supply of $1/\nu = 0.5$ and $\lambda = 1.5$ implies a range of deviations where reference dependent preferences are relevant of 0.135 (roughly ± 0.0675). Alternatively, assuming an elasticity of labor supply of $1/\nu = 0.5$ and $\lambda = 2.5$ implies a range of 0.3 (roughly ± 0.15). A larger labor supply elasticity would lead to a larger range.

Even the narrower of the bounds I calculate here (± 0.0675) are quite substantial relative to the observed range of unanticipated variation in daily average hourly earnings. Only 101 of the 1826 days had unanticipated variation in the average log wage outside this range. In other words, virtually all days (about 94.5 percent) saw unanticipated wage variation small enough to imply reference dependence and target earnings behavior at these reasonable parameter values. Intuitively, since most of the observed unanticipated variation in the wage is likely within the range where target earnings behavior is relevant, strong hints of such behavior ought to be observed in the full sample estimates, and the fact that the full-sample estimates of the labor supply elasticity are strongly positive suggests that reference dependence is not playing a large role in taxi driver labor supply.

These calculations notwithstanding, I investigate how the estimated elasticity varies with the daily level of unanticipated wage variation by estimating separate labor supply functions for days where the absolute deviation between average log hourly earnings and expected log average hourly earnings is very small or is larger. I split the sample into three subgroups: 1) days with an absolute deviation in the bottom 25 percent of days (an absolute log wage deviation less than 0.009), 2) days with an absolute deviation in the second quartile of days

³⁰ See, for example, Tversky and Kahneman (1991), Abdellaoui Bleichrodt and Paraschiv (2007), and Tovar (2009). A recent field study comparing taxes and bonuses as incentives for the use of reusable bags at supermarkets (Homonoff, 2014), finds a much higher value for λ of 5 or more in this low stakes setting.

Table 6: Wage Elasticity, IV Regression of Average Log Daily Hours, by Shift Subsamples by Deviation of Average Log Daily Wage from Expected Value

Sample	Deviation Percentile	Elasticity	Elasticity	Elasticity
		All Shifts	Day Shifts	Night Shifts
(1)	0-25	0.1398	-0.1045	0.1095
	($N = 1287951$)	(0.0668)	(0.1040)	(0.0751)
(2)	25-50	0.4249	0.1463	0.5625
	($N = 1274650$)	(0.0310)	(0.0414)	(0.0384)
(3)	50-100	0.5809	0.3784	0.6268
	($N = 2484742$)	(0.0102)	(0.0121)	(0.0135)

Note: The 25th percentile and median across days of the absolute deviation of the average log daily wage from its expected value are 0.00896 and 0.01834 respectively. The expected value is the predicted value from a regression of average log hourly earnings on indicators for day of week, week of year, year, the high fare period (on or after September 4, 2012), and major holiday. Each estimated elasticity is from a separate IV regression. The instrument for average log hourly earnings is the average of average log hourly earnings for a non-overlapping sample of drivers on the same day. “Elasticity” is the estimated coefficient of log average hourly earnings from a regression of log shift duration which additionally includes the set of variables listed in the note to table 5. The listed sample sizes are for the “All Shifts” samples based on the underlying sample of 5,047,343 shifts for 8,802 drivers from 2009-2013. Robust standard errors clustered by driver are in parentheses.

(an absolute log wage deviation between 0.009 and 0.01834), and 3) days with an absolute deviation above the median of days (an absolute log wage deviation larger than 0.01834).³¹ My view is that absolute deviations from the expected wage in the first two groups are so small relative to the calculated bounds that they should provide the reference dependent preference model a fair chance to exhibit predictive power.

Table 6 contains IV estimates of labor supply elasticities for subsamples of the sample of shifts I have been using broken down by days where the absolute deviation the average of log average hourly earnings from its daily expected value is lower or higher than the threshold described above. These estimates are calculated first using all shifts and then separately for day shifts and night shifts. The results in the first row of the table, for days on which averages wages are relatively close to their predicted value (the 25 percent of days with smallest absolute deviation), show small elasticities that are not significantly different from zero for either day shifts or night shifts. The estimated elasticities on days with intermediate levels of difference between average wages and their predicted values (from

³¹ The distribution considered uses the date the shift started as the operative date. No adjustment is made for the fact that some shifts span calendar days, and the same distribution is used for all shifts, regardless of whether they are day or night shifts.

the 25th to the 50th percentile of the distribution across days), contained in the second row of the table, are positive, larger, and statistically significant. Finally, the third row of the table contains estimates for the for days with relatively larger deviations of average wages from their predicted values (above the median), and these elasticities are quite large, ranging from 0.38 for day shifts to 0.63 for night shifts. As before, elasticities are uniformly larger on night shifts than day shifts.

The finding that elasticities are smaller on days with small deviations than on days with large deviations suggests that there may be some reference dependence on days with small deviations of the wage from its expected value. However, the fact that the smaller estimated elasticities are apparent only for values of the absolute deviation of the average log earnings from its expected value smaller than (optimistically) the median value of 0.18 implies that the coefficient of loss aversion would be very small. Returning to the expression for the range of relevant deviations in equation 13, a range for $\ln(W^{**}) - \ln(W^*)$ of 0.036 (twice 0.018) and a labor supply elasticity of 0.5 implies value of the coefficient of loss aversion of $\lambda = 1.11$. This is smaller than existing estimates of the coefficient of loss aversion in other settings (footnote 30), implying about an 11 percent increase in marginal utility at the reference point.

The results of this analysis suggests a tension between the strength of reference dependent preferences (the coefficient of loss aversion) and the range of wage variation over which such preferences are relevant. The estimates in this section suggest that the elasticity is smaller on days when the absolute deviation between average wages and the expected wage is small. While this is consistent with reference dependence, the magnitudes are such that there cannot be a particularly sharp kink in the utility function at the reference point. In other words, to the extent there is reference dependence in this setting, the coefficient of loss aversion is quite small, which limits that the importance of these preferences for labor supply behavior.

