A Semiparametric Life Cycle Labor Supply Model with Non-Additive Fixed Effects *

Elizabeth Schroeder†

July 2010

Abstract

This paper estimates the Frisch elasticity of labor supply, which represents the intertemporal elasticity of substitution. Estimation of this elasticity has previously required assuming that utility is either separable between consumption and leisure or quasi-homothetic with respect to leisure. These restrictions are required to generate hours equations in which individual effects, representing the marginal utility of wealth, enter additively and can be differenced out. Using PSID data, I relax this assumption for a sample of prime-age men. I estimate a semiparametric labor supply equation, using a control function to control for the individual effects. This strategy allows fixed effects to be both non-additive and correlated with the regressors. The average structural function and average partial effects of wages on hours are identified and estimated.

JEL Classification: J22
Keywords:

---

*

†Department of Economics, Georgetown University, Washington DC 20057 USA. eas38@georgetown.edu

1 Introduction

The intertemporal substitution elasticity is used in many macroeconomic models and is essential to the analysis of tax and benefit policies. This elasticity is often estimated with household data using marginal utility of wealth-constant labor supply functions. The wage elasticity in these equations is known as the Frisch elasticity, and measures the change in hours of labor supplied over time in response to receiving different wages at different periods of the life-cycle. It is interpreted as the response of labor supply to anticipated changes in the wage.

Frisch labor supply equations are typically derived from utility functions that are separable in consumption and leisure, which generate labor supply equations that are additive in marginal utility of wealth. At a minimum, estimation of Frisch, or "lambda-constant", labor supply equations has previously required utility to be quasi-homothetic, in which case marginal utility of can be treated as additive even if preferences are not separable. These assumptions are made so that marginal utility, which is unobserved, can be treated as a fixed effect. The assumption that utility is separable between consumption and leisure is widely recognized in the literature as unlikely to be true. Quasi-homotheticity is also a strong and potentially misleading assumption. If these assumptions are false, estimates of the Frisch elasticity used for policy analysis could be severely biased.

I examine the severity of this bias by estimating a life-cycle labor supply equation for married men that allows for a more general form of utility, and thus does not require additive fixed effects. Specifically, I allow the fixed effect to enter the hours equation in an unspecified, nonlinear, way. The hours equation is estimated using the double-index semiparametric least squares estimator of Ichimura and Lee (1988). A control function is employed to account for the fixed effect, which contains the marginal utility of wealth. This strategy allows the individual effects to be both non-additive and correlated with other variables.

2 Literature

Modern life-cycle labor supply estimation began with Heckman and MaCurdy (1980) and MaCurdy (1981 and 1983). These papers formulate the labor supply decision in a given period as a function of current state variables, including wages and household characteristics. An individual solves the following problem.

$$V(a_{it}, t) = \max\left[U\left(c_{it}, h_{it}\right) + \beta E_t V\left(a_{i,t+1}, t+1\right)\right] \tag{1}$$

$$\text{s.t. } a_{i,t+1} = \left(1 + r_{t+1}\right)\left(a_{it} + w_{it}h_{it} - c_{it}\right) \tag{2}$$

The first order conditions, assuming an interior solution for consumption, are:

$$U_c(c_{it}, h_{it}) = \lambda_{it} \tag{3}$$

$$U_h(c_{it}, h_{it}) \geq \lambda_{it} w_{it} \tag{4}$$

where λ_{it} is the marginal utility of wealth, $\frac{\delta V}{\delta a_{it}}$. Demands for quantities of goods and leisure can then be written as functions of current prices and the marginal utility of wealth. These demands are known as Frisch demands. A Frisch labor supply equation takes the following form.

$$h_{it} = f_{it}(w_{it}, \lambda_{it}) \tag{5}$$

Here, $f_{it}()$ may be a function of preference-shifting variables such as household or individual characteristics. Using this equation to estimate labor supply elasticities has two benefits. First, equation (5) does not include consumption, and so can be estimated without consumption data. Second, past and future realizations of wages and any preference variables enter the hours decision only through their effect on current marginal utility of wealth. Marginal utility is unobserved, but the solution to the agent's optimization problem keeps expected marginal utility constant over the life cycle. Assuming rational expectations and perfect capital markets, marginal utility evolves according to the following Euler equation.

$$\lambda_{it} = E[\beta(1 + r_{t+1})\lambda_{i,t+1}] \tag{6}$$

Estimation typically employs a log-approximation of the Euler equation, which breaks λ_{it} into distinct components.

$$\ln \lambda_{it} = \mu_t + \ln \lambda_{i0} + \varepsilon_{it} \tag{7}$$

Individuals determine their marginal utility of wealth at the beginning of the life cycle, setting λ_{i0}. Marginal utility in each subsequent period differs from this initial level by a time effect μ_t, which is a function of the common discount rate and interest rate, and an idiosyncratic forecast error ε_{it}. The Frisch labor supply equation can now be estimated, given proper treatment of λ_{i0}.

