Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers

Henry S. Farber

Princeton University

The labor supply of taxi drivers is consistent with the existence of intertemporal substitution. My analysis of the stopping behavior of New York City cabdrivers shows that daily income effects are small and that the decision to stop work at a particular point on a given day is primarily related to cumulative daily hours to that point. This is in contrast to the analysis of Camerer et al., who find that the daily wage elasticity of labor supply of New York City cabdrivers is substantially negative, implying large daily income effects. This difference in findings is due to important differences in empirical methods and to problems with the conception and measurement of the daily wage rate used by Camerer et al.

I. 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 whereas 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, al-

I thank 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. This paper was written while I was a visiting scholar at the Russell Sage Foundation.

[*Journal of Political Economy,* 2005, vol. 113, no. 1]

ternatively, can select a job with the optimal wage-hours combination from a dense joint distribution of jobs. Evidence that neither of these is a credible assumption is that the distribution of hours is quite lumpy, with a substantial fraction of workers reporting usual weekly hours of precisely 40.[1]

There is an emerging literature on labor supply that is not subject to this criticism because it investigates labor supply responses in settings in which workers are free to set their hours of work. By and large, this literature finds substantial positive labor supply elasticities and evidence consistent with the standard neoclassical labor supply model. In this study, I contribute to this emerging literature by providing a new analysis of the labor supply of New York City taxi drivers. My analysis shows that daily income effects are small, as one would expect in a standard intertemporal labor supply model, and that the decision to stop work at a particular point on a given day is primarily related to cumulative daily hours to that point.

My findings are in direct contrast to those of Camerer et al. (1997), who also study the labor supply of New York City cabdrivers. They find that the daily wage elasticity of labor supply of New York City cabdrivers is substantially negative, implying large daily income effects that could be interpreted as target earnings behavior. This difference in findings is due to important differences in empirical methods and to problems with the conception and measurement of a daily wage rate used by Camerer et al.

A. *Hours Constraints in the Traditional Labor Supply Literature*

There is a substantial literature demonstrating that workers appear not to be free to set hours and providing some estimates of the importance of this restriction on 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 35–40 percent of workers would prefer to work more hours at their current wage rate, with a smaller fraction preferring to work 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 workers' annual hours is larger

[1] This is based on tabulation of the 2002 merged outgoing rotation group files from the Current Population Survey. It is the case that over 50 percent of workers report 40 hours as their usual weekly work hours. While some of this heaping at 40 hours may be due to the natural tendency of respondents to round off, it is clear that there is substantial bunching of hours.

for workers who change jobs than for those who remain on the same job. Dickens and Lundberg (1993) estimate a model of labor supply in which 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.

B. *The Emerging Literature on Labor Supply in Settings with Flexible Hours*

In this subsection, I critically review four recent studies that analyze labor supply responses in settings in which 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 intertemporal labor supply elasticity.

1. 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 in this case 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 noncooperative multivendor participation problem that, in equilibrium, has more vendors show up to work games with larger expected attendance. Oettinger concludes that there is a substantial positive intertemporal labor supply elasticity implicit in the daily participation decisions of the vendors he studies.

2. 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 receive a fixed compensation per message delivered and are free to set their own hours. The experiment consisted of dividing the company's messengers into two groups (A and B). Group A received a significantly enhanced fee per message 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 intertemporal 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 the daily labor supply is inconsistent with the standard neoclassical model, and they argue that the decline is due to a combination of (1) messengers being loss-averse relative to a fixed daily benchmark or target and (2) a 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 in which drivers want to supply substantially more labor, it is efficient for the messengers to work more days but work fewer hours each day.

3. Taxi Drivers: Camerer et al. (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 prespecified 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 the 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. As in 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 in which it is easy to make money (pick low-hanging fruit, so to speak), the drivers quit early, whereas on days in which fares are scarce, drivers work longer hours. If workers can substitute labor for leisure intertemporally across days, then they should work more on days with higher wage rates relative to other days. 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 effect on daily labor supply 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" (1997, 408), and they report 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. Fare opportunities vary dramatically and unpredictably over the course of a day. In my analysis, I present evidence that within-day variation swamps between-day variation in accounting for hourly wage variation. In this context, characterizing a day by the average income per hour earned that day clearly makes little sense.

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" in which, if there is any misspecification or measurement error, there will be a negative bias on the coefficient of the wage. Both Camerer et al. and Chou address this concern through the use of an instrumental variable estimation in which the instrument is the average daily wage of other workers on the same calendar date. However, if there are calendar date effects on the wage that are also correlated with labor supply conditional on the wage, 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. I estimate a model of the decision to stop work or continue driving at the conclusion of each fare. Estimates based on this approach, using new data on New York City taxi drivers, show that the primary determinant of the decision to stop work is cumulative hours worked on that day. There are no sub-

[2] Indeed, the literature on intertemporal substitution in labor supply as it relates to macroeconomics typically assumes no income effects of shocks to annual income on labor supply (the so-called λ-constant assumption). See, e.g., MaCurdy (1981) and Pencavel (1986).

stantial income effects, and the labor supply behavior of the cabdrivers is consistent with the standard neoclassical model.

Camerer et al. graciously made their TRIP data available to me, and my reanalysis of their data using my framework verifies my finding that cumulative hours worked are the primary determinant of the decision to stop work. Additionally, I have applied their approach to my new data, and I am able to reproduce their estimate of a negative labor supply elasticity. Taken together, the pair of findings reported in this paragraph strongly imply that the difference in our results is due to the different econometric and conceptual frameworks rather than to differences in data.

II. Conversations with Cabdrivers

While they were not conducted in a systematic fashion, I have had informal conversations with cabdrivers in New York City and elsewhere when traveling for the past few years. The information I have gained is not meant as evidence to test competing models. However, it does provide some information on what cabdrivers 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 most 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 11 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 two exceptions (out of about 25 drivers interviewed), 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?

Interestingly, the two drivers who said they had a target both owned their cabs. These drivers clearly explained to me what their target was and how it was derived on the basis of their expenses. 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? One driver answered that (1) it usually took 10–11 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." This driver did not, in fact, appear to be a target earner. The other driver responded that (1) it usually took eight to nine hours, (2) he would continue driving to reach the target, and (3) he would stop when the target was reached early and spend more time with his family. This single driver did, in fact, appear to be a target earner.

My impression from these interviews taken together is that drivers do not consciously behave as though 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 at which 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.

