An Empirical Analysis of the Daily Labor Supply of Stadium Vendors

Gerald S. Oettinger

University of Texas at Austin

This paper analyzes the daily labor supply behavior of food and beverage vendors at a single stadium over an entire baseball season. This labor market is attractive for the study of labor supply both because the vendors unilaterally decide whether to participate on each game date and because changes in product demand conditions across days are large and highly predictable, generating exogenous game-to-game variation in the vendor "wage." I exploit the observable shifts in product demand conditions across games to estimate the labor supply (participation) elasticity of stadium vendors. Estimates that recognize that demand conditions and vendor labor supply decisions simultaneously determine the vendor wage always find substantial labor supply elasticities, typically in the .55–.65 range. In contrast, estimates that ignore the endogeneity of the vendor wage yield severely downward-biased labor supply elasticities. These results highlight the importance of using demand shift instruments to identify labor supply elasticities in specific labor markets.

I. Introduction

How responsive is work effort to transitory wage changes? Most estimates of life cycle models of labor supply using individual panel data (e.g., MaCurdy 1981; Browning, Deaton, and Irish 1985; Altonji

I am very grateful to Chris Ciora for providing the data and answering questions. I thank Steve Bronars for very useful discussions and an anonymous referee for especially helpful comments on an earlier draft. I also thank Henry Farber, Daniel Hamermesh, Alan Krueger, Kenneth McLaughlin, George Neumann, Cecilia Rouse, and seminar participants at Princeton, City University of New York Graduate School, and the 1998 Society of Labor Economists meetings for comments. All remaining errors are my own.

[*Journal of Political Economy*, 1999, vol. 107, no. 2]

1986) imply that this elasticity is very small. However, these studies relate annual changes in hours worked to annual changes in average hourly earnings, and it seems doubtful that the measured annual wage changes are either fully anticipated or purely transitory. As a result, the observed wage change may be correlated with an unobserved (but possibly large) change in expected lifetime wealth, and therefore the estimate of the intertemporal substitution elasticity (as well as other labor supply elasticities of interest) is likely to be biased downward.[1]

Recently, Camerer et al. (1997) analyzed the daily labor supply behavior of New York City taxicab drivers. They argue that this is a promising group to study because the daily demand for taxi services is subject to large but transitory shocks due to weather conditions, conventions, day of the week effects, and so forth. As a result, there is exogenous variation in the cabdriver "wage" across days, but the very temporary nature of these wage changes implies that any wealth effects should be negligible. In sharp contrast to the prediction from dynamic labor supply models, however, Camerer et al. find significantly *negative* wage elasticities of hours worked; the cabdrivers in their samples work fewer hours on high-wage days. Inclusion of individual fixed effects and estimation by instrumental variables methods do not alter this result. Camerer et al. argue that their findings suggest that cabdrivers have very short (1-day) horizons and fixed daily income targets. They observe more generally, however, that their results raise questions about the empirical relevance of the idea that workers substitute work effort toward times when the return to work is high.

While the motivation for studying the effects of daily wage changes on labor supply is compelling, the source of the daily variation in cabdriver wages is not entirely clear. Although Camerer et al. posit that shifts in demand for taxi services are the driving force, they have no data on obvious demand shifters that allow them to check this assumption. Likewise, they have no aggregate quantity data, which could shed light on the source of the wage movements across days. But if the observed daily wage fluctuations result partly from unobserved shifts in the labor supply curves of cabdrivers, then the estimated labor supply elasticities reported by Camerer et al. are inconsistent.[2] Clearly, it would be desirable to have data on exogenous

[1] MaCurdy (1985) and Card (1990) develop the argument in detail. Card also surveys the literature on intertemporal labor supply.

[2] Although Camerer et al. directly control for weather conditions and the day and time of the driver's shift, some potentially important supply shifters—e.g., daily demand shocks in industries in which many cabdrivers hold second jobs—are not measured.

demand shifters that could be used as instruments for the observed daily wage.[3]

The present paper follows in the spirit of Camerer et al. (1997) by analyzing the daily labor supply behavior of a different group of workers: stadium vendors at major-league baseball games. Like those authors, I seek to measure how labor supply responds to day-to-day variation in the wage. To address this question, I obtained complete participation and earnings data for every vendor at every game at a single stadium during the 1996 baseball season. As shown below, the wage varied substantially across games. But, as with cabdrivers, this wage variation could result from either demand shifts or supply shifts. However, in contrast to cabdrivers, the key shifter of demand for vendor services—game attendance—and a number of good ex ante predictors of attendance are readily observable. These demand shift variables can serve as instruments for the wage and allow credible estimates of the labor supply elasticity to be obtained. In its focus on a labor market subject to large and observable demand shifts, this paper resembles Carrington's (1996) analysis of the Alaskan labor market during the era of construction of the Trans-Alaska Pipeline System. Unlike Carrington, though, I have panel data on individuals, which allow both individual-level and aggregate analyses of vendor labor supply behavior to be undertaken.

Briefly summarizing the empirical results, I always find that the wage elasticity of labor supply (participation) for stadium vendors is positive and substantial. Typical elasticity estimates are in the .55–.65 range, although the results vary a bit depending on the exact model specification and on whether the estimation uses individual-level panel data or time-series data on aggregate vendor participation and average game earnings. The aggregate analyses additionally reveal that treating the daily "wage" as exogenous in the labor supply equation (i.e., assuming that demand shifts are the sole source of unexplained game-to-game wage variation) causes a large downward bias in the estimated labor supply elasticity. Thus one must instrument for the vendor wage using observable demand shifters to consistently estimate the vendor labor supply elasticity. This last finding suggests that estimates of labor supply elasticities for specific labor markets, where unobserved individual labor supply shocks plausibly have an important common component, must be interpreted very cautiously unless demand shifters are used to instrument for the wage.

[3] Camerer et al. do present instrumental variable estimates, but their instruments (summary statistics of the daily wages earned by *all* cabdrivers in their sample) address only the problem of measurement error in individual cabdriver wages and not the conceptually distinct problem of endogeneity of the wage because of shifts in labor supply.

The remainder of the paper consists of five parts. Section II describes the relevant features of the labor market for vendors. Section III develops a theoretical model of the labor supply behavior of vendors. Section IV outlines the econometric model and the estimation strategy. Section V presents the empirical analysis. Finally, Section VI puts the results in perspective and offers some concluding observations.

II. Background Information on the Vendor Labor Market

This paper analyzes the labor supply of stadium vendors using the complete set of participation and earnings histories of all the vendors who worked at any of the games played at a single stadium during the 1996 major-league baseball season. In this section, those aspects of the vendor labor market that are relevant either for modeling vendor behavior or for analyzing data on vendor earnings and labor supply are described briefly.

During the 1996 season, the food and beverage vending business (i.e., the sale of food and beverage products by mobile vendors circulating through the stadium) was operated by a firm that hereafter is referred to as "the vending subcontractor." The vending subcontractor, which took over the vending business at the stadium in 1995, paid the primary concessions contractor a share of the vendors' sales revenue in return for the exclusive right to operate the vending business.[4] The vending subcontractor in turn hired vendors to walk through the stands and sell food and beverage products at the games. The stadium vendors sold six different products: beer, cotton candy, lemon ice, peanuts, popcorn, and soda. The selection of products sold by the vendors and their prices (which include sales tax) were decided by the primary concessions contractor prior to the season and did not change during the season.

The vending subcontractor paid its vendors a straight commission on their dollar sales (net of sales tax) at each game during the 1996 season. Both the decision to adopt a pure piece-rate compensation scheme and the choice of commission rate levels were made solely by the vending subcontractor. At the request of the vending subcontractor, I do not disclose the exact schedule of vendor commission rates or the share of vendor sales revenues that the subcontractor paid to the primary concessions contractor. However, commission rates were lower for vendors with less seniority at the start of the 1996

[4] The primary concessions contractor continued to operate the fixed concession stands located in the stadium.

season and for vendors who sold beer. Specifically, with product held constant, the commission rate for vendors in the highest seniority category exceeded the commission rate for vendors in the lowest seniority category by .02. With date of hire held constant, the commission rate for beer sales was .05 lower than the commission rate for sales of other products. This pattern of commission rates prevailed throughout the 1996 season. Thus vendor commission rates were not adjusted in response to vendor sales performance during the 1996 season.