In order to handle the unobserved marginal utility of wealth, past studies impose a labor supply function of the following form (Browning 1986).

$$g(h_{it}) = \psi_{it}(w_t) + \phi_i(\lambda_t) \tag{8}$$

The estimate of $\frac{\partial h_{it}}{\partial w_{it}}$ gives an estimate of the Frisch elasticity, the effect of a change in wages holding λ constant. If the function $\phi_i()$ is the natural log, equation (7) can be substituted in for the last term, and the time-invariant individual effect, $\ln \lambda_{i0}$, can be treated as a fixed effect.

The obvious benefit of this framework is that accounting for the fixed effect controls for the influence of all past and future time periods on the current hours choice. The cost, however, is that generating a labor supply equation in which the marginal utility term is additive or log-additive requires restrictions on preferences. A common strategy follows Heckman and MaCurdy (1980) and Macurdy (1981), and is summarized by Blundell and MaCurdy (1999). A log specification is generated by a utility function that is separable in consumption and labor.

$$U_{it} = g(c_{it}, Z_{it}) + \exp(-Z_{it}\rho - v_{it})(h_{it})^{\sigma} \tag{9}$$

Which gives the first order condition:

$$\ln h_{it} = \frac{1}{1-\sigma}\left(\ln w_{it} + \rho Z_{it} + \ln \lambda_{it} - \ln \sigma + v_{it}\right) \tag{10}$$

$$\ln h_{it} = \delta \ln w_{it} + \alpha_{i0} + \mu_t + \beta Z_{it} + e_{it} \tag{11}$$

where $\alpha_{i,t} = \delta(\ln \lambda_{i,t} - \ln \sigma)$ and α_{i0}, the individual effect, comes from substituting in the updating process for $\lambda_{i,t}$.[1] The individual effect contains time zero marginal utility of wealth and is thus theoretically correlated with w_{it} and Z_{it}. Since wages in time t affect wealth, and the preference variables contained in Z affect utility, the wages and preference

[1] Here, $\delta = \frac{1}{\sigma-1}$, $\beta = \delta\rho$, and $e_{it} = \delta v_{it}$.

variables in all time periods can be expected to be correlated with the marginal utility of wealth. The marginal utility term is therefore treated as a fixed effect, and the hours equation is estimated in first differences.

$$\Delta \ln h_{it} = \delta \Delta \ln w_{it} + \beta \Delta Z_{it} + \theta_t + \Delta e_t \tag{12}$$

The estimate of δ is an estimate of the Frisch elasticity.

MaCurdy (1981) estimates an equation of this form and finds the Frisch elasticity to be 0.23, which is reduced to 0.10 when time dummies are included to control for the interest rate effects μ_t. Since then, estimates have tended to fall within or close to this range, include those of Altonji (1986) and Ham (1986), despite different specifications for preferences and different instruments used to control for remaining time-varying endogeneity of the wage.

There is ample evidence, however, rejecting intratemporal separability between consumption and leisure, which is assumed by both MaCurdy and Altonji. Altonji estimates an equation similar to (12), but then tests the separability assumption by adding terms for cross substitution between consumption and hours. He concludes that the assumption of separability is unlikely to be true. Browning and Meghir (1991) devise a methodology for testing for weak separability. Using a system of conditional demand functions for household commodities in the UK Family Expenditure Survey, they test whether these demands depend on labor supply. The authors find that variation in labor supply variables is important in explaining variation in budget shares. They conclude that separability of demand for goods from both hours of work and labor force participation is rejected, for both males and females. Blundell, Browning and Meghir (1993) use this dataset to extend the idea of conditioning on labor supply to the marginal utility of wealth-constant framework. They specify a set of preferences that allow them to exploit results on two-stage budgeting

(Gorman 1959), first estimating within-period preferences, and then aggregating cohorts and estimating intertemporal preferences. They find that labor market variables have significant effects on consumption growth, which is a further rejection of separability.

Blundell, Fry and Meghir (1990) show that relaxing additive separability in a log-hours labor supply equation can only be done by imposing homothetic preferences. An alternative specification is found in Browning, Deaton and Irish (1985). The authors relax additive separability by defining an individual's profit function and deriving the corresponding duel problem to utility maximization. Here, the transformation for g() in equation (5) is linear. The authors show, however, that treatment of λ or $\ln(\lambda)$ as additive in the hours equations implies intra-period quasi-homotheticity. (This point is also discussed in Browning (1986) and Nickell (1988)). Preferences of this type restrict hours of work and expenditures to be linearly related to within-period full income. Browning, Deaton and Irish estimate this model using a pseudo panel created with cohort means. While the elasticity is not parameterized as it is in the MaCurdy-type specification, they find the intertemporal elasticity at the mean of hours to be around 0.4 when allowing nonseparability.

