Revisiting the Role of Home Production in Life-Cycle Labor Supply

December 2008

R. Jason Faberman
Federal Reserve Bank of Philadelphia*

Abstract

This paper revisits the argument, posed by Rupert, Rogerson, and Wright (2000), that estimates of the intertemporal elasticity of labor supply that do not account for home production are biased downwards. I use the American Time Use Survey, a richer and more comprehensive data source than those used previously, to replicate their analysis and explore how other factors interact with household and market work hours to affect the labor supply elasticity. In a straightforward replication of their research, I find an elasticity of about 0.4, somewhat larger than previously estimated. In a richer cohort analysis that accounts for demographics, household characteristics, and the endogeneity of wages, however, I obtain an elasticity estimate that is essentially zero.

JEL Codes: D13, D91, J22
Keywords: household production, intertemporal elasticity of labor supply, life cycle time use

---

* Ten Independence Mall, Philadelphia, PA, 19106. Jason.Faberman@phil.frb.org. I thank Randy Wright for comments on this draft and Nathan Brownback for his research assistance. The views presented here are my own and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1. Introduction

The intertemporal elasticity of labor supply plays a critical role in business cycle theory, yet there exists still a debate on its magnitude over the long run. Early real business cycle models, such as that of Kydland and Prescott (1982), postulate a relatively high elasticity, while micro labor studies of supply decisions over individuals' life cycles (MaCurdy, 1981; Altonji, 1986) find elasticity estimates that are positive but economically small. Rupert, Rogerson, and Wright (1995, 2000) argue that the estimates from these micro studies are biased downwards because they neglect the role of household production in the labor supply decision and its behavior over the life cycle. When they include time spent on non-market work in their analysis, they find labor supply elasticities that are considerably larger, though still lower than those posited by macroeconomists.

This paper revisits the role of household production in the labor supply decision over the life cycle. It does so using data from the American Time Use Survey (ATUS). Previous studies generally appealed to either individual panel data (e.g., MaCurdy, 1981; Altonji, 1986) or time use data (Rupert, Rogerson, and Wright, 2000, henceforth RRW). The panel data do not have information on time spent outside of market work, while the earlier time use data come from relatively small surveys with limited labor market information. The ATUS, on the other hand, is considerably larger and since it is drawn from the Current Population Survey (CPS), it contains a wealth of labor market information.

I replicate the RRW analysis using the ATUS data from 2003 through 2007, which involves estimating the elasticity of labor supply using data on male workers

aggregated into synthetic age cohorts. I obtain estimates that are slightly higher than those found by Rupert, Rogerson, and Wright. Including household work hours in their study produces labor supply elasticity estimates between 0.22 and 0.34. Including household work hours with the ATUS data produces estimates between 0.34 and 0.42. In every specification, the inclusion of household work hours produces a considerably higher estimate than the regressions where they are ignored.

The RRW analysis is not ideal, however. For one, they consider the supply of market and non-market labor as perfect substitutes, and consequently estimate an elasticity of the supply of both market and non-market labor. In practice, though, it only the elasticity of *market* labor supply that is of macroeconomic importance. In addition, like most of the labor supply studies that precede them, they focus on only male workers. Finally, their wage measures come from an outside data source, introducing the potential for measurement error. The first issue is just a matter of specification, and I find that it leads to only a small reduction in the estimated elasticity. The latter two issues are driven primarily by data constraints. Using the ATUS relaxes these constraints. I first disaggregate the data into various demographic categories (gender, marital status, education, race) and estimate the elasticity separately for each group. I then disaggregate the age cohort data into finer demographic detail and estimate the elasticity through a fixed effects approach. The first approach produces similar estimates (about 0.3) across all groups. The second approach however, yields significantly lower estimates. Including household work when using the disaggregated cohorts produces elasticity estimates that are about 0.10 to 0.13.

Finally, I explore how much household characteristics and the endogeneity of wages may affect the elasticity estimates. One might think that at least some substitution between home work and market work over the life cycle is driven by household characteristics, such as the number of children in the household. One might also think that the biases that arise from endogenous wages at the individual level may not wash out at the aggregate level. For example, the accumulation of human capital over the life cycle may alter the entire life cycle wage path. I replicate the analysis for the age cohorts and the more disaggregated cohorts and include the number of household children and the real wage and hours worked of any cohabitating partners. Doing so reduces the effect of including household work hours on the estimated labor supply elasticity further. Using my preferred specification, the estimate using only age cohorts falls from 0.41 to 0.22, and the estimate using the finer cohorts falls from 0.13 to a statistically insignificant 0.06, which is the same estimate one obtains if they ignore both household work and the household characteristics. I address the endogeneity of wages by instrumenting them with the locational information of the individual (state of residence, whether they live in an urban area). It is well known that wages vary substantially across regions for a variety of reasons (cost of living, local amenities, etc.), but there is no reason to believe that hours should exhibit geographic variation, at least on the intensive margin. I aggregate the first stage wage estimates from this IV approach into age cohorts and the more disaggregated cohorts, and replicate the previous analyses. When accounting for all factors, (household work hours, household characteristics, and endogenous wages), I obtain labor supply elasticity estimates of about 0.04, about the same as those without the IV specification, and from an economic standpoint, essentially zero.