By the start of the 1996 baseball season, the vending subcontractor had recruited an extensive pool of vendors. The vendors are independent contractors, not employees, and therefore are not covered by minimum-wage legislation and are not subject to employee payroll taxes. More important for present purposes, each vendor also decides freely whether or not to work at any given game. To help ensure that "enough" vendors show up at each game, the vending subcontractor asks vendors to sign up in advance for the games they intend to work on a publicly posted work schedule. If the schedule suggests that vendor turnout for a particular game will be "low," given anticipated attendance, the vending subcontractor will often contact prospective vendors who have not signed up and try to enlist their services. The important point, however, is that the schedule is not binding in any way; unscheduled vendors who show up to work a game are never turned away and scheduled vendors who miss a game are never disciplined. Moreover, vendors whom the subcontractor tries to enlist on short notice are under no obligation to work and receive no extra compensation if they choose to work.[5]

At the stadium, a vendor's work routine is straightforward. Vendors arrive about one hour before the game begins. After arrival, the subcontractor assigns each vendor a product (or products) to sell at that game. To the extent possible, the subcontractor accommodates the preferences of individual vendors when making product assignments. The subcontractor also assigns each vendor to work at either the field level or the mezzanine level. Vendors assigned to the field level are not supposed to sell on the mezzanine level, and vice versa, but there is a third main seating area in the stadium, the upper level, where *all* vendors are allowed to sell. Around game time, vendors begin walking through the stands and selling their

[5] The vending subcontractor did not set strict work schedules for the vendors because doing so would have threatened the vendors' independent contractor status. Maintaining the vendors' independent contractor status also might explain why the subcontractor paid the vendors a pure piece rate, even though this compensation scheme generally is suboptimal (Lazear 1986). Details of the rules governing employee vs. independent contractor classification of a worker can be found in U.S. Department of Treasury (1996).

products. Vendors are allowed to sell only "in the seats"; selling along the stadium concourses and, in particular, in front of the concession stands is prohibited. Vendors who have sold out of their product return to the supply room on their assigned level and "buy" a new batch of the product with the proceeds from the previous load. Vendors continue selling until the end of the seventh inning, when they return to their assigned supply rooms and settle their daily sales accounts.

III. A Model of Vendor Labor Supply

The quantity of labor supplied by each vendor on each game date depends on both a participation decision and a choice of effort (e.g., how fast to circulate through the stands, how loud and frequently to yell). In contrast to the standard model of labor supply, however, the vendors make no hours decision (given participation); instead, each participating vendor works from the start of the game through the seventh inning.[6] In addition, because effort choices are unobservable, this paper focuses exclusively on participation decisions.[7] The absence of an hours margin and my focus on the participation decision have two implications that the reader should keep in mind when perusing the empirical results. First, unlike most papers on labor supply that estimate the elasticity of hours worked with respect to the wage in some broadly defined labor market, this paper estimates the elasticity of participation in a very narrowly defined labor market with respect to the wage in that labor market. Second, to the extent that participating vendors supply more effort when the return to effort is high, the participation elasticity understates the overall labor supply response to shifts in demand.

Vendor i is assumed to participate at game t if and only if his expected earnings equal or exceed his opportunity cost. Let C_{it} denote the opportunity cost for vendor i on date t, and let $G(\cdot|\mathbf{Z}_{it})$ be the distribution from which C_{it} is drawn. The term $\mathbf{Z}_{it}$ is a vector of opportunity cost shifters that might vary across both individuals and dates, so there is no presumption that all vendors draw from the same op-

[6] In principle, a vendor could quit selling early, but the subcontractor claims that this does not occur. The apparent reason is that the cashiers who settle the vendors' daily sales accounts are not available for this task until after the seventh inning, so a vendor who quits selling early still would have to wait until the end of the seventh inning before logging out for the day. This behavior, if it did occur, could be thought of as a reduction in effort.

[7] In an earlier version of the paper (Oettinger 1997), I model the effort decision and conduct empirical analyses that suggest that effort elasticities with respect to both game attendance and the commission rate might be positive and substantial. However, these conclusions rely on an untestable assumption about the form of the production function mapping unobserved effort into observed earnings.

portunity cost distribution. Given the set of realized opportunity costs, aggregate vendor participation at game t as a function of expected earnings per vendor, y, is $N_t(y) = \sum_{i=1}^{N_t^p} 1[C_{it} \leq y]$, where N_t^p denotes the number of potential vendors on date t and $1[\cdot]$ is an indicator function for the condition inside the brackets. The term $N_t(y)$ is a positive step function in y, and therefore the aggregate labor supply (participation) curve is upward-sloping. Taking the expectation of $N_t(y)$ over the random opportunity cost realizations yields the expected or average aggregate participation curve, $E(N_t(y)) = \sum_{i=1}^{N_t^p} G(y|\mathbf{Z}_{it})$. If opportunity costs are continuous random variables, $E(N_t(y))$ is a smooth upward-sloping labor supply curve, and changes in components of $\mathbf{Z}_{it}$ that are common to all vendors will cause systematic shifts in the aggregate labor supply curve.[8]

Vendor i's (latent) earnings at game t are assumed to be $Y_{it} = F(N_t, \mathbf{X}_t) + e_{it}$, where N_t denotes the total number of vendors working at game t, $\mathbf{X}_t$ is a vector of product demand shifters at game t, and e_{it} captures random determinants of earnings. Allowing for a vendor-specific earnings component is straightforward, and vendor fixed effects are included in the individual-level empirical analysis. This extension does not alter any of the model's predictions, however, and therefore I present the case in which vendors are identical to simplify the exposition. With identical vendors, each anticipates the same earnings conditional on product demand conditions and aggregate vendor participation, namely $y \equiv E(Y_{it}|N_t, \mathbf{X}_t) = F(N_t, \mathbf{X}_t)$. When $\mathbf{X}_t$ is held fixed, the function $F(N_t, \mathbf{X}_t)$ can be viewed as an (expected) average product of labor curve. If some of the marginal vendor's sales come from customers who are diverted from the inframarginal vendors, then vendors compete for sales and the average product of labor curve (eventually) slopes downward. Changes in product demand conditions alter average vendor productivity for any given level of aggregate vendor participation and therefore shift the average product curve.

Each potential vendor participates at game t as long as expected earnings at least cover the opportunity cost of participation. Thus

[8] If vendor i observes his own opportunity cost realization, but not the opportunity cost realizations of the other potential vendors, before deciding whether to participate at game t, then the expected aggregate labor supply curve from vendor i's perspective is

$$E_i(N_t(y)) = 1[C_{it} \leq y] + \sum_{j \neq i} G(y|\mathbf{Z}_{jt}).$$

Clearly, $E_i(N_t(y))$ need not equal $E_j(N_t(y))$ for $i \neq j$, and therefore expected earnings at game t could differ for identical vendors. However, as long as the opportunity cost distributions are common knowledge, the difference between $E_i(N_t(y))$ and $E_j(N_t(y))$ for $i \neq j$ will become trivial as N_t^p becomes large.

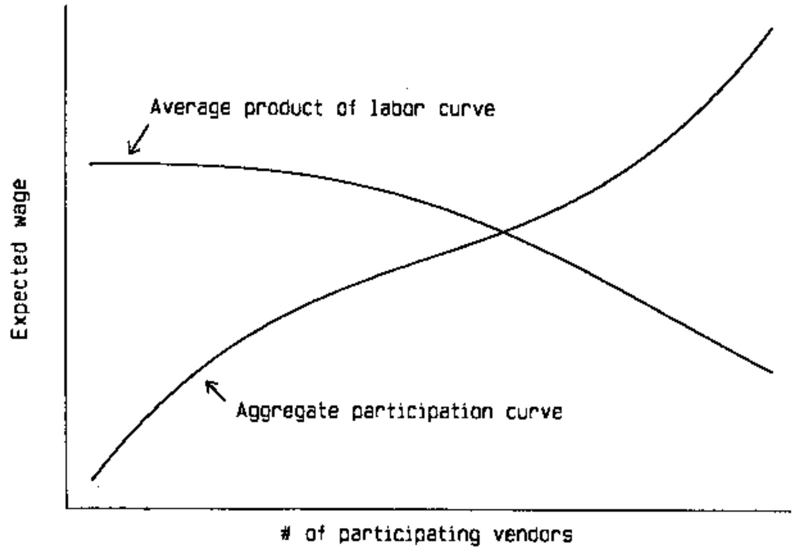

FIG. 1.—Vendor labor market equilibrium

the equilibrium number of participating vendors and the equilibrium level of earnings per vendor at game *t* are determined by the requirement that the marginal vendor's opportunity cost of participation equals expected vendor earnings. This condition holds where the aggregate labor supply and average product of labor curves cross, as shown in figure 1. Clearly, positive product demand shocks raise the probability of participation, aggregate participation, and earnings per vendor in equilibrium, all else equal. Positive opportunity cost shocks also raise equilibrium earnings per vendor but reduce the probability of participation and aggregate participation, other things equal. Of course, since both curves can shift from game to game, simple regressions of participation measures on vendor earnings measures will not yield a consistent estimate of the labor supply (participation) elasticity in general. Instead, obtaining consistent estimates requires instrumental variable–type estimation. The next section describes two approaches, one using individual panel data and the other using the time-series data on aggregate participation and average vendor earnings, for consistently estimating the elasticity of vendor labor supply.