Blundell, Fry and Meghir discuss the limitations of quasi-homothetic preferences. As expenditures increase, preferences become linear Leontief, implying that the rich have zero within-period substitution effects. In addition, the intertemporal elasticity of substitution tends to zero as expenditures rise. The authors conclude that these may be strong restrictions to impose, and a high price to pay for relaxing separability between consumption and leisure.

An alternative approach to estimating the intertemporal elasticity of substitution is to parameterize utility in such a way that preference parameters can be estimated in two stages. First, within-period preferences can be estimated using the first order conditions for consumption and labor. Next, a suitably parameterized intertemporal Euler equation can be estimated to recover intertemporal parameters. This approach requires specification of intertemporal preferences, but has the advantage of removing the restrictions on utility discussed above, as it does not require marginal utility to be additive in the first order condition for hours. It does require data on consumption, however, which can be difficult to obtain at the individual level, since purchases are often made at the household level. The two-stage approach also requires dealing with the endogeneity of consumption in the hours decision.

I relax the restrictions on utility that are needed to estimate an additive fixed effects Frisch labor supply equation, but do not require data on consumption. I estimate a log-hours equation in which the marginal utility term is allowed to enter the equation in an unspecified manner, allowing for nonseparability and nonhomotheticity. I identify and estimate the average structural function, which gives the expected value of the hours function at a given X, averaged over the marginal utility of wealth.

This methodology also has the advantage of allowing the Frisch elasticity to vary at different points in the data. Models that imply a constant Frisch elasticity are typically rejected, and the literature has found significant differences in elasticity among wealth quantiles. For example, using a two-stage approach similar to that discussed above, Ziliak and Kniesner (1999) find that the Frisch elasticity rises with wealth, so that the hours response to a wage change is about 40% higher for the wealthiest quartile of men than for the poorest quartile. The authors conclude that examining only average elasticities obscures the distributional effects of tax policy. I therefore also estimate the distribution of the wage elasticity, both averaged over the individual effect and unconditionally.

3 Estimation and identification

I relax the assumptions that estimating equation (8) imposes on utility by allowing for an arbitrary relationship between the marginal utility of wealth and the observed variables that determine labor supply. This is achieved by estimating an hours equation that allows the individual effect to enter in an unspecified way. Denote the observed variables $X_{it} = [\ln w_{it}, Z_{it,}\mu_t]$. Let $\eta_{it} = [\lambda_{i0}, \varepsilon_{it}]$. The equation of interest is

$$\ln h_{it} = h^*(X_{it}, \eta_{it}) + e_{it} \tag{13}$$

Here, e_{it} is a zero-mean error term, assumed to be uncorrelated with X and η. The function $h^*()$ is unspecified, and allows for interactions between its two arguments. Life-cycle theory tells us that the marginal utility of wealth is a function of an individual's wages and other characteristics in every period of the life cycle. The unobserved λ_{i0} is therefore correlated with the variables in X_{it}, and correlation between the two arguments of $h^*()$ must be taken into account. The objective is to estimate structural effects using equation (13). In order to identify the effects of changes in the X variables on hours, I make the following three assumptions. Maurer, Klein and Vella (2007) apply a similar approach in a semiparametric binary choice model that is estimated by maximum likelihood.

Define X_i as individual i's realization of X for all time periods; $X_i = X_{i,1}$, $X_{i,2}$,..., $X_{i,T}$

.

Assumption 1:

$h_{i,t}$ depends on X_i and the error term only through contemporaneous components.

$$h_{i,t} \perp X_i, \eta_i | X_{i,t}, \eta_{i,t}$$

This assumption follows directly from life-cycle theory and is the driving intuition behind the Frisch labor supply equation. Past and future wages and taste-shifters affect labor supply for individual i at time t only by changing the value of marginal utility of wealth. After controlling for the unobserved heterogeneity, λ_{i0}, the variables in X_{it} and the idiosyncratic forecast error ε_{it} have no effect on hours in other time periods.