III. A Model of Taxi Driver Daily Labor Supply

The standard employment arrangement of New York City cabdrivers 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 economist's 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 in which 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 intertemporal utility function for a cabdriver 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

$$U_t = a(x_t) + b(l_t), \tag{1}$$

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

$$U = \sum_{t=0}^{T} (1 + \rho)^{-t}[a(x_t) + b(l_t)], \tag{2}$$

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

$$Y_0 + \sum_{t=0}^{T} (1 + r)^{-t} y_t(1 - l_t) = \sum_{t=0}^{T} (1 + r)^{-t} x_t, \tag{3}$$

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

The Lagrangian expression for constrained maximization of this utility function is

$$V = \sum_{t=0}^{T} (1 + \rho)^{-t}[a(x_t) + b(l_t)] + \lambda\left\{Y_0 + \sum_{t=0}^{T} (1 + r)^{-t}[y_t(1 - l_t) - x_t]\right\}, \tag{4}$$

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

$$\frac{\partial V}{\partial x_t} = a'(x_t) - \lambda\theta^t = 0, \tag{5}$$

$$\frac{\partial V}{\partial l_t} = b'(l_t) - \lambda\theta^t y_t'(1 - l_t) = 0, \tag{6}$$

and

$$\frac{\partial V}{\partial \lambda} = Y_0 + \sum_{t=0}^{T} (1 + r)^{-t}[y_t(1 - l_t) - x_t] = 0, \tag{7}$$

where $\theta = (1 + \rho)/(1 + r)$. Solving equations (5) and (6) for $\lambda\theta^t$ yields the result that

$$y_t'(1 - l_t) = \frac{b'(l_t)}{a'(x_t)}. \tag{8}$$

A. *Income Effects and the Shape of the Labor Supply Function*

Equation (8) implies that hours are selected so that the marginal wage from working an additional increment of 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 $y_t'(1 - l_t)$, which I call the marginal wage, would equal the fixed wage rate, and equation (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 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 the rate of time preference is much larger than the market interest rate (θ is large), then the individual is impatient relative to the market. In this case, the relevant time horizon is short and measurable daily income effects on labor supply are possible. In contrast, if θ is smaller, implying that the rate of time preference is not substantially larger than the market interest rate, then individuals will smooth their consumption of goods and leisure over time, and there will not be large daily income effects on labor supply.

Given that taxi drivers, like virtually all workers, make consumption commitments that span many days (e.g., apartment rental), it seems clear that they are able and desire to smooth consumption across days. In terms of the model, taxi drivers have small values of θ at the daily level. The clear prediction is that daily hours worked by taxi drivers are positively related to transitory variation in the marginal wage. Daily income effects are inconsequential.

A more permanent shift in earnings opportunities, such as the one that likely occurred in New York after September 11, 2001, and that occurs regularly in recessions, can have important income effects.[3] 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.

The prediction of the intertemporal model with regard to transitory changes in the marginal wage stands in stark contrast to the prediction of daily target earnings behavior. 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 intertemporal labor supply model, this is an extreme case of a large value of θ 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.

B. *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 taxicabs varies during the day. Thus it would not necessarily be optimal for drivers to quit on the basis of 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 them.

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 can end the shift. The theory outlined here has several sharp predictions for this modeling approach.

1. 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.

[3] 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.

2. 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 intertemporal nature of daily labor supply and the resulting small daily income effect.
3. 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.

IV. Empirical Models of Taxi Driver Labor Supply

A. *The Discrete-Choice Stopping Model*

As I noted in the previous section, I estimate a model of taxi driver daily labor supply as a survival time model in which 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

$$R_{idc}(\tau) = \gamma_1 h_\tau + \gamma_2 y_\tau + \boldsymbol{X}_{idc}\boldsymbol{\beta} + \mu_i + \epsilon_{idc\tau}, \tag{9}$$

where i indexes the particular driver, d indexes the date, and c indexes hour of the day. The quantity h_τ measures hours worked on the shift at τ, y_τ measures income earned on the shift at τ, and $\boldsymbol{X}$ measures other factors affecting the determination of the optimal stopping time and the comparison with τ. Elements of the vector $\boldsymbol{X}_{idc}$ include measures of weather and sets of fixed effects for hour of the day, day of the week, and location within New York City. These measures are included to capture variation in earnings opportunities from continuing to drive. The quantity ϵ is a random component with a standard normal distribution. The individual stops driving at τ if $R_{idc}(\tau) \geq 0$, and this implies a standard probit specification based on the latent variable defined in equation (9).

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.

B. Camerer et al.'s Target Earnings Model

Camerer et al. 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

$$\ln H_{it} = \eta \cdot \ln W_{it} + \boldsymbol{X}_{it}\boldsymbol{\beta} + \epsilon_{it}, \tag{10}$$

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 $\boldsymbol{X}_{it}$ are other factors affecting labor supply. The parameter η is meant to represent the elasticity of labor supply, and Camerer et al.'s estimates of η are strongly negative.

An important conceptual problem with this model is that it relies on there being significant exogenous transitory day-to-day variation in the average wage. This is the variation that drives the estimate of η in equation (10). However, as I demonstrate below, there is not significant transitory interday variation in the average wage. Nor is there significant autocorrelation in the hourly wage on a particular day. Thus it is hard to see a source of legitimate variation in the average hourly wage that would drive the estimate of the labor supply elasticity.

There is also an important econometric problem with this approach that is recognized by Camerer et al. There is an inherent division bias that can lead to a negative bias in the estimate of η. This bias arises because 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 instrumental variables approach. One potential problem with this approach is that there might be day-specific factors that affect both the wage and hours conditional on the wage of all drivers to some degree. To the extent that this is the case, their instrument will not purge their estimates of η of their inherent negative bias. Another potential problem with this instrumental variables approach is that calendar dates on which there is only one driver in the sample cannot be used in the analysis.

I present ordinary least squares (OLS) estimates of models like equation (10). However, my data do not have sufficient numbers of drivers on any particular date to replicate Camerer et al.'s instrumental variables analysis.

V. Data and Preliminary Statistics

The data necessary to carry out my analysis 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 details 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. After contacting more than 70 leasing companies, one was found that was still in business and was willing to provide trip sheets. 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 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 for the same driver 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 interday relationships in labor supply.

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 important, there are financial incentives working in favor of a complete listing of trips. In my informal interviews, I have asked drivers in New York about their trip sheet practices, and most told me that they are careful about filling out the sheets, some because of fines levied as a result of incomplete trip sheets. Apparently, taxicabs are stopped by New York City police officers or by Taxi and Limousine Commission inspectors, either randomly or for cause.[4] 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 performed several regularity checks to ensure that the trip sheets are internally consistent, and where they are not, I cleaned the data using a set of reasonable rules. These rules are outlined in detail in Appendix A.

I coded the starting and ending locations on the trip sheets into 11 categories. These are Downtown Manhattan (below Fourteenth Street), Midtown Manhattan (Fourteenth Street to Fifty-ninth Street), Uptown Manhattan (above Fifty-ninth 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 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 data on hourly rainfall at LaGuardia Airport.