7 The Discrete-Choice Stopping Model

While estimating the slope of the relationship between hours of work and average hourly earnings is informative about models of labor supply, modeling the number of hours worked on a particular day in this way is made difficult by the fact that the available wage at any point in time not constant (or monotone decreasing or increasing) within the day due to variation in demand during the day. In this situation it may be appropriate to think of the end of each fare is a decision point for the driver: The driver can continue to work or the driver can end the shift. This is a dynamic discrete choice problem based on a driver's

comparison of the marginal utility of stopping with the marginal utility of continuing and making the same assessment after successive fares. A reduced form version of this comparison at any point in the shift should account for variation in prospective earnings opportunities (to account for variation in the benefit of continuing), hours worked to that point (to account for increasing disutility of work during a shift), and other variables that could affect preferences for work (e.g., hour of day). To the extent there are reference dependent preferences or important daily income effects, the reduced form comparison should also include a measure of income earned to that point on the shift.

The neoclassical labor supply model and the model with income reference dependence have sharp and contrasting predictions for this modeling approach. The neoclassical model implies that the probability of ending a shift after a given trip will be positively related to accumulated hours (conditional on accumulated income) and unrelated to accumulated income (conditional on accumulated hours and assuming income effects are zero with regard to daily wage variation). In contrast, the model with preferences that are dependent on an income reference level implies that the probability of ending a shift after a given trip will be positively related to accumulated income (conditional on accumulated hours) and less strongly related to accumulated hours (conditional on accumulated income).³²

Following Farber (2005), a reasonable approximate solution to the dynamic stopping problem can be implemented empirically as a simple discrete choice problem. At any point τ during the shift, a driver can calculate the forward-looking expected optimal stopping point, τ^* . The optimal stopping point may be a function of many factors including hours worked so far on the shift and expectations about future earnings possibilities. If daily income effects are important, the optimal stopping point may also be a function of income earned so far on the shift. A driver will stop at τ if $\tau \geq \tau^*$ so that $\tau - \tau^* \geq 0$.

A reduced form representation of $R(\tau) = \tau - \tau^*$ on a given shift is

$$R_{it}(\tau) = \gamma_1 h_\tau + \gamma_2 y_\tau + X_{it}\beta + \mu_i + \epsilon_{it\tau}, \quad (14)$$

where i indexes the particular driver, t indexes time (hour of day, day of week, week of year, year), μ_i is a driver fixed effect, and ϵ_i is a random component. The quantity h_τ is a vector of indicators for ranges of hours worked at τ on the shift, y_τ is a vector of indicators for ranges of income earned at τ on the shift, and X measures other factors affecting the

³² The dual reference point approach suggested by Koszegi and Rabin (2006) and implemented by Crawford and Meng (2011) suggests that behavior will be governed by both an income target (like the standard reference dependent preferences model) and by an hours target (with implications more like the neoclassical model). I discuss the implications of my estimates in the context of this model later in this section.

determination of the optimal stopping time and the comparison with τ . Elements of the vector X_{it} include sets of fixed effects for hour of the day by day of the week (168), week of the year (52), and year (5) as well as indicators for the period subsequent to the September 4, 2012 fare increase and major holiday. These are included to capture systematic variation in earnings opportunities and utility of leisure from continuing to drive.

The standard discrete choice model implies that the individual stops driving at τ if $R_{it}(\tau) \geq 0$, and, with appropriate distributional assumptions this implies a typical probit or logit specification based on the latent variable defined in equation 14. However, the presence of thousands of individual driver fixed effects makes estimation of a latent variable model impractical, and I proceed using a linear probability model.³³

I estimate the model separately for day and night shifts for the 2/15 sample of drivers used earlier. I drop the shortest one percent and longest one percent of shifts. The resulting sample contains 51,021,936 trips for 6312 drivers on 2,201,443 day shifts and 49,661,892 trips for 6380 drivers on 2,116,675 night shifts.³⁴ Both day and night shifts average about 23 trips, with day shifts averaging 8.8 hours and night shifts averaging 8.2 hours. Drivers earned an average of \$244 per day shift and \$261 per night shift.

Table 7 contains estimates of the marginal effects of accumulated income and hours on the probability of a shift ending after a specific trip relative the base (lowest) categories from the linear probability model specified above. Columns 1 and 2 contain estimates of the effects of accumulated income on day and night shifts respectively. As shift income accumulates (conditional on hours) during day shifts (column 1), the probability of stopping increases slowly by 6 percentage points at \$300 relative to less than \$100. However, there is no increase in the probability of stopping as income accumulates during night shifts (column 2), with the stopping probability at \$300 no different than the stopping probability at less than \$100 (conditional on hours).

Columns 3 and 4 of table 7 contain the analogous results for accumulation of hours (conditional on income). These show sharply increasing increases in the probability of stopping with accumulated hours on both day and night shifts. On day shifts (column 3), as hours

³³ The problem with the logit or probit model in this case is not the usual problem of lack of sufficient observations within each group (driver) for consistency. The problem is simply the scale of estimating a probit or logit model with tens of millions of observations and thousands of coefficients. I have estimated probit versions of the model without the driver fixed effects but with all of the other variables, and the results (in terms of marginal effects of the variables on the probability of stopping) are very close to those found using a linear probability model with the same variables.

³⁴ Of the 7,578 drivers in the sample, 5,114 have driven both day and night shifts and so are represented in both samples.

Table 7: Marginal Effects of Income and Hours on Probability of Ending Shift

Linear Probability Model					
Income (\$)	(1) Day Shift	(2) Night Shift	Hours	(3) Day Shift	(4) Night Shift
100-149	0.0001 (0.0003)	-0.0045 (0.0003)	3-5	0.0020 (0.0004)	-0.0049 (0.0003)
150-199	0.0044 (0.0006)	-0.0077 (0.0005)	6	0.0001 (0.0007)	0.0007 (0.0006)
200-224	0.0157 (0.0010)	-0.0062 (0.0007)	7	0.0034 (0.0011)	0.0223 (0.0010)
225-249	0.0264 (0.0013)	-0.0046 (0.0008)	8	0.0281 (0.0017)	0.0536 (0.0016)
250-274	0.0389 (0.0017)	-0.0042 (0.0011)	9	0.0750 (0.0025)	0.0897 (0.0022)
275-299	0.0506 (0.0020)	-0.0033 (0.0013)	10	0.1210 (0.0035)	0.1603 (0.0031)
300-349	0.0596 (0.0024)	-0.0027 (0.0017)	11	0.1236 (0.0050)	0.2563 (0.0051)
350-399	0.0607 (0.0028)	0.0011 (0.0024)	12	0.1004 (0.0078)	0.2573 (0.0142)
≥ 400	0.0702 (0.0034)	0.0101 (0.0035)	≥ 13	0.1093 (0.0050)	0.2406 (0.0063)

Note: Based on estimates of two linear probability models for the probability of stopping: day shifts (columns 1 and 3) and night shifts (columns 2 and 4). The base category for income is \$0-\$99 and the base category for hours is 0-2. Both models additionally include sets of fixed effects for driver, hour of the day by day of the week (168), week of the year (52), and year (5) as well as indicators for the period subsequent to the September 4, 2012 fare increase and major holiday. Robust standard errors clustered by driver are in parentheses. See text for information on sample size and composition.

accumulate (conditional on income), the probability of stopping increases by 12 percentage points from early in the shift to hour 11. The effect of hours is even larger on night shifts (column 4), where the probability of stopping increases by 25 percentage points from early in the shift to hour 11.