I interpret these results as suggesting that, as RRW argue, accounting for hours worked at home is an important part of estimating the intertemporal labor supply elasticity over the life cycle, but other factors, such as the number of household children (which determine, in part, the amount of household work required) and the labor market experiences of other household members (which can act as a substitute for an individual's own labor market participation), are at least as important. Based on my estimates, these latter factors completely mitigate the effect substitution between market work and home work has on estimates of the labor supply elasticity. Accounting for both household work and household characteristics produces an elasticity estimate that is essentially zero. Accounting for the endogeneity of wages does not alter the result.

The next section illustrates a standard model of an individual's labor supply decision that includes home production. The following section describes the data and presents some basic evidence on wages and time use. Section 4 presents the empirical analysis, and Section 5 concludes.

2. Model

2.A. Labor Supply Theory

I motivate my analysis with a standard model of lifetime labor supply. The model appears in various forms in numerous studies of the intertemporal elasticity of labor supply, including the RRW study.[1] In the model, individuals maximize lifetime utility subject to an intertemporal budget constraint. Like RRW, I assume that consumers receive utility from a composite consumption good, c_{it}, and household goods and

[1] The seminal examples include Lucas and Rapping (1969), Ghez and Becker (1975),MaCurdy (1981), and Altonji (1986).

services, Z_{it}. The consumption good is purchased in the market and the household goods and services (whose required amount is determined exogenously) can either be purchased in the market or produced at home. At age t, individual i allocates her time between market work, h_{it}^{m}, which earns a real wage w_{it}, household work, h_{it}^{H}, and leisure. The individual solves

$$\max_{c_{it}, x_{it}, h_{it}^{m}, h_{it}^{H}} \sum_{t=1}^{T} \beta^{t} \left[u(c_{it}, Z_{it}) - v(h_{it}^{m}, h_{it}^{H}) \right],$$

subject to

i.) $$\sum_{t=1}^{T} (1+r)^{-t} \left(c_{it} + p_t x_{it} \right) = A_0 + \sum_{t=1}^{T} (1+r)^{-t} w_{it} h_{it}^{m} ,$$

ii.) $H \geq h_{it}^{m} + h_{it}^{H} \geq 0$, and

iii.) $Z_{it} = x_{it} + g(h_{it}^{H})$,

where x_{it} is the amount of household services purchased in the market, p_t is their price, and A_0 is the initial level of assets. The total per period time endowment is H, the discount rate is β, and the real interest rate is r. As is typical in the literature, I assume that utility is additively separable in hours and consumption.

Given this formulation and assuming an interior solution, the first order conditions are

(1.1) $$\beta^{t} U_{c,it} = (1+r)^{-t} \lambda_{it} ,$$

(1.2) $$\beta^{t} U_{Z,it} = (1+r)^{-t} \lambda_{it} p_t ,$$

(1.3) $$\beta^{t} v_{m,it} = (1+r)^{-t} \lambda_{it} w_{it} , \text{ and}$$

(1.4) $$\beta^{t} \left(g' U_{Z,it} - v_{H,it} \right) = 0 .$$

Equation (1.3) usually forms the basis for estimating the intertemporal elasticity of labor supply over the life cycle. For example, if as an RRW, one lets $v(\cdot) = \phi\left(h_{it}^{m} + h_{it}^{H}\right)^{\gamma}$, one can express (1.3) in logs as

(2) $$(\gamma - 1)\ln\left(h_{it}^{m} + h_{it}^{H}\right) = \ln\lambda_{it} - \ln\gamma\phi - t\ln\beta(1+r) + \ln w_{it},$$

where λ_{it} is the Lagrange multiplier for the utility maximization problem, which represents the marginal utility of wealth. The intertemporal elasticity of labor supply in this case would be $1/(\gamma - 1)$, which could theoretically be estimated through (2) using OLS. In the empirical section, I appeal to a more general form of $v(\cdot) = \phi h_{it}^{m\,\gamma_1} h_{it}^{H\,\gamma_2}$, which allows for imperfect substitution between market work and household work. The first order condition now becomes

(3) $$(\gamma_1 - 1)\ln h_{it}^{m} = \ln\lambda_{it} - \ln\phi\gamma - t\ln\beta(1+r) + \ln w_{it} - \gamma_2 \ln h_{it}^{H}.$$

Again, one can recover the labor supply elasticity, now measured as $1/(\gamma_1 - 1)$ and clearly interpreted as an elasticity of *market* labor supply, from (3) using OLS. The key point of the RRW study is that the omission of hours of household work from the estimating equation will introduce a downward bias in the estimate of the labor supply elasticity.