A. *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. Appendix table B1

[4] Das Gupta (2002, 23) notes that "the rules governing drivers have become more elaborate and punitive" and that tickets are "zealously issued."

contains average 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 nine drivers and more than 20 shifts 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 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^2 from a regression of daily hours on a set of driver fixed effects is 0.181 with a residual root mean squared error (RMSE) of 2.30. Figure 1*a* 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.[5] 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 income is within-driver variation across days. The standard deviation of daily income is \$59.57. The R^2 from a regression of daily income on a set of driver fixed effects is 0.169 with a residual RMSE of \$55.27. A labor supply model in which drivers had fixed but potentially different targets would imply that more of the variation in income would be accounted for by driver fixed effects. Figure 1*b* contains a kernel density estimate of daily income.[6]

Column 8 of Appendix table B1 contains the daily average for each driver of his hourly wage rate (total income divided by working hours). These averages show less interdriver 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 since the R^2 from a regression of the daily wage on a set of driver fixed effects is 0.089 with a residual RMSE of \$4.35. Figure 1*c* contains a kernel density estimate of the shift average hourly wage.

B. Trip-Level Summary Statistics

Figure 2 contains kernel density estimates of the distributions of trip times, waiting times, and fares for the 13,464 trips in my sample. I have

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

[6] All kernel density estimates in this study use the Epanechnikov kernel. The bandwidths are listed in the figures.

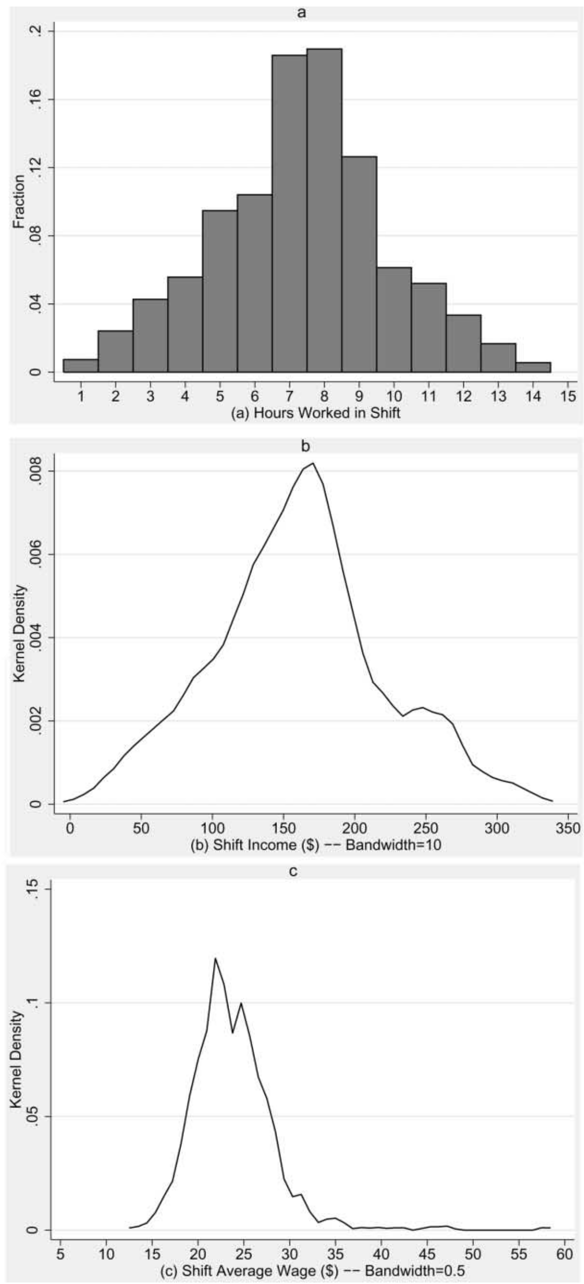

FIG. 1.—Distributions of hours, income, and average wage by shift: *a*, hours worked in a shift; *b*, shift income; *c*, shift average hourly wage.

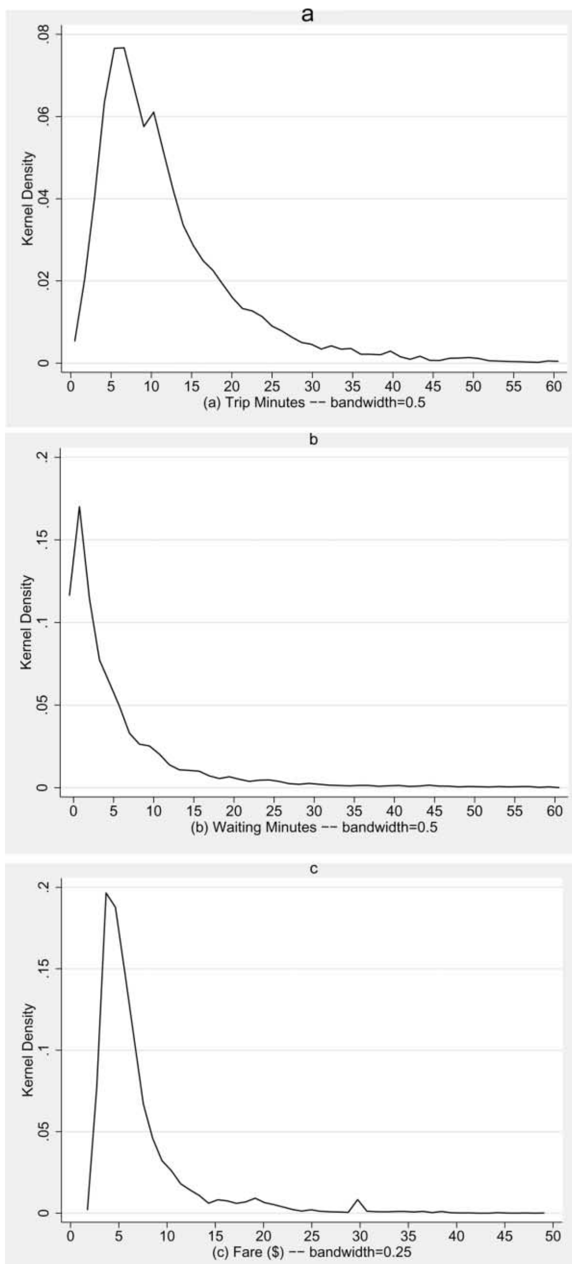

FIG. 2.—Kernel density estimates of *a*, trip times, *b*, waiting times, and *c*, fares

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 are essentially zero.[7]

Figure 2*a* contains a kernel density estimate of the distribution of trip times. Median time 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 percent at LaGuardia Airport, 0.8 percent at Kennedy Airport, 0.9 percent in Brooklyn, and 0.5 percent in Queens).

Figure 2*b* contains a kernel density estimate of the distribution of waiting times before each trip. Median waiting time is three minutes, and the mean waiting time is 5.9 minutes. There are 2,892 trips with zero waiting time. As I describe in Appendix A, I reclassified 316 long waiting times between fares as break times, and fares after these breaks are classified as having preceding waiting times of zero.