Assumption 2:

There exists a single index $X_{i,t}\beta$ such that $h_{i,t}$ and $X_{i,t}$ are conditionally independent given $X_{i,t}\beta$ and η_{it}

$$h_{it} \perp X_{it} | X_{it}\beta, \eta_{it}$$

This assumption is a dimensionality reduction, or index restriction, which states that the effect of a change in X_{it} on h_{it} can be summarized through a single index. An index restriction is not required theoretically, but is required for the function to be well identified using a reasonable sample size of data. Note that wages and household characteristics have been included in the same index. Ideally, one might estimate a model with three indices, to allow arbitrary interactions of the wage, the characteristics in Z, and the individual effect. This is not feasible given the assumptions needed for the current estimator, however. Thus I assume that the relationship between the individual effect and the X variables can be summarized by the index restriction.

The remaining barrier to estimating the hours equation is that a conventional orthogonality condition is violated. As described above, theory indicates that X_{it} is correlated with the unobserved marginal utility of wealth effect, λ_{i0}. A control function is therefore employed to restore the desired orthogonality conditions. Once the marginal utility of wealth is controlled for, $X_{i,t}$ is uncorrelated with the remaining error term. The control function assumption is stated as follows.

Assumption 3:

There exists a control function V_i such that η_{it} and $X_{i,t}$ are conditionally independent given V_i

$$\eta_{it} \perp X_{i,t}|V_i$$

This assumption allows for identification of what are known in the semiparametric literature as structural effects. By requiring $X_{i,t}$ and η_{it} to move separately in the data, conditional on the control function, the effects of a change in X while holding λ constant can be identified.

The choice of an appropriate control function is guided by the theory of life-cycle labor supply. Marginal utility of wealth depends on wages, taste-shifters, and any non-wage income that contributes to wealth, in all periods of an individual's life. A linear combination, specifically the average, of these observed variables is employed here to control for the part of the error term that is correlated with X_{it}. Conditioning on this function, X_{it} and λ_{i0} are independent.

Let V_i be a vector of the time means of each X variable for individual i, as well as the mean of household income other than his own wage, for the individual over time. Other household income provides a natural exclusion restriction, entering the control function but not $X_{i,t}$. Income obviously affects wealth, and thus marginal utility, but does not enter the hours equation once marginal utility has been controlled for. In addition, age is left out of the control function, assuming that the men in the sample have the same expected life span and thus the same average age over time. The control function and hours equation to be estimated become:

$$\widetilde{X}_i = \frac{1}{T}\sum_{t=1}^{T}\widetilde{X}_{it}$$

$$V_i = \widetilde{X}_i\gamma \tag{14}$$

$$h_{i,t} = h(X_{i,t}\beta, \widetilde{X}_i\gamma) + e_{i,t} \tag{15}$$

This approach is similar to Chamberlain's (1982) "correlated random effects" model, in which the dependence of the individual effect on the X variables is modeled as a combination of past and future Xs. It is also related to the identification strategy of Altonji and Matzkin (2005), who use additional external variables as controls in a nonlinear panel data model. The approach used here is much more practical to implement, however, because of the index restrictions.

The final estimator is the double-index semiparametric least squares estimator of Ichimura and Lee (1988). Let $\theta = [\beta \quad \gamma]$. The estimate of the parameters is

$$\widehat{\theta} = \min_{\theta}\sum_{i,t}\tau_{it}\left(h_{it} - \widehat{E}\left[h|X_{it}\beta, \widetilde{X}_i\gamma\right]\right)^2 \tag{16}$$

Here, $\widehat{E}$ is a nonparametric expectation. The indices are orthogonalized, then the joint density is estimated as the product of two normal kernels. Local smoothing is used as a bias-reduction technique, following Klein and Vella (2006). They find that using local smoothing, rather than Ichimura and Lee's suggestion of higher-order kernels, significantly improves the finite sample performance of the double-index estimator. A trimming function, τ_{it}, is employed, placing zero weight on observations that have index values below the fifth or above the 95th percentile of the distributions.

The semiparametric estimation described above allows estimation of the average structural function (ASF) suggested by Blundell and Powell (2000, 2003). The ASF describes how the structural function, $h^*(X_{it}\beta, \eta_{it})$, averaged over the unobserved individual heterogeneity $\eta_{i,t}$, depends on X. This is an important object to estimate in the present application, as the relationship between the structural index, $X_{it}\beta$, and hours of work depends on the marginal utility of wealth. For example, individuals with a lower level of wealth (and thus higher marginal utility), might be less responsive to a change in the index variables if they need to keep working to maintain a minimum level of income. Households with greater wealth may have the ability to be more flexible when preference variables change. In particular, this means that the Frisch elasticity may depend on on an individual's level of marginal utility.