2.B. Empirical Considerations

Even if one correctly includes a measure of household work hours in the estimation of (3), other issues make it difficult to obtain consistent estimates of the elasticity using available data sources without strong assumptions about the evolution and aggregation of wages, and the marginal utility of wealth. The first issue is that while the standard labor supply model takes wages as given, in reality, they are endogenous. In

addition, the accumulation of human capital over an individual's lifetime will likely affect the entire wage profile over the life cycle, as Imai and Keane (2004) note. The second issue is the endogeneity and unobservability of the marginal utility of wealth, which, given the life cycle behavior of asset accumulation, will not only vary systematically across individuals but also likely change over the life cycle.

There have generally been two empirical approaches to estimating the intertemporal labor supply elasticity. The first, used prominently by MaCurdy (1981), Altonji (1986), and more recently by Imai and Keane (2004) and Chang and Kim (2006), involves appealing to longitudinal household microdata, such as the Panel Study of Income Dynamics (PSID) or National Longitudinal Survey of Youth (NLSY). The use of longitudinal data is appealing because it allows a first differencing of the data and hence an estimation of the labor supply elasticity that accounts for unobservable individual characteristics. In addition, these data include plausibly valid proxies for the marginal utility of wealth, such as food expenditures (PSID) or assets (NLSY).

The data have shortcomings, though. For one, earnings and work hours are measured at an annual frequency. As Altonji (1986) and others acknowledge, the ratio of reported annual earnings to annual hours produces a wage measure that suffers from some degree of measurement error. Since this wage measure uses the dependent variable from (3) in its denominator, any measurement error in hours will be negatively correlated with the measurement error in the wage, creating a downward bias in the elasticity estimate. The proxies for the marginal utility of wealth may also be inadequate. Recent work by Aguiar and Hurst (2005) suggests that food expenditures may be a poor proxy for consumption in this setting because individuals tend to substitute time for money in

food consumption over the life cycle.[2] The asset measure in the NLSY79 is a suitable measure of wealth, but annual data is only available for individuals up to 36 years of age in the NLSY79, forcing Imai and Keane (2004), who use the data in their study, to use simulated data for later years of the life cycle. This is problematic because one would expect the wealth effects to matter most in these later years. The biggest issue with the longitudinal data from the perspective of this study, though, is their lack of information on time spent on anything besides market work, including hours of household work.

The second empirical approach involves aggregating cross-sectional household data into synthetic age cohorts and estimating the intertemporal labor supply elasticity using the cohort observations. This was the approach used in early studies by Ghez and Becker (1975) and Smith (1977) and the household production study by Rupert, Rogerson, and Wright (2000). It is also the approach used in this study. The use of age cohorts also has its shortcomings, but given that time use data are not longitudinal, the approach is necessary if one wants to include data on time outside of market work in the analysis. One obvious issue is aggregation, namely whether unobservable characteristics and idiosyncratic shocks at the individual level aggregate in a way that is either mean-zero or uncorrelated across age cohorts. For example, human capital accumulation early in life likely affects the entire wage profile over the life cycle. Another issue is whether the marginal utility of wealth, λ_t, interpreted as the marginal utility for a representative agent when using age cohorts, is constant over the life cycle. The RRW study deals with this issue by assuming complete markets, so that the agents can perfectly smooth their

[2] Specifically, with an additive log utility specification for $u(c,Z)$, equation (1.1) implies that consumption is equal to the inverse of the marginal utility of wealth. With detailed time use and expenditure data, Aguiar and Hurst show that individuals tend to spend more time shopping for bargains and preparing food later in their life cycle, reducing their food expenditures but keeping their food consumption essentially constant.

consumption over the life cycle. If one assumes that $U(c) = \log(c)$, then by equation (1.1), it is straightforward to see that perfectly smooth consumption will lead to a constant λ over the life cycle.