In order to quantify the extent to which shocks to income and hours affect the probability of shift ending, I used the estimates of the models to predict the probability that a shift continues (the survival probability) for each category of income and hours.³⁵ I then

³⁵ The survival probability is constructed by taking one minus the predicted probability of stopping after each trip and creating the running product of this quantity during each shift.

Figure 7: Effect of 20 percent Income Shock on Shift Survival

compute the survival probabilities for each of two counterfactuals. In the first counterfactual, I shock income by increasing the observed fare on each trip by 20 percent (holding trip time constant). I then recalculate accumulated income and compare the predicted survival probability by hour with that predicted by the base (observed) data. It is important to note that this is a shock to income but not to the expected wage. The expected wage is accounted for by the date-time variables and other controls in the model. To the extent that reference dependence, and perhaps income targeting, is important, survival probabilities at given hours should be lower due to the income shock. Without reference dependence (or important daily income effects), survival probabilities should be unaffected by the income shock.

Figure 7 contains separate plots for day and night shifts of the effect on shift survival of the 20 percent uniform income shock. This shows a small reduction in the survival probability throughout day shifts (about 2.4 percentage points by 12 hours) and virtually no reduction (0.4 percentage points by 12 hours) on night shifts.

In the second counterfactual, I shock hours by increasing the observed time on and between each trip by 20 percent (holding fare income constant). I then recalculate accumulated hours and compare the predicted survival probability by income category with that predicted by the base (observed) data. Note that this is a shock to hours worked but not to the expected wage, which is accounted for by the date-time variables and other controls in the model. With reference dependence on daily income resulting in income targeting, the prediction is that the increase in hours worked at a given income level will not substantially

Figure 8: Effect of 20 percent Hours Shock on Shift Survival

affect the likelihood of a shift surviving to earn that amount. The neoclassical model predicts that the increase in hours required to earn a given amount will be associated with reduced likelihood of a shift surviving to earn that amount.

Figure 8 contains separate plots for day and night shifts of the effect on shift survival of the 20 percent uniform hours shock. This shows a substantial reduction in the survival probability throughout both day and night shifts. The survival probability at \$300-349 is about 7 percentage points lower on day shifts and 22 percentage points lower on night shifts when hours are 20 percent higher.

The conclusion from this analysis complements the results of the estimation of the slope of the labor supply function. The evidence is largely consistent with the neoclassical model. Hours are the important driver in the stopping decision, with some evidence of a marginal effect of income (and, hence, income reference points) on the stopping decision on day shifts only.

In order to investigate how accumulated hours and income affect drivers' stopping decisions in a less constrained way, I estimated a version of the stopping model that interacts the accumulated hours and income indicators. In other words, rather than the 10 income and 10 hours categories listed in table 7, this version of the model included the full set of interactions. The model additionally includes driver fixed effects as well as the other controls listed in the note to table 7.³⁶ Figure 9 contains separate plots for day shifts (left panel)

³⁶ In order to limit the computational burden, this model is estimated using a 1/15 random sample of

Figure 9: Marginal Effects of Hours and Income on Shift End Probability

and night shifts (right panel) of the estimated marginal effect of each hours-income category on the probability of a shift ending. For each hours-income category, the size of the square marker is proportional to the magnitude of the estimated marginal effect on the stopping probability relative to the lowest income-hours category (less than \$100, 0-2 hours).³⁷

The estimates for day shifts, in the left panel of figure 9, show that the probability of stopping rises with hours within income categories. The probability of stopping also rises somewhat with income within hours categories but less sharply than the rise with hours. This suggests that there may be some role for daily income reference points on day shifts but that the central determinant of a driver's stopping decision is accumulated hours. The estimates for night shifts, in the right panel of the figure, show a clearer pattern. The probability of stopping on night shifts rises sharply with hours within income categories but the stopping probability is roughly unaffected by income within hours categories. These patterns suggest a limited role, at best, for income reference dependence.

drivers rather than the 2/15 random sample of drivers used in the earlier estimation. In order to investigate whether any differences I might find are due to the difference in sample, I also re-estimated the model shown in table 7 using this smaller sample, and the results are virtually identical to those shown in the table.

³⁷ The missing points (5 each on day shifts and night shifts) are for hours-income combinations that are observed for ten or fewer trips (out of a total of approximately 25 million trips on each of day shifts and night shifts). These are all for unlikely combinations of hours and income (very low income with long hours or very high income with short hours).

Crawford and Meng’s (2011) analysis of taxi driver labor supply posits reference points in both daily income and daily hours, with loss aversion to both income below an income reference point and hours above an hours reference point. The idea, based on Koszegi and Rabin (2006), is that there are two “domains of losses”: 1) if income is below the income reference point then the individual is within the domain of losses for income and has a higher marginal utility of income and 2) if hours are above the hours reference point then the individual is within the domain of losses for hours and has a higher marginal disutility of work (a higher marginal utility of leisure). While the theory does not have unambiguous predictions on how the decision to end a shift depends on income and hours, the model does suggest in rough terms that the likelihood of stopping 1) will be lowest when income is below the income reference point and hours are below the hours reference point (within the domain of losses in income and outside the domain of losses in hours), 2) will be highest when income is above the income reference point and hours are above the hours reference point (outside the domain of losses in income and within the domain of losses in hours), and 3) will be intermediate in the other cases.

With regard to the estimates shown in figure 9, this is weakly consistent with the pattern for day shifts (left panel) where the probability of stopping increases from southwest to northeast as both income and hours increase. However, as I noted earlier, the strongest relationship is in the hours dimension (south to north) so that the probability of stopping is increasing strongly in hours and weakly in income. The pattern for night shifts (right panel) does not provide support for dual reference point model because, while the stopping probabilities increase from south to north, indicating a strong increase with hours at all income levels, the stopping probabilities are flat going west to east, indicating no relationship between the stopping probability and accumulated income at each level of hours.

