

WORKING PAPER #473
PRINCETON UNIVERSITY
INDUSTRIAL RELATIONS SECTION
MAY 2003
Revised: May 1, 2003

Is Tomorrow Another Day?
The Labor Supply of New York City Cab Drivers *

Henry S. Farber
Princeton University

Abstract

I model the labor supply of taxi drivers as the result of optimization based on an inter-temporal utility function. Since income effects in response to temporary fluctuations in daily earnings opportunities are likely to be small, cumulative hours will be much more important than cumulative income in the decision to stop work on a given day. However, if these income effects are large due to very high discount and interest rates, then labor supply functions could be backward bending, and, in the extreme case where the wage elasticity of daily labor supply is minus one, drivers could be target earners. Indeed, Camerer, Babcock, Lowenstein, and Thaler (1997) and Chou (2000) find that the daily wage elasticity of labor supply of New York City cab drivers is substantially negative and conclude that it is likely that cab drivers are target earners. I conclude from my empirical analysis, based on new data, of the stopping behavior of New York City cab drivers that, when accounting for earnings opportunities in a reduced form with measures of clock hours, day of the week, weather, and geographic location, cumulative hours worked on the shift is a primary determinant of the likelihood of stopping work while cumulative income earned on the shift is weakly related, at best, to the likelihood of stopping work. This is consistent with there being inter-temporal substitution and inconsistent with the hypothesis that taxi drivers are target earners.

*Industrial Relations Section, Firestone Library, Princeton University, Princeton, NJ 08544-2098. email: farber@princeton.edu. The authors thanks Orley Ashenfelter, David Autor, Avinash Dixit, Danny Kahneman, Alan Krueger, Robert Solow, and participants in numerous workshops for helpful comments and discussion. Susan Merino, Gregory Evans, Julia Stahl, and Hannah Pierce provided able research assistance.

1 Introduction

There is a very large literature in economics estimating the wage elasticity of labor supply. This literature has been surveyed exhaustively (Killingsworth and Heckman, 1986; Pencavel, 1986; Blundell and MaCurdy, 1999), and a reasonable summary of the findings is that labor supply elasticities for men are very small and often not significantly different from zero while labor supply elasticities for women are somewhat larger (though considerably less than one). One criticism of this literature is that the standard neoclassical model assumes that workers are free to set their hours in response to changes in the wage or, alternatively, can select a job with the optimal wage-hours combination from a dense joint distribution of jobs. Evidence that neither of these are credible assumptions is that the distribution of hours is quite lumpy, with a substantial fraction of workers reporting usual weekly hours of precisely forty.¹

There is an emerging literature on labor supply that is not subject to this criticism because it investigates labor supply responses in settings where workers are free to set their labor supply. In this study, I contribute to this emerging literature by providing a new analysis of the labor supply of New York City taxi drivers. In contrast to earlier work on taxi driver labor supply (Camerer, Babcock, Loewenstein, and Thaler, 1997; Chou 2000), which concluded that taxi drivers were target earners, I find no evidence of target earnings behavior among the taxi drivers I studied, and I conclude that their labor supply behavior is consistent with a standard neoclassical inter-temporal labor supply model.

1.1 Hours constraints in the traditional labor supply literature

There is a substantial literature that demonstrates that workers appear not to be free to set hours and provide some estimates of the effect of this restriction on the estimation of the hours distribution. Some of this work uses information, available in some surveys, regarding workers preferred hours of work relative to actual work hours. Kahn and Lang (1991) and Dickens and Lundberg (1993) report that 35 to 40 percent of workers would prefer to work more hours at their current wage rate with a smaller fraction preferring to work

¹ Based on tabulation of the the 2002 merged outgoing rotation group files from the Current Population Survey, it is the case that over fifty percent of workers report 40 hours as their usual weekly work hours.

shorter hours. Ham (1982) estimates separate labor supply functions for constrained and unconstrained workers and finds that they differ substantially. Altonji and Paxson (1988) use data on hours of work in longitudinal data to demonstrate that the temporal variation in a workers' annual hours is larger for workers who change jobs than for workers who remain on the same job. Dickens and Lundberg (1993) estimate a model of labor supply where workers choose among a finite set of alternative jobs with fixed wage-hours combinations. They find that this model fits the observed hours distribution quite well.

1.2 The Emerging Literature on Labor Supply in Settings with Flexible Hours

In this section, I critically review four recent studies that analyze labor supply responses in settings where workers are free to set their hours of work. Two of the studies conclude that workers are target earners with negative elasticities of labor supply, and two find that workers have a substantial positive inter-temporal labor supply elasticity.

1.2.1 Taxi Drivers: Camerer, Babcock, Loewenstein, and Thaler (1997) and Chou (2000)

Camerer et. al. (1997) investigate the daily hours of work of New York City taxi drivers. These drivers lease their cabs for a pre-specified period (day, week, or month) for a fixed fee, are responsible for fuel and some maintenance, and keep 100 percent of their fare income after paying fixed costs. They are free to drive as much or as little as they want during the lease period. This leasing arrangement is close to the incentive-theorist's first-best solution to the firm-worker principal-agent problem of selling the firm to the worker.

The core of their analysis consists of computing a daily wage rate as the ratio of daily income to daily hours. They then regress logarithm of daily hours on the logarithm of this wage rate and find a significant and substantial *negative* elasticity of labor supply. They conclude that this is consistent with a target earnings model, in which drivers stop working after reaching their target daily income. They argue further that this is inconsistent with a standard neoclassical model of labor supply.

Chou (2000) carries out an analysis of the labor supply of taxi drivers in Singapore that closely follows that of Camerer, et. al. Like the earlier paper, he finds a significant negative relationship between log hours worked and the log wage rate calculated as the ratio of daily income to daily hours. Chou concludes that drivers appear to set targets over a short horizon.

I am puzzled by these findings for both economic and econometric reasons. Economically, target earning implies that, on days when it is easy to make money (pick low-hanging fruit, so to speak), the drivers quit early, while on days when fares are scarce, drivers work longer hours. If workers can substitute labor for leisure inter-temporally across days, then they should work more on days with low-hanging fruit relative to other days. Effectively, the target earnings model is a very strong version of a model with a backward-bending labor supply curve. This implies very strong effects of daily income on daily labor supply; so strong as to overwhelm any substitution effect. A finding that daily income, which is a small fraction of income over reasonable longer periods (monthly, annual) has such a strong income effect demands careful scrutiny.²

A second source of concern is the assumption that there is a wage rate characterizing a day that a driver uses parametrically to determine his hours of work. Camerer, et. al. state that the wages of taxi drivers are “. . . relatively constant within a day” (Camerer, et. al., 1997, p. 408), and they present evidence showing substantial positive autocorrelations in the hourly wage available within a given day. In contrast, I do not find significant autocorrelations of the hourly wage within a shift. I find that fare opportunities vary dramatically and unpredictably over the course of a day. In my analysis, I present evidence that within-day variation is much more important than between-day variation in accounting for hourly wage variation.

An important econometric concern, one that is recognized by the authors, is that they are regressing hours on a wage measure that is computed using the reciprocal of hours. This leads to a “division bias” where, if there is any misspecification or measurement error, there will be a negative bias on the coefficient of the wage. Both Camerer et. al. (1997) and Chou

² Indeed, the literature on inter-temporal substitution in labor supply as it relates to macro-economics typically assumes no income effects of shocks to annual income on labor supply (the so-called λ -constant assumption). See, for example, MaCurdy (1977) and Pencavel (1986).

(2000) address this concern through the use of an instrumental variable estimation where the instrument is the average wage of other workers on the same calendar date. However, if there are calendar-date effects on labor supply independent of the wage rate, this instrument will be ineffective in purging the estimated labor supply elasticity of bias.

I propose an alternative approach to estimating taxi drivers' work hours that is not subject to the same criticisms. Estimates based on this approach, using new data on New York City Taxi drivers, do not show evidence of the income effects, large or small, that are necessary for a target earnings interpretation. Since using these new data I am able to reproduce the negative labor supply elasticity estimated by Camerer, et al. and by Chou, the difference in our results are likely due to the different econometric frameworks rather than to differences in data.

1.2.2 Stadium Vendors: Oettinger (1999)

Oettinger (1999) investigates the days of work of stadium vendors at baseball games. The stadium vendors that Oettinger studies are hired in the sense that they are approved to sell at games. The vendors are free to work or not work at any particular game without notice to their employer, and they receive a fixed commission rate on sales. Their hours are fixed for any game for which they show up to work to the extent that they are not supposed to leave early. The interesting labor supply margin is the number of days particular vendors choose to work. Oettinger carefully models the factors that make certain games more lucrative for vendors (e.g., larger crowds due to factors such as quality of opponent and day of the week). His analysis accounts for the multi-vendor participation problem which, in equilibrium, has more vendors show up to work games with larger expected attendance. Oettinger concludes that there is a substantial positive inter-temporal labor supply elasticity implicit in the daily participation decisions of the vendors he studies.

1.2.3 Bicycle Messengers: Fehr and Goette (2002)

Fehr and Goette (2002) investigate the hours per day and days per month of bicycle messengers in Zurich through implementation and analysis of a very interesting experiment. Bicycle messengers employed by the company studied by Fehr and Goette received a fixed compen-

sation per message delivered and are free to set their own hours. The experiment consisted of dividing the companies messengers into two group (A and B). Group A received a significantly enhanced per-message fee for one month after which the fee declined to normal. Group B received the enhanced fee in the second month. Their results are clear. Monthly labor supply of the group with the enhanced fee increased significantly relative to that of the group with the normal fee during the same month. The conclusion of Fehr and Goette is that there is a large positive inter-temporal elasticity of labor supply.

The increase in labor supply took the form of an increase in the number of days worked during the month that was partially offset by a decrease in labor supply on any particular day. Fehr and Goette argue that the decline in daily labor supply is inconsistent with the standard neo-classical model, and they argue that the decline is due a combination of 1) messengers being “loss averse” relative to a fixed daily benchmark or target and 2) lower likelihood of failing to reach the daily benchmark in the high fee regime.

It may be that workers have benchmark earnings levels that they would like to meet, but I disagree with the assertion that the reduction in daily hours is inconsistent with the standard neoclassical model. It may be that, in the high fee month where drivers want to supply substantially more labor, it is efficient for the messengers to work more days but work less each day.