Given the lack of data on wealth or assets in the time use survey, I cannot control for potential life-cycle variations in λ. Consequently, I use many of the same assumptions of the RRW study. Note, though, that the main purpose of this study is to reevaluate the importance of home production in estimating the intertemporal elasticity of labor supply using more robust and exhaustive data. The issues I outline above may affect estimates of the true elasticity the true elasticity, but one can still study how the elasticity changes in relative terms when home production is included. In future work, I hope to explore obtaining an estimate of the labor supply elasticity that accounts for the noted issues.

3. Data

I use the time diary data from the American Time Use Survey (ATUS) produced by the Bureau of Labor Statistics (BLS). I pool the data for all individuals in the survey years 2003 through 2007. The pooled panel contains demographic, labor market, and time use information for 72,922 individuals. This is a large increase in the number of observations used in previous studies, in some cases by an order of magnitude.[3] Individuals are sampled from the outgoing rotation groups of the Consumer Population Survey (CPS). Each respondent keeps a detailed time diary for one full day, and their activities are classified into a wide range of categories. Respondents are also re-interviewed about their current labor market situation, with the interview asking many of

[3] For example, MaCurdy uses 5,130 person-year observations and Altonji uses between 3,269 and 10,036 person-year observations from the PSID, while RRW use either 799 or 1,165 individual observations pooled from three different time use surveys.

the same questions as the standard CPS survey. The survey also collects basic information about the household, and includes the responses form the final month in the CPS from all household members. The significance of this information is that I can use it to get measures of labor market information for other household members.[4]

Respondents report their time diaries for a single day, which can prove problematic for identifying employed individuals. Namely, full-time workers who report on a weekend will likely report zero work hours would incorrectly be identified as non-employed if one were to use the time-use work hours to determine employment status. Luckily, respondents also report their "usual" weekly hours at their current jobs. I use this information to identify employed individuals. This results in a sample of 16,787 males and 18,834 females between the ages of 22 and 62 that have positive "usual" hours and positive reported wages.

Given the potential for day-of-week bias in the response of time use, I aggregate respondents into synthetic weeks prior to my creation of the synthetic cohorts. The aggregation groups respondents into those reporting on a weekday, Saturday, or Sunday, with holidays treated as Sundays. It then calculates the (sample-weighted) mean of the time devoted to each activity by day-of-week group and age cohort and aggregates these means to a weekly total. Even though I identify employed individuals through the usual work hours measure, I use the work hours derived from this synthetic week aggregation in the regression analysis so that they are consistent with the measured time spent in other activities.

[4] The final CPS interviews occur 2 to 5 months prior to the ATUS interview. I note potential measurement issues with this timing below.

The work hours measure that I employ includes all time spent working on the job (including multiple jobs), including time in work-related activities (e.g., business-related outings), and down time at work (e.g., lunch breaks). It does not include time spent commuting or searching for work. I also replicate my analyses using various other definitions of work hours, including just the "core" time spent working at all jobs, a total work time measure that includes commuting and job search, and the total usual work hours measure. All measures produce very similar results. I measure household work hours as the total time spent in house work, child and adult care, pet care, vehicle care, shopping for goods, and purchasing services. They are aggregated into synthetic weeks then age cohorts in the same manner as work hours.

Real wages are measured as total weekly earnings (deflated by the CPI) divided by total usual work hours. Earnings are reported as part of the re-interview information. The measure is a considerable improvement over those used in previous studies. First, the information is up to date and of high enough frequency to be a reliable measure of the current price paid for market work. The longitudinal data used in the earlier studies (MaCurdy, 1981; Altonji, 1986) must rely on annual earnings and hours data, where the propensity for measurement error is high. Second, the wages reported correspond to the individuals reporting the time use data. The RRW study must use wages from an outside source—they match CPS wage data to the time use data at the age cohort level—which is likely also fraught with measurement error.

Table 1 reports the average wage and time spent by individuals aged 22 to 62 in various activities. It reports the estimates for all workers (male and female), and married

male workers.[5] Male workers have higher wages and considerably more hours of market work than female workers, 41.8 to 34.9 but women have more hours in household work, 26.5 to 17.4. The difference in household work is higher for home work, child care and purchasing goods and services. Combined, women perform about two hours more total work. Given that women workers also spend over two hours more with personal care, they also engage in 5.4 hours less leisure and socializing than male workers, 34.4 to 39.8.

Figure 1 shows the behavior of wages, market work, and household work for workers over the life cycle. I show the trends of each using a fourth-order polynomial. Again, the differences between male and female workers are clear. Aside from the obvious differences in levels, there are also differences in the life cycle patterns. The market work hours of male workers increase during their twenties and thirties, while the hours of female workers are mostly flat until their late fifties. Female workers have greater changes in their household work hours over the life cycle, though. Male workers also have a wage profile that rises throughout their life cycle, while the wages of female workers flattens out in their late thirties.