I also used the estimates from the model with income-hours interactions to recalculate the income and hours counterfactuals shown earlier, and the results (not shown here) are virtually identical to those shown in figures 7 and 8. Shift survival probabilities, and, by extension, shift durations are unaffected by a 20 percent income shock while income earned is substantially lower when hours at a given income level are 20 percent higher. These patterns provide additional evidence consistent with the conclusions I drew from figures 7, 8, and 9.

To summarize, the pattern of results for variation in the stopping probability with hours and income and for the two counterfactuals is consistent with the predictions of the neo-classical model and with my findings based on estimating the labor supply elasticity. The

pattern of results is not consistent with the basic income reference dependent preferences model in that 1) the stopping probability varies only a small amount with income conditional on hours on day shifts and does not vary at all with income conditional on hours on night shifts and 2) there is virtually no effect of the counterfactual income shock on labor supply. The pattern of results is consistent with reference dependence in daily hours in that 1) the stopping probability increases substantially with hours conditional on income on both day shifts and night shifts and 2) income is lower on shifts where counterfactually hours are higher at a given income level. However, it is difficult to see how this differs from the pattern predicted by the neoclassical model.

8 Do Different Drivers Use Different Models?

Economists typically assume a single model of behavior applies to all agents and estimate the parameters of said model. The scale of the data available here makes it feasible to estimate separate labor supply models for individual drivers. It may be that some drivers exhibit reference dependent preferences (substantial negative labor supply elasticities) while others are optimizers (positive labor supply elasticities).³⁸

In this section, I return to estimating the slope of the relationship between log hours worked and the log wage and estimate separate labor supply models by driver for the large number of drivers who are observed on a substantial number of shifts. Given the difference in decision making margins and estimated elasticities for day shifts and night shifts, I make a distinction between “day shift drivers” and “night shift drivers.” I classify drivers who work at least 750 day shifts between 2009 and 2013 as day shift drivers and drivers who work at least 750 night shifts between 2009 and 2013 as night shift drivers. There are 1267 day shift drivers and 1205 night shift drivers in the 2/15 random sample. While most drivers are observed working different shifts over the period, drivers do tend to specialize. For example, there are only two drivers who are classified as both day shift and night shift. The mean number of night shifts for day shift drivers is 46.6 and the mean number of day shifts for night shift drivers is 55.8. I proceed estimating separate labor supply models for the day shifts of the day shift drivers and for the night shifts of the night shift drivers.

The left panel of figure 10 contains kernel density estimates of the distributions across

³⁸ Doran (2014) discusses some estimates that allow for a driver-specific relationship accumulated between income and the probability of stopping for a sample of 66 drivers who are observed for at least 500 trips (about 20 shifts). He finds that about half of these drivers respond to a positive transitory income shock with an increase in the likelihood of stopping, implying a negative wage elasticity of labor supply.

Figure 10: Kernel Density Estimates of Distribution of Estimated Elasticities Over Individual Drivers, Separately for Day Shift and Night Shift Drivers

day shift and night shift drivers of the wage elasticity of labor supply estimated from separate IV regression models for each driver.³⁹ The density estimates are weighted by the inverse sampling variance of the individual estimated elasticities. The right panel of the figure contains the cumulative distribution function implied by the kernel density estimates. It is clear from the figure that there is substantial variation across drivers in estimated labor supply elasticities which reinforces the value of having sufficient data and variation to estimate separate models.⁴⁰

Examination of the distributions of elasticities summarized in Figure 10 yield two strong conclusions. First, there is little evidence of a substantial number of individual drivers having the strongly negative labor supply elasticities implied by reference dependence. Only 0.12 percent of day shift drivers and 0.04 percent of night shift drivers are estimated to have elasticities less than -0.5 and less than 2.1 percent of day shift drivers and 0.24 percent of

³⁹ Each model, a regression for each driver of log shift hours on log average shift hourly earnings, also includes indicators for day of week, week of year, year, the period subsequent to the September 4, 2012 fare increase, and major holiday. The instrument used is as described earlier and used in Table 5: the average of average hourly earnings for a non-overlapping sample of drivers on the same day and shift.

⁴⁰ The variance of these estimated distributions are overstated as measures of the variance of the distribution of the underlying elasticities due to the fact that the sampling errors in estimating the elasticities is included. Given the positive mean of the distribution of elasticities, this implies that the fraction negative is overstated.

night shift drivers are estimated to have elasticities less than -0.25. About 25 percent of day shift drivers and 11 percent of night shift drivers are estimated to have negative elasticities.⁴¹ The average elasticity (weighted by inverse sampling variance) is 0.105 for day shift drivers and 0.321 for night shift drivers.

The second conclusion from figure 10 is that night shift drivers generally have larger elasticities than day shift drivers. In fact, the plot of the two CDFs indicates that the distribution of elasticities for night shift drivers stochastically dominates the distribution for day shift drivers. The generally larger elasticities for night shift drivers is consistent with these drivers being more likely to be able to adjust hours to unanticipated changes in earnings opportunities revealed during the shift. Note that with reference dependence, these same considerations would imply that elasticities would be more strongly negative for night shift drivers as they could work longer when wages are unexpectedly low in order to reach the daily earnings reference point while drivers on day shifts may have more limited ability to extend their shifts.

9 Do Drivers Learn to Optimize?

It is clear that optimizing behavior dominates reference dependence in the sense that individuals who set labor supply optimally will earn more money working fewer hours than individuals who set labor supply as a result of reference dependence. Drivers who are optimizers will work relatively more hours on days when the wage is high, while drivers with reference dependent preferences will work relatively more on days when the wage is low. It may be that experience helps taxi drivers take advantage of high wage days, by teaching them not only to behave as optimizers rather than target earners but also to take better advantage of earnings opportunities by modifying their driving strategies.⁴²

Given the importance of labor income and the time commitment of taxi drivers, the value of learning how to optimize well is potentially very large in this industry.⁴³ In order

⁴¹ These proportions are calculated weighted by the inverse sampling variance of the estimated elasticities.

⁴² Learning to be an optimizer is not a new idea. Camerer, et al (1997) present evidence suggesting that the more experienced taxi drivers in their sample do not exhibit target earnings behavior. List (2003) studied two markets for collectibles and found that experience resulted in a reduction in the importance of the endowment effect. This is related to reference dependence in that both the endowment effect and reference dependence flow from loss aversion.