2 Conversations with Cab Drivers

While not conducted in a systematic fashion, I have been having informal conversations with cab drivers in New York City and elsewhere when traveling for the past couple of years. The information I have gained is not meant as evidence to test competing models. However, it does provide some information on what cab drivers are thinking about in deciding on their labor supply. I worked hard to avoid asking leading questions regarding their decision making.

I began by asking drivers about their contracting arrangements, and in all cases they leased their cabs, sometimes on a daily basis but usually on a weekly basis. I then probed how much the drivers worked, with most responding that they worked eight to eleven hours

in a shift for six days per week. When I asked how drivers decided when it was time to stop for the shift, most said that they got tired after some period of time and stopped. Several elaborated by saying that they would stop if fares seemed scarce or if they got a fare that took them near the garage, often in Queens. Some said they were constrained by the need to get the taxi back to the garage at shift end or, in the case of longer term leases, to a set meeting place to turn the cab over to another driver with whom they were sharing the car.

I then asked if they had an income target that they needed to meet before they quit. With one exception, the drivers denied having a target and many reiterated that they quit when tired. I would then ask what would happen if it were a particularly good day or a particularly bad day. The answer was generally that you never know what will happen tomorrow, so why worry much about a single day. The one driver who said he had a target, told me what it was and on what it was based. I then probed by asking 1) how many hours it generally took to reach the target, 2) what happened if he got to that point and was short of the target, and 3) what happened if he reached the target substantially earlier than the usual hours. The answers were 1) it usually took ten to eleven hours, 2) he would stop if he was short at that point because he was tired, and 3) he would continue to drive after reaching the target because he “might as well.”

Taken together, my impression from these interviews is that drivers do not consciously behave as if they are target earners. The reasoning they articulate is consistent with a standard neoclassical model with small daily income effects. Effectively, saying that you stop when you are tired is equivalent to saying that you quit because the marginal utility of leisure increased to the point where it was optimal to stop. Of course, there may be a difference between how drivers say they are behaving and how they actually behave. For that reason, I turn to the systematic theoretical and empirical analysis.

3 A Model of Taxi Driver Daily Labor Supply

The standard employment arrangement of New York City cab drivers is that the driver leases the cab for a fixed period, usually a 12-hour shift, a week, or a month. The driver pays a fixed fee for the cab plus fuel and certain maintenance costs, 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 economists first-best solution to the agency problem. In a manner of speaking, the employer has “sold the firm to the worker.”

A fully optimizing model of taxi driver daily labor supply is based on the solution of a dynamic programming problem where a driver at a given point in his shift (economically, geographically, and temporally) compares his utility if he stops working with his expected utility from continuing to work. While I do not formulate and solve this model, I do sketch its main components.

Consider a simple inter-temporal utility function for a cab driver with utility derived each day from consumption of goods and leisure. Let this utility function be additively separable in utility between periods and in goods and leisure within a day. On this basis, the utility of a driver on day t is

$$U_t = a(x_t) + b(\ell_t), \quad (3.1)$$

where x_t is daily goods consumption and ℓ_t is daily leisure consumption. The inter-temporal utility function, defined over some undefined set of T periods is

$$U = \sum_{t=1}^T (1 + \rho)^{-t} [a(x_t) + b(\ell_t)], \quad (3.2)$$

where ρ is the rate of time preference and $a(\cdot)$ and $b(\cdot)$ have positive first derivatives and negative second derivatives. The lifetime budget constraint is

$$Y_0 + \sum_{t=1}^T (1 + r)^{-t} y_t (1 - \ell_t) = \sum_{t=1}^T (1 + r)^{-t} x_t, \quad (3.3)$$

where the price of consumption goods is normalized to one, r is the discount rate, Y_0 represents initial wealth, $1 - \ell_t$ represents work hours (the complement of leisure time), and $y_t(1 - \ell_t)$ represents daily earnings generated as a function of work time. The first derivative of $y_t(1 - \ell_t)$ is assumed to be positive.

The Lagrangian expression for constrained maximization of this utility function is

$$V = \sum_{t=1}^T (1 + \rho)^{-t} [a(x_t) + b(\ell_t)] + \lambda \left[Y_0 + \sum_{t=1}^T (1 + r)^{-t} [y_t(1 - \ell_t) - x_t] \right], \quad (3.4)$$

where λ is interpreted as the marginal utility of initial wealth. This expression is maximized with respect to x_t and ℓ_t , and the first order conditions are

$$\frac{\partial V}{\partial x_t} = a'(x_t) - \lambda\theta^t = 0, \quad (3.5)$$

$$\frac{\partial V}{\partial \ell_t} = b'(\ell_t) - \lambda\theta^t y'_t(1 - \ell_t) = 0, \quad (3.6)$$

and

$$\frac{\partial V}{\partial \lambda} = \left[Y_0 + \sum_{t=1}^T (1+r)^{-t} [y_t(1 - \ell_t) - x_t] \right] = 0, \quad (3.7)$$

where $\theta = (1+r)/(1+\rho)$. Solving equations 3.5 and 3.6 for $\lambda\theta^t$ yields the result that

$$y'_t(1 - \ell_t) = \frac{b'(\ell_t)}{a'(x_t)}. \quad (3.8)$$

3.1 Income Effects and the Shape of the Labor Supply Function

Equation 3.8 implies that hours are selected so that the marginal earnings from working an additional increment time is equal to the marginal rate of substitution of leisure for goods within a single period. If taxi drivers were hourly employees earning a fixed wage rate, then $y'_t(1 - \ell_t)$, which I call the marginal wage, would equal the fixed wage rate, and equation 3.8 would imply the standard labor supply result that hours are selected to equate the fixed wage rate and the marginal rate of substitution of leisure for goods.

The labor supply function implicit in the solution of this problem is depends centrally on the marginal utility of wealth (λ) and the relative discount factor (θ). The marginal utility of wealth is a function of initial wealth, presumably minimal for taxi drivers, and the general level of earnings opportunities (the scale of y_t) over the relevant time horizon. If the relevant time horizon is short, then short-run fluctuations in earnings opportunities will have strong effects on λ , and income effects on labor supply could be important. If the relevant time horizon is longer, then short-run fluctuations in earnings opportunities will not have strong effects on λ , and income effects on labor supply are not likely to be important. The time horizon is crucially determined by the relative discount factor, θ . If θ is large, implying that the interest rate in the budget constraint is larger than the rate of time preference, then short-run fluctuations in income will have larger income effects on labor supply. In the

extreme case, if the interest rate goes to infinity, so that there is no savings or borrowing and income and consumption are equal on a daily basis, then income effects are likely to be very important. At the other extreme, if the interest rate is small relative to the rate of time preference, then individuals will smoother their consumption of goods and leisure over time, and there will not be large income effects from daily fluctuations in income.

Given that taxi drivers, like virtually all workers, make consumption commitments that span days (e.g., apartment rental), it seems clear that they are able to smooth consumption across days. Thus, transitory variation in the marginal wage will be directly reflected in labor supply in that period. As long as there are many periods, as is surely the case when considering a period to be a single day, there are not likely to be significant income effects. A transitory increase in the marginal wage (an increase in $y'_t(1 - \ell_t)$) implies a decrease in leisure and an increase in labor supply due to the negative second derivative of $b(\ell_t)$. Leisure becomes more valuable as more hours are worked.

In the context of daily taxi driver labor supply, the clear prediction is that hours worked are positively related to transitory variation in the marginal wage. Daily income effects are inconsequential. This stands in stark contrast to the predictions of daily target earnings behavior. Essentially, daily target earnings behavior implies that income effects dominate substitution effects so that the elasticity of hours with respect to changes in the marginal wage rate is minus one. In the context of the inter-temporal labor supply model, this is an extreme case of a very high interest rate coupled with 1) a marginal utility of goods consumption that is very large until some target level of goods consumption and low thereafter and 2) very low marginal disutility of leisure until the target is reached.

3.2 Modeling Daily Hours of Work

Modeling the number of hours worked on a particular day is made difficult by the fact that the marginal wage function is likely not monotonic in hours worked. In other words, the second derivative of $y(\cdot)$ can change sign. In fact, this is quite likely as the demand for taxi cabs varies during the day. Thus, it would not necessarily be optimal for drivers to quit based on time-specific lulls in traffic during the day. This implies that there can be multiple local maxima that satisfy the second order conditions, and the driver is assumed to be aware

of this and select the global maximum from among these.

One approach to modeling hours worked is to consider the problem to be consistent with a survival time (hazard) model. 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. The theory outlined here has several sharp predictions for this modeling approach.

- The likelihood of quitting for the day is positively related to the number of hours already worked. This is due to the monotonically increasing marginal utility of leisure with hours worked.
- The likelihood of quitting for the day conditional on the number of hours worked so far should not be substantially related to income already earned during the day. This is due to the inter-temporal nature of daily labor supply and the resulting small daily income effect.
- The likelihood of quitting for the day conditional on the number of hours worked so far should be negatively related to further earnings opportunities on that day. This includes within-day variation in the marginal wage as well as day-specific transitory earnings effects.

All three of these predictions are inconsistent with the predictions of a target earnings model, where a worker is expected to quit when income on that day reaches the target level. The target model predicts that: 1) the likelihood of quitting for the day is not substantially related to the number of hours already worked, 2) the likelihood of quitting for the day is centrally determined by income already earned during the day, and 3) the likelihood of quitting for the day conditional on the number of hours worked so far on that day should be positively related to day-specific earnings opportunities as the daily income target is likely to be reached after fewer hours.

4 Empirical Models of Taxi Driver Labor Supply

As I suggest in the previous section, I estimate a model of taxi driver daily labor supply as a survival time model where quitting can occur at discrete points in time corresponding to the

ends of fares. Without deriving the full dynamic solution to the optimal stopping problem, I can derive a reasonable approximate solution that I can implement 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 $R(\tau) = \tau - \tau^*$ is

$$R_{idq}(\tau) = \gamma_1 h_\tau + \gamma_2 y_\tau + X_{idq}\beta + \mu_i + \delta_d + \alpha_q + \epsilon_{idh\tau}, \quad (4.1)$$

where i indexes the particular driver, d indexes the date (including measures of the day of the week and the month), q indexes the hour of the day interacted with the day of the week, h_τ measures hours worked on the shift at τ , y_τ measures income earned on the shift at τ , and X measures other factors affecting the determination of the optimal stopping time and the comparison with τ . The quantity ϵ is a random component with a standard normal distribution. The sets of fixed effects for date, hour/day-of-week, and location are included to capture variation in earnings opportunities from continuing to drive. The individual stops driving at τ if $R_{idq}(\tau) \geq 0$, and this implies a standard probit specification based on the latent variable defined in equation 4.1.