Figure 2*c* 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 fares (which exclude tips) are so small because most trips are intra-Manhattan. The average intra-Manhattan fare is $5.90, whereas 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.

C. 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.

Figure 3*a* contains the average hourly wage by clock hour for the driver-hours in my sample along with the average ± 1.96 standard errors. There is clearly variation over the day in the hourly wage, with the hourly

[7] There are 41 trips with times greater than 60 minutes and 26 waiting times greater than 60 minutes. There are eight fares greater than $50.

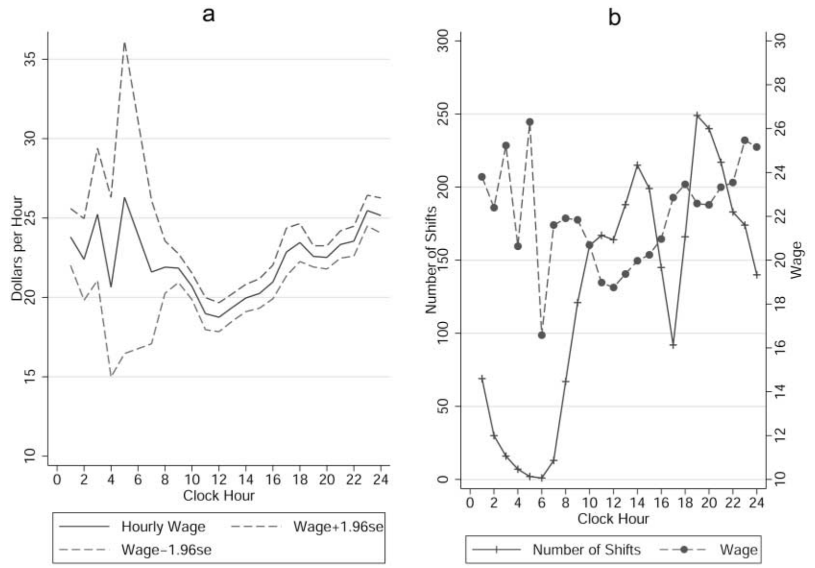

FIG. 3.—Hourly wage, by clock hour: *a*, hourly wage; *b*, number of shifts and wage

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 early morning hours is relatively large because of the smaller number of driver-hours in my sample during that part of the day.

Figure 3*b* overlays the plot of the average hourly wage by clock hour with a plot of the number of driver-hours in my sample by clock hour. While it is not the case that I have a random selection of cabs on the streets of Manhattan at any point during 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 6:00 a.m., after which the number of driver-hours increases through 1:00 p.m. The number of driver-hours drops sharply at 4:00 p.m. and 5:00 p.m., likely reflecting the change of shifts, before increasing to the daily maximum at 7:00 p.m. Subsequently, the number of driver-hours drops consistently through 6:00 a.m.

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 one at 6:00 a.m. to 249 at 7:00 p.m., whereas the average wage varies from \$18.75 at noon to \$25.47 at 11:00 p.m.[8] This pattern is

[8] The wage is even higher at \$26.30 at 5:00 a.m., but this is based on only two driver-hours.

TABLE 1
ANALYSIS OF VARIANCE OF HOURLY WAGE

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Driver identification		x		x	x		x	x
Hour of day			x	x	x		x	x
Day of week			x	x	x		x	x
Hour of day × day of week			x	x	x		x	x
Weather					x		x	x
Date						x	x	x
Date × driver identification								x
Degrees of freedom used	0	20	141	161	165	315	470	701
R^2	.00	.05	.15	.17	.17	.15	.28	.35
RMSE	6.67	6.51	6.31	6.25	6.24	6.49	6.17	6.12

NOTE.—The statistics in the table are based on linear regressions with dummy variables included for the indicated categories in each column. The sample used 3,025 hours for which all 60 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, and an indicator for maximum temperature greater than or equal to 80 degrees.

consistent with the number of drivers adjusting to daily fluctuations in the pattern of demand.[9] The lower wage at midday may reflect short lunch breaks not recorded as such by drivers.

Table 1 contains the analysis of variance results from a series of regressions of the hourly wage on a sequence of sets of variables. Column 1 refers to a regression with only a constant. The RMSE of this regression (the standard deviation of the wage) is 6.67. Controlling for driver fixed effects (col. 2) accounts for only 5 percent of the variation in the hourly wage, and the RMSE is reduced slightly to 6.51. Controlling for the hour of the day, the day of the week, and their interaction (col. 3) accounts for 15 percent of the variation in the hourly wage, and controlling driver fixed effects along with the hour of the day, the day of the week, and their interaction (col. 4) accounts for 17 percent of the hourly wage variation.

I controlled additionally for the weather (four variables: [1] hourly rainfall, [2] daily snowfall, [3] daily low temperature less than 30 degrees Fahrenheit, and [4] daily high temperature greater than or equal to 80 degrees Fahrenheit) in column 5. These variables do not improve the R^2 substantially, but the hourly wage is significantly related to the weather measures ($p = .033$). Specifically, the hourly wage is $1.04 lower

[9] Oettinger (1999) models the labor supply (participation) 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 include the labor supply of all vendors. The analogous participation 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
Autocorrelations of Hourly Wage

Period	(1)	(2)	Observations
0	1.0	1.0	3,025
1	.0687*	−.0369	2,232
2	.0984	.0241	1,652
3	.0625*	−.0072	1,312
4	.0261	−.0351	1,034
5	.0644	.0079	785
6	.1354*	.1297*	539
7	−.0082	−.0290	334

Note.—Col. 1 contains autocorrelations of hourly wages indexed by hour on shift. Col. 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 set of driver fixed effects and a set of hour-of-day fixed effects.
* Significantly different from zero at the .05 level.

on hot days ($p = .007$), \$0.49 lower on cold days ($p = .158$), and, surprisingly, \$0.94 lower for each 0.1 inch of hourly rain ($p = .103$).

Controlling only for the 316 distinct calendar dates in my sample (col. 6) accounts for 15 percentage points of the hourly wage variation. Taken together (col. 7), the calendar date and other variables account for 28 percent of the variation in the hourly wage with an RMSE of 6.17. Overall, more than 70 percent of the variation in the hourly wage is unsystematic within-driver, within-day, interhour variation.