IV. Econometric Methodology

A. *Individual-Level Labor Supply*

The theory developed above assumes that each potential vendor compares opportunity costs and expected earnings when making the participation decision. Thus the empirical model specifies equations for the log of the opportunity cost of participation and the log of earnings from vending for potential vendor i on date t:

$$\ln C_{it} = \mathbf{Z}_{it}\gamma + \theta_i + u_{it} \tag{1}$$

and

$$\ln Y_{it} = \mathbf{X}_{it}\beta + \alpha_i + \epsilon_{it} + \nu_{it}. \tag{2}$$

Of course, the econometrician does not observe $\ln C_{it}$ and observes $\ln Y_{it}$ only for participating vendors.

Equation (1) specifies opportunity costs for vendor i on date t as a function of observable cost shifters ($\mathbf{Z}_{it}$), a vendor-specific fixed effect (θ_i), and an independent random shock (u_{it}). The idiosyncratic disturbance, u_{it}, although unobserved by the econometrician, is assumed to be known to the potential vendor prior to the participation decision. The vendor fixed effect captures any fixed (over the season) differences across vendors in their nonvending opportunities. The vector $\mathbf{Z}_{it}$ includes indicators for the day, time, and season of game t and interactions of these variables with some basic vendor demographic characteristics (dummies for age category, race, and sex).[9] The game time indicators capture any common systematic variation in opportunity costs across times of the day, days of the week, and seasons of the year, and the interactions capture any systematic differences across demographic categories in the temporal pattern of opportunity costs. The vector $\mathbf{Z}_{it}$ also contains measures of weather conditions on date t (the daytime high temperature and a dummy for 24-hour rainfall in excess of a quarter inch). Finally, $\mathbf{Z}_{it}$ sometimes includes indicators for the opposing team and measures of the home and visiting teams' current positions in the standings. These specifications allow for the possibility that vendors are themselves baseball fans whose participation decisions are influenced by the quality of the opposing team or the game's importance to the pennant race.

Equation (2) specifies vendor i's earnings on date t, given participation, as a function of observable predictors of earnings ($\mathbf{X}_{it}$), a vendor-specific fixed effect (α_i), and two unobserved (by the econometrician) random components (ϵ_{it} and ν_{it}). This is a reduced-form earnings equation, and therefore $\mathbf{X}_{it}$ consists only of variables that vendors observe before making participation decisions. In the empirical analysis, $\mathbf{X}_{it}$ always includes measures of when game t is played (day of week, time of day, and season), weather conditions on date

[9] The time-invariant vendor demographic characteristics are not included in $\mathbf{Z}_{it}$ because of the presence of the vendor fixed effects in (1).

t, whether game t is a special promotional date (e.g., "free bat day"), the opposing team, the current place in the standings of both the home and visiting teams, and the number of other *potential* vendors on date t.[10] In addition, $\mathbf{X}_{it}$ includes actual attendance at game t in some specifications. These specifications implicitly assume that vendors have perfect foresight about game attendance, which is not an unreasonable approximation if vendors have important prior information about game attendance that the econometrician does not observe.[11]

The vendor fixed effect in (2) captures permanent (over the season) earnings differences among vendors resulting, say, from differences in vending ability or commission rates. The two transitory unobserved earnings components distinguish conceptually between those components that the vendor observes prior to the participation decision (ϵ_{it}) and those that he does not (ν_{it}). For example, vendor i's energy level on date t would be part of ϵ_{it}, whereas vendor i's random sales luck on date t would be part of ν_{it}. The unforecastable part of product demand conditions also would belong to ν_{it}. These random shocks are assumed to be mutually independent and independent of all the other variables in the model.

As noted earlier, potential vendor i participates on date t only if expected earnings equal or exceed opportunity costs. Thus, when $\Omega_{it} \equiv \{\mathbf{Z}_{it}, \mathbf{X}_{it}, \theta_i, \alpha_i, u_{it}, \epsilon_{it}\}$ denotes the information set at the time of the participation decision, vendor i participates only if $E(Y_{it}|\Omega_{it}) \geq C_{it}$ or, equivalently, if $\ln E(Y_{it}|\Omega_{it}) \geq \ln C_{it}$. Taking the expectation of (2) conditional on Ω_{it} gives $E(\ln Y_{it}|\Omega_{it}) = \mathbf{X}_{it}\beta + \alpha_i + \epsilon_{it}$. Given the independence assumptions on the error terms, $E(\ln Y_{it}|\Omega_{it})$ and $\ln E(Y_{it}|\Omega_{it})$ differ by only a constant (namely, $\ln E(e^{\nu})$), and therefore only the estimated constant term is affected if one assumes that participation is determined by a comparison of $\ln C_{it}$ with $E(\ln Y_{it}|\Omega_{it})$ instead of $\ln E(Y_{it}|\Omega_{it})$. Thus, when P_{it} is the indicator for participation by vendor i on date t, it is assumed that $P_{it} = 1$ if and only if $E(\ln Y_{it}|\Omega_{it}) \geq \ln C_{it}$.

Estimation of the individual-level model of vendor participation proceeds in several steps, essentially following the methodology in Lee (1978) and Willis and Rosen (1979). First, the idiosyncratic er-

[10] The number of other potential vendors on date t might affect vendor i's earnings on date t by affecting the aggregate participation of other vendors. The construction of measures of the number of other potential vendors on each date t is described in Sec. V of the paper.

[11] The R^2 from a regression of game attendance on all the ex ante observable predictors of attendance is .76. Thus, from the econometrician's perspective, there exists nontrivial unexplained game-to-game variation in attendance, and this variation might be (partially) forecastable by potential vendors.

ror terms are assumed to be normally distributed, and a reduced-form probit model for participation is estimated. The explanatory variables consist of observable opportunity cost shifters ($\mathbf{Z}_{it}$), observable predictors of earnings ($\mathbf{X}_{it}$), and a complete set of individual vendor dummies.[12]

Next, a selectivity-corrected log earnings equation is estimated. The adjustment for selection is necessary because earnings are observed only for vendors with $E(\ln Y_{it}|\Omega_{it}) \geq \ln C_{it}$, and this self-selected participation induces a correlation between the idiosyncratic errors and the explanatory variables in the sample of participating vendors. In particular, expected log earnings conditional on the observables, the vendor fixed effect, and participation are

$$E(\ln Y_{it}|\mathbf{X}_{it}, \mathbf{Z}_{it}, \alpha_i, P_{it} = 1) = \mathbf{X}_{it}\beta + \alpha_i + E(\epsilon_{it}|\epsilon_{it} + u_{it} \geq \mathbf{Z}_{it}\gamma - \mathbf{X}_{it}\beta + \theta_i - \alpha_i). \quad (3)$$

The conditional expectation on the right-hand side of (3) represents the bias from self-selected participation. However, including an inverse Mills ratio term, derived from the reduced-form participation probit estimates, as an additional explanatory variable in (2) corrects for this bias (Heckman 1976). The selection correction term is identified if some variables affect opportunity costs (and hence participation) but not earnings. I assume that the interactions between vendor i's demographic characteristics and the day, time, and season of game t satisfy this condition. This assumption says that the temporal variation in vendor opportunity costs differs systematically across demographic categories but the temporal variation in vendor expected earnings (productivity) does not.[13]

Given the earnings equation estimates, predicted log earnings are constructed using (3). This yields an uncensored sample of "wage" observations for all the potential vendors on every game date, regardless of actual participation decisions, and allows estimation of the structural probit model of participation. This last model explains participation decisions as a function of opportunity cost shifters ($\mathbf{Z}_{it}$), predicted log earnings ($\widehat{\ln Y_{it}}$), and vendor fixed effects. The

[12] In general, estimates of nonlinear panel data models with individual fixed effects are consistent only as the panel grows arbitrarily long (Chamberlain 1980). Fortunately, the panel used here is sufficiently long—the average number of observations per vendor for the participation probit model is nearly 50—that the asymptotic consistency result is likely to be approximately valid.