⁴³ Haggag, McManus, and Paci (2014) find substantial learning among New York City taxi drivers about how to find fares efficiently.

to investigate the importance of learning, I analyze the labor supply of new taxi drivers and how it changes as they accumulate experience. The first step is to identify new drivers. This would be relatively straightforward if the actual hack license numbers of the drivers were available since these numbers are assigned in sequence. However, the hack license numbers are available only in encrypted form. As a result, I identify drivers as new if they were not observed driving for some period from January 2009 forward.

In order to determine a reasonable period of non-observation to consider a driver a new entrant, I analyzed the labor supply patterns over time in the 2/15 random subsample that I relied on above. It turns out that there is a fair amount of entry, exit, and reentry among taxi drivers. For example, about 27 percent of the 8802 drivers in my sample did not drive for at least one three-month period and then returned to driving. About 14 percent did not drive for a six-month period and returned to driving. After not driving for a year, about 5 percent of drivers returned, after 18 months of not driving about 3 percent of drivers returned, and after two year of not driving, about 1.5 percent of drivers returned.

Conservatively, I define new drivers to be those are not observed driving for a full year at the start of my data. In other words, while my data start in January 2009, I use the 2009 data only to identify drivers who are observed driving at some point in the 2010-2013 period but did not drive in 2009. Without a better method to identity new drivers, I live with the likelihood that about 5 percent of drivers I classify as new are, in fact, experienced to a greater or lesser extent. The restriction to drivers who did not have any shifts in 2009 eliminates about 64 percent of drivers and about 87 percent of shifts. In order to ensure a sufficient sample, I use the full sample of drivers who did not have any shifts in 2009 for this analysis, rather than only those in the 2/15 random subsample I have been using.

The sample of new entrants contains 4,814,278 shifts for 24,114 drivers. However, fully 7,038 (29 percent) of these drivers are observed for only a single shift, and 4,555 of these single shifts have only a single trip. While it is possible that 29 percent of new drivers go through the procedure of getting a hack license only to quit after a single trip or shift, it may also be that these are simply data anomalies. In what follows, I delete the 7,038 drivers and their associated 7,038 shifts from the analysis. The remaining analysis sample contains 4,807,240 shifts for 17,076 drivers.

I define experience based on number of shifts driven rather than by the passage of time, and I divide experience into ten categories. I then estimate separate labor supply models for each of the ten categories. Table 8 contains average characteristics of shifts by experience category. There are some interesting patterns. Average hours worked increases early with

Table 8: Mean Shift Characteristics, by Experience Level

Experience	Hours	Income	AHE	Day	Night	N
Week 1	9.06 (0.01)	227.08 (0.34)	24.97 (0.02)	0.53 (0.00)	0.39 (0.00)	112387
Week 2	9.55 (0.01)	252.85 (0.36)	26.56 (0.02)	0.53 (0.00)	0.41 (0.00)	102625
Week 3	9.61 (0.01)	258.80 (0.37)	27.05 (0.02)	0.51 (0.00)	0.42 (0.00)	98241
Week 4	9.63 (0.01)	262.16 (0.37)	27.38 (0.02)	0.50 (0.00)	0.43 (0.00)	95217
Month 2	9.61 (0.01)	266.28 (0.19)	27.87 (0.01)	0.49 (0.00)	0.43 (0.00)	358312
Month 3	9.52 (0.01)	268.54 (0.18)	28.42 (0.01)	0.47 (0.00)	0.45 (0.00)	406779
Month 4-6	9.44 (0.00)	271.84 (0.12)	29.05 (0.01)	0.45 (0.00)	0.46 (0.00)	891679
Month 7-12	9.29 (0.00)	273.80 (0.10)	29.77 (0.01)	0.44 (0.00)	0.47 (0.00)	1243343
Year 2	9.03 (0.00)	277.16 (0.10)	31.01 (0.01)	0.44 (0.00)	0.46 (0.00)	1216563
$\geq$ Year 3	8.70 (0.01)	275.25 (0.22)	31.97 (0.01)	0.46 (0.00)	0.43 (0.00)	282094

Note: Means from the set of all shifts from 2010-2013 for drivers first observed in 2010 or later. Standard deviations of means in parentheses. Means for Day and Night are fraction of shifts that are day shifts and night shifts respectively. The column headed AHE contains mean average hours earnings. The last column (N) contains the number of shifts for each experience level. I define Month 2 as 5-8 weeks, Month 3 as 9-13 weeks, Months 4-6 as 27-52 weeks, and Year 2 as 53-104 weeks, and $\geq$ Year 3 as at least 105 weeks. A week is defined as seven shifts, rounding up.

experience (from weeks 1 to 3) then declines after 6 months. Income per shift increases sharply between weeks 1 and 2 and continues to increase as experience accumulates. Average hourly earnings increases sharply with experience, by about 22 percent between the first week and year 3 or later. Finally, it appears that new drivers are relatively more likely than more experienced drivers to be on day shifts and, conversely, more experienced drivers are relatively more likely than less experienced drivers to be on night shifts. As the data in figure 6 show, average hourly earnings are higher on night shifts, but further analysis (not shown here) suggests that movement toward night shifts accounts for very little of the growth in

average hourly earnings with experience.⁴⁴

What I am interested in here is whether drivers learn to be better optimizers as they gain experience. Evidence for this would be labor supply elasticities that are less likely to be negative and are on average larger positive values as experience accumulates. I estimated separate IV regressions for each of the 10 experience categories using the model with controls described in the note to table 5 and without driver fixed effects.⁴⁵ The top panel of Figure 11 contains a plot of the estimated elasticities for each experience group with the associated 95 percent confidence intervals. The estimated elasticity in the first week is positive and grows steadily from 0.3 in week 2 to about 0.6 after 6 months and to about 0.7 in the third and later years.

I showed earlier that there are substantial differences in labor supply elasticities by shift. The second and third panels of Figure 11 contains plots of the estimated labor supply elasticities by level of experience separately for day and night shifts. The second panel of the figure shows small positive elasticities on day shifts that grow from about 0.1 early on to about 0.45 by year 2. The pattern for night shifts, in the third panel of the figure is more dramatic. The estimated elasticity is about 0.15 in the first week but increases to a significantly positive 0.35 by week 3 and grows further to 0.75 by year 2 and to 0.8 in year 3 or later. This confirms the earlier finding that elasticities are larger on night shifts as drivers can adjust hours more easily after learning about earnings opportunities.