Column 8 of table 1 additionally controls for the interaction of driver identification 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 interactions, 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 identification 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 1 of table 2 contains the autocorrelations of hourly wages for the first eight hours of shifts. The autocorrelations are surprisingly small, suggesting that there is not important transitory daily variation in the hourly wage. And it may be that part of this autocorrelation is related to time-of-day and driver fixed effects, which do not represent transitory daily variation in the hourly wage. I investigate this by examining the autocorrelations of residuals from a regression of the hourly wage on hour-of-day and driver fixed effects These autocorrelations are contained in column 2 of table 2, and, with the exception of an anomalous significant, positive sixth-order autocorrelation, they are very small, neg-

TABLE 3
LABOR SUPPLY FUNCTION ESTIMATES: OLS REGRESSION OF LOG HOURS

Variable	(1)	(2)	(3)
Constant	4.012	3.924	3.778
	(.349)	(.379)	(.381)
Log(wage)	−.688	−.685	−.637
	(.111)	(.114)	(.115)
Day shift	…	.011	.134
		(.040)	(.062)
Minimum temperature	…	.126	.024
< 30		(.053)	(.058)
Maximum temperature	…	.041	.055
≥ 80		(.055)	(.064)
Rainfall	…	−.022	−.054
		(.073)	(.071)
Snowfall	…	−.096	−.093
		(.036)	(.035)
Driver effects	no	no	yes
Day-of-week effects	no	yes	yes
R^2	.063	.098	.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.

ative, and not significantly different from zero. This suggests little or no role for transitory daily shocks to the hourly wage.

While there is significant day-to-day variation in the hourly wage, the results in tables 1 and 2 suggest that most variation in the wage appears to be nonforecastable 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.

VI. 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 by Chou (2000) and formulated in equation (10). I then implement the probit model of the probability of stopping based on the latent variable defined in equation (9).

A. *Estimation of the Log Hours Function*

Table 3 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 based on 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 based on specifications that include additional regressors. The specification in column 2 includes controls for shift, day of the week, and four measures of the daily weather. The estimated wage elasticity is virtually unchanged in this specification. The results indicate no difference in hours worked between the day shift and the night shift. The estimates of the coefficients of the weather measures suggest that shifts are longer on cold days and shorter on snowy days. Rainfall is not statistically significantly related to shift length. The estimates in column 3 additionally control for driver fixed effects. There are significant differences across drivers in hours worked ($p < .0005$), but accounting for these differences does not have a substantial effect on the estimated negative wage elasticity. Interestingly, when driver fixed effects are accounted for, the day shift indicator is significantly negative, suggesting that, when a specific driver moves from the night shift to the day shift, hours increase by about 13 percent and vice versa. Finally, controlling for the weather along with driver fixed effects in column 4 suggests that, while hours are unrelated to temperature extremes or rainfall, drivers do work fewer hours when it has snowed.

The key consistent finding is that there appears to be a substantial negative elasticity of labor supply, as is found by Camerer et al.[10] However, 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 and is not highly autocorrelated hour to hour, it is not likely that the wage measured this way can be considered parametric to the labor supply decision. Thus I do not consider estimates of a labor supply elasticity derived using this daily regression approach to be reliable, even if a convincing instrument for the wage were available.

B. Estimation of the Probit Optimal Stopping Model

The probability that a driver ends his shift increases sharply as hours and income accumulate. Panel A of table 4 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 25 percent 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 in the table. Panel B of table 4 contains the simple empirical hazard of a driver

[10] The similarity of our findings is important because it implies that any differences in results that I find using other empirical models are due to differences in analytic approach rather than to differences in the data used.

TABLE 4
EMPIRICAL HAZARD OF STOPPING BY HOURS AND BY INCOME

A		B	
Hour	Hazard	Income ($)	Hazard
≤ 2	.0050	< 25	.0016
3–5	.0244	25–49	.0078
6	.0517	50–74	.0149
7	.0992	75–99	.0228
8	.1467	100–124	.0338
9	.1147	125–149	.0596
10	.1609	150–174	.1111
11	.2450	175–199	.1477
≥ 12	.2636	200–224	.1283
		≥ 225	.1977

NOTE.—Based on the sample of 13,461 trips in 584 shifts for 21 drivers. These figures 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.

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 IV, I outlined a 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, weather, and fixed effects for driver, calendar date, hour of the day, and day of the week. The latent variable defined in equation (9) forms the basis for estimation of this model. Table 5 contains estimates of this model in which income and hours are constrained to enter linearly.[11] The normalized probit estimate is $\hat{\boldsymbol{\beta}} \cdot \phi(\boldsymbol{X}^*\boldsymbol{\beta})$, where $\phi(\cdot)$ is the standard normal density. Given that the probability that a shift will end 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.[12] 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 eight total hours (5.5 trip hours, 2.5 waiting hours, 0.5 break hour) on a dry day with moderate temperatures in midtown Manhattan at 2:00 p.m.

The estimates in column 1 of table 5 contain only the hours measures

[11] This constraint is relaxed in table 6.

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

TABLE 5
Hazard of Stopping after Trip: Normalized Probit Estimates

Variable	X^*	(1)	(2)	(3)	(4)	(5)
Total hours	8.0	.013	.037	.011	.010	.010
		(.009)	(.012)	(.005)	(.005)	(.005)
Waiting hours	2.5	.010	−.005	.001	.004	.004
		(.010)	(.012)	(.006)	(.006)	(.005)
Break hours	.5	.006	−.015	−.003	−.001	−.002
		(.008)	(.011)	(.005)	(.005)	(.005)
Shift income ÷ 100	1.5	.053	.036	.014	.016	.011
		(.022)	(.030)	(.015)	(.016)	(.015)
Minimum temperature < 30	.0	…	…	…	.002	.002
					(.007)	(.006)
Maximum temperature ≥ 80	.0	…	…	…	−.016	−.016
					(.007)	(.007)
Hourly rain	.0	…	…	…	.005	.015
					(.124)	(.113)
Daily snow	.0	…	…	…	.007	.007
					(.006)	(.005)
Downtown	.0	…	…	…	…	.000
						(.006)
Uptown	.0	…	…	…	…	.000
						(.005)
Bronx	.0	…	…	…	…	.075
						(.057)
Queens	.0	…	…	…	…	.044
						(.038)
Brooklyn	.0	…	…	…	…	.078
						(.036)
Kennedy Airport	.0	…	…	…	…	.056
						(.031)
LaGuardia Airport	.0	…	…	…	…	.061
						(.029)
Other	.0	…	…	…	…	.134
						(.071)
Driver (21)		no	yes	yes	yes	yes
Day of week (7)		no	no	yes	yes	yes
Hour of day (19)	2:00 p.m.	no	no	yes	yes	yes
Log likelihood		−2,039.2	−1,965.0	−1,789.5	−1,784.7	−1,767.6

Note.—The sample includes 13,461 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 5:00 a.m. to 10:00 a.m. have a common fixed effect. The evaluation point is after 5.5 driving hours, 2.5 waiting hours, and 0.5 break hour in a dry hour on a day with moderate temperatures in midtown Manhattan at 2:00 p.m. Robust standard errors accounting for clustering by shift are reported in parentheses.

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, on the basis of the improvement in the log likelihood, there are significant differences across drivers in their probability of quitting. Interestingly, after one accounts for interdriver differences, the probability of quitting 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 hours on trips.