[13] To take a concrete example, the assumption is that the relative opportunity costs of participating on a weekday afternoon vs. a Sunday differ on average between a 16-year-old and a 30-year-old (because of systematic differences across age categories in school enrollment status, marital status, the presence of young children in the household, etc.) but that relative expected earnings from vending at these two times are the same for both vendors.

model is identified as long as some of the predictors of vendor i's earnings at game t do not also directly influence vendor i's opportunity cost of participation. Many of the product demand shifters—for example, attendance at game t, whether game t is a promotional date, or (if one assumes that the vendors are not baseball fans) the identity of the opposing team at game t—plausibly satisfy this condition.

B. *Aggregate Labor Supply*

The theory developed earlier also suggests an empirical model for aggregate outcomes in the vendor labor market. Although an aggregate analysis does not use all of the variation in the data and must ignore the issues of unobserved vendor heterogeneity and self-selected participation, it does offer several advantages. First, the aggregate model can be estimated in one step and is therefore much simpler. Second, testing alternative identifying assumptions is more straightforward for the aggregate model. Third, and most important, one can assess the extent of the bias that results from ignoring the potential endogeneity of the daily vendor wage by a simple comparison of instrumental variable and ordinary least squares (OLS) estimates. In contrast, no simple comparison is possible for the individual-level specification because earnings are observed only for vendors who choose to work.

I specify the aggregate empirical model

$$\ln N_t = \mathbf{Z}_t\tilde{\gamma} + \delta \ln Y_t + \tilde{u}_t \tag{4}$$

and

$$\ln Y_t = \mathbf{X}_t\tilde{\beta} + \pi \ln N_t + \tilde{v}_t, \tag{5}$$

where N_t and Y_t denote aggregate vendor participation and average vendor earnings on date t, respectively. Equation (4) is the empirical counterpart of the aggregate labor supply curve derived in Section III. Consistent with the theory, aggregate participation is assumed to depend on average vendor earnings on date t (Y_t), observable measures of date t opportunity costs ($\mathbf{Z}_t$), and unobserved components of date t opportunity costs ($\tilde{u}_t$). The vector $\mathbf{Z}_t$ consists of indicators for the day, time, and season of game t, weather conditions on date t, the log of the number of *potential* vendors on date t, and, in some specifications, indicators for the opposing team and current place in the standings of the home and visiting teams.

Equation (5) is the empirical counterpart of the average product of labor curve derived in Section III. Again consistent with the theory, average vendor productivity (earnings) is assumed to depend on aggregate vendor participation on date t (N_t), observable measures of date t product demand ($\mathbf{X}_t$), and unobserved components of date t product demand ($\tilde{v}_t$). The vector $\mathbf{X}_t$ consists of the ex ante predictors of both attendance and the demographic composition of the crowd (e.g., the time of game dummies, the promotional date dummy, the opponent dummies, the weather variables, etc.) and, in some specifications, the actual realization of log attendance.

Theory suggests that δ is positive (aggregate labor supply slopes upward) and that π is probably negative (vendors' average product slopes downward). However, because of the simultaneous relationship in (4) and (5), OLS estimation of (4) will produce a downward-biased estimate of δ, the wage elasticity of aggregate vendor participation, if the true signs of δ and π are as predicted by the theory. Intuitively, OLS attributes all unexplained wage variation to demand shifts, and therefore the estimated labor supply curve will be too steep if labor supply shifts are the source of some unexplained wage variation. One can estimate δ consistently, however, by estimating (4) by two-stage least squares (2SLS), using the product demand shifters that do not directly influence vendor opportunity costs (i.e., attendance or certain of its ex ante predictors) as instruments for vendor earnings.

V. Empirical Analysis

A. *Descriptive Statistics*

Before turning to the econometric analysis, I first summarize the cross-section data on vendor demographic characteristics, the time-series data on game characteristics, and the panel data on participation and earnings. A total of 127 vendors worked at one or more of the 81 home games during the season, and aggregate participation over the entire season was 3,580 vendor-games. Of course, to analyze participation decisions, one needs to know not only which vendors actually participate on each date but also which vendors *could* have participated. If both the hire date and the quit date were observed for every vendor, the set of potential vendors on each date could be computed simply as the number of vendors who had already been hired and who had not yet quit. Unfortunately, although hire dates are available for all the vendors, quits are not explicitly observed. As a result, a quit imputation rule must be specified before one can calculate the number of "active" vendors on each date (i.e., the number of vendors who potentially could have worked).

All the empirical analyses are performed for two alternative quit imputation rules. The first rule imputes a quit on the date on which a spell of nonparticipation reaches 31 days. Thus, under this rule, a vendor is defined as active on date t if and only if she has worked or was hired within the last 30 days.[14] Under this rule, there are 6,182 active-status vendor-games during the season. One could argue, however, that this imputation rule is not sufficiently conservative since some vendors do in fact participate after nonparticipation spells of more than 30 days.[15] Thus the second rule ensures that no quits are imputed in error by simply assuming that vendors never quit and hence are active on all dates after the hire date. Under this rule, the number of active-status vendor-games rises to 8,712. For obvious reasons, the two definitions of active status are referred to as "narrow" (vendors who have worked or were hired in the last 30 days) and "broad" (all previously hired vendors), respectively. Although I believe that the narrow definition of active status corresponds more closely to the set of vendors actually at risk of participating, the estimated labor supply elasticities are always fairly similar for both definitions.

Table 1 summarizes the cross-sectional variation in the data. Since participation and earnings vary across games for any given vendor, the table reports cross-sectional means and standard deviations of the time-averaged earnings and time-averaged (or cumulative) participation over the season for all vendors. Both unweighted summary statistics and summary statistics weighted by the number of games worked (or, where noted, by the number of active-status games) are presented. The unweighted statistics reflect average characteristics of a randomly sampled vendor who *ever* worked during the season, whereas the weighted statistics represent average characteristics of a randomly sampled working vendor.

Panel A of the table summarizes vendor participation behavior. Obviously, vendors are active for fewer games and have correspondingly higher participation rates (defined as games worked divided by games active) under the narrow definition of active status. The average vendor works at about half of the games during his active spell, but some vendors work at every game whereas others work once and then never again. For the narrow definition of active status, the average participation rate is much smaller when not weighted

[14] This rule seems consistent with evidence from (unreported) hazard model estimates, which shows that the probability of participation on date t given a current nonparticipation spell of L days is substantial for $L < 15$ but is very close or equal to zero for all $L > 30$.

[15] While only 0.56 percent of the participation observations (20 out of 3,580) are immediately preceded by a nonparticipation spell in excess of 30 days, 15 percent of the vendors (19 out of 127) have a nonparticipation spell of at least 30 days followed by participation at some point during the season.

TABLE 1

Summary Statistics on Vendor Participation Behavior, Average Vendor Earnings, and Vendor Demographic Characteristics

Variable	Unweighted Mean	Unweighted Standard Deviation	Unweighted N	Weighted Mean	Weighted Standard Deviation	Weighted N	Minimum	Maximum
	A. Vendor Participation Behavior							
Narrow definition of active status:								
Number of active games	48.68	25.27	127	61.69	21.62	6,182	6	81
Participation rate	.474	.287	127	.576	.270	6,182	0	1
Broad definition of active status:								
Number of active games	68.60	18.53	127	73.56	13.70	8,712	6	81
Participation rate	.408	.321	127	.411	.326	8,712	.012	1
	B. Vendor Average Earnings							
Per game	35.71	15.80	127	43.81	17.37	3,580	.83	106.16
Per hour of actual vending time	15.92	7.09	127	19.58	7.75	3,580	.36	47.69
	C. Vendor Characteristics							
Total games worked by vendor	28.19	24.57	127	49.43	21.69	3,580	1	81
Male	.803		127	.820		3,580		
Female	.197		127	.180		3,580		
African-American	.528		127	.535		3,580		
Hispanic	.284		127	.279		3,580		
Other	.189		127	.187		3,580		
16–19 years old	.252		127	.208		3,580		
20–24 years old	.213		127	.189		3,580		
25–39 years old	.472		127	.513		3,580		
40–54 years old	.063		127	.089		3,580		

Note.—The columns labeled N list the number of vendors for the unweighted sample statistics and the number of vendor-game observations (the sum of the sample weights) for the weighted sample statistics. The weighting variable is the number of games worked by each vendor except in panel A, where the weighting variable is the number of games in which each vendor was active. Under the narrow definition of active status, a vendor is included in the pool of potential vendors on date t only if he last worked or was hired in the previous 30 days. Under the broad definition of active status, a vendor is included in the pool of potential vendors on any date t after the hire date. The vendor participation rate is defined as the ratio of games worked to games active for each vendor. Standard deviations and extreme values are not reported for dichotomous variables.

by the number of games active because vendors who quit shortly after being hired have both the shortest active spells (under the narrow definition) and the lowest participation rates during their active spells. In contrast, there is no difference between the unweighted and weighted average participation rates under the broad definition of active status, which indicates that participation propensities do not vary systematically by date of hire.