At a fixed realization of X, X_0, the ASF is defined as

$$\mu(X_0) = \int h^*(X_0\beta, \eta_{it}) dF_{\eta_{it}} \tag{17}$$

This gives the average value of the hours function at X_0, with the average taken over the marginal density of the individual heterogeneity. Employing the control function assumption, this expression becomes.

$$\begin{aligned} \mu(X_0) &= \int\int h^*(X_0\beta, \eta_{it}) dF_{\eta_{it}|\overline{X}\gamma} dF_{\overline{X}\gamma} \\ &= \int h(X_0\beta, \overline{X}\gamma) dF_{\overline{X}\gamma} \end{aligned} \tag{18}$$

The estimates of the index parameters, $\widehat{\beta}$ and $\widehat{\gamma}$, and the function $\widehat{h}()$ above, allow computation of the predicted $\widehat{h}$ at any X_{it} and $\widetilde{X}_i$ combination. The average structural function is computed at a given X_0 as

$$\widehat{\mu}(X_0) = \frac{1}{N}\sum_{i=1}^{N}\widehat{h}(X_0\widehat{\beta}, \overline{X_i}\gamma) \tag{19}$$

The average is taken over the marginal distribution of the estimated control function.

The estimation above also allows computation of the average partial effects (APE), which give the change in hours with respect to a change in X, averaged over the marginal distribution of the individual effect. In particular, denote the wage elasticity at a given X_0 as h_w, and the APE as $\delta(X_0)$.

$$\begin{aligned} h_w(X_0\beta, \eta_{it}) &= \frac{\partial h(X_0\beta, \eta_{it})}{\partial w} \\ \delta(X_0) &= E_\eta\left[h_w(X_0\beta, \eta_{it})\right] \end{aligned} \tag{20}$$

Given assumption 3, the control function assumption, this expression can be rewritten as

$$\delta(X_0) = E_{\overline{X_i}\gamma}\left[h_w(X_0\beta, \overline{X_i}\widehat{\gamma})\right] \tag{21}$$

(Wooldridge 2002). This function gives estimates of the Frisch elasticity for different values of the structural index, averaged over the marginal utility of wealth. To estimate $\delta(X_0)$, first the wage elasticity is estimated for each individual, at different values of the control function index, by a local linear regression. The wage elasticity at each combination of indices is computed as $\frac{\partial h(X_0\widehat{\beta}, \overline{X_i}\widehat{\gamma})}{\partial\left(X_0\widehat{\beta}\right)} * \widehat{\beta}^{wage}$. Next, the estimated APE of the wage for a given X_0 is the average of these elasticities over the control function.

$$\widehat{\delta}(X_0) = \frac{1}{N}\sum_{i=1}^{N}\widehat{h}_w(X_0\widehat{\beta}, \overline{X_i}\widehat{\gamma}) \tag{22}$$

Given the interest in the literature in the variation of the Frisch elasticity with respect to different levels of wealth, I also compute unconditional wage elasticities. Using the local linear regression estimates of the Frisch elasticity for each individual, without averaging out the individual effect, I examine variation in the elasticity with respect to the control function. This provides an illustration of how different values of marginal utility of wealth impact an individual's responsiveness to changes in the wage.

4 Results

The data are from the Michigan Panel Study of Income Dynamics (PSID) from the years 1984 to 1994. The sample was chosen to most closely match the samples used in the standard papers on life-cycle labor supply, and therefore includes prime-age men, aged $25 - 55$, who were employed during each period in the sample. The hours variable used is an individual's annual hours of work. Characteristics in X include marital status, self-reported health status on a scale of 1 to 5, the numbers of total children and young children (under age 6) present in the household, age, age-squared, education, and an interaction between age and education.

Although the control function accounts for endogeneity of the wage due to time-invariant heterogeneity, there is still an important potential source of correlation between the wage term and the error. The typical labor-supply equation uses a measure of hourly earnings that is computed by dividing labor income by the number of hours worked. If hours of work are measured with error, a negative correlation is induced between the measurement error of hours and the measurement error of wages (see Altonji 1986). To avoid this problem, I use data only for workers who report an hourly wage rate. This strategy has the disadvantage of limiting the sample to workers who earn an hourly wage, rather than an annual salary. As these workers are not likely to be a random sample of employees, the results and conclusions below can only be interpreted as statements about prime-age men who work for an hourly wage.

The first column of Table 1 presents the results of estimating the hours equation by OLS. The wage coefficient is -0.011 and significantly different from zero. The third column presents the results of fixed effects estimation, which provides an estimate of the standard Frisch labor supply model, controlling for unobserved marginal utility of wealth with the fixed effect. The signs and significance levels of several of the coefficients change, indicating that the individual effects are in fact correlated with the regressors. The wage coefficient is still negative, but smaller in magnitude than the OLS coefficient and not significantly different from zero. Life-cycle theory predicts that the intertemporal elasticity of substitution, captured here by the wage coefficient, must be positive. Thus the simple fixed effects model seems to fail in this dataset.