The estimates in column 3 additionally control for hour of day and day of 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 fact that shift start and end times tend to be concentrated in a relatively narrow range of clock hours. 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 points 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.[13]

Finally, the estimates in column 5 include controls for geographic location at the end of the trip. The omitted location is midtown Manhattan (between Fourteenth Street and Fifty-ninth 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. The reason may be that drivers are closer to home in the outlying boroughs. Controlling for location has an insubstantial effect on the coefficients of the key hours and income measures.

A potential problem with the estimates in table 5 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

[13] The results reported in the discussion of table 1 suggest that the hourly wage is lower in rainy weather. This is not consistent with an increase in demand. One possible explanation might be that demand does increase when it is raining, but there is also a rain-induced slowdown in traffic that increases trip times for a given distance traveled and fare. Simple tabulation of the trip-level data shows that dollars earned per minute with a fare in the cab is statistically significantly lower by about $0.04 when it is raining.

sharp increase in the hazard of a shift ending as the shift progresses suggests that a more flexible specification is appropriate. Table 6 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 10-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 5 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 probability of quitting is calculated for a driver in his eighth hour on the shift having earnings between $150 and $174 on a dry day with moderate temperatures in midtown Manhattan at 2:00 p.m.

The estimates in column 1 of table 6 are based on a model that includes only the unconstrained hours effects. They 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 column 2 are based on a model that includes only the unconstrained income effects. There is 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, is a strong predictor 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 6, 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 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 = .008$).

When hours are controlled for, income may be picking up intertem-

TABLE 6
Hazard of Stopping after Trip: Normalized Probit Estimates

Variable	(1)	(2)	(3)	(4)
Hour:				
≤ 2	−.142	…	−.137	−.042
	(.014)		(.017)	(.012)
3–5	−.122	…	−.108	−.028
	(.014)		(.017)	(.010)
6	−.095	…	−.084	−.026
	(.015)		(.018)	(.009)
7	−.048	…	−.041	−.012
	(.016)		(.018)	(.008)
9	−.032	…	−.034	−.006
	(.021)		(.022)	(.010)
10	.014	…	.011	.031
	(.025)		(.033)	(.020)
11	.098	…	.086	.085
	(.038)		(.051)	(.036)
≥ 12	.117	…	.109	.119
	(.063)		(.075)	(.060)
Income ($):				
< 5	…	−.109	−.121	−.036
		(.010)	(.022)	(.014)
25–49	…	−.103	−.067	.005
		(.010)	(.027)	(.022)
50–74	…	−.096	−.063	−.002
		(.010)	(.021)	(.014)
75–99	…	−.088	−.065	−.010
		(.010)	(.020)	(.011)
100–124	…	−.077	−.053	−.009
		(.011)	(.018)	(.009)
125–149	…	−.052	−.035	−.007
		(.011)	(.017)	(.008)
175–199	…	.037	.014	.011
		(.018)	(.021)	(.010)
200–224	…	.017	−.022	.006
		(.020)	(.025)	(.013)
≥ 225	…	.087	.009	.015
		(.019)	(.032)	(.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
Weather (4)	no	no	no	yes
p-value:				
Hours = 0	.000	…	.000	.000
Income = 0	…	.000	.008	.281
Log likelihood	−2,028.9	−2,058.7	−2,016.5	−1,753.1

Note.—The sample includes 13,461 trips in 584 shifts for 21 drivers. Probit estimates are normalized to reflect the marginal effect at $\boldsymbol{X}^*$ of $\boldsymbol{X}$ on the probability of stopping. The normalized probit estimate is $\hat{\boldsymbol{\beta}} \cdot \phi(\boldsymbol{X}^*\boldsymbol{\beta})$, where $\phi(\cdot)$ is the standard normal density. The values of $\boldsymbol{X}^*$ chosen for the fixed effects are equally weighted for each day of the week and driver. The hours from 5:00 a.m. to 10:00 a.m. have a common fixed effect. The evaluation point is at eight total hours with income of $150–$174 in a dry hour on a day with moderate temperatures in midtown Manhattan at 2:00 p.m. Robust standard errors accounting for clustering by shift are reported in parentheses.

poral and interdriver differences in the value of continuing to drive. In order to address this possibility, the estimates in column 4 of table 6 are derived from a model that includes additional controls for driver, day of the week, hour of the day, weather, and geographic location. As I showed in table 5, these variables are all important determinants of the probability of stopping. Interestingly, hours worked remains 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–$174 category. All other estimates on the income categories are smaller along with not being significantly different from zero. The hypothesis that all the income coefficients are zero cannot be rejected at conventional levels ($p = .28$).

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

C. *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 the six drivers for whom I have more than 40 trip sheets.

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 difficult the interpretation of estimates of the effect of cumulative

TABLE 7
DRIVER-SPECIFIC HAZARD OF STOPPING AFTER TRIP: NORMALIZED PROBIT ESTIMATES

	Driver					
Variable	4	10	16	18	20	21
Hours	.073	.056	.043	.010	.195	.198
	(.060)	(.047)	(.015)	(.007)	(.045)	(.030)
Income÷100	.178	.039	.064	.048	−.160	−.002
	(.167)	(.059)	(.041)	(.020)	(.123)	(.150)
Number of shifts	40	45	70	72	46	46
Number of trips	884	912	1,754	2,023	1,125	882
Log likelihood	−124.1	−116.0	−221.1	−260.6	−123.4	−116.9

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 at eight hours with income of $150 in a dry hour of moderate temperature in Manhattan during a night shift. Robust standard errors accounting for clustering by shift are reported in parentheses. All specifications include controls for day of week, weather (high and low temperatures, rainfall, and snowfall), location (in or out of Manhattan), and shift (day or night).

hours and cumulative income on the probability of stopping. 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 not withstanding, table 7 contains estimates of probit models of the probability of stopping after a trip for each of the six drivers. The set of control variables is pared down because of a 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 the four weather variables (hourly rainfall, snowfall, low temperature, and high temperature). Fixed effects for day of the week are included in all specifications. Cumulative hours and cumulative income are controlled for linearly.

The results provide support for the importance of cumulative hours as the primary factor in determining the probability of stopping. The probability of stopping is significantly positively related to hours worked at less than the .05 level for three of the six drivers (drivers 16, 20, and 21). For the remaining three drivers (drivers 4, 10, and 18), the probability of stopping is significantly positively related to hours worked at the .11, .12, and .08 levels, respectively. For only one of the six drivers (driver 18) is the probability of stopping significantly positively related to cumulative income at conventional levels, and for two of the six drivers (drivers 20 and 21), the estimated marginal effect of income is insignificantly negative.

Overall, the stopping rules for these six individual drivers are consistent with the neoclassical model of labor supply and are not consistent with the target earnings model.