Panel B of the table summarizes how average earnings per game and per hour of actual vending time vary across vendors who worked during the season. Average earnings, like participation rates, exhibit a great deal of heterogeneity across vendors: the lowest-earning vendor (who worked at only one game) had average earnings of less than \$1.00 per game worked, whereas the highest-earning vendor had average earnings in excess of \$100 per game worked. Participation-weighted average vendor earnings are much larger than unweighted average vendor earnings, reflecting the more frequent participation of relatively high-earning vendors. Finally, panel C of the table summarizes vendor demographic characteristics, revealing that the vendors are disproportionately male, nonwhite, and young.

Table 2 cuts the data in a different way to summarize their variation in the time dimension. Since participation and earnings vary across vendors at any given game, the table focuses on the time-series variation by first calculating average earnings and average (or cumulative) participation across all vendors for each game and then reporting the means and standard deviations of these game averages. As in table 1, both unweighted and weighted summary statistics are reported, although the weights are now the number of vendors who worked at the game (or, where noted, the number of vendors who were active on the game date). In contrast to table 1, there are few notable differences between the unweighted and weighted statistics.

Panel A of table 2 describes participation behavior across games. The main result of interest is that the game participation rate (defined as the ratio of participating vendors to active vendors at a given game) varies considerably between games, regardless of how active status is defined. Thus, on some dates, most of the active vendors participate, whereas on other dates, most of the active vendors do not. There are no games, however, in which all the active vendors choose to participate.

Panel B of the table shows average game earnings, both per vendor and per hour of actual vending time, and also characterizes the extent of game-to-game variation in earnings. Average daily earnings over all participating vendors are around \$43 at the typical game

TABLE 2

SUMMARY STATISTICS ON GAME PARTICIPATION BEHAVIOR, AVERAGE GAME EARNINGS, AND GAME CHARACTERISTICS

	Unweighted			Weighted				
Variable	Mean	Standard Deviation	N	Mean	Standard Deviation	N	Minimum	Maximum
	A. Game Participation Behavior							
Narrow definition of active status:								
Number of active vendors	76.32	5.87	81	76.77	5.60	6,182	62	85
Participation rate	.574	.113	81	.576	.112	6,182	.238	.831
Broad definition of active status:								
Number of active vendors	107.56	17.96	81	110.52	16.82	8,712	76	127
Participation rate	.421	.108	81	.411	.106	8,712	.118	.684
	B. Game Average Earnings							
Per vendor	43.39	10.45	81	43.81	10.29	3,580	26.55	73.15
Per hour of actual vending time	19.39	4.84	81	19.58	4.91	3,580	11.15	34.19

	C. Game/Date Characteristics							
Total vendors working at game	44.20	9.93	81	46.40	8.98	3,580	15	66
Attendance	24,394	8,953	81	25,189	9,053	3,580	12,700	49,674
Time of first 7 innings (min.)	135.4	16.3	81	135.6	16.1	3,580	98.8	184.3
Weekday (Monday–Thursday) day game	.111		81	.075		3,580		
Weekday (Monday–Thursday) night game	.407		81	.413		3,580		
Friday (night) game	.161		81	.182		3,580		
Saturday (night) game	.161		81	.173		3,580		
Sunday (day) game	.161		81	.158		3,580		
Promotional date	.271		81	.277		3,580		
Before Memorial Day	.321		81	.295		3,580		
Summer	.518		81	.584		3,580		
After Labor Day	.161		81	.121		3,580		
Home team in first place	.506		81	.515		3,580		
Home team games out of first	1.22	(1.83)	81	1.09	(1.67)	3,580	0	7
Opponent in first place	.210		81	.241		3,580		
Opponent games out of first	6.64	(7.48)	81	6.01	(6.81)	3,580	0	29
Daytime high temperature	88.84	(6.96)	81	89.34	(7.02)	3,580	64	99
24-hour rainfall > .25 inch	.099		81	.095		3,580		

NOTE.—The columns labeled N list the number of games for the unweighted sample statistics and the number of vendor-game observations (the sum of the sample weights) for the weighted sample statistics. The weighting variable is the number of vendors working at each game except in panel A, where the weighting variable is the number of active vendors on each date. The game participation rate is defined as the ratio of participating vendors to active vendors for each date. Standard deviations and extreme values are not reported for dichotomous variables. See the note to table 1 for the definitions of narrow and broad active status.

but fluctuate considerably across games, ranging from a low of $26 to a high of $73. The coefficient of variation is roughly the same for average earnings per hour of actual vending time, so between-game differences in actual time spent vending are not the source of this variation. Rather, there clearly are "high-wage" and "low-wage" games, although the summary statistics cannot reveal the source of this wage variation. Finally, panel C of the table reports summary statistics for game-specific characteristics. The most interesting finding is large game-to-game variation in both attendance and aggregate vendor participation.

B. *Individual-Level Estimation Results*

Estimation results for the individual-level empirical model are presented in tables 3–5. Because there might be random unobserved date or game effects, the estimated covariance matrices in these tables always allow for arbitrary correlation among vendors' transitory error components *within* each date t.[16] In addition, the estimated covariance matrices in tables 4 and 5 are adjusted to account for the fact that the associated empirical models include an explanatory variable that is itself estimated. These adjustments follow the methodology described in Murphy and Topel (1985), appropriately modified for the case of a robust covariance matrix.

Table 3 presents the estimates of the reduced-form participation probit model for both the narrow and broad definitions of active status. Estimates both with and without the log of game attendance included as an explanatory variable are reported. As I pointed out earlier, including the attendance measure is justified if vendors have perfect foresight over attendance (or at least have significant prior information that the econometrician cannot observe). Panel A of the table shows estimated coefficients and standard errors for selected explanatory variables. Panel B shows p-values and degrees of freedom for χ^2 tests of the joint significance of (1) the vendor dummies, (2) the opposing team dummies, and (3) the interactions between the dummies for vendor demographic characteristics and the time of game dummies. Because interactions between the demographic indicators and the time of game dummies are included, the estimated coefficients in panel A describe the participation behavior of 25–39-year-old white males at different game times.

All the specifications tell the same basic story. Relative to the reference game (a nonpromotional Sunday game), the probability of par-

[16] It is assumed, however, that the transitory error components are uncorrelated *across* dates.

TABLE 3

ESTIMATES OF REDUCED-FORM PROBIT MODEL FOR PARTICIPATION

	Definition of Active Status			
	Narrow		Broad	
	(1)	(2)	(3)	(4)
	A. Coefficient Estimates and Standard Errors			
Monday–Thursday day game	−.7242	−.6370	−.6893	−.5952
	(.1363)	(.1383)	(.1314)	(.1379)
Monday–Thursday night game	−.0700	.0311	−.0012	.1074
	(.1604)	(.1647)	(.1591)	(.1663)
Friday (night) game	.4002	.4187	.3850	.4137
	(.1990)	(.2028)	(.1839)	(.1881)
Saturday (night) game	.3125	.3180	.3005	.3038
	(.1623)	(.1580)	(.1556)	(.1508)
Promotional date	.2624	.1919	.2590	.1950
	(.0753)	(.0717)	(.0651)	(.0588)
Opponent in first place	.2057	.1373	.2303	.1701
	(.0762)	(.0853)	(.0872)	(.0925)
Home team games out of first	−.0511	−.0465	−.0518	−.0480
	(.0283)	(.0269)	(.0237)	(.0220)
Daytime high temperature	.0127	.0114	.0071	.0052
	(.0036)	(.0036)	(.0065)	(.0063)
24-hour rainfall > .25 inch	.0017	−.0056	−.0109	−.0211
	(.0824)	(.0815)	(.0684)	(.0671)
Log of attendance	...	.2589	...	.2533
		(.1055)		(.0976)
	B. χ^2 Statistic p-Values and Degrees of Freedom			
Individual vendor dummies	<.0001	<.0001	<.0001	<.0001
	[123]	[123]	[124]	[124]
Opponent dummies	<.0001	<.0001	<.0001	<.0001
	[12]	[12]	[12]	[12]
Vendor demographic indicators × day/time/season dummies	<.0001	<.0001	<.0001	<.0001
	[36]	[36]	[36]	[36]
Observations	6,029	6,029	8,601	8,601
Log likelihood	−2,905.4	−2,902.9	−3,248.1	−3,245.2