The three clear predictions of the theory outlined above hold for this probit model: 1) The probability of quitting will be positively related to hours worked ($\gamma_1 > 0$), 2) the probability of quitting will be unrelated to income earned ($\gamma_2 = 0$) unless daily income effects are important, and 3) the probability of quitting will be negatively related to further earnings opportunities as captured here by the day-of-week effects, hour-of-day effects, and other factors.

4.1 A Target Earnings Model of Taxi Driver Labor Supply

A direct empirical implementation of a target earnings model is straightforward. Define the target as

$$T_{it} = X_{it}\beta + \mu_i + \epsilon_{it}, \quad (4.2)$$

where i indexes the driver and X_{it} are driver characteristics, X_{it} is a set of characteristics describing shift t (e.g., day shift vs. night shift, weather, day of week), μ_i is a driver-specific fixed factor, and ϵ_{it} is a normally distributed random component. It is somewhat problematic to include measures that allow substantial inter-day variation in the target for a particular driver. To the extent that the target can vary from day to day, the usefulness of the target as a construct for predicting hours of work is severely limited.

Because fares come in discrete packages, the target itself is not observed directly, but it is known that the target lies between the value of total fares received before the final fare (Y_{K-1}) and the value of total fares received including the final fare (Y_K). This is

$$Y_{K-1} < T_{it} < Y_K. \quad (4.3)$$

Using the definition of the target in equation 4.2 and the assumption of normality of ϵ_{it} , the probability that the target is bounded as defined in equation 4.3 is

$$Pr(Y_{K-1} < T_{it} < Y_K) = \Phi\left(\frac{Y_K - X_{it}\beta - \mu_i}{\sigma}\right) - \Phi\left(\frac{Y_{K-1} - X_{it}\beta - \mu_i}{\sigma}\right), \quad (4.4)$$

where σ^2 is the variance of ϵ . The probability in equation 4.4 forms the basis of a double truncation likelihood function that can be maximized in a straightforward fashion.

A very close approximation to the maximum likelihood estimates resulting from estimation of equation 4.4 can be obtained by estimating equation 4.2 by OLS where T_{it} is measured as total shift income minus one-half of the final fare. Comparison of estimates obtained by maximum likelihood and by OLS shows for all specifications I use that the estimates and standard errors are virtually identical. This is not surprising given that the average fare is small relative to total daily income so that the target is very close to total daily income. As a result, I present the OLS estimates.

There is no explicit test of the target earnings hypothesis that can be based on the specific parameter estimates in this model. An indirect test is based on variation the difference between total daily fare income and the predicted target by date for a particular driver. Some variation will naturally occur because the total daily income will exceed the daily target by a value that bounded from above by the size of the last fare. We know that the target lies between total daily income less the last fare and total daily income. If the target

is uniformly distributed in this interval, as it is reasonable to assume, then our best direct measure of the target is total daily income minus one-half of the last fare.³ To the extent that the target construct is useful, there should not be substantial within-driver inter-day variation of the difference between this target measure and the predicted target. In the extreme case, if drivers set their target independently each day then the target has no predictive value for labor supply.

4.1.1 Camerer et. al.'s Empirical Implementation of the Target Earnings Model

Camerer, et. al. (1997) and Chou (2000) use the prediction of the target earnings model, that daily hours worked will be negatively related to hourly earnings opportunities for that day, as a test of the model. They measure hourly earnings opportunities as a fixed daily wage rate computed as total fare income divided by hours worked. They then estimate a regression, with one observation for each shift, of the form

$$\ln H_{it} = \eta \cdot \ln W_{it} + X_{it}\beta + \epsilon_{it}, \quad (4.5)$$

where H_{it} represents the hours worked by driver i on day t , $W_{it} = Y_{it}/H_{it}$, Y_{it} is the total fare income of driver i on day t , and X_{it} are other factors affecting labor supply. The parameter η is meant to represent the elasticity of labor demand, and their estimates of η are strongly negative.

The authors recognize the inherent division bias that can lead to a negative bias in the estimate of η . This bias is due to the fact that the wage rate is computed using the dependent variable in the denominator. If the model is not perfectly specified or if there is any measurement error, the estimate of η will be biased downward. They address this problem directly through the use of an instrumental variables estimator. The instrument they use is the wage computed for other drivers on the same calendar date, and they find similar, though somewhat weaker, results with their IV approach. One potential problem with this approach is that there might be day-specific factors that affect the hours of all

³ I have also estimated the parameters of equation 4.2 using OLS by assuming that the observed target is total daily income minus the last fare. The estimates are virtually identical to the maximum likelihood estimates of the double truncation model.

drivers to some degree. To the extent that this is case, their instrument will not purge their estimates of η of their inherent negative bias.

I will present OLS estimates of models like equation 4.5. However, my data do not have sufficient numbers of drivers on any particular date to replicate their IV analysis.

A mathematically equivalent representation of the hours function in equation 4.5 can be derived due to the fact that both hours and the computed wage enter the model in logarithmic form. Equation 4.5 can be written as

$$\ln H_{it} = \eta \cdot \ln Y_{it} - \eta \cdot \ln H_{it} + X_{it}\beta + \epsilon_{it}. \quad (4.6)$$

Solving equation 4.6 for $\ln H_{it}$ yields

$$\ln H_{it} = \eta^* \cdot \ln Y_{it} + X_{it}\beta^* + \epsilon_{it}^*, \quad (4.7)$$

where log hours are a function of log income rather than the log wage.⁴ The coefficient on log income in equation 4.7 is related mathematically to the coefficient on the log wage in equation 4.5 such that

$$\eta^* = \left(\frac{\eta}{1 + \eta} \right). \quad (4.8)$$

Of course, the OLS estimates of η and η^* will not be related as in equation 4.8 because of the different orthogonality conditions enforced by OLS.

The underlying variation used to estimate η^* in equation 4.7 comes only from variation in $\ln Y_{it}$, and it is difficult to know why there is variation in income if it is meant to reflect the target income level. Clearly, inter-day transitory variation in earnings opportunities should not affect income. It should only affect hours. Variation in the target could come from systematic differences across drivers in their target levels, but, as I noted above, the predictive power of the target for hours is limited if one allows the target to vary substantially over time for a particular driver. Since the two formulations of the hours functions in equations 4.5 and 4.7 are equivalent, this represents an anomaly in the context of the target earnings model.

⁴ Since income and hours are jointly determined, OLS estimates of equation 4.6 are biased.

5 Data and Preliminary Statistics

The data necessary to carry out this study are available on “trip sheets” that drivers fill out during each shift. Each trip sheet lists the driver’s name, hack number, and date, along with detail on each trip. The information for each trip includes the start time, start location, end time, end location, and fare. In order to obtain a sample of trip sheets, in the Summer of 2000 my research assistants created a list of taxi leasing companies from the current edition of the New York City Yellow Pages. They then contacted these companies by phone or in person. Of the seventy leasing companies contacted, twenty-one were out of business, twelve refused to talk, and three were no longer leasing companies. Six were responsive to our inquiries, and they received personal visits. However, only one provided trip sheets, and these proved unusable because the company redacted information on the drivers’ names and the dates.

Another company called was responsive to our inquiries and our contact mailed us a set of trip sheets.⁵ We were sent 244 trip sheets for 13 drivers covering various dates over the period from June 1999 through May 2000. We contacted the leasing company again in the Summer of 2001, and, while our original contact was no longer with the firm, we were sent an additional 349 trip sheets for 10 drivers covering various dates over the period from June 2000 through May 2001. Two of the drivers appear in both groups so that I have a total of 593 trip sheets for 21 drivers over the period from June 1999 through May 2001. A few of these trip sheets refer to common dates so that I have data on 584 shifts. The drivers in my sample lease their cabs weekly for a fee of \$575. Each driver pays for his own fuel and keeps all of his fare income and tips.

An unfortunate consequence of receiving the trip sheets in an unsystematic fashion is that I have no information on the number of shifts worked. If a trip sheet is not available for a specific driver on a given day, I cannot determine if that driver did not work on that day or if the trip sheet was simply not provided. This prevents me from examining in any conclusive way inter-day relationships in labor supply.

⁵ The person contacted was a basketball fan who particularly liked Bill Bradley so that he was willing to cooperate with work done at Princeton.

Completeness of the trip sheets is a concern. Unfortunately, I do not have the shift summary printed by the meter after each shift, which lists the total number of trips, in order to verify the completeness of the trip sheets. As a result I cannot do the kind of careful *ex post* checking that Camerer, et. al. were able to perform by comparing the trip sheets to the daily summary printed by the meter. However, for several reasons it is likely that the trip sheets are relatively complete. First, there is no particular disincentive for drivers to avoid listing trips on their trip sheets. The trip sheets are not used for tax or other financial purposes. More importantly, there are financial incentives working in favor of complete listing of trips. In my informal interviews, I have asked drivers in New York about their trip-sheet practices, and they have told me that they are careful about filling out the sheets because of fines levied due to incomplete trip sheets. Apparently, taxi cabs are stopped by New York City police officers or by Taxi and Limosine Commission inspectors, either randomly or for cause.⁶ When stopped, drivers are asked for their trip sheet and a printout of the meter summary to that point. The driver can be fined a substantial amount for each fare that is a shortfall between the number of fares listed on the meter summary and the number of fares listed on the trip sheet. Additionally, from time to time, police request trip sheets as part of the investigation of a crime. In the end, there is no way to ensure that the trip sheets are complete, and I proceed under the assumption that they are.

I coded the starting and ending locations on the trip sheets into eleven categories. These are Downtown Manhattan (below 14th street), Midtown Manhattan (14th Street to 59th Street), Uptown Manhattan (above 59th Street), The Bronx, Queens, Brooklyn, Staten Island, Kennedy Airport, LaGuardia Airport, Newark Airport, and other. Almost all trips (92 percent) started and ended in Manhattan.

I additionally collected data from the National Atmospheric and Oceanic Administration (NOAA) on temperature and rainfall in New York City. I collected daily average, minimum, and maximum temperatures, and total daily rainfall and snowfall in Central Park. I also collected hourly rainfall data at LaGuardia Airport.