D. Reanalysis of the Camerer et al. TRIP Data

Given the sharp contrast between my conclusion that cabdriver work time on a given day is consistent with a neoclassical intertemporal labor supply model and inconsistent with a target earnings model and the opposite conclusion of Camerer et al., it is important to understand the basis of the difference. Clearly, the statistical and conceptual models used are very different, and these differences in approach likely account for the difference in findings. However, the samples of trip sheets used also differ, and it is possible that the difference in data could account for the difference in results.

I have already applied the Camerer et al. regression approach to my data by estimating log hours regressions by OLS (table 3), and the results show the same significantly negative labor supply elasticity found in the earlier study. The remaining question is whether, when applying my statistical model to the Camerer et al. data, I find results consistent with the estimates reported in tables 5 and 6.

Camerer et al. supplied me with their TRIP data, a sample of 70 trip sheets for 13 drivers that can be used to estimate my probit stopping model. They report substantial first- and second-order autocorrelations of the hourly wage of approximately 0.5, which is evidence consistent with there being significant daily fluctuations in the hourly wage. This stands in sharp contrast to my calculation of the first- and second-order autocorrelations of the hourly wage in my data of 0.07 and 0.10, respectively (table 2), which are not consistent with there being significant daily fluctuations in the hourly wage. I recalculated the hourly wage in the TRIP data using the same method I describe in Section V.*C*, and I computed the autocorrelations. My calculations yield autocorrelations that differ substantially from those presented by Camerer et al. for the same data. I find first- and second-order autocorrelations of the hourly wage of 0.17 and 0.22, respectively.[14] While these autocorrelations are somewhat larger than those found in my data, they are not nearly as large as those reported by Camerer et al., and they do not provide support for important daily fluctuations in the hourly wage.

Next, I estimate the probit stopping model using the Camerer et al. TRIP data. Of the eight drivers in the TRIP sample, five are observed for only a single shift. Since I include driver fixed effects in the speci-

[14] As in my data, autocorrelations of hourly wage residuals from the Camerer et al. data are insignificantly negative.

TABLE 8
HAZARD OF STOPPING AFTER TRIP: CAMERER ET AL. TRIP DATA: SIXTH OR LATER HOUR
NORMALIZED PROBIT ESTIMATES

Variable	All (1)	<Median Experience (2)	≥Median Experience (3)	Low Experience (4)	Medium Experience (5)	High Experience (6)
Total hours	.101	.003	.102	.001	.307	.098
	(.030)	(.004)	(.044)	(.004)	(.112)	(.049)
Shift income÷100	.055	.016	−.136	.063	−.515	−.027
	(.119)	(.018)	(.101)	(.122)	(.342)	(.116)
Driver fixed effects	yes	yes	yes	yes	yes	yes
Number of drivers	8	3	5	1	4	3
Number of shifts	65	26	39	10	33	22
Number of trips	554	239	315	60	319	175
Log likelihood	−144.2	−55.7	−83.4	−20.0	−70.9	−40.6

NOTE.—Probit estimates in the bottom rows are normalized to reflect the marginal effect at X^* (eight hours of work and $125 in earnings) of a change in 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 reported standard errors (in parentheses) are computed by applying the delta method to robust standard errors of β that account for clustering by driver/shift.

fication, I drop these five drivers from my reanalysis. With only 65 shifts for eight drivers, there is not substantial variation in shift starting time. All shifts started between 1:00 p.m. and 8:00 p.m., and 56 of the 65 shifts started between 4:00 p.m. and 7:00 p.m. This makes it difficult to identify hours worked on a shift separately from clock hour. None of the 65 trips ended until more than six hours had been worked. As a result of these factors, I estimate the stopping model including only the 554 trips that began after working six hours, and I measure hours and income linearly without any control for hour of the day or day of the week. In addition to the linear measures of hours worked and shift income, the model includes only driver fixed effects.

Column 1 of table 8 contains normalized probit estimates of the stopping model using the Camerer et al. TRIP data. There is a significant positive relationship between the probability of stopping after a trip and cumulative hours. The point estimate suggests that each additional hour of work increases the likelihood of stopping after a trip by 10 percentage points. There is no significant relationship between cumulative shift income and the probability of stopping. This pattern of evidence is consistent with estimates of the stopping model based on my data and with a standard neoclassical intertemporal labor supply model. There do not appear to be strong income effects in the TRIP data.

The role of experience in the Camerer et al. TRIP data.—Camerer et al. suggest that their finding of a negative wage elasticity of labor supply might be the result of income targeting by less experienced drivers, with more experienced drivers behaving in a manner consistent with the neoclassical model. They investigate this by breaking their sample by experience level using the hack license number assigned consecutively by the New York City Taxi and Limousine Commission. They break the sample at the median hack number, ranging in their sample from 50,893 to 610,069, with higher numbers indicating less experience. When Camerer et al. reestimate their OLS model separately for the high- and low-experience samples, they find a significant negative labor supply elasticity (−0.841) for the low-experience group and a positive labor supply elasticity for the high-experience group (0.613).

I reestimated the probit model using the TRIP data separately for each of the two experience groups defined by Camerer et al. The results of this estimation are in columns 2 and 3 of table 8. They show no relationship between the probability of stopping and shift income for either experience group. Interestingly, there is a significant relationship between hours worked and the probability of stopping for the more experienced group but not for the less experienced group. This suggests that there might be some experience-related learning about optimal labor supply behavior.

Using the median hack number to break the sample may not yield a sharp distinction in experience given that there is not a sharp break in the sample at the median (450,079). Fully nine of the 13 drivers have hack numbers in the 400,000 range. There are three more experienced drivers with hack numbers less than 400,000 (50,893, 96,013, and 318,867) and one less experienced driver with a hack number of 610,069. All five of the drivers with single shifts in the sample have hack numbers in the 400,000 range. The eight drivers with multiple shifts can be more naturally divided into three groups: (1) low experience (≥ 500,000), (2) medium experience (400,000–499,999), and (3) high experience (< 400,000).

I reestimated the probit model using the TRIP data separately for each of these three experience groups. The results of this estimation are in columns 4–6 of table 8. The results are interesting. There is a significant positive relationship between hours worked and the probability of stopping after a trip for both the medium- and high-experience groups. There are negative and insignificant estimates of the relationship between cumulative income and the probability of stopping for both of these groups of drivers. For the least experienced driver (alone in the low-experience group), there is no significant relationship of either hours or income with the probability of stopping.

The overall pattern of evidence by experience level suggests indirectly that Camerer et al.'s experience result is due to the single very low experienced driver (hack number of 610,069, more than 130,000 higher than the next more experienced driver). On the basis of a single driver, I am reluctant to conclude that taxi drivers learn from experience to behave optimally whereas less experienced drivers are more likely to be target earners.