NOTE.—The estimated covariance matrix allows for an arbitrary error covariance structure across vendors at any given game but assumes independent errors across games, after allowing for vendor fixed effects. The sample sizes are slightly smaller than the total number of active observations in tables 1 and 2 because the inclusion of vendor fixed effects eliminates vendors who either always participated or never participated. All the specifications also include as explanatory variables the log of the number of (other) active vendors, the number of games the opposing team is out of first place, and indicators for the season (before Memorial Day or after Labor Day) and for whether the home team was in first place.

ticipation is higher for promotional dates and on Friday and Saturday nights but is much lower for weekday afternoon games. The probability of participation also is higher for games against first-place teams and when the home team is closer to first place in its own division. In some specifications, the participation probability rises with the daily high temperature, but the magnitude of this effect is always very small. Finally, when attendance is included as an explanatory variable, the participation probability rises with crowd size. The hypothesis tests in panel B of the table reveal highly significant differences in participation probabilities across individual vendors and across opposing teams. The interactions between the demographic indicators and the time of game indicators also are jointly highly significant, which I interpret as evidence of important differences in the temporal pattern of opportunity costs across demographic categories. This interpretation assumes, plausibly, that the effects of temporal variation in product demand are the same for vendors from all demographic groups and therefore are captured by the time of game main effects. The (unreported) estimated coefficients on the interaction terms generally support this interpretation; for example, teenage vendors, who are likely to be enrolled in school, have much lower participation probabilities on weekday afternoons and in the spring and fall than vendors from the reference category.

Table 4 presents estimates of the reduced-form log earnings equation. The dependent variable is earnings per hour of vending time, but the results are qualitatively identical if one uses daily earnings instead.[17] For each of the specifications of the log earnings equation in table 4, the inverse Mills ratio (selection correction) term is derived from the reduced-form participation probit in the corresponding column of table 3. The selection term is identified off of the interactions between the vendor demographic characteristics and the time of game dummies, which were highly significant in the participation model but are excluded from the earnings model. There is some evidence that participation is positively selected since the estimated coefficient on the inverse Mills ratio is always positive and is statistically significant for the narrow definition of active status. This result says that a vendor is less likely to participate when the unobserved transitory earnings component is low (e.g., when the vendor is fatigued). Positive selection is precisely what one would expect if the transitory shocks to daily earnings and daily opportunity costs are independent.

The coefficient estimates are essentially identical for specifications that differ only in how active status is defined. In contrast, though, whether one controls for game attendance has a large effect on many of the other estimated coefficients because attendance is highly correlated with many of these explanatory variables. For example, without controls for attendance, earnings are about 30 percent lower at weekday night games and about 15 percent higher on

[17] The daily earnings measure is probably a noisier (and hence inferior) measure of the vendor wage since random variation in game length causes some of the variation in daily earnings.

promotional dates, other things equal. However, these effects basically disappear once one controls for attendance. Apparently, the lower earnings on weekday afternoons and higher earnings on promotional dates are driven by swings in attendance. Even after one controls for attendance, though, some of the other explanatory variables remain correlated with earnings because of true temporal effects in product demand, "crowd composition" effects, or aggregate participation effects. For instance, when attendance is held constant, earnings are about 11 percent higher at Saturday night games. This might reflect that any given individual spends more on a Saturday night (a true temporal effect), that high spenders are a larger fraction of the crowd on Saturday nights (a compositional effect), or that aggregate vendor participation does not increase proportionately with attendance on Saturday nights. Finally, there is very strong evidence in all specifications of systematic individual vendor effects and opposing team effects in earnings.

Table 5 presents estimates of the main model of interest, the structural participation probit. The key explanatory variable is predicted log hourly earnings, and for each structural probit specification, this variable is constructed from the log earnings equation in the corresponding column of table 4. In columns 1 and 3, attendance is not part of the earnings model, and therefore the participation model is identified by assuming that the opposing team indicators, the measures of the home and visiting teams' places in the standings, and the promotional date indicator affect earnings but do not influence participation except through their effects on earnings. Excluding the opposing team dummies and the team performance measures from the structural probit amounts to an assumption that vendors are *not* baseball fans whose participation decisions depend partially on the quality of the opponent or the importance of the game. In contrast, in columns 2 and 4, attendance is part of the earnings model, and the model is identified by excluding only attendance and the promotional date indicator from the opportunity cost equation. Thus these specifications allow for the possibility that vendors are also fans.

The coefficient on predicted log earnings is positive and highly significant in all the specifications. Thus, all else equal, an increase in expected earnings raises the probability of participation. Moreover, the estimated coefficient is basically the same under either set of identifying assumptions described above. Likewise, the results are qualitatively identical if one reestimates the earnings models using the log of *daily* earnings as the dependent variable and then uses predicted log daily earnings in the structural participation probit. Many of the other coefficient estimates are quite similar to those

TABLE 4

ESTIMATES OF REDUCED-FORM LOG EARNINGS EQUATION

	DEFINITION OF ACTIVE STATUS			
	Narrow		Broad	
	(1)	(2)	(3)	(4)
	A. Coefficient Estimates and Standard Errors			
Monday–Thursday day game	−.0565	.1435	−.0492	.1550
	(.0689)	(.0477)	(.0672)	(.0428)
Monday–Thursday night game	−.3058	−.0607	−.3095	−.0645
	(.0517)	(.0455)	(.0548)	(.0465)
Friday (night) game	−.0312	.0480	−.0280	.0463
	(.0582)	(.0406)	(.0594)	(.0406)
Saturday (night) game	.1117	.1152	.1091	.1115
	(.0458)	(.0357)	(.0460)	(.0369)
Promotional date	.1550	.0266	.1702	.0393
	(.0533)	(.0342)	(.0565)	(.0375)
Opponent in first place	.0692	−.0556	.0582	−.0602
	(.0658)	(.0490)	(.0640)	(.0503)

Home team games out of first	−.0404	−.0347	−.0305	−.0260
	(.0248)	(.0150)	(.0220)	(.0132)
Daytime high temperature	.0069	.0047	.0106	.0071
	(.0027)	(.0018)	(.0036)	(.0029)
24-hour rainfall > .25 inch	.1242	.1084	.1247	.1086
	(.0643)	(.0470)	(.0685)	(.0469)
Log of attendance	. . .	.5680	. . .	.5600
		(.0606)		(.0635)
Inverse Mills ratio (selectivity correction)	.1736	.1523	.1051	.0818
	(.0715)	(.0712)	(.0669)	(.0656)
	B. χ^2 Statistic p-Values and Degrees of Freedom			
Individual vendor dummies	<.0001	<.0001	<.0001	<.0001
	[125]	[125]	[126]	[126]
Opponent dummies	<.0001	<.0001	<.0001	.0002
	[12]	[12]	[12]	[12]
Observations	3,579	3,579	3,580	3,580
R^2	.650	.670	.649	.669

Note.—The estimated covariance matrix allows for an arbitrary error covariance structure across vendors at any given game but assumes independent errors across games, after allowing for vendor fixed effects. One earnings observation is lost under the narrow definition of active status because there is one vendor who participated at only one game, which took place more than 30 days after the date of hire. All the specifications also include as explanatory variables the log of the number of (other) active vendors, the number of games the opposing team is out of first place, and indicators for the season (before Memorial Day or after Labor Day) and for whether the home team was in first place.