The conclusion from this analysis is that drivers become more responsive to earnings opportunities as they accumulate experience but there is no evidence, even early, of the negative elasticities implied by reference dependence. The general pattern of elasticities increasing with experience suggests that drivers do learn to become better optimizers.

10 Selection: Do Inefficient Drivers Quit Driving Taxis?

Sample selection is an alternative explanation for the patterns of variation of outcomes with experience found in the previous section. Drivers who do not have a substantial positive labor supply elasticity with respect to unanticipated wage changes will find it harder to earn money

⁴⁴ Haggag, McManus, and Paci (2014) use the 2009 TPEP data to investigate learning by New York City taxi drivers, and their results suggest that experienced drivers are more productive because they are able to earn more after ending a trip in an area and at a time where new earnings opportunities are less promising.

⁴⁵ I do not include driver fixed effects when estimating separately by experience category results because inclusion of driver fixed effects in the models whose estimates are shown in Table 5 had only marginal effects on the estimated elasticities.

Figure 11: Wage Elasticity of Labor Supply, IV Estimates by Experience

as a taxi driver. Such drivers (those drivers with small positive labor supply elasticities as well as those drivers with reference dependent preferences yielding negative labor supply elasticities) may be relatively likely to quit the business (stop being taxi drivers). The result will be a sample of drivers that becomes progressively composed of drivers with substantial positive labor supply elasticities as experience accumulates while no particular driver has a growing labor supply elasticity.

The key question is whether or not drivers who exit early have negative or smaller positive labor supply elasticities than do drivers who do not exit early, and I turn now to an analysis of this question. I define exit understanding that a gap in driving does not necessarily imply exit. I use the same algorithm I used earlier in determining who is a new driver to calculate the fraction of new drivers with a gap of various sizes who were observed driving earlier (before the gap). After dropping the 7,038 new drivers who are observed only for a single shift, about 25.5 percent of the remaining 17,076 drivers in my sample of new drivers did not drive for at least one three-month period and then returned to driving. About 12.1 percent did not drive for a six-month period and returned to driving. After not driving for a year, 4.8 percent of drivers returned. Based on these tabulations, in order to minimize errors in determining who is a short-term driver, I require a full year of observation subsequent to the last observed shift for a driver to classify that driver as having exited. For this reason, I drop from the analysis 4,399 drivers who entered the industry on or after January 1, 2013. There are 12,677 drivers who entered between January 2010 and December 2012, of whom 7,064 are observed driving in December 2013. Of the 12,677 new drivers, 1,352 (10.7 percent) drove 7 or fewer shifts, and 2,268 (17.9 percent) drove 28 or fewer shifts.

Table 9 contains IV estimates of labor supply elasticities from a model that interacts log average hourly earnings with a set of indicators for longevity of the driver (total number of shifts recorded for the driver). In order to focus on differences in elasticity at the start of the driving career, this analysis uses only shifts early in the career even for drivers who are observed much longer. The first three columns contain estimates of elasticities on the first 12 shifts recorded for each driver.⁴⁶ These elasticities reflect the responsiveness of labor supply to unanticipated wage variation at the very beginning of a driver's experience. The estimates show that drivers who will quit driving with 12 or fewer shifts have a statistically significantly lower elasticity than drivers in any category with a greater number of total shifts (p -value ≤ 0.01 for all tests). These patterns hold for day shifts and night shifts considered separately

⁴⁶ Drivers who are observed for fewer than 12 shifts necessarily have fewer than 12 observations.

Table 9: Wage Elasticity, IV Regression of Average Log Daily Hours
By Longevity (Total Number of Shifts as Taxi Driver)

Longevity (Total Shifts)	First 12 Shifts			First 30 Shifts		
	All Shifts	Day Shifts	Night Shifts	All Shifts	Day Shifts	Night Shifts
2-12	0.0807 (0.0441)	0.0361 (0.0477)	0.0965 (0.0499)	0.1752 (0.0304)	0.1027 (0.0336)	0.1719 (0.0367)
13-30	0.2036 (0.0416)	0.1063 (0.0457)	0.2065 (0.0474)	0.2912 (0.0282)	0.1674 (0.0321)	0.2771 (0.0347)
151-300	0.2262 (0.0411)	0.1240 (0.0452)	0.2381 (0.0466)	0.3188 (0.0276)	0.1903 (0.0313)	0.3114 (0.0337)
31-60	0.2101 (0.0419)	0.1086 (0.0462)	0.2213 (0.0469)	0.3021 (0.0284)	0.1713 (0.0323)	0.2953 (0.0340)
61-90	0.2111 (0.0420)	0.1123 (0.0466)	0.2201 (0.0471)	0.3025 (0.0284)	0.1734 (0.0326)	0.2957 (0.0339)
91-150	0.2299 (0.0411)	0.1270 (0.0450)	0.2367 (0.0467)	0.3195 (0.0277)	0.1898 (0.0313)	0.3092 (0.0337)
≥ 301	0.2287 (0.0412)	0.1245 (0.0451)	0.2375 (0.0468)	0.3228 (0.0276)	0.1921 (0.0311)	0.3124 (0.0337)
Number Drivers	12677	8722	7318	12677	9159	8091
Number Shifts	139690	74253	54878	331434	171103	135871

Note: Each column presents estimated elasticities from a single IV regression. The elasticities the coefficients of the interaction of $\ln W$ with a set of indicators for total shifts observed for each driver. The instrument set is the interaction for average hourly earnings is the average of average hourly earnings for a non-overlapping sample of drivers on the same day with the set of indicators for total shifts observed for each driver. All models include a set of indicators for day of week (6), calendar week (51), year (3), the period subsequent to the September 4, 2012 fare increase (1), and major holiday (1). Robust standard errors clustered by driver are in parentheses.

as well as for all shifts. Note further that considering up to the first 12 shifts, elasticities are larger on night shifts than on day shifts at every longevity, but these differences by shift are not significantly different from zero at conventional levels.

The next three columns of table 9 columns contain estimates of elasticities on up to the first 30 shifts recorded for each driver.⁴⁷ The pattern of results is qualitatively identical to those for the first 12 shifts. Drivers who will quit driving with 12 or fewer shifts have a statistically significantly lower elasticity on than drivers in any category with a greater level of total shifts (p -value ≤ 0.01 for all tests). When considering up to the first 30 shifts, night

⁴⁷ Again, drivers who are observed for fewer than 30 shifts necessarily have fewer than 30 observations.