⁶ Das Gupta (2002) notes that “. . . the rules governing drivers have become more elaborate and punitive.” and that tickets are “. . . zealously issued.”

5.1 Procedures for Cleaning Data

I perform several regularity checks to insure that the sheets are internally consistent, and where they are not, I clean the data using a set of reasonable rules. I outline that process here. The main consistency checks are that 1) trips must start no earlier than the end of the previous trip and 2) trips must start before they end (be of positive duration). I also develop a set of rules to impute starting or ending times for trips where such times are missing and fares where fares are missing. For the most part, the trips on my trip sheets satisfy these rules and have non-missing time and fare data. In the small number of cases where these rules are not satisfied or time or fare data are missing (about one percent of trips), I use the following procedures.

- Impute new trip times for 38 trips of zero duration. I start by computing the length of time between the end of the previous trip and the start of the next trip. Clearly, the time of the “current” trip is no greater than this available gap. I compute a predicted trip length based on regression of trip length on the fare and its square (R-squared = 0.68). In almost all cases, this predicted trip length is less than the available gap. I then center the predicted trip length in the available gap and compute the implied adjusted start and end times for the trip. In the handful of cases where the predicted trip length is longer than the gap (never by more than a few minutes), I set the trip length equal to the available gap and compute the adjusted start and end times for the trip accordingly.
- Impute new trip times for 52 cases where one or both of the starting and ending times are missing. Most of these are imputed using the same predicted trip length defined above as a function of the fare. If the start time is missing, it is computed as the end time minus the predicted trip length. Analogously, if the end time is missing, it is computed as the start time plus the predicted trip length. Where both times are missing, they are computed as defined above for trips of zero duration. There are a handful of cases where one or the other of the times are missing and the fare is missing. In these cases, I predict trip length using a complete interaction of starting and ending locations of the trip (R-squared = 0.44).

- Adjust times where they are out of sequence. There are 83 cases where a trip apparently started before the previous trip ended. The mean overlap was 2.7 minutes, and the maximum overlap was 9 minutes. I adjust these by reducing the end time of the previous trip by one-half of the overlap and increasing the start time of the current trip by one-half of the overlap.
- Impute fares for 45 cases where the fare is missing. I impute fares in these cases using predicted values from a regression of fare on trip length and a complete set of interactions of starting and ending locations (R-squared = 0.83). In the 7 cases where location is missing I predict based on a model using only trip length (R-squared = 0.64).
- Classify as breaks (non-work time during the day) certain long wait times between fares. I classify as breaks 1) 195 periods of at least 30 minutes between a fare that ended in Manhattan and a fare that started in Manhattan, 2) 112 periods of at least 60 minutes between fares that either started or ended outside Manhattan but did not end at either LaGuardia or Kennedy Airport, and 3) 9 periods of at least 120 minutes that ended at LaGuardia or Kennedy Airport.
- Impute starting locations to 144 trips and ending locations to 223 trips with missing or illegible locations. Where the starting location is missing, I assume that the starting location is the same as the ending location of the previous trip. Where the ending location is missing, I assume that the ending location is the same as the starting location of the next trip. In cases where the necessary information for imputation by these rules is not available, I examined the individual strip sheets and made an imputation based on the “nearest” available locations, the size of the fares, and elapsed time of the trip. Since over 96 percent of trips start in Manhattan and over 95 percent of trips end in Manhattan, almost all of the allocated locations are in Manhattan.

Table 1: Summary Statistics, by Driver

Driver	# of Shifts	Average Trips	Working Hours	Driving Hours	Waiting Hours	Break Hours	Total Income	Average Wage
Driver 1	39	23.56	6.85	4.32	2.53	0.90	157.58	23.16
Driver 2	14	12.29	3.89	2.78	1.11	2.41	97.10	25.11
Driver 3	6	23.67	6.66	4.61	2.05	0.74	163.42	25.56
Driver 4	40	22.10	6.28	4.52	1.76	0.39	147.51	23.89
Driver 5	23	16.52	6.46	3.98	2.48	2.11	144.96	23.65
Driver 6	6	29.33	8.62	6.48	2.14	2.42	205.00	24.02
Driver 7	24	22.29	6.47	4.42	2.05	0.74	160.71	25.59
Driver 8	37	25.32	7.78	5.13	2.64	0.86	172.44	22.54
Driver 9	19	25.58	7.17	5.47	1.70	0.54	162.02	23.23
Driver 10	45	20.27	6.35	3.90	2.45	1.65	133.19	21.46
Driver 11	6	24.00	7.15	5.22	1.93	0.71	182.81	26.57
Driver 12	13	19.46	6.15	4.03	2.13	0.55	157.95	25.78
Driver 13	10	26.50	7.03	4.71	2.31	0.53	154.19	22.55
Driver 14	17	21.29	7.06	4.49	2.57	0.64	165.84	23.59
Driver 15	8	39.62	10.82	7.64	3.17	0.19	228.26	21.12
Driver 16	70	25.06	6.84	4.56	2.28	0.93	172.01	25.62
Driver 17	10	23.00	5.88	3.71	2.17	0.54	144.57	23.35
Driver 18	72	28.10	8.53	5.84	2.69	0.60	203.05	24.01
Driver 19	33	17.06	6.91	4.63	2.29	0.97	163.51	23.91
Driver 20	46	24.46	7.10	4.80	2.30	0.67	156.23	22.00
Driver 21	46	19.17	5.32	3.66	1.66	0.24	128.97	24.72
All	584	23.05	6.90	4.64	2.26	0.86	161.33	23.79

5.2 Shift-Level Summary Statistics

There are a total of 13,464 trips listed for the 584 shifts on the 593 trip sheets for the 21 drivers in the cleaned sample. Table 1 contains total statistics by shift for each driver. I have data on an average of 27.8 shifts per driver. I have more than 30 shifts for 9 drivers and more than 20 sheets for 11 drivers. Hours worked per day is defined as the sum of driving time (the sum over trips of the time between the trip start time and the trip end time) and waiting time (the sum over trips of the time between the end of the last trip and the start of the current trip). Waiting time is substantial, accounting for 33 percent of working time on average. Break time averages about 52 minutes per shift.

There is substantial variation across drivers in average hours worked per day, with means

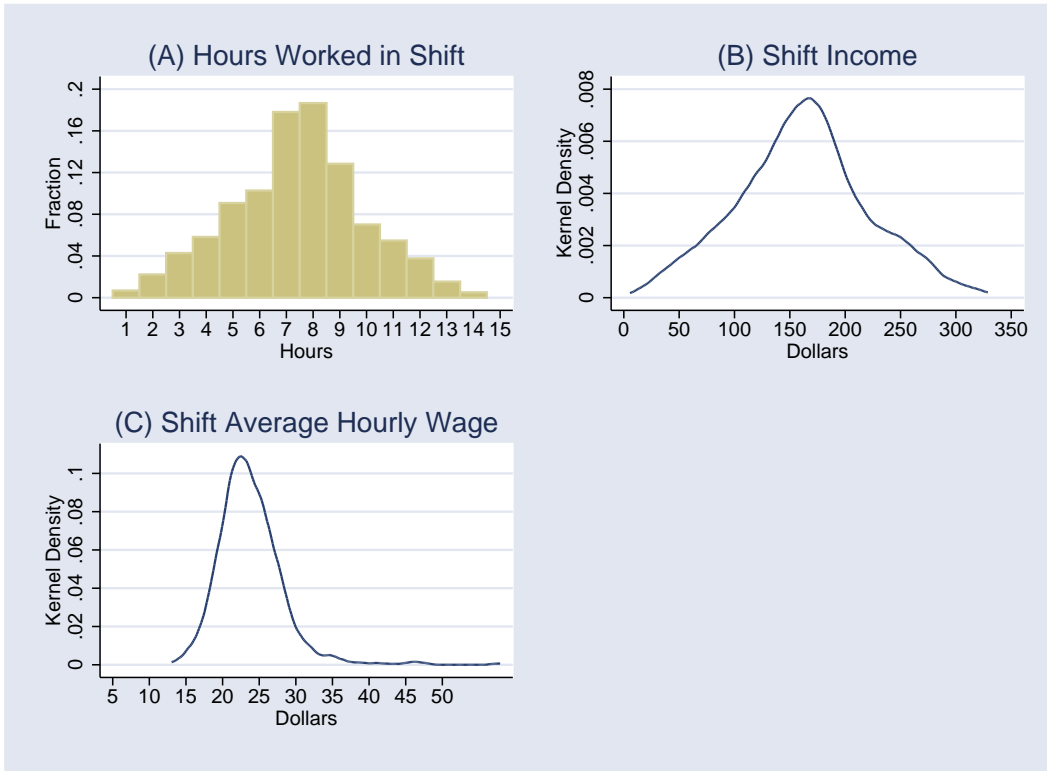


Figure 1: Distributions of Hours, Income, and Average Wage by Shift

ranging from 3.89 to 10.82. Still, the majority of the variation in daily work hours is within-driver variation across days. The standard deviation of daily work hours is 2.50. The R-squared from a regression of daily hours on a set of driver fixed effects is 0.181 with a residual root mean squared error of 2.30. Panel A of figure 1 contains a histogram of hours worked for the 584 shifts. The distribution is single-peaked, with the mode at eight hours.

There is also substantial variation across drivers in total fare income per day, with means ranging from \$97.10 to 228.26.⁷ Not surprisingly, daily income covaries strongly with daily hours with a simple correlation of 0.91. As with hours, the majority of the variation in daily work hours is within-driver variation across days. The standard deviation of of daily income is 59.57. The R-squared from a regression of daily income on a set of driver fixed effects is 0.169 with a residual root mean squared error of 55.27. A labor supply model where drivers had fixed but potentially different targets would imply that more of the variation in income

⁷ Income per day is the sum of fares. Tip income is not measured or accounted for.

would be accounted for by driver fixed effects. Panel B of figure 1 contains a kernel density estimate of daily income.

The last column of table 1 contains the daily average for each driver of their hourly wage rate (total income divided by working hours). These show less inter-driver variation, ranging from a low of \$21.12 to a high of \$26.57. The standard deviation of the daily wage rate is \$4.48. Most of this is within-driver variation as the R-squared from a regression of the daily wage on a set of driver fixed effects is 0.089 with a residual root mean squared error of \$4.35. Panel C of figure 1 contains a kernel density estimate of the shift average hourly wage.