4. Evidence on the Labor Supply Elasticity

4.A. Replication of Previous Findings

The main purpose of this paper is to compare my estimates of the intertemporal labor supply elasticity over the life cycle to earlier estimates when home production is included. As such, the natural starting point is to replicate the earlier research with the ATUS data. I focus on the results from the RRW study, since they are only other analysis

[5] For a more detailed analysis of individual time use patterns over time, see Aguiar and Hurst (2007). I draw my distinctions for types of market and household work in part from their study.

that looks at the role of home production. They pool their time use data and aggregate them into age cohorts. They then estimate the regression in equation (2) for male workers aged 22 to 62 using the synthetic age cohorts as their observations. They use several alternative specifications for the $v(h)$ function, which include $v(h) = \phi h_{it}^{\gamma}$, $v(h) = -\phi\left(112 - h_{it}^{\gamma}\right)$, $v(h) = \phi\left(168 - s_{it} - h_{it}^{\gamma}\right)$, and $v(h) = \phi \exp\left\{\gamma h_{it}\right\}$, where $h_t = h_{it}^{m} + h_{it}^{H}$, $H = 112$ is a measure of total hours that deducts 56 hours per week for sleep, and s_{it} is a measure of actual time reported spent sleeping and on personal care. They estimate their regressions weighting by cohort size and weighting by the variance of work hours.

Table 2 reports the results of my replication of the RRW analysis using the ATUS data alongside the original estimated labor supply elasticities from their study. I only report results weighted by cohort size since the variance-weighted results are nearly identical.[6] In their study, they find that ignoring household work produces labor supply elasticity estimates between 0.09 and 0.13, while accounting for it produces substantially higher elasticities, on the order of 0.22 to 0.34. Using the ATUS data, I find somewhat higher elasticities when ignoring home production for all but the specification that accounts for variations in sleep and personal care. The other three elasticity estimates range from 0.13 to 0.23. When I include household production in the regressions, all four specifications produce higher elasticity estimates than the case when household work is excluded and produce higher elasticity estimates than those found by RRW. The

[6] As RRW do, I also replicate the results for 22-45 year olds. The results are qualitatively similar to those reported in Table 2. In addition, I replicate the results using the usual reported work hours, "core" work hours, and all work time, including commuting and job search. The usual hours measure produces slightly higher estimates, which may be the result of measurement error. The other measures of work hours produce nearly identical estimates to those reported.

elasticity estimates when household work is included range from 0.34 to 0.42. In addition, the estimates imply larger biases from the exclusion of household work than those found by RRW. Including household work in their study increases the point estimate of the elasticity between 10 and 19 basis points, while including household work with the ATUS data increases the point estimate between 17 and 31 basis points.

If anything, the replication of the RRW study using a more robust source of time use data reinforces their finding that accounting for household production is important for consistently estimating the intertemporal elasticity of labor supply. This direct replication of their specifications produces estimates that are near the upper bound of those found in previous studies (e.g., Ghez and Becker, 1975; McCurdy, 1981; and Altonji, 1986) who find estimates ranging between -0.06 and 0.45.

4.B. Accounting for Demographics and Household Characteristics

The data used in earlier studies often did not allow a thorough analysis of the labor supply elasticity by demographics. Analyzing the behavior of women proved particularly difficult. For one thing, the data sets used, whether they were panel data or time use data, were relatively small, and produced very small sample sizes of both male and female workers. In addition, the time periods over which most of these studies occur (the 1960s and 1970s) are when women's labor force participation was substantially lower than it is now. Since the literature uses estimating equations similar to equation (3), derived based on an interior solution to the utility maximization problem, it requires an estimation that uses individuals with positive work hours. Thus, much of the earlier work (Ghez and Becker, 1975, and Smith, 1977, are notable exceptions) focuses on either male workers or married male workers.

Given the number of observations in my sample and the higher labor force participation rates for women in 2003-07, I am able to study the role of various demographics in estimating the labor supply elasticity. I do so in two ways. In the first approach, I split the data by various demographic dimensions and estimate equation (3) using age cohort data for each demographic group. This is informative as to how the inclusion of home production affects various groups, but it does not allow an accounting of variations of other characteristics within each group. Therefore, in my second approach, I disaggregate the age cohorts by various demographic characteristics and estimate the labor supply elasticity using a fixed effects regression that controls for these characteristics.