TABLE 5

ESTIMATES OF STRUCTURAL PROBIT MODEL FOR PARTICIPATION

	Definition of Active Status			
	Narrow		Broad	
	(1)	(2)	(3)	(4)
	A. Coefficient Estimates and Standard Errors			
Predicted log hourly earnings	.7644	.7282	.6125	.6045
	(.1990)	(.2173)	(.1819)	(.1934)
Monday–Thursday day game	−.6815	−.7347	−.6258	−.6897
	(.1716)	(.1404)	(.1612)	(.1494)
Monday–Thursday night game	.1624	.0638	.1869	.0966
	(.1735)	(.1882)	(.1665)	(.1942)
Friday (night) game	.4105	.3842	.3783	.3629
	(.2094)	(.2111)	(.1803)	(.1901)
Saturday (night) game	.2923	.2927	.2739	.2729
	(.1714)	(.1581)	(.1539)	(.1462)
Opponent in first place	...	.1203	...	.1504
		(.1022)		(.1015)
Home team games out of first	...	−.0173	...	−.0321
		(.0268)		(.0233)

Daytime high temperature	−.0031	.0078	−.0066	.0002
	(.0054)	(.0042)	(.0060)	(.0039)
24-hour rainfall > .25 inch	−.2690	−.0860	−.2500	−.0853
	(.1288)	(.1087)	(.1239)	(.0883)
	B. χ^2 Statistic p-Values and Degrees of Freedom			
Individual vendor dummies	<.0001	<.0001	<.0001	<.0001
	[123]	[123]	[124]	[124]
Opponent dummies	. . .	<.0001	. . .	<.0001
		[12]		[12]
Vendor demographic indicators × day/time/season dummies	<.0001	<.0001	<.0001	<.0001
	[36]	[36]	[36]	[36]
Sample average elasticity of participation with respect to hourly earnings	.568	.546	.757	.759
Elasticity of participation with respect to hourly earnings at covariate sample means	.511	.480	.677	.675
	(.132)	(.143)	(.199)	(.216)
Observations	6,029	6,029	8,601	8,601
Log likelihood	−2,953.1	−2,907.5	−3,319.9	−3,250.9

NOTE.—The estimated covariance matrix allows for an arbitrary error covariance structure across vendors at any given game but assumes independent errors across games, after allowing for vendor fixed effects. The sample sizes are slightly smaller than the total number of active observations in tables 1 and 2 because the inclusion of vendor fixed effects eliminates vendors who either always participated or never participated. All the specifications include as explanatory variables indicators for the season (before Memorial Day or after Labor Day). The specifications in cols. 2 and 4 also include as explanatory variables the number of games the opposing team is out of first place and an indicator for whether the home team is in first place.

reported in table 3. In particular, there is very strong evidence that opportunity costs of time vary greatly across individual vendors and across times of the day and days of the week.

Expressing the estimated coefficients on log hourly earnings as elasticities illustrates more clearly the magnitude of the effect of expected earnings on participation. Two distinct but similar measures of the average elasticity of participation with respect to expected earnings are reported at the bottom of table 5: (1) the sample mean of the estimated individual elasticities and (2) the estimated elasticity at the sample mean values of the covariates.[18] The first elasticity measure seems preferable given that most of the covariates are dummy variables, but computing a standard error is less cumbersome for the second measure. The sample average estimated participation elasticities are approximately .55 and .75 for the narrow and broad definitions of active status, respectively, and are always highly significant. Thus stadium vendors appear to supply labor quite elastically. The estimates imply that, when one starts at the mean vendor wage, a one-standard-deviation increase in the wage would raise the probability of participation by about .08 for a vendor with the mean participation probability. This translates into an increase in aggregate vendor participation of about six vendors when the pool of potential vendors is at its average size.

C. *Aggregate Estimation Results*

Table 6 presents estimates of the aggregate participation model. As discussed earlier, both product demand shifts *and* labor supply shifts are potential sources of variation in the daily vendor wage, and therefore the wage measure must be treated as endogenous in the aggre-

[18] In particular, let $\hat{\beta}_y$ denote the estimated coefficient on predicted log earnings (reported in the first row of table 5), let $\hat{\boldsymbol{\beta}}$ denote the entire estimated coefficient vector, let $\mathbf{X}_{it}$ denote the vector of covariates for potential vendor i on date t, and let $\bar{X}$ denote the sample mean value of $\mathbf{X}_{it}$. Then the estimated elasticity of the probability of participation for vendor i on date t is given by $\hat{\beta}_y\phi(\mathbf{X}_{it}\hat{\boldsymbol{\beta}})/\Phi(\mathbf{X}_{it}\hat{\boldsymbol{\beta}})$, where $\phi(\cdot)$ denotes the standard normal probability density function and $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. The first elasticity measure reported at the bottom of table 5 is

$$\frac{\sum_{t=1}^{T}\sum_{i=1}^{N_t^p}\hat{\beta}_y\phi(\mathbf{X}_{it}\hat{\boldsymbol{\beta}})/\Phi(\mathbf{X}_{it}\hat{\boldsymbol{\beta}})}{\sum_{t=1}^{T}N_t^p},$$

where N_t^p denotes the number of active vendors on date t, and the second elasticity measure reported at the bottom of table 5 is $\hat{\beta}_y\phi(\bar{X}\hat{\boldsymbol{\beta}})/\Phi(\bar{X}\hat{\boldsymbol{\beta}})$.

TABLE 6

ESTIMATES OF THE AGGREGATE PARTICIPATION MODEL

Dependent Variable: Log of Aggregate Participation

	OLS		2SLS		
	(1)	(2)	(3)	(4)	(5)
	Coefficient Estimates and Standard Errors				
Log of average hourly earnings of participating vendors	.2378	.0858	.5346	.6209	.6457
	(.0986)	(.1107)	(.1508)	(.1525)	(.2064)
Monday–Thursday day game	−.3764	−.4024	−.3640	−.3604	−.3997
	(.0650)	(.0596)	(.0692)	(.0718)	(.0724)
Monday–Thursday night game	.0870	−.0086	.1838	.2120	.1587
	(.0594)	(.0580)	(.0723)	(.0742)	(.0847)
Friday (night) game	.1772	.1515	.2040	.2118	.2114
	(.0586)	(.0514)	(.0630)	(.0653)	(.0646)
Saturday (night) game	.0735	.0841	.0408	.0312	.0286
	(.0587)	(.0508)	(.0635)	(.0657)	(.0636)
Opponent in first place	...	.0410	...	...	.0272
		(.0613)			(.0745)
Home team games out of first	...	−.0586	...	...	−.0313
		(.0212)			(.0269)
Daytime high temperature	.0008	.0057	−.0002	−.0005	.0041
	(.0028)	(.0028)	(.0029)	(.0031)	(.0034)
24-hour rainfall > .25 inch	−.1080	.0027	−.1520	−.1648	−.0734
	(.0621)	(.0613)	(.0679)	(.0703)	(.0774)
	Included as Controls?				
Opponent indicators	no	yes	no	no	yes
Measures of team quality	no	yes	no	no	yes
	Exclusion Restrictions (Instruments for Log Earnings)				
Promotional date indicator	...	...	yes	yes	yes
Log attendance	...	...	no	yes	yes
Opponent indicators	...	...	yes	no	no
Measures of team quality	...	...	yes	no	no
	Overidentification Test				
p-value	...	...	.060	.021	.400
Degrees of freedom	...	...	16	1	1
	Test of Joint Significance of Instruments in First Stage of Regression				
p-value	...	...	.0013	<.0001	<.0001
Degrees of freedom	...	...	17	2	2
Observations	81	81	81	81	81
R^2	.727	.847	.692	.669	.774

NOTE.—All the specifications also include as explanatory variables the log of the total number of active vendors and indicators for the season (before Memorial Day or after Labor Day).

gate participation equation. Thus an instrumental variables estimation procedure is needed to obtain consistent estimates of the participation elasticity. To assess the importance of recognizing this endogeneity of vendor earnings—in other words, the importance of labor supply shifts as a source of unexplained earnings variation—table 6 reports OLS estimates of the aggregate participation equation in addition to the more appropriate 2SLS estimates. Theory predicts that OLS estimates of the participation elasticity will be downward biased because OLS implicitly assumes that demand shifts are the source of all unexplained wage variation. The wage measure used in the estimates in table 6 is average earnings per hour of actual vending time calculated over all participating vendors on each date; the results are qualitatively the same if one instead uses average *daily* earnings. Likewise, table 6 reports estimates only for the case in which the pool of potential vendors on each date is calculated using the narrow definition of active status; the results are qualitatively similar when the potential vendor pool is calculated using the broad definition of active status.