Figure 2: Distributions Starting and Ending Hours of Shifts

Figure 2 contains graphs of the distributions of starting and ending times of shifts. There is substantial concentration in starting times, with 33.1 percent of shifts starting in the three-hour period from 6AM to 8:59AM and 34.9 percent of shifts starting in the two-hour period from 4PM to 5:59PM. No shifts start between 11PM and 3AM. There is slightly lower concentration in shift ending times, with 29.5 percent of shifts ending in the two-hour period from 3PM to 4:59PM and 31.4 percent of shifts ending between 11PM and 12:59AM. After a lull following the afternoon spike in the shift-ending rate, the ending rate increases throughout the evening. Almost no shifts end between 5AM and 10AM.

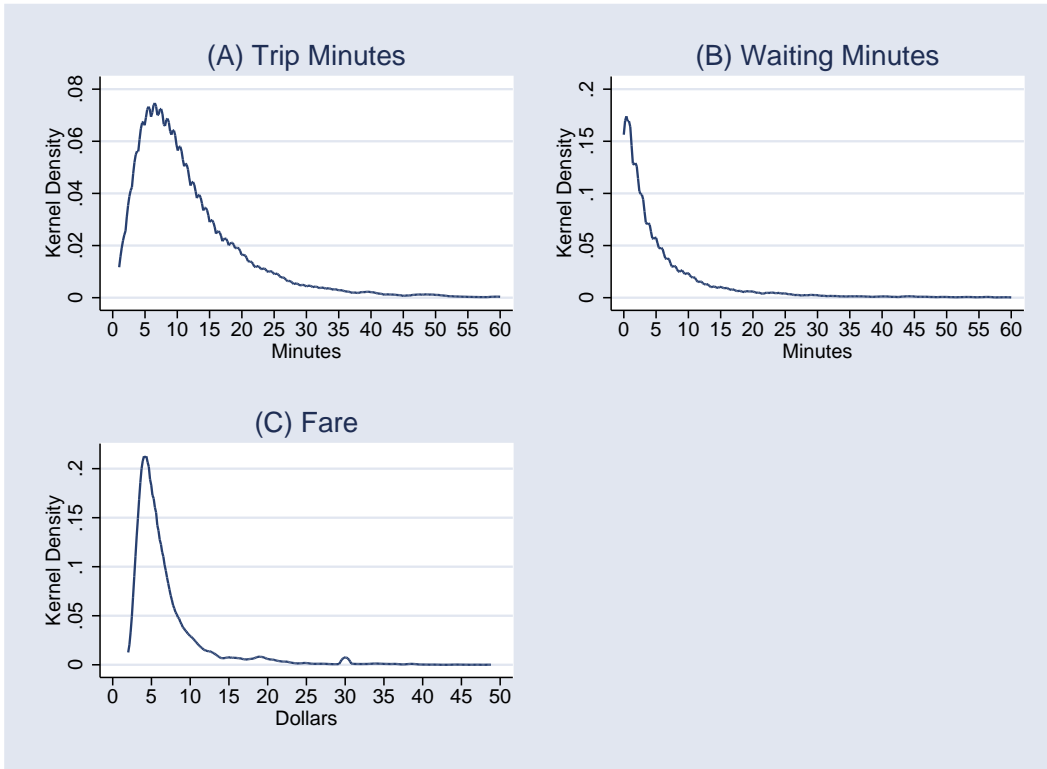


Figure 3: Kernel Density Estimates of Trip Times, Wait Times, and Fares

5.3 Trip-Level Summary Statistics

Figure 3 contains kernel density estimates of the distributions of trip minutes, waiting minutes, and fares for the 13,464 trips in my sample. I have truncated these distributions at 60 minutes for trip and waiting times and at \$50 for fares. These truncations do not change the shape of the distribution since the kernel density estimates in these upper tails is essentially zero.⁸

Panel A of figure 3 contains a kernel density estimate of the distribution of trip minutes. Median minutes per trip is 10 minutes and the mean is 12.1 minutes. That trips are this short reflects the fact that 92.6 percent of the trips in my sample begin and end in Manhattan. Only 3.8 percent of trips originate outside Manhattan (2.6 percent at LaGuardia Airport and 1.0 percent at Kennedy Airport). Only 4.3 percent of fares end outside Manhattan (1.6

⁸ There are 41 trips with times greater than 60 minutes and 26 waiting times greater than 60 minutes. There are 8 fares greater than \$50.

percent at LaGuardia Airport, 0.8 percent at Kennedy Airport, 0.9 percent in Brooklyn, 0.5 percent in Queens).

Panel B of figure 3 contains a kernel density estimate of the distribution of waiting minutes before each trip. Median waiting minutes is 3 minutes and the mean waiting time is 5.9 minutes. There are 2892 trips with zero wait time. Recall that I reclassified 316 long wait times between fares as break times, and fares after these breaks are classified as having preceding wait times of zero.

Panel C of figure 3 contains a kernel density estimate of the distribution of fares. The median fare is \$5.30, and the mean fare is \$7.00. Once again, the small size of the fares (which exclude tips) is due to the fact that most trips are intra-Manhattan. The average intra-Manhattan fare is \$5.90, while the average fare that starts or ends outside Manhattan is \$20.73. The small blip at \$30 represents the flat rate from Kennedy Airport to Manhattan in force during my sample period.

5.4 Variation in the Hourly Wage

Variation in the hourly wage rate of taxi drivers both within particular shifts and between driving days is central to understanding taxi driver labor supply. In order to examine this variation, I computed an hourly wage measure 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, 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.

Panel A of figure 4 contains the average hourly wage by clock hour for the driver-hours in my sample along with the average ± 1.96 standard deviations. There is clearly variation over the day in the hourly wage, with the hourly wage rising from noon through midnight and falling between midnight and noon with a temporary peak during the morning rush hour. The variation around the average wage for the late night hours is relatively large due to the smaller number of driver-hours in my sample during that part of the day. Panel B of figure 4 overlays the plot of the average hourly wage by clock hour with a plot of the number of driver

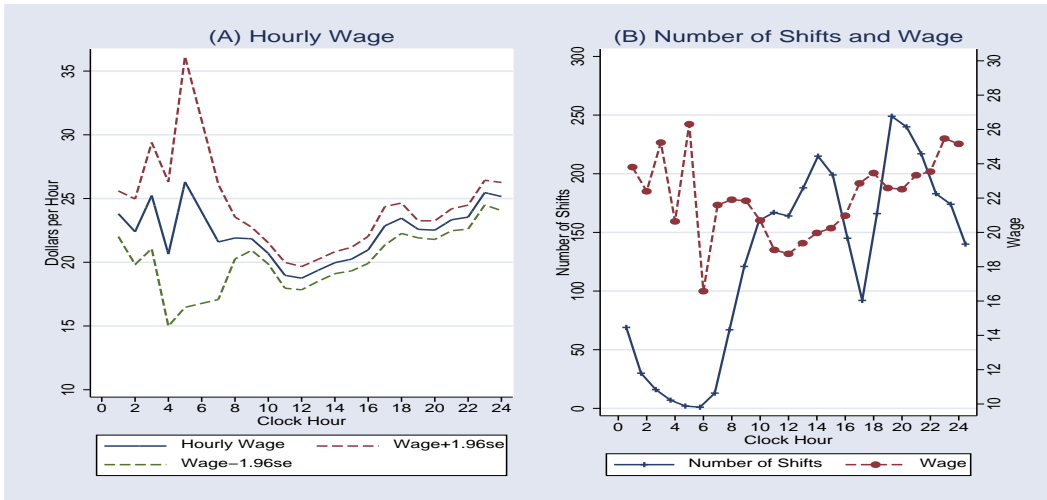


Figure 4: Hourly Wage and Number of Trips, by Clock Hour

hours in my sample by clock hour. While it is not the case that I have a random selection of cabs across the day, my trip sheets do provide some evidence on variation over the day in the number of cabs on the street. The minimum is at 6 AM, after which the number of driver-hours increases through 1 PM. The number of driver-hours drops sharply at 4 PM and 5 PM, likely reflecting the change of shifts, before increasing to the daily maximum at 7 PM. Subsequently, the number of driver-hours drops consistently through 6 AM.

It is interesting that the hourly wage is much less variable over the day than the supply of driver-hours. The number of driver hours ranges from 1 at 6 AM to 249 at 7 PM while the average wage varies from \$18.75 at noon to \$25.47 at 11 PM.⁹ This pattern is consistent with the number of drivers adjusting to daily fluctuations in the pattern of demand.¹⁰ The low wage at noon may reflect short lunch breaks not coded as such by drivers.

Table 2 contains the analysis of variance results from a series of regressions of the hourly wage on a sequence of sets of variables. The first column refers to a regression with only a constant. The root mean squared error of this regression (the standard deviation of the

⁹ The wage is even higher at \$26.30 at 5 AM, but this is based on only two driver hours.

¹⁰ Oettinger (1999) models the labor supply decisions of stadium vendors across days as a function of predictable fluctuations in demand. His model takes into account the facts that other vendors are making similar decisions and that these decisions affect own income. Oettinger's data includes the labor supply of all vendors. The analogous data for New York City taxi drivers would include the labor supply of all drivers. I do not have access to such data.

Table 2: Analysis of Variance of Hourly Wage

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Driver ID (ID)		x			x	x	x	x	x
Hour of Day (HoD)			x		x	x	x	x	x
Day of Week (DoW)				x	x	x			
HoD*DoW								x	x
Weather						x	x	x	x
Date							x	x	x
Date*ID									x
DF used	0	20	23	6	49	53	335	470	701
R-Squared	0.00	0.05	0.09	0.01	0.12	0.12	0.18	0.28	0.35
Root MSE	6.67	6.51	6.40	6.65	6.49	6.40	6.40	6.17	6.12

Note: The statistics in this table are based on linear regressions with dummy variables included for the indicated categories in each column. The sample used contains 3025 hours for which all sixty minutes are either part of a trip or waiting time between trips. Hours that include a break are excluded. The Weather variables include hourly rainfall, daily snowfall, an indicator for minimum temperature below 30 degrees F, and an indicator for maximum temperature greater than or equal to 80 degrees F.

wage) is 6.67. Controlling for driver fixed effects (column 2) accounts for only 5 percent of the variation in the hourly wage, and the root MSE is reduced slightly to 6.51. Controlling for the hour of the day (column 3) accounts for 9 percent of the variation in the hourly wage, and controlling for the day of the week (column 4) explains only 1 percent of the wage variation. Controlling for all three categories simultaneously (column 5) account for 12 percent of the variation in the hourly wage.