Table 10: Mean Shift Characteristics, by Experience Level
Drivers with more than 52 Weeks Total Driving Experience

Experience	Hours	Income	AHE	Day	Night	N
Week 1	9.31 (0.02)	222.65 (0.57)	23.90 (0.03)	0.55 (0.00)	0.39 (0.00)	35597
Week 2	9.69 (0.02)	243.89 (0.57)	25.30 (0.03)	0.54 (0.00)	0.41 (0.00)	35597
Week 3	9.72 (0.02)	248.99 (0.57)	25.75 (0.03)	0.52 (0.00)	0.42 (0.00)	35597
Week 4	9.81 (0.02)	253.76 (0.57)	26.06 (0.03)	0.51 (0.00)	0.43 (0.00)	35597
Month 2	9.75 (0.01)	256.73 (0.29)	26.53 (0.02)	0.50 (0.00)	0.43 (0.00)	150388
Month 3	9.61 (0.01)	258.36 (0.26)	27.07 (0.02)	0.48 (0.00)	0.44 (0.00)	187985
Month 4-6	9.48 (0.01)	262.87 (0.16)	27.96 (0.01)	0.47 (0.00)	0.45 (0.00)	488761
Month 7-12	9.29 (0.00)	270.78 (0.11)	29.45 (0.01)	0.44 (0.00)	0.46 (0.00)	977522
Year 2	9.03 (0.00)	277.16 (0.10)	31.01 (0.01)	0.44 (0.00)	0.46 (0.00)	1216563
$\geq$ Year 3	8.70 (0.01)	275.25 (0.21)	31.97 (0.01)	0.46 (0.00)	0.43 (0.00)	282094

Note: Means from set of all shifts from 2010-2013 for drivers with more than 52 weeks total driving experience first observed in 2010 or later. Standard deviations of means in parentheses. See note to table 8 for detailed definitions.

shift elasticities are significantly greater than day shift elasticities at every longevity (p -value < 0.05 for all tests). While not shown here, I repeated this analysis for the first 60 shifts, and the pattern of results is the same.

To summarize, there is no evidence of the negative elasticities associated with reference dependence even for drivers who quit after a small number of shifts, but it does appear that drivers who quit the business quickly are less responsive to unanticipated wage changes.

Next I investigate whether the patterns of means by experience found in table 8 and the growth of labor supply elasticity with experience shown in figure 11 are due to systematic exit of drivers at low levels of experience. Table 10 contains mean shift characteristics by experience for the subsample of drivers who are first observed in 2010 and have more than 52 weeks total driving experience. This subsample has 3,453,701 shifts for 5,371 drivers. The

general pattern of variation of each outcome with experience for this sample of long-term drivers is identical to that for the sample of all drivers shown in table 8. This is evidence that systematic exit from the industry is not driving the patterns seen in that table, and the fact that average hourly earnings grow with experience suggests that drivers do learn.

The estimates in table 9 show that the estimated elasticities for drivers who remain for the long term, while significantly positive, are relatively small (in the 0.1 to 0.25 range) early in their experience. This suggests that the selective exit of low-elasticity new drivers is not an important factor in the growth of elasticity with experience shown in figure 11. To get at this directly, I computed IV estimates of the wage elasticities by experience for a sample longer-term drivers. To identify the longer-term drivers, I note that most exit happens very quickly, with 12 percent driving 7 or fewer shifts, 15.6 percent driving 12 or fewer shifts, and 22 percent driving 28 or fewer shifts.⁴⁸ The rate of exit slows down substantially after 28 shifts, and I use this value to demarcate longer-term drivers.

The top panel of figure 12 shows the IV estimates of wage elasticities by experience for longer-term drivers, defined as those who are observed to drive more than 28 shifts. These estimates shows the same pattern of increasing labor supply elasticity as was seen in figure 11. The second and third panels of figure 12 contain plots of the estimated labor supply elasticities by level of experience separately for day and night shifts. The general pattern is the same as is shown in figure 11 for the combination of long-term and short-term drivers. Elasticities grow with experience for both day shift drivers and night shift drivers, but they grow more rapidly for night shift drivers and wind up at a much higher level. These estimates support the conclusion that drivers, in fact, learn to optimize and that the increase in elasticity with experience is not an artifact of selection.

To summarize, there is clear evidence that drivers with low labor supply elasticity tend to exit the industry early. Short-term drivers exhibit significantly lower elasticity of labor supply than do long-term drivers at the same (low) level of experience. The evidence on longer-term drivers, that the labor supply elasticity grows with experience over the first few years of driving, is consistent with the view that these drivers learn to optimize over time by learning (how) to adjust their daily labor supply positively in response to unanticipated variation in daily earnings opportunities. There is no evidence in these data that the labor supply decisions of drivers, either those who exit early or those who are longer-term, are substantially influenced by income reference dependence.

⁴⁸ This tabulation is based on the 17,076 new drivers who drive for more than a single shift. The 7,038 drivers who are observed for a single shift are not considered.

Figure 12: Wage Elasticity of Labor Supply, IV Estimates by Experience
(New Drivers With More Than 28 Shifts)

11 Final Remarks

Following on the work of Camerer, Babcock, Loewenstein, and Thaler (1997), I evaluated the role of reference dependence vs neoclassical optimizing behavior in determining daily hours decisions of taxi driver in New York City using the complete driving records of all NYC taxi drivers over the 2009-2013 period. The high-level conclusion is that there is little evidence that reference dependence is an important factor determining the labor supply of NYC taxi drivers. This conclusion rests on several pieces of evidence:

- Koszegi and Rabin's (2006) model of expectations-based reference points implies that reference dependence is relevant only for unanticipated wage movements in the wage while individuals respond neoclassically to anticipated wage movements. However, only about one-eighth of total hourly wage variation is unanticipated. Thus, most variation in wages (the seven-eighths that is anticipated) affects labor supply neoclassically.
- Instrumental variable estimates of labor supply elasticities, based on a very large sample of about 5 million shifts for 8,802 drivers, are strongly positive and range from 0.4 to 0.8, depending on shift. There is no evidence of negative labor supply elasticities or target earnings behavior (a prediction of the reference dependence model).
- The model further predicts that reference dependence will be relevant when unanticipated transitory wage variation is small and not relevant when this variation is large. While I do find that the elasticity of labor supply is smaller on shifts (particularly day shifts) with very little unanticipated wage movement, the range of unanticipated wage variation over which these smaller elasticities are found is so narrow that it implies only a very limited degree of loss aversion (a coefficient of loss-aversion of about 1.1).
- Estimation of a discrete choice model of the probability of stopping after a trip (ending a shift) yields the clear result that the stopping probably is at best weakly correlated with accumulated income but strongly positively related to accumulated hours worked.
- Allowing individual drivers to have their own labor supply elasticities yields a range of estimated elasticities of labor supply, but the large majority of elasticities are estimated to be positive, and only a tiny fraction of drivers are estimated to have substantial negative elasticities.