VII. Final Remarks: Tomorrow Is Another Day

The results of the probit stopping models I estimate, whether applied to my taxi data or to the Camerer et al. TRIP data, show consistently that the probability of stopping daily work after a particular trip is strongly related to hours worked to that point and not significantly related to cumulative income earned. This pattern is consistent with a conventional neoclassical intertemporal labor supply model with weak or nonexistent income effects due to transitory changes to the marginal daily wage. There is no evidence either for important income effects or for target earnings behavior. These findings are in sharp contrast to the finding of negative daily wage elasticities of labor supply reported by Camerer et al. (1997).

While I estimated a model of labor supply, I did not estimate a wage elasticity of labor supply. In order to estimate such an elasticity, there would need to be an exogenous permanent or transitory shift in the general level of earnings opportunities. Some natural experiment (such as Oettinger [1999]) or randomized experiment (such as Fehr and Goette [2002]) that yielded changes to daily earnings opportunities would provide the right sort of variation.

A complete analysis of the labor supply response to changes in earnings opportunities needs to examine both daily hours responses and responses in the number of days worked (the participation margin). Owing to data limitations, I cannot investigate the daily labor supply of taxi drivers on the participation margin, and my analysis is restricted to the daily hours margin.[15] Despite this limitation, the fact that I find no evidence of a daily income effect on hours worked leads me to conclude that the labor supply behavior of taxi drivers can be characterized by a conventional neoclassical intertemporal labor supply model.

Appendix A

Procedures for Cleaning Data

I perform several regularity checks to ensure that the trip 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 (have a 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 nonmissing time and fare data. In the small number of cases in

[15] Oettinger (1999) examines labor supply response on the participation margin in a setting in which daily hours are fixed. Fehr and Goette (2002) examine labor supply response on both margins.

which these rules are not satisfied or time or fare data are missing (about 1 percent of trips), I use the following procedures.

1. I 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 a regression of trip length on the fare and its square ($R^2 = 0.68$). 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 in which 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.
2. I impute new trip times for 52 cases in which one or both of the starting and ending times are missing. Most of these times 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 in which one or the other of the times is 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 ($R^2 = 0.44$).
3. I adjust times where they are out of sequence. There are 83 cases in which a trip apparently started before the previous trip ended. The mean overlap was 2.7 minutes, and the maximum overlap was nine minutes. I adjust these times 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.
4. I impute fares for 45 cases in which 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 ($R^2 = 0.83$). In the seven cases in which location is missing, I predict on the basis of a model using only trip length ($R^2 = 0.64$).
5. I classify as breaks (nonwork time during the day) certain long waiting 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) nine periods of at least 120 minutes that ended at LaGuardia or Kennedy Airport.
6. I 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 in which the necessary information for imputation by these rules is not available, I examined the individual trip 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 the allocated locations are in Manhattan.

Appendix B

TABLE B1
SHIFT-LEVEL SUMMARY STATISTICS, BY DRIVER

Driver	Number of Shifts (1)	Average Trips (2)	Working Hours (3)	Driving Hours (4)	Waiting Hours (5)	Break Hours (6)	Total Income (7)	Average Wage (8)
1	39	23.56	6.85	4.32	2.53	.90	157.58	23.16
2	14	12.29	3.89	2.78	1.11	2.41	97.10	25.11
3	6	23.67	6.66	4.61	2.05	.74	163.42	25.56
4	40	22.10	6.28	4.52	1.76	.39	147.51	23.89
5	23	16.52	6.46	3.98	2.48	2.11	144.96	23.65
6	6	29.33	8.62	6.48	2.14	2.42	205.00	24.02
7	24	22.29	6.47	4.42	2.05	.74	160.71	25.59
8	37	25.32	7.78	5.13	2.64	.86	172.44	22.54
9	19	25.58	7.17	5.47	1.70	.54	162.02	23.23
10	45	20.27	6.35	3.90	2.45	1.65	133.19	21.46
11	6	24.00	7.15	5.22	1.93	.71	182.81	26.57
12	13	19.46	6.15	4.03	2.13	.55	157.95	25.78
13	10	26.50	7.03	4.71	2.31	.53	154.19	22.55
14	17	21.29	7.06	4.49	2.57	.64	165.84	23.59
15	8	39.62	10.82	7.64	3.17	.19	228.26	21.12
16	70	25.06	6.84	4.56	2.28	.93	172.01	25.62
17	10	23.00	5.88	3.71	2.17	.54	144.57	23.354
18	72	28.10	8.53	5.84	2.69	.60	203.05	24.01
19	33	17.06	6.91	4.63	2.29	.97	163.51	23.91
20	46	24.46	7.10	4.80	2.30	.67	156.23	22.00
21	46	19.17	5.32	3.66	1.66	.24	128.97	24.72
All	584	23.05	6.90	4.64	2.26	.86	161.33	23.79

References

Altonji, Joseph G., and Christina H. Paxson. 1988. "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs." *J. Labor Econ.* 6 (April): 254–76.

Blundell, Richard, and Thomas E. MaCurdy. 1999. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, vol. 3A, edited by Orley Ashenfelter and David Card. Amsterdam: North-Holland.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Q.J.E.* 112 (May): 407–41.

Chou, Yuan K. 2000. "Testing Alternative Models of Labor Supply: Evidence from Taxi Drivers in Singapore." Research Paper no. 768, Dept. Econ., Univ. Melbourne.

Das Gupta, Monisha. 2002. "Of Hardship and Hostility: The Impact of 9/11 on Taxi Drivers in New York City." Manuscript, Ethnic Studies Dept., Univ. Hawaii at Manoa.

Dickens, William T., and Shelley J. Lundberg. 1993. "Hours Restrictions and Labor Supply." *Internat. Econ. Rev.* 34 (February): 169–92.

Fehr, Ernst, and Lorenz Goette. 2002. "Do Workers Work More if Wages Are

High? Evidence from a Randomized Field Experiment." Working Paper no. 125, Inst. Empirical Res. Econ., Univ. Zurich.

Ham, John C. 1982. "Estimation of a Labour Supply Model with Censoring Due to Unemployment and Underemployment." *Rev. Econ. Studies* 49 (July): 335–54.

Kahn, Shulamit, and Kevin Lang. 1991. "The Effect of Hours Constraints on Labor Supply Estimates." *Rev. Econ. and Statis.* 73 (November): 605–11.

Killingsworth, Mark R., and James J. Heckman. 1986. "Female Labor Supply: A Survey." In *Handbook of Labor Economics*, vol. 1, edited by Orley Ashenfelter and Richard Layard. Amsterdam: North-Holland.

MaCurdy, Thomas E. 1981. "An Empirical Model of Labor Supply in a Life-Cycle Setting." *J.P.E.* 89 (December): 1059–85.

Oettinger, Gerald S. 1999. "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors." *J.P.E.* 107 (April): 360–92.

Pencavel, John H. 1986. "Labor Supply of Men: A Survey." In *Handbook of Labor Economics*, vol. 1, edited by Orley Ashenfelter and Richard Layard. Amsterdam: North-Holland.