I controlled additionally for the weather (4 variables: 1) hourly rainfall, 2) daily snowfall, 3) daily low temperature less than 30 degrees F, and 4) daily high temperature greater than or equal to 80 degrees F.) in column 6. These variables do not improve the R-squared substantially, but the hourly wage is significantly related to the weather measures (p -value = 0.016). Specifically, the hourly wage is \$1.15 lower on hot days (p -value = 0.003), \$0.53 lower on cold days (p -value = 0.121), and, surprisingly, \$0.94 lower for each 0.1 inch of hourly rain (p -value = 0.104).

Controlling for the 316 distinct calendar dates in my sample (column 7) accounts for an additional 6 percentage points of the hourly wage variation. Interacting the day of the

Table 3: Autocorrelations of Hourly Wage

Periods	(1)	(2)	N
0	1.0	1.0	3025
1	0.0687	-0.0508	2232
2	0.0984	0.0292	1652
3	0.0625	-0.0130	1312
4	0.0261	-0.0362	1034
5	0.0644	-0.0020	785
6	0.1354	0.1326	539
7	-0.0082	-0.0022	334

Note: Column 1 contains autocorrelations of hourly wages indexed by hour on shift. Column 2 contains autocorrelations of residuals of hourly wages indexed by hour on shift. The residuals are calculated from a regression of the hourly wage on a complete set of interactions of hour of the day and day of the week.

week with the hour of the day (column 8) accounts raises the R-squared to 0.28 with a root MSE of 6.17. This illustrates the fact that hour-specific demand for taxi services varies over the week. Overall, more than 70 percent of the variation in the hourly wage is within-day inter-hour variation.

The last column of table 2 controls for the interaction of driver ID and calendar date. Essentially, each driver is allowed to have a day-specific mean hourly wage. While somewhat more of the variation in the hourly wage is accounted for with these interaction, it is still the case that almost two-thirds of the variation in the hourly wage cannot be accounted for by the interaction of driver id with calendar date and the interaction of the day of the week with the hour of the day.

Understanding the time-series properties of the hourly wage during a shift is important in understanding what drivers might infer from the current wage about earnings from continuing to drive. Toward this end, column A of table 3 contains the autocorrelation matrix of hourly wages for the first eight hours of shifts. The first order autocorrelation is surprisingly small at 0.07, suggesting that wages are not highly serially correlated over time.¹¹ And it may be that

¹¹ These autocorrelations are strikingly different from those presented by Camerer, et. al. (1997), who report first and second autocorrelations of the hourly wage of about 0.5.

part of this autocorrelation is related to time-of-day and day-of-week effects. I investigate this by examining the autocorrelations of residuals from a regression of the hourly wage on a complete set of interactions of hour-of-day and day-of-week effects. These autocorrelations are contained in column 2 of table 3, and they are small and negative suggesting a week negative inter-hour relationship in wage rates.

While there is significant day-to-day variation in the hourly wage, the results in tables 2 and 3 suggest that most variation in the wage appears to be non-forecastable within-day variation. On this basis, predicting hours of work with a model that assumes a fixed hourly wage rate during the day does not seem appropriate.

6 Estimation of the Labor Supply Models

I begin the presentation of estimates of labor supply models by estimating the labor supply model used by Camerer, et al (1997) and Chou (2000) and formulated in equation 4.5. I then present estimates of the direct target earnings model formulated in equations 4.2-4.4. Finally, I implement the probit model of the probability of stopping based on the latent variable defined in equation 4.1.

6.1 Estimation of the Log Hours Function

Table 4 contains estimates of OLS regressions of log hours on the log wage rate computed as the log of the ratio of daily income to daily hours for the 584 shifts in my sample. The estimates in column 1 of the table are for a specification that includes only the log wage as a regressor, and the estimated elasticity of labor supply is significantly negative at -0.688. The successive columns of the table are for specifications that include additional regressors. The specification in column 2 includes driver fixed effects. While there are clearly significant differences across drivers in hours worked, the estimated elasticity is virtually unchanged. Controlling for shift and the day of the week improves the fit of the model, with drivers on the night shift working about 13 percent more hours than drivers on the day shift. Once again, the estimated wage elasticity of labor supply is virtually unchanged. Finally, controlling for the weather in column 4 suggests that, while hours are unrelated to temperature extremes or

Table 4: Labor Supply Function Estimates
 OLS Regression of log Hours

Variable	(1)	(2)	(3)	(4)
Constant	4.012 (0.349)	3.882 (0.354)	3.776 (0.379)	3.778 (0.381)
log(wage)	-0.688 (0.111)	-0.647 (0.112)	-0.636 (0.115)	-0.637 (0.115)
Night Shift	---	---	0.128 (0.062)	0.134 (0.062)
Min Temp < 30	---	---	---	0.024 (0.058)
Max Temp \geq 80	---	---	---	0.055 (0.064)
Rainfall	---	---	---	-0.054 (0.071)
Snowfall	---	---	---	-0.093 (0.035)
Driver Effects	No	Yes	Yes	Yes
Day-of-Week Effects	No	No	Yes	Yes
R-squared	0.063	0.162	0.185	0.198

Note: The sample includes 584 shifts for 21 drivers. The dependent variable is log hours worked (driving time plus time between fares excluding declared breaks and breaks between fares one hour or longer). The mean of the dependent variable is 1.84. Standard errors are in parentheses.

rainfall, drivers do drive fewer hours when it has snowed. The estimated elasticity is again unchanged.

There is strong potential for negative bias in the estimated elasticity because the wage is computed mechanically using the inverse of hours worked. Additionally, since the wage is not constant over the working day, it is not likely that the wage measured this way can be considered parametric to the labor supply decision. Absent a convincing exclusion restriction from the hours function to provide an instrument for the wage, I do not have a solution to this problem.

My finding of a negative elasticity of labor supply using this method is consistent with the findings of Camerer, et al (1997) and Chou (2000). In fact, the magnitude of my estimates is comparable to theirs. The consistency of our findings is important because it implies that any differences in results that I find using other empirical models is due to differences in

analytic approach rather than to differences in the data used.

6.2 OLS Estimation of the Target Earnings Function

Table 5 contains OLS estimates of the target function in equation 4.2 where the target is measured for each shift as the average of total shift income and total shift income minus the last fare. There are clearly significant differences across drivers in their targets. Based on the estimates in column 1 of the table, the standard deviation of the average target values across drivers is \$59.02. However, there remains substantial unexplained variation within drivers in their targets. The R-squared of the model in column 2 is only 0.183 and the root mean squared error after accounting for driver fixed effects is \$54.30. This is not consistent with a target earnings model.

Table 5: OLS Estimates of Target Function

Variable	(1)	(2)	(3)	(4)
Constant	156.937 (2.444)	152.966 (8.695)	---	157.814 (8.556)
Night Shift	---	---	8.637 (7.348)	9.600 (7.337)
Min Temp < 30	---	---	---	-0.248 (7.004)
Max Temp \geq 80	---	---	---	13.707 (7.747)
Rainfall	---	---	---	-3.073 (8.535)
Snowfall	---	---	---	-7.751 (4.270)
Driver Effects	No	Yes	Yes	Yes
Day-of-Week Effects	No	No	Yes	Yes
R-squared	0.000	0.183	0.214	0.224
Root MSE	59.02	54.30	53.59	53.43

Note: The sample includes 583 shifts for 21 drivers. One shift is not included because it ended after one trip. The dependent variable is average of total earnings for the day and total earnings for the day minus the last fare. The mean of the dependent variable is \$156.94. Standard errors are in parentheses.

Controlling for some shift-specific factors that could affect targets does not change this conclusion. Adding a night-shift indicator and fixed effects for day of the week improves the

fit but reduces the root mean squared error only to \$53.59. Adding controls for temperature and precipitation in column 4 does not improve the fit significantly. Drivers appear to set lower targets when it snows and higher targets when it is very hot, but these results are of marginal statistical significance.

There is no firm test of the target earnings model implicit in the direct estimation of the target function. However, the fact that the target shows substantial variation within drivers across days, even after accounting for the day of the week, shift, and the weather suggests that if drivers are target earners then their targets are highly variable. Such variable targets severely limits the value of the target earnings construct as a tool for understanding labor supply.

6.3 Estimation of the Probit Optimal Stopping Model

The probability that a driver ends his shift increases sharply as hours and income accumulate. The left side of table 6 contains the simple empirical hazard of a driver stopping after a trip ending in a given time interval since the start of the shift. The likelihood that a driver will stop increases from 0.5 percent in the first two hours to 15 percent in hour 8 and to over 0.25 by hour 12. Note that since there are multiple trips ending in any given hour, the probability that a driver stops in a given hour is considerably larger than the single-trip hazards listed

Table 6: Empirical Hazard of Stopping by Hours and by Income

Hour	Hazard	Income (\$)	Hazard
<=2	0.0050	<25	0.0016
3-5	0.0244	25-49	0.0078
6	0.0517	50-74	0.0149
7	0.0992	75-99	0.0228
8	0.1467	100-124	0.0338
9	0.1147	125-149	0.0596
10	0.1609	150-174	0.1111
11	0.2450	175-199	0.1477
>=12	0.2636	200-224	0.1283
		>=225	0.1977

Note: Based on the sample of 13461 trips in 584 shifts for 21 drivers. These are the probability of stopping after a trip in the designated hour of the shift or with total income in the indicated range conditional on not having stopped earlier.

in the table. The right side of table 6 contains the the simple empirical hazard of a driver stopping after a trip ending with shift income in the given interval. The likelihood that a driver will stop increases from 0.16 percent when the driver has earned less than \$25 to 11 percent when the driver has earned \$150-174. Analogously to the case of hours, there are likely to be multiple trips ending in any of these intervals so that the probability that a driver stops with income in any particular interval is considerably larger than the single-trip hazards listed in the table.