The results of the first method are in Table 3. I estimate elasticities separately by gender, race, and education. I also include married men and white males, since these subgroups are often used in the literature. The results for male workers are very similar to the ATUS results in Table 2. Including household work hours on the right-hand side of the regression reduces the elasticity estimate only slightly from 0.41 to 0.35. White men have a stronger response to the inclusion of household work hours than all men. The elasticity rises from 0.23 to 0.47. Married men, on the other hand, have a similar change in their estimated elasticity (about a 15 basis point increase in the point estimate), but much lower elasticity estimates overall. Even when household work is included, the I find an elasticity of labor supply for married men of 0.10, which is statistically insignificant.

As one might expect, the intertemporal elasticity of women's labor supply is more sensitive to the inclusion of household work hours, rising by 23 basis points, versus the

13 basis point increase men experience. Overall, however, women have a lower estimated elasticity, 0.25 compared to 0.35 for men. Estimating the elasticity for all workers combined produces an increase of about 22 basis points in the labor supply elasticity, to 0.38.

Turning to the results by race, white workers have estimated elasticities similar to those for all workers. The results for Hispanic workers are similar as well, though they have a smaller increase in the estimated elasticity when including home production (about 10 basis points) and the point estimate is slightly lower, 0.31. The results for black workers show considerably lower elasticities and a smaller effect of including household work. The estimated elasticity only rises 6 basis points to 0.17.

Finally, save for the results for the college educated, splitting the data by education produces similar estimates for all groups. Including household work in the estimation increases the estimated elasticity by 10 to 23 basis points across the four education groupings (high school or less, some college, college, and post-graduate), and for all but the college educated, produces an estimated elasticity between 0.29 and 0.32. For those with a college degree, it is 0.19.

For my second method, I disaggregate the age cohorts two ways. Ideally, I would create one highly disaggregated synthetic cohort panel divided by the various demographic characteristics but the need to aggregate the time use data into a synthetic week within each cell limits my ability to do so. I create cohorts for age-gender-marital status-education groups and age-gender-marital status-race groups. For the first, I use three education groups: high school or less, some college, and college or more, and for the second, I use three race groups: Hispanic, black (non-Hispanic), and all others (which

is predominantly white.) Because of small cell sizes in some categories among the very young (e.g., married individuals with college degrees), I restrict the sample to those aged 25 to 62. These panel regressions are similar to the previous regressions except that they now use $38 \times 2 \times 2 \times 3 = 456$ observations and use fixed effects to control for gender, marital status, and either education or race.

The results are in Table 4. The first two columns list the OLS results for all workers grouped by age cohort using only 25 to 62 year olds. It provides a baseline for evaluating the panel data results. Comparing these results to those for all workers in Table 3 shows that removing the first three years leads to somewhat lower elasticity estimates (about 5 basis points in both specifications), but otherwise produces similar results. The inclusion of household work in the baseline specification produces a labor supply elasticity estimate of 0.33. The next two columns present the results for synthetic cohorts that include education with and without household production hours, respectively. The use of finer cohort detail reduces the elasticity estimates dramatically. The estimate that includes household work is 0.13, and represents only a 7 basis point increase over the case where household work is not included. Using cohorts based on race rather than education produces nearly the same results, with an estimated elasticity of 0.10 when household work is included.

Besides demographics, household composition could also affect the estimation of the labor supply elasticity. Indeed, the model in this paper assumes an exogenous requirement of household goods and services, which the individual must either produce at home or purchase in the market. The number of household children can potentially drive exogenous variations in this requirement over time. While many children are the result of

planning by couples, there often remain uncertainties about the timing and success of an actual birth, not to mention the potential for having more than one child. Thus, for a life cycle study such as this one, it is plausible to assume that the number of household children is exogenous to our estimating equation. In addition, while my model is for a single individual, one can easily envision a model of couples who choose their market and household work hours to jointly maximize utility.

The ATUS includes information on the number of household children younger than 18, $child_{it}$. It also includes a wide array of demographic information on the spouses or unmarried partners who live in the household through their most recent CPS interview. I use this data to create a measure of the spouse/partner's real hourly wage, w_{it}^{s} (in the same way that I create the respondent's real wage), and use their reported total "usual" hours as a measure of their market work hours, $h_{it}^{m,s}$.[7] Unfortunately, since only one respondent per household completes the time diary, I do not have household work hours for the spouse/partner. I include the broader category of spouses and unmarried partners in the household because one need not be married to substitute their market and household work hours with their cohabitant to jointly maximize utility.