Reference dependence and target earnings behavior leads to inefficient labor supply decisions. Responding positively to earnings opportunities, as the neoclassical model implies,

is efficient in the sense that more money is earned when the return to working (the wage) is higher. This raises two questions regarding efficiency of work decisions that I have addressed here. First, do new drivers learn to take advantage of strong earnings opportunities by working more on high wage days and less on low wage days? Second, do new drivers who start with negative or small positive labor supply elasticities quit the business? The answer to both questions is yes. The estimated labor supply elasticity grows substantially with experience, and new drivers with small labor supply elasticities are more likely to quit. Those who remain have growing elasticity with experience. In other words, new drivers who continue in the industry learn how to be better optimizers (respond more positively to wage variation). There is no evidence of negative labor supply elasticities for new drivers even at the start.

To summarize, the overall pattern is clear. Drivers tend to respond positively to both anticipated and unanticipated increases in earnings opportunities. This is consistent with the neoclassical optimizing model of labor supply, and I conclude that reference dependence does not play an important role in determining the labor supply of taxi drivers in New York City.

References

- Abdellaoui, Mohammed, Han Bleichrodt, and Corina Paraschiv. "Loss Aversion under Prospect Theory: A Parameter-Free Measurement," *Management Science* 53(10) (October 2007) pp. 1659-1674.
- Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman. "Reference Points and Effort Provision," *American Economic Review* 101(2) (2011) pp. 470-92
- Agarwal, Sumit; Mio Diao; Jessica Pan; and Tien Foo Sing, "Labor Supply Decisions of Singaporean Cab Drivers," mimeo, University of Singapore, October 2013.
- Ashenfelter, Orley, Kirk Doran, and Bruce Schaller. "A Shred of Credible Evidence on the Long-run Elasticity of Labour Supply," *Economica* 77 (2010) pp. 637-650.
- Blundell, Richard and Thomas MaCurdy. "Labor Supply: A Review of Alternative Approaches," in *Handbook of Labor Economics*, O. Ashenfelter and D. Card, eds. Amsterdam, North Holland, vol 3a. 1986: pp. 1559-1695.
- Camerer, Colin, "Taxi Drivers and Beauty Contests." *Engineering & Science*, California Institute of Technology, 1997, pp. 10-19.
- Camerer, Colin."Prospect Theory in the Wild: Evidence from the Field," in *Choices, Values, and Frames*. D. Kahneman and A. Tversky, eds. Cambridge UK: Cambridge University Press, 2000, pp. 288-300.
- Camerer, Colin; Linda Babcock; George Loewenstein; and Richard Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics* 112 (May 1997), pp. 407 – 441.
- Chang, Tom and Tal Gross. "How many pears would a pear packer pack if a pear packer could pack pears at quasi-exogenously varying piece rates?." *Journal of Economic Behavior and Organization* 99 (2014) pp. 1-17.
- Chou, Yuan K. "Testing Alternative Models of Labor Supply: Evidence from Taxi Drivers in Singapore." *The Singapore Economic Review* 47 (2002).
- Crawford, Vincent P. and Juanjuan Meng. "New York City Cab Drivers Labor Supply Revisited: Reference-Dependent Preference with Rational Expectations Targets for Hours and Income," *American Economic Review* 101 (2011) pp. 1912-1932.
- Doran, Kirk. "Are Long-Term Wage Elasticities of Labor Supply More Negative than Short-Term Ones?" *Economic Letters* 122 (2014) pp. 208-210.

- Farber, Henry S. "Is Tomorrow Another Day? The Labor Supply of New York City Cab Drivers," *Journal of Political Economy* 113 (February 2005), pp. 46-82.
- Farber, Henry S. "Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers," *American Economic Review* 98 (June 2008), pp. 1069-1082.
- Fehr, Ernst and Lorenz Goette. "Do Workers Work More if Wages are High? Evidence from a Randomized Field Experiment," *American Economic Review*, 97(1) (2007) pp. 298-317.
- Haggag, Kareem, Brian McManus, and Giovanni Paci. "Learning by Driving: Productivity Improvements by New York City Taxi Drivers." June 2014. Available at SSRN: <http://ssrn.com/abstract=2456733>.
- Haggag, Kareem and Giovanni Paci. "Default Tips." *American Economic Journal: Applied Economics* 6(3) (July 2014): pp. 1-19.
- Homonoff, Tatiana. "Can Small Incentives Have Large Effects? The Impact of Taxes versus Bonuses on Disposable Bag Use." Working Paper No. 575, Industrial Relations Section, Princeton University, March 2013.
(Online at <http://arks.princeton.edu/ark:/88435/dsp014q77fr47j>).
- Kahneman, Daniel and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47 (March 1979), 263-292.
- Koszegi, Botond and Matthew Rabin. "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics* 121 (November 2006), pp. 1133-1165.
- List, John A. "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics* 118 (February 2003), pp. 41 – 71.
- MaCurdy, Thomas E. "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy* 89 (December 1981): 1059-1085.
- Nguyen, Quang and Pingsun Leung. "Revenue Targeting in Fisheries: The Case of Hawaii Longline Fishery," MPRA Paper No. 13846, March 7, 2009.
(Online at <http://mpa.ub.uni-muenchen.de/13846>).
- Oettinger, Gerald S. "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors," *Journal of Political Economy* 107 (April 1999), pp. 360-392.

- Ordonez, Lisa; Maurice E. Schweitzer; Adam D. Galinsky, and and Max H. Bazerman “Goals Gone Wild: The Systematic Side Effects of Over-Prescribing Goal Setting,” *Academy of Management Perspectives* (February 2009), pp. 6-16.
- Tovar, Patricia. “The Effects of Loss Aversion on Trade Policy: Theory and Evidence,” *Journal of International Economics* 78 (2009), 154-167.
- Tversky, Amos and Daniel Kahneman. “Loss Aversion in Riskless Choice: A Reference-Dependent Model,” *Quarterly Journal of Economics* 106 (November 1991), 1039-1061.