In section 4, I outlined probit model of the probability that a driver stops after a given trip as a function of hours worked to that point, income earned to that point, current location, and fixed effects for driver, calendar date, and the interaction of hour of the day with day of the week. The latent variable defined in equation 4.1 forms the basis for estimation of this model. Table 7 contains estimates of this model where income and hours are constrained to enter linearly.¹² The normalized probit estimate is $\hat{\beta} \cdot \phi(X^*\beta)$, where $\phi(\cdot)$ is the standard normal density. Given that the probability of a shift ending early in the shift is very small and there are more trips early than late, evaluating the density at the mean for this normalization could be misleading.¹³ As a result, I selected reasonable values for the key variables. I evaluate the marginal effect of a variable on the probability of quitting after 8 total hours (5.5 trip hours, 2.5 waiting hours, 0.5 break hours) on a dry day with moderate temperatures in mid-town Manhattan at 2PM.

The estimates in the first column of table 7 contain only the hours measures and income. Total hours is defined as the sum of hours spent on trips, hours “waiting” between trips, and hours on break. The probability of stopping is significantly positively related to income and not significantly related to time worked. The magnitude of the income effect is that a \$10 increase in income implies a 0.53 percentage point increase in the probability of stopping work. The estimates in column 2 include driver fixed effects, and, based on the improvement in the log-likelihood, there are significant differences across drivers in their probability of quitting. Interestingly, after accounting for inter-driver differences, the probability of quitting

¹² This specification constraint is relaxed in table 8.

¹³ The probit specification naturally accounts for this nonlinearity as the marginal effect of any given variable on the probability of quitting is smaller in the tails of the distribution.

Table 7: Hazard of Stopping after Trip: Normalized Probit Estimates

Variable	X^*	(1)	(2)	(3)	(4)	(5)
Total Hours	8.0	0.013 (0.009)	0.037 (0.012)	0.011 (0.005)	0.010 (0.005)	0.010 (0.005)
Waiting Hours	2.5	0.010 (0.010)	-0.005 (0.012)	0.001 (0.006)	0.004 (0.006)	0.004 (0.005)
Break Hours	0.5	0.006 (0.008)	-0.015 (0.011)	-0.003 (0.005)	-0.001 (0.005)	-0.002 (0.005)
Shift Income x100	1.5	0.053 (0.022)	0.036 (0.030)	0.014 (0.015)	0.016 (0.016)	0.011 (0.015)
Min Temp < 30	0.0	---	---	---	0.002 (0.007)	0.002 (0.006)
Max Temp \geq 80	0.0	---	---	---	-0.016 (0.007)	-0.016 (0.007)
Hourly Rain	0.0	---	---	---	0.005 (0.124)	0.015 (0.113)
Daily Snow	0.0	---	---	---	0.007 (0.006)	0.007 (0.005)
Downtown	0.0	---	---	---	---	0.000 (0.006)
Uptown	0.0	---	---	---	---	0.000 (0.005)
Bronx	0.0	---	---	---	---	0.075 (0.057)
Queens	0.0	---	---	---	---	0.044 (0.038)
Brooklyn	0.0	---	---	---	---	0.078 (0.036)
Kennedy Airport	0.0	---	---	---	---	0.056 (0.031)
LaGuardia Airport	0.0	---	---	---	---	0.061 (0.029)
Other	0.0	---	---	---	---	0.134 (0.071)
Driver (21)		No	Yes	Yes	Yes	Yes
Day-of-Week (7)		No	No	Yes	Yes	Yes
Hour-of-Day (19)	2PM	No	No	Yes	Yes	Yes
Log L		-2039.2	-1965.0	-1789.5	-1784.7	-1767.6

Note: The sample includes 13461 trips in 584 shifts for 21 drivers. Probit estimates are normalized to reflect the marginal effect at X^* of X on the probability of stopping. The normalized probit estimate is $\hat{\beta} \cdot \phi(X^* \beta)$, where $\phi(\cdot)$ is the standard normal density. The values of X^* chosen for the fixed effects are equally weighted for each day of the week and for each driver. The hours from 5AM to 10AM have a common fixed effect. The evaluation point is after 5.5 driving hours, 2.5 waiting hours, 0.5 break hours in a dry hour on a day with moderate temperatures in mid-town Manhattan at 2PM. Robust standard errors accounting for clustering by shift are reported.

is not significantly related to income but is significantly related to total hours. A one hour increase in total time is associated with a 3.7 percentage point increase in the probability of quitting. Given that the probability of stopping for trips ending in the eighth hour is 0.14, this is a substantial effect. In neither column 1 nor column 2 do waiting hours or break hours have a significantly different effect on the probability of stopping work than do hours on trips.

The estimates in column 3 additionally control for the hour of the day and the day of the week. The marginal effect of hours on the probability of stopping is much lower in this specification, with a one hour increase in total time increasing the probability of quitting by 1.1 percentage points. This is not surprising, given the bunching of shift start times and ending times shown in figure 2. The magnitude and statistical significance of the marginal effect of income on the probability of stopping is reduced further by the additional control variables.

The estimates in column 4 add the weather measures. The only statistically significant finding is that the probability of stopping work is 1.6 percentage point lower on hot days. Perhaps surprisingly, given anecdotal reports of the difficulty of finding taxis in rainy weather, there is no relationship between hourly rainfall and the probability of stopping. This would suggest that reported difficulty in finding a taxi in rainy weather is due to increased demand. However, the results reported in table 2 and the related discussion suggest that the hourly wage is lower in rainy weather, implying that there has not been an increase in demand relative to supply.

Finally, the estimates in column 5 include controls for geographic location at the end of the trip. The omitted location is midtown Manhattan (between 14th Street and 59th Street). The general pattern is clear. Drivers are substantially more likely to stop working after a fare that takes them outside Manhattan. For example, drivers have a 6.1 percentage point higher probability of stopping work after a fare to LaGuardia Airport, located in Queens. This may be due to drivers being closer to home in the outlying boroughs. Controlling for location appears to have an insubstantial effect on the coefficients of the key hours and income measures.

A potential problem with the estimates in table 7 is that hours and income are constrained

to affect the underlying latent variable linearly. While the probit function introduces a particular form of nonlinearity that reduced the marginal effects in the tails of the distribution, the sharp increase in the hazard of a shift ending as the shift progresses suggests that a more flexible specification is appropriate. Table 8 contains normalized probit estimates of the probability of stopping after a trip estimated using a nine-category classification of total hours since starting the shift and a ten-category classification of fare income earned since the start of the shift. I use total hours since the start of the shift based on the results in table 7 that hours between trips, either waiting for a fare or on break, do not appear to have different effects than driving hours with fares. The marginal effect of a variable on the on the probability of quitting is calculated for a driver in his eighth 8 hour on the shift having earnings between \$150 and \$174 on a dry day with moderate temperatures in mid-town Manhattan at 2PM.

The estimates in the first column of table 8 are for a model that includes only the unconstrained hours effects. These show the sharply increasing probability of ceasing work after a trip as hours accumulate. A driver is 9.5 percentage points less likely to stop work after a trip ending in the sixth hour than after one ending in the eighth hour. On the other end, a driver is 9.8 percentage points more likely to stop work after a trip ending in the eleventh hour than after one ending in the eighth hour. The estimates in the second column are for a model that includes only the unconstrained income effects. These also a sharply increasing probability of ceasing work after a trip as income accumulates. A driver is 9.6 percentage points less likely to stop work after a trip ending with income of \$50-74 than after one ending with income of \$150-174. On the other end, a driver is 8.7 percentage points more likely to stop work after a trip ending with income greater than or equal to \$225 than after one ending with income of \$150-174.

The fact that either hours or income, taken alone, are strong predictors of the likelihood of stopping work is not surprising given their strong positive correlation within a shift ($\rho = 0.94$). When both sets of variables are included in the model in column 3 of table 8, the hours effects remain strong, with a driver being 8.4 percentage points less likely to stop work after a trip ending in the sixth hour than after one ending in the eighth hour and 8.6 percentage points more likely to stop work after a trip ending in the eleventh hour than

Table 8: Hazard of Stopping after Trip: Normalized Probit Estimates

Variable	(1)	(2)	(3)	(4)
Hour ≤ 2	-0.142 (0.014)	---	-0.137 (0.017)	-0.042 (0.012)
Hour 3-5	-0.122 (0.014)	---	-0.108 (0.017)	-0.028 (0.010)
Hour 6	-0.095 (0.015)	---	-0.084 (0.018)	-0.026 (0.009)
Hour 7	-0.048 (0.016)	---	-0.041 (0.018)	-0.012 (0.008)
Hour 9	-0.032 (0.021)	---	-0.034 (0.022)	-0.006 (0.010)
Hour 10	0.014 (0.025)	---	0.011 (0.033)	0.031 (0.020)
Hour 11	0.098 (0.038)	---	0.086 (0.051)	0.085 (0.036)
Hour ≥ 12	0.117 (0.063)	---	0.109 (0.075)	0.119 (0.060)
Income < 25	---	-0.109 (0.010)	-0.121 (0.022)	-0.036 (0.014)
Income 25-49	---	-0.103 (0.010)	-0.067 (0.027)	0.005 (0.022)
Income 50-74	---	-0.096 (0.010)	-0.063 (0.021)	-0.002 (0.014)
Income 75-99	---	-0.088 (0.010)	-0.065 (0.020)	-0.010 (0.011)
Income 100-124	---	-0.077 (0.011)	-0.053 (0.018)	-0.009 (0.009)
Income 125-149	---	-0.052 (0.011)	-0.035 (0.017)	-0.007 (0.008)
Income 175-199	---	0.037 (0.018)	0.014 (0.021)	0.011 (0.010)
Income 200-224	---	0.017 (0.020)	-0.022 (0.025)	0.006 (0.013)
Income ≥ 225	---	0.087 (0.019)	0.009 (0.032)	0.015 (0.018)
Driver (21)	No	No	No	Yes
Day-of-Week (7)	No	No	No	Yes
Hour-of-Day (19)	No	No	No	Yes
Location (9)	No	No	No	Yes
p -value Hours = 0	0.000	---	0.000	0.000
p -value Income = 0	---	0.000	0.008	0.281
Log L	-2028.9	-2058.7	-2016.5	-1753.1

Note: The sample includes 13461 trips in 584 shifts for 21 drivers. Probit estimates are normalized to reflect the marginal effect at X^* of X on the probability of stopping. The normalized probit estimate is $\hat{\beta} \cdot \phi(X^* \beta)$, where $\phi(\cdot)$ is the standard normal density. The values of X^* chosen for the fixed effects are equally weighted for each day of the week and driver. The hours from 5AM-10AM have a common fixed effect. The evaluation point is at 8 total hours with income of \$150-174 in a dry hour on a day with moderate temperatures in mid-town Manhattan at 2PM. Robust standard errors accounting for clustering by shift are reported.

after one ending in the eighth hour. However, the income effect is attenuated, particularly at the high end. A driver is 6.3 percentage points less likely to stop work after a trip ending with income of \$50-74 than after one ending with income of \$150-174. However, there is no significant difference in the probability of stopping after a fare with income greater than or equal to \$225 relative to one ending with income of \$150-174. Still, the hypothesis that the income variables all have zero coefficients can be rejected at conventional levels (p -value = 0.008).