Table 5 presents the results that include the number of household children and the wage and work hours of the partner. The first two columns in Table 5 correspond in specification to the first two columns of Table 4, where I only disaggregate the data in age cohorts. Including the household information raises the intertemporal labor supply

[7] The CPS data precede the ATUS interview data by 2-5 months. Thus, there exists the potential for measurement error that arises from potential changes in wages, hours, or labor force status during that time. CPS gross worker flow estimates that about 6 percent of workers leave their job each month (including those who immediate find new work, see Fallick and Fleischmann, 2004), so the issue is non-trivial. Unfortunately, I have no suitable alternative measure of partners' employment behavior.

elasticity by 8 basis points when I exclude household work, most likely because the number of household children acts a proxy for hours of household work. When I add the household information to the specification that includes household work, though, the labor supply elasticity declines by about 10 basis points. I interpret this as evidence that variations in the amount of household production that the number of household children requires over the life cycle accounts for at least some of the relationship between household work hours and the market wage.

The next four columns of Table 5 replicate the finer cohort disaggregation specifications from the third and fourth columns of Table 4, which use the education groups.[8] Columns (3) and (4) are an exact replication of the Table 4 specifications to show that the shift to using all partners and not just spouses has a negligible effect on the results. Columns (5) and (6) replicate the results with the household information. In this case, its inclusion yields a slight drop in the estimated labor supply elasticity. In fact, when both household work hours and the household characteristics are included, the estimated elasticity is no different from what one estimates from the fixed effects specification that includes only the real wage. Statistically, the point estimate (0.06) is only marginally significant, and economically, it is essentially zero.

The results suggest that the inclusion of household work hours clearly addresses a downward bias in estimates of the intertemporal elasticity of labor supply. At the same time, however, it is also clear that demographics and household characteristics matter as well. Theory dictates that household characteristics should matter because the number of children helps determine the amount of home production a household needs, while the

[8] The results by race are very similar, so for the sake of space, I do not report them.

earnings and work hours of an individual's partner provide substitutability alternatives to the individual the same way that household work does.

4.C. Accounting for Endogenous Wages

Without longitudinal household data that include information on either consumption or wealth and time use, I am unable to address all of the econometric concerns outlined in Section 2. For instance, even with the inclusion of demographic, household, and time use information, my estimates must still rely on the strong assumption of a constant marginal utility of wealth (λ) over time. I can, however, address the endogeneity of wages, provided I can obtain a valid instrument from the ATUS data.

I instrument wages with geographic variations in the real wage at the individual level. Wages have large, persistent differences across geographic areas. In addition, there is no reason to believe that an individual's location is correlated with hours of work. One might think that variations in labor laws across U.S. states might make location variables poor instruments, but in the U.S., these variations in these laws generally involve differences in the minimum wage or unemployment insurance compensation, which are relevant primarily for the extensive margin of labor participation. Since my entire analysis is restricted to individuals who work positive hours, these laws are not an issue.

I use data on the state of residence and a dummy variable for whether an individual lives in a metropolitan area, since the existence of an urban wage premium is well-known (e.g., Glaeser and Mare, 2001). The data come from the CPS interview portion of the ATUS data.[9] In the first stage, I regress the real wage of each individual on

[9] As with the spousal information, the geographic data from the CPS are 2-5 months old at the time of the CPS interview. In this case, however, there is no concern about measurement error in the location because

a set of fixed effects for *j* states, the urban dummy variable, and the exogenous regressors used in the second stage. I then aggregate the predicted wage from this regression to the cohort level of interest, and the use the log of the predicted wage in the second stage.

The IV results are in Table 6. I replicate my analysis for age cohorts and for age-gender-partner present-education cohorts for all workers aged 25 to 62.[10] The first four columns report the results for the age cohort data. The first column shows that the IV estimates produce a much larger elasticity (0.57) when one ignores all other factors, though given the large standard errors, the estimate is not statistically significant. The second column shows the estimate when household work hours are included. The elasticity increases substantially to 1.34. When one includes the number of household children and the spousal information, however, the elasticity drops substantially. Including both household work hours and household characteristics produces an elasticity of 0.04, which is considerably lower than the OLS estimate reported in the second column of Table 5 and more in line with the elasticity estimates from the results in Table 5 that use finer cohort disaggregation.

The last four columns of Table 6 present the IV results when using the finer cohorts. In short, instrumenting for wages does little to change these elasticity estimates. In fact, it is striking how little the estimates differ from their counterparts in Table 5. The estimated elasticities range from -0.07 to 0.09, though none of these are statistically significant. When both household work hours and the household characteristics variables

---

the CPS is a household-based, not an individual-based survey. In other words, they only re-interview individuals who lived in the household at the time of both interviews.