Figure 1 presents the estimate of the ASF. The ASF is upward sloping over most of the support of the structural index. It increases from a value of 7.58 to a value of 7.95. The dependent variable is log hours, so these values correspond to annual hours of work ranging from 1,939 to 2,853. Since an increase in the structural index is associated with an increase in the ASF, variables that increase the index can be interpreted as increasing the expected number of hours of work, averaged over the distribution of the marginal utility of wealth.

Table 2 presents the parameter estimates of the structural index. Semiparametric estimates using kernels are identified only up to location and scale. A constant is therefore excluded from each index. The coefficient on age is normalized to one in the structural index, and the coefficient on the average wage is normalized to one in the control function index. Given this normalization, the remaining coefficients can be interpreted in relative terms. All variables have been standardized to have mean zero and standard deviation of one. The log wage, education, number of children, and marital status all enter the index with positive coefficients, and have t-statistics greater than 2. A one standard deviation increase in education, however, has a five times greater impact on the index than a one standard deviation increase in the log wage. A standard deviation increase in the number of children present impacts the index twice as much as education, and a standard deviation increase in marital status contributes by far the most to the index. The coefficient on self-reported health status is also positive, but not significantly different from zero. In contrast to the number of children, which increases the index, the number of young children impacts the index negatively, with a standard deviation increase in each leading to about the same magnitude of impact on hours. The interaction between age and education is also negative and significant, indicating a drop-off to the impact of education as the individual ages. Several of the time dummy variables are significant as well.

Table 3 presents the parameter estimates for the control function index. The mean of income, the mean numbers of children and young children, and mean health status all impact the index negatively and have coefficient estimates that are strongly significantly different from zero. Only the mean of marital status enters with the opposite sign. Standard deviation changes in the means of health and number of children have the greatest impact, with about the same magnitude. The number of young children has about 75% of the impact of the overall number of children, and in this case both move the index in the same direction. The mean of education has the smallest coefficient, and is not significantly different from zero.

The Frisch elasticity is estimated as the Average Partial Effect of the wage on the structural function, and is presented in Figure 2. The estimated elasticity ranges from -0.025 up to 0.0042. While starting out negative and initially increasing, it remains relatively flat over most of the support of the structural index. At the median of the structural index, the Frisch elasticity is estimated to be 0.0005826. This number is lower than MaCurdy's estimate of 0.1, but generally in keeping with the findings in panel data that the Frisch elasticity is positive but small. The semiparametric model is an improvement over the fixed effects model, in that the wage elasticity is positive, as predicted by life-cycle theory. This result suggests that the control function is successfully capturing the impact of the unobserved marginal utility of wealth term, and that allowing the individual heterogeneity to enter the hours equation non-additively significantly improves the estimation.

While there is not a great deal of variation in the wage elasticity with respect to the structural index, it is also instructive to see how it varies with the control function. Figure 3 presents the estimates of the wage elasticity over the support of the control function index, as described above. The wage elasticity decreases as the index increases. The control function captures the individual's marginal utility of wealth, so wealth is decreasing as the index increases. The interpretation of Figure 3 is therefore that wealthier individuals have a greater Frisch elasticity, indicating that they are more flexible in responding to changes in the wage. This result is in agreement with the findings of Ziliak and Kniesner, who also document a higher Frisch elasticity for wealthier men.

5 Conclusion

This paper contributes to the literature on life-cycle labor supply by estimating a Frisch labor supply equation without requiring additive fixed effects. Allowing fixed effects to enter the hours equation nonlinearly allows for a non-separable, non-quasihomothetic utility function. The Frisch elasticity is found to be positive, at 0.00058, and significantly different from zero. This value is not outside the range of estimates found in the existing literature, although it is quite close to zero. Prime-age men who work for an hourly wage are therefore found to respond very little to changes in the wage when making labor supply decisions. The difference between the linear fixed effects estimate and the semiparametric estimate suggests that an hours equation with additive fixed effects represents a misspecification of the labor supply function. This result can be interpreted as a rejection of the assumptions on utility that are necessary to generate a labor supply equation with additive fixed effects.

In addition to parameter estimates, the shape of the hours function and its wage derivatives are explored. The Average Structural Function and Average Partial Effects are identified and estimated, with the ASF found to be increasing in the structural index while the APE remains relatively flat. In addition, the wage elasticity is found to be decreasing in the marginal utility of wealth.

References

[1] Altonji, Joseph G. (1986): "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *Journal of Political Economy*, Vol. 94, No. 3, Part 2, pp. S176 - S215.