Controlling for hours, income may be picking up inter-temporal and inter-driver differences in the value of continuing to drive. In order to address this possibility, the estimates in column 4 of table 8 are derived from a model that includes additional controls for driver, day of the week, hour of the day, and geographic location. As I showed in table 7, these variables are all important determinants of the probability of stopping. Interestingly, hours remain an important factor in the stopping decision, but the pattern is changed somewhat. The differences are much smaller early in the shift, with only a 2.5 percentage point reduction in the likelihood of stopping in the sixth hour relative to the eighth hour. However, the differences late in the shift are unchanged, with an 8.5 percentage point increase in the likelihood of stopping in the eleventh hour relative to the eighth hour.

Income is no longer a factor. Only the coefficient on income less than or equal to \$25 is significantly different from zero, and its magnitude suggests only a 3.6 percentage point reduction relative to the \$150-174 category. All other estimates on the income categories are smaller along with not being significantly different from zero. The hypothesis that all of the income coefficients are zero cannot be rejected at conventional levels (p -value = 0.28).

The clear conclusion from this analysis is that hours of work is a central determinant of the stopping decision but that daily income is not related to the stopping decision. Thus, there are not important daily income effects, and the evidence does not support a conclusion that taxi drivers are target earners.

7 Labor Supply Models for Specific Drivers

A restriction implicit in the analysis of the previous section is that drivers are assumed to follow a common behavioral model aside from an intercept shift in the underlying function determining the probability of stopping after a trip. This could be misleading, particularly if some drivers are target earners and other are not or if drivers are target earners but have heterogeneous targets. In order to investigate the importance of driver heterogeneity, I estimate separate stopping models for five drivers for whom I have a large number of shifts.

This analysis is complicated by the fact that drivers are fairly consistent across days in their starting times, and, while their ending times show more variation as a result of the labor supply decision, individual drivers do not tend to “cover the clock.” Shifts for particular drivers tend to end in a fairly narrow time window. As a result, it is not feasible to control for clock hour in the analysis. This is a problem because clock hour is an important component of drivers’ calculations of the value of continuing to work, the inability to control directly for clock hour makes interpretation of estimates of the effect of cumulative hours and cumulative income on the probability of stopping difficult. Since, for any particular driver, shift starting times are fairly standard across days, the cumulative hours measure is accounting both for changes in the marginal utility of leisure as hours accumulate and for changes in the value of continuing to drive as the clock hour advances. Given the dual role that cumulative hours are playing and the substantial correlation of cumulative income with cumulative hours, the interpretation of the effect of cumulative income on the probability of stopping work is also affected.

These problems of interpretation notwithstanding, table 9 contains estimates of probit models of the probability of stopping after a trip for each of five drivers. The set of control variables is pared down due to lack of variability within drivers. There are no controls for clock hour. A single geographic location control (Manhattan vs. Non-Manhattan) is included along with a single weather variable (hourly rainfall). Fixed effects for day of the week are included in all specifications.

The results are somewhat mixed but generally provide support for the importance of cumulative hours as the primary factor in determining the probability of stopping. The

Table 9: Driver-Specific Hazard of Stopping after Trip: Normalized Probit Estimates

Variable	Driver 4	Driver 10	Driver 16	Driver 18	Driver 20
Hour \leq 4	-0.290 (0.107)	-0.131 (0.059)	-0.206 (0.063)	-0.006 (0.011)	-0.423 (0.059)
Hour 5-7	-0.198 (0.095)	-0.055 (0.067)	-0.169 (0.064)	-0.004 (0.009)	-0.397 (0.056)
Hour \geq 10	0.256 (0.107)	0.143 (0.100)	0.107 (0.096)	0.093 (0.042)	---
Income < 100	-0.314 (0.072)	-0.088 (0.061)	-0.140 (0.052)	0.012 (0.026)	-0.162 (0.097)
Income 100-149	-0.241 (0.065)	0.015 (0.065)	-0.122 (0.049)	0.016 (0.020)	0.003 (0.102)
Income \geq 200	0.023 (0.165)	0.021 (0.085)	0.087 (0.082)	0.027 (0.014)	-0.154 (0.112)
Manhattan	-0.308 (0.128)	-0.203 (0.132)	-0.122 (0.076)	-0.015 (0.016)	-0.317 (0.110)
Hourly Rain	-0.220 (0.762)	0.841 (0.660)	-1.075 (0.588)	-0.036 (0.089)	0.237 (0.722)
Day-of-Week (7)	Yes	Yes	Yes	Yes	Yes
p -value Hours = 0	0.003	0.000	0.000	0.000	0.000
p -value Income = 0	0.001	0.105	0.020	0.255	0.048
Number of Shifts	40	45	70	72	46
Number of Trips	884	912	1754	2023	1125
Log L	-124.2	-143.0	-226.1	-250.2	-127.9

Note: Probit estimates are normalized to reflect the marginal effect at X^* of X on the probability of stopping. The normalized probit estimate is $\hat{\beta} \cdot \phi(X^*\beta)$, where $\phi(\cdot)$ is the standard normal density. The values of X^* chosen for the fixed effects are equally weighted for each day of the week. The evaluation point is after 8 or 9 total hours with income of \$150-199 in a dry hour in Manhattan. Robust standard errors accounting for clustering by shift are reported. Driver 20 worked ten hours twice and each time stopped after one trip in that range. These trips are aggregated with the 8-9 hour category.

hypothesis that the cumulative hours variables have zero coefficients can be rejected at conventional levels for all drivers, and the pattern of hours effects consistently implies lower probabilities of quitting earlier in shifts and higher probabilities of quitting late in shifts. The evidence is mixed with regard to the importance of cumulative income. For drivers 10 and 18 the hypothesis that the cumulative income variables have zero coefficients cannot be rejected at conventional levels. The income variables have a marginally significant effect for driver 20, but the pattern of coefficients is not consistent. Driver 20 appears less likely to stop when income is both below \$100 and above \$200 relative to intermediate levels. The

income variables have a consistent statistically significant effect for drivers 4 and 16, with lower probabilities of stopping when income is below \$150.

8 Final Remarks

The inter-temporal labor supply model nests the conventional upward-sloping labor supply model and the target earnings model. The conventional model results from a case where discount and interest rates are sufficiently small that idiosyncratic daily fluctuations in income are not of consequence in determining daily labor supply. There is inter-temporal substitution in this case. The target earnings model results from a case where discount and interest rates are so high that idiosyncratic daily fluctuations in have strong effects on daily labor supply. There is no inter-temporal substitution in this case, and this implies that expenditures and income must match on a daily basis.

Note that a more permanent shift in earnings opportunities, such as that which likely occurred in New York after September 11, 2001 and that which occurs regularly in recessions can have important income effects.¹⁴ It is certainly possible in this case that the labor supply schedule could be backward-bending in response to these long-run changes where there is not the possibility of a high wage tomorrow.

I have focused in my analysis on the daily labor supply decision, which I have modeled as being embedded in an inter-temporal labor supply model. However, largely due to data limitations, I have not considered relationships across days in individual labor supply decisions. A fuller analysis requires a complete time-series of drivers' labor supply over a substantial period of time. Such an analysis would allow estimation of income effects in response to long-run changes in earnings opportunities.

I conclude from my empirical analysis of the stopping behavior of New York City cab drivers that, when accounting for earnings opportunities in a reduced form with measures of clock hours, day of the week, the weather, and geographic location, cumulative hours worked on the shift is a primary determinant of the likelihood of stopping work while cumulative

¹⁴ Das Gupta (2002) documents the substantial negative effect that the events of 9/11 had on taxi driver income in New York City. She also provides some evidence that hours worked per shift increased slightly.

income earned on the shift is weakly related, at best, to the likelihood of stopping work. This is consistent with there being substantial opportunity for inter-temporal substitution in labor supply and no support for the hypothesis that taxi drivers are target earners.

9 References

- Altonji, Joseph and Christina Paxson, "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs" *Journal of Labor Economics*, Vol. 6, No. 2. (Apr., 1988), pp. 254-276.
- 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, 1999, vol 3a, pp. 1559-1695.
- 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*, May 1997, pp. 407 – 441.
- Chou, Yuan K. "Testing Alternative Models of Labor Supply: Evidence from Taxi Drivers in Singapore," University of Melbourne, Department of Economics Research Paper No. 768, 2000.
- Das Gupta, Monisha. "Of Hardship and Hostility: The Impact of 9/11 on Taxi Drivers in New York City," University of Hawaii at Manoa, undated unpublished mimeo, 2002.
- Dickens, William and Shelley Lundberg. "Hours Restrictions and Labor Supply," *International Economic Review*, 34 (February 1993), pp. 169-192.
- Fehr, Ernst and Lorenz Goette. "Inter-temporal Substitution at Work? Evidence from a Field Experiment," mimeo, April 2002.
- Ham, John. "Estimation of a Labor Supply Model with Censoring Due to Unemployment and Underemployment," *The Review of Economic Studies* 49 (1982), pp. 335-354.
- Kahn, Shulamit and Kevin Lang. "The Effect of Hours Constraints on Labor Supply Estimates," *The Review of Economics and Statistics* 93 (1991), 605-611.
- Killingsworth, Mark R. and James J. Heckman. "Female Labor Supply: A Survey," in *Handbook of Labor Economics*, O. Ashenfelter and R. Layard, eds. Amsterdam, North Holland, 1986, vol 1, pp. 103-204.
- MaCurdy, Thomas E. "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, December 1981, pp. 1059-1085.
- Oettinger, Gerald S. "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors," *Journal of Political Economy* (April 1999), pp. 360-392.
- Pencavel, John. "Labor Supply of Men: A Survey," in *Handbook of Labor Economics*, O. Ashenfelter and R. Layard, eds. Amsterdam, North Holland, 1986, vol 1, pp. 3-102.