[10] In unreported results, I also replicate the age cohort results including the 22-24 year olds. The estimated elasticities are considerably higher, on the order of 30 basis points, across the board. Unfortunately, it is difficult to assess whether this implies a bias in the reported results for two reasons. First, the standard errors on the IV elasticity estimates are huge, so even a 30 basis point difference is not statistically large, and second, the data for 22-24 year olds by finer cohorts is too sparse to examine whether these large differences persist in the more disaggregated data.

are included, it produces essentially the same labor supply elasticity estimate (about 0.04) as the IV specification in the fourth column of Table 6, which uses the age cohorts, and the OLS specification that uses the finer cohorts in the last column of Table 5. Statistically and economically speaking, these estimates are essentially zero. The IV estimates do little to alter the earlier interpretation of the results.

5. Conclusions

This paper revisits the role of home production in the estimation of the intertemporal elasticity of labor supply over the life cycle. Rupert, Rogerson, and Wright (2000) argue that ignoring the home production when estimating this elasticity will lead to a downward bias because changes in the wage over the life cycle may occur concurrent with changes in household work. Using a new, more robust data source on individual time use, I replicate the analysis of Rupert, Rogerson, and Wright and find somewhat higher labor supply elasticity estimates than they do. I then replicate the analysis for various demographic groups, and while the point estimates vary widely, the main finding is that a higher estimated elasticity when household work hours is included persists.

I also explore the role of demographics using more disaggregated synthetic cohorts, which lets me control for demographics through a fixed effects specification. The estimated labor supply elasticity declines substantially. When I include the number of household children and the work hours and wage of an individual's partner, the estimate is essentially zero. Finally, I control for the endogeneity of wages using an instrumental variables approach. This produces estimates that are highly sensitive to the specification used in the age cohort analysis but are generally consistent across specifications in the analysis of the more disaggregated cohorts. Once household work

and household characteristics are accounted for in the instrumental variables approach, I again obtain an elasticity estimate that is essentially zero.

The finding of a labor supply elasticity of zero is consistent with the notion that the income and substitution effects of labor supply roughly cancel out over the life cycle, leading to a flat long run labor supply curve. The finding comes even though I control for hours of household work. I argue that this is because much of the life-cycle relationship between wages and household work are due to changes in household characteristics over the life cycle.

Much of my analysis is made possible by the use of a larger, richer data set than previous studies. The data, however, are still not ideal. First, the data are not longitudinal, so one cannot control for unobserved characteristics that may be correlated with wages and hours, nor can one evaluate the response of hours to changes in the wage at the individual level. Second, the data contain no information on consumption or wealth, which forces me maintain the strong assumption of a constant marginal utility of wealth over the life cycle throughout this study. Future work should explore estimation approaches that can incorporate time use and either wealth or consumption data within a panel data framework.

References

Aguiar, Mark and Hurst, Erik (2005) "Consumption versus Expenditure," *Journal of Political Economy* 113(5): 919-48.

---------- (2007). "Measuring Trends in Leisure: The Allocation of Time Over Five Decades," *Quarterly Journal of Economics* 122(3): 969-1006.

Altonji, Joseph G. (1986). "Intertemporal Substitution in Labor Supply: Evidence from Micro Data." *Journal of Political Economy* 94(3, Part 2): S176-S215.

Chang, Yongsung and Kim, Sun-Bin (2006). "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy," *International Economic Review* 47(1): 1-27.

Fallick, Bruce and Fleischmann, Charles A. (2004). "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Federal Reserve Board of Governors, Finance and Economics Discussion Series Working Paper No. 2004-34.

Ghez, Gilbert and Becker, Gary S. (1975). *The Allocation of Time and Goods Over the Life Cycle*. New York: Columbia University Press (National Bureau of Economic Research).

Glaeser, Edward L. and Mare, David C. (2001). "Cities and Skills," *Journal of Labor Economics* 19(2): 316-42.

Imai, Susumu and Keane, Michael P. (2004). "Intertemporal Labor Supply and Human Capital Accumulation," *International Economic Review* 45(2): 601-641.

Kydland, Finn E. and Prescott, Edward C. (1982). "Time to Build and Aggregate Fluctuations," *Econometrica* 50(6): 1345-70.

Lucas, Robert E, and Rapping, Leonard A. (1969). "Real Wages, Employment, and Inflation," *Journal of Political Economy* 77(5); 721-54.

MaCurdy, Thomas E. (1981) "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy* 89(5): 1059-85.

Rupert, Peter, Rogerson, Richard and Wright, Randall (1995) "Estimating Substitution Elasticities in Household Production Models," *Economic Theory* 6(1): 179-93.

----------, (2000). "Home Work in Labor Economics: Household Production and