[2] Altonji, J. and R. Matzkin (2005): "Cross Section and Panel Data Estimators for Nonseparable Models with Endogenous Regressors." *Econometrica* 73, 1053-1102.

[3] Altug, S. and R. Miller (1990): "Household Choices in Equilibrium." *Econometrica* Vol. 58, No. 3. pp. 543-570.

[4] Blundell, R., M. Browning and C. Meghir (1994). "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." The Review of Economic Studies, Vol. 61, No. 1 (Jan., 1994), pp. 57-80

[5] Blundel, R., V. Fry and C. Meghir. (1990) "Preference Restrictions in Life-Cycle Models." Microeconometrics: Surveys and Applications. Ed. Florens, J. M. Ivaldi, J Laffont, and F. Laisney. Basil Blackwell.

[6] Blundell and T. MaCurdy (1999): "Labor Supply: A Review of Alternate Approaches," Handbook of Labor Economics, Volume 3.

[7] Blundel, R., and C. Meghir. (1991) "The Effects of Male and Female Labor Supply on Commodity Demands." *Econometrica* Vol.59, No. 4, pp. 925-951

[8] Blundell, R. and J. Powell (2003): "Endogeneity in Nonparametric and Semiparametric Regression Models." in M. Dewatripont, L. Hansen and S.Turnovsky (eds.): Advances in Economics and Econometrics: Theory and Applications, Eigth World Congress, Vol. II, 312-357.

[9] Browning, Martin (1986). "The Costs of Using Frisch Demand Functions That Are Additive in the Marginal Uility of Expenditure." Economic Letters 21. pp 205-207.

[10] Browning, M., A. Deaton and M. Irish (1985): "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle." *Econometrica* Vol. 53, No. 3. pp. 503-544.

[11] Chamberlain, G. (1982): "Multivariate Regression Models for Panel Data." *Journal of Econometrics* 18, 5 - 46.

[12] Heckman, J. and T. MaCurdy (1980): "A Life Cycle Model of Female Labor Supply." *The Review of Economic Studies,* Vol. 47, No.1, Econometrics Issue. pp. 47 - 74.

[13] Ichimura, H. and L.F. Lee (1991): "Semiparametric least squares (SLS) and weighted SLS estimation of multiple index models: Single equation estimation", in Nonparametric and Semiparametric Methods in Econometrics and Statistics, ed. W. Barnett, J. Powell and G. Tauchen, Cambridge University Press.

[14] Klein, R. and F. Vella (2006): "A Semiparametric Model for Binary Response and Continuous Outcomes Under Index Heteroscedasticity." IZA working paper 2383.

[15] MaCurdy, T. (1981). "An Empirical Model of Labor Supply in a Life Cycle Setting." *Journal of Political Economy* 89, 1059–1085.

[16] —— (1983): "A Simple Scheme for Estimating an Intertemporal Model of Labor Supply and Consumption in the Presence of Taxes and Uncertainty." *International Economic Review*, Vol. 24, No. 2. pp. 265-289.

[17] Maurer, Jurgen, Roger Klein and Francis Vella (2006). "Health, Demographics and Retirement: Evidence from a Semiparametric Binary Choice Model with Nonadditive Correlated Fixed Effects."

[18] Wooldridge, Jeffrey M. *Econometric Analysis of Cross Section and Panel Data.* MIT Press, 2002.

Table 1: OLS estimation.

	OLS	t-stat	Fixed Effects	t-stat
ln wage	-0.011	(-1.020)	-0.007	(-0.350)
education	-0.032	(-2.740)	0.012	(0.660)
kids	0.009	(2.620)	0.003	(0.550)
youngkid	-0.018	(-2.010)	-0.003	(-0.390)
age	-0.029	(-3.780)	0.002	(0.210)
age2	0.000	(2.990)	0.000	(0.220)
age*ed	0.001	(2.950)	0.000	(-0.720)
health	-0.008	(-1.950)	0.005	(1.160)
married	0.044	(3.740)	0.035	(2.270)
constant	8.400	(42.220)	7.543	(29.500)

Table 2: structural index.

	index coefficient	t-stat
ln wage	0.291	(2.492)
education	1.586	(2.670)
children	3.280	(4.050)
young children	-3.629	(-3.463)
age squared	-0.070	(-0.449)
age*education	-7.075	(-3.193)
health	0.261	(1.439)
married	13.265	(3.257)
year 1985	0.660	(1.665)
year 1986	1.233	(3.076)
year 1987	2.756	(4.605)
year 1988	3.722	(4.624)
year 1989	4.422	(4.422)
year1990	5.398	(4.660)
year 1991	6.689	(4.402)
year 1992	7.496	(4.138)
year 1993	0.729	(1.572)
year 1994	2.466	(5.199)