Columns 1 and 2 of table 6 present OLS estimates of the aggregate participation equation. The specification in column 1 assumes that the opponent dummies and the measures of home and visiting team performance do not directly influence labor supply decisions. For this specification, the estimated participation elasticity of .24 is significantly different from zero but is much smaller than the estimates from the individual-level model. In column 2, when the identity of the opposing team and the quality of both teams are allowed to directly affect labor supply, the OLS estimate of the participation elasticity falls to .09 and is no longer statistically significant. Thus, on the basis of the OLS estimates, one would conclude that participation responds weakly to expected earnings if at all.

The 2SLS estimates in the remainder of table 6 lead to a very different conclusion, however. The specifications in columns 3 and 4 both assume that vendor participation decisions are not directly influenced by the identity of the opponent or the home and visiting teams' positions in the standings, but they use different sets of instruments for the log of average hourly earnings. The resulting participation elasticity estimates of .53 and .62 are much higher than the OLS estimates and are very similar to the elasticity estimates obtained in the individual-level analysis. However, using the test described in Newey (1985), one rejects the model's overidentifying restrictions at conventional significance levels (p-values of .060 and .021 in cols. 3 and 4, respectively). Thus the specification in column 5 once again allows the opponent dummies and the variables measuring team quality to directly affect participation, and it uses realized game attendance and the promotional date indicator as instruments for average hourly earnings. The participation elasticity estimate is hardly changed at .65, but now the model's one overidentifying restriction cannot be rejected (p-value = .400). Table 6 also reports hypothesis tests for each 2SLS specification that show that the instruments are highly correlated with the log of average earnings in the first-stage regression.[19] Finally, it is worth noting that the signs and significance of the estimated coefficients on the other explanatory variables generally match the results from the individual-level analysis in table 5.

In summary, the evidence from the aggregate labor supply analysis confirms the conclusion from the individual-level analysis that vendor participation decisions are quite responsive to expected earnings. In particular, the aggregate data suggest that the participation elasticity is between .55 and .65. Perhaps more important, the aggregate analysis also reveals that the estimated participation elasticity is severely downward biased if one ignores the endogeneity of vendor earnings that arises through game-to-game shifts in the vendor labor supply curve. This finding suggests more generally that, at least in analyses of specific labor markets in which unobserved individual labor supply shocks are likely to have an important common component, one must find plausible demand shift instruments to obtain credible estimates of labor supply elasticities.

VI. Conclusion

This paper has analyzed the daily labor supply behavior of stadium vendors and has found that vendor participation decisions are quite responsive to expected earnings. The evidence from both individual-level and aggregate analyses suggests that the elasticity of participation with respect to hourly earnings is around .6. These estimates are always highly significant, although the confidence intervals are rather wide at conventional significance levels. At the same time, aggregate estimates of the participation elasticity that ignore the endogeneity of vendor earnings are dramatically downward biased. This finding suggests that day-to-day shifts in the aggregate labor supply curve of vendors, driven by common shocks to opportunity costs of participation, are an important source of wage variation across games. It also highlights the importance of finding demand shifters to use as instruments in labor supply analyses, especially in

[19] As one would expect, the (unreported) first-stage regressions always reveal a strong positive relationship between the log of average hourly earnings per vendor and the demand shift instruments (game attendance or the promotional date indicator).

studies of specific labor markets in which unobserved individual supply shocks might be expected to be highly correlated.

While I believe that I have provided strong evidence that vendors are more likely to work when expected rewards are high, the study has a few limitations that deserve mention. First, the labor market for stadium vendors is just one particular labor market, and given the youth and independent contractor status of the participants and the decidedly part-time nature of the employment, it is a rather unusual one at that. Thus one should be cautious in generalizing the results found here to other labor markets.

Second, the analysis presented here has ignored the effort margin of vendor labor supply and therefore probably does not capture the total labor supply response of vendors to daily wage changes. For example, if the supply of effort conditional on participation also responds positively to product demand conditions, then the overall labor supply elasticity exceeds the participation elasticity measured here.[20] Since the observed game-to-game variation in aggregate participation does not come close to equalizing the wage across games, effort incentives probably do vary across games. However, since a higher wage per unit of effort has a direct impact on earnings in addition to the indirect effect operating through effort choice, the effort elasticity can be identified only by imposing an (arbitrary) assumption about the form of the production function that maps observable demand conditions and unobservable effort into earnings.

As for its place in the labor supply literature, this paper adds to a small and somewhat disparate set of studies that use high-frequency data on labor supply and wages to try to learn something about how work effort responds to transitory wage changes. For example, Carrington (1996) analyzes quarterly data on industry employment, hours, and earnings in Alaska during the period of construction of the Alaskan oil pipeline and concludes that labor was supplied quite elastically on both the extensive and intensive margins. Treble (1996), exploiting data on a one-time temporary (2-week duration) change in the piece rates paid to miners working at a particular British coal mine in the 1890s, estimates a large and positive elasticity of miners' output with respect to the wage, which implies a positive effort elasticity. Finally, as already discussed,

[20] A positive effort elasticity could help explain why the subcontractor does not maintain a much larger vendor pool. A larger vendor pool would increase both recruiting costs and nonvendor labor costs (e.g., more cashier hours need to be hired if more vendors work at games) but would not increase aggregate vendor sales revenue by much if higher aggregate participation induced a sufficient reduction in each vendor's effort level for any given state of demand.

Camerer et al. (1997) study the daily hours decisions of cabdrivers and, surprisingly, find negative wage elasticities of hours worked.

While these studies and the present one are similar in their use of high-frequency data, they do not all focus on the same dimension of labor supply. As a result, the elasticities estimated in the different papers are conceptually distinct and are not directly comparable. Nevertheless, the present paper delivers the general message that estimated labor supply elasticities for specific labor markets can be seriously biased and therefore must be interpreted very cautiously, unless demand shift variables are used to instrument for the endogenous wage. Because the dynamic labor supply literature has focused mainly on hours decisions, an analysis of high-frequency individual-level data from a labor market with *both* an important hours margin on the supply side and large and observable shocks on the demand side could make a valuable contribution.

References

Altonji, Joseph G. "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *J.P.E.* 94, no. 3, pt. 2 (June 1986): S176–S215.

Browning, Martin J.; Deaton, Angus; and Irish, Margaret. "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle." *Econometrica* 53 (May 1985): 503–43.

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

Card, David. "Intertemporal Labor Supply: An Assessment." Working Paper no. 269. Princeton, N.J.: Princeton Univ., Indus. Relations Sec., 1990.

Carrington, William J. "The Alaskan Labor Market during the Pipeline Era." *J.P.E.* 104 (February 1996): 186–218.

Chamberlain, Gary. "Analysis of Covariance with Qualitative Data." *Rev. Econ. Studies* 47 (January 1980): 225–38.

Heckman, James J. "The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Simple Estimator for Such Models." *Ann. Econ. and Soc. Measurement* 5 (Fall 1976): 475–92.

Lazear, Edward P. "Salaries and Piece Rates." *J. Bus.* 59 (July 1986): 405–31.

Lee, Lung-Fei. "Unionism and Wage Rates: A Simultaneous Equations Model with Qualitative and Limited Dependent Variables." *Internat. Econ. Rev.* 19 (June 1978): 415–33.

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

———. "Interpreting Empirical Models of Labor Supply in an Intertemporal Framework with Uncertainty." In *Longitudinal Analysis of Labor Market Data*, edited by James J. Heckman and Burton Singer. Cambridge: Cambridge Univ. Press, 1985.

Murphy, Kevin M., and Topel, Robert H. "Estimation and Inference in

Two-Step Econometric Models." *J. Bus. and Econ. Statis.* 3 (October 1985): 370–79.

Newey, Whitney K. "Generalized Method of Moments Specification Testing." *J. Econometrics* 29 (September 1985): 229–56.

Oettinger, Gerald S. "Daily Labor Supply of Stadium Vendors: An Empirical Analysis of Participation and Effort Decisions." Working Paper no. 383. Princeton, N.J.: Princeton Univ., Indus. Relations Sec., 1997.

Treble, John G. "Intertemporal Substitution of Effort: Some Empirical Evidence." Manuscript. Bangor, U.K.: Univ. Wales, Dept. Econ., 1996.

U.S. Department of Treasury. Internal Revenue Service. "Independent Contractor or Employee? Training Manuals." Washington: Dept. Treasury, 1996. Web link: http://www.irs.ustreas.gov/prod/bus_info/training.html

Willis, Robert J., and Rosen, Sherwin. "Education and Self-Selection." *J.P.E.* 87, no. 5, pt. 2 (October 1979): S7–S36.