Table 3: control function index

	index coefficient	t-stat
mean income	-0.771	(-7.088)
mean education	-0.142	(-0.828)
mean children	-2.381	(-9.588)
mean young children	-1.839	(-13.569)
mean health	-2.644	(-12.271)
mean married	1.168	(12.490)

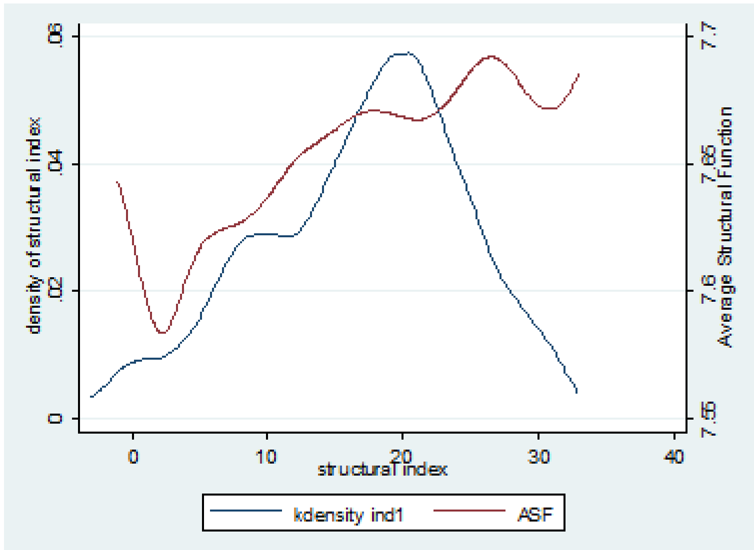

Figure 1

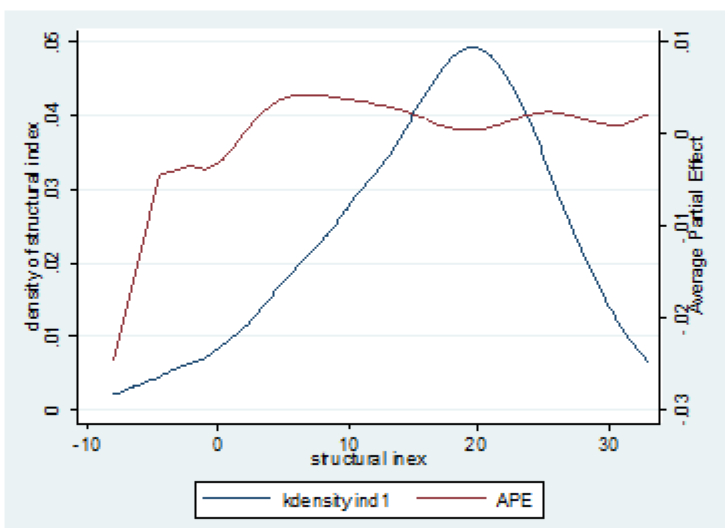

Figure 2

Figure 3

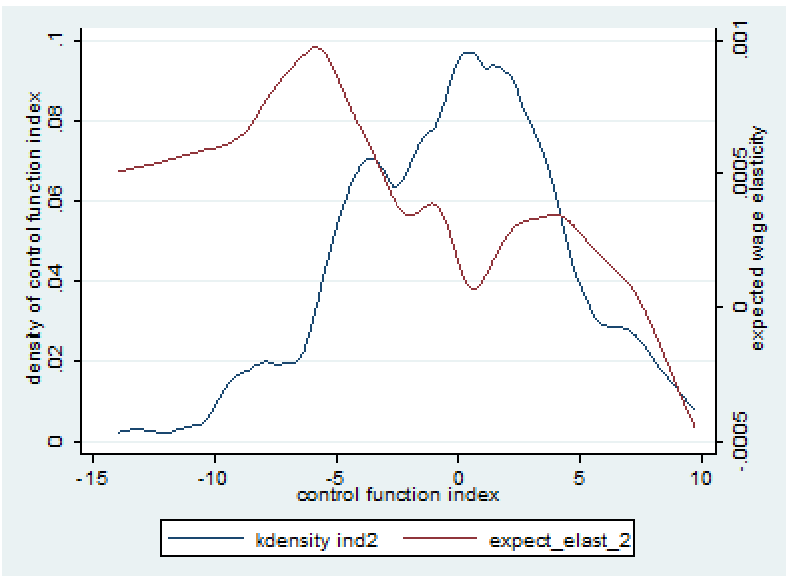

density of control function index
0
.02
.04
.06
.08
.1
-15
-10
-5
0
5
10
control function index
.001
.0005
0
-.0005
expected wage elasticity
kdensity ind2
expect_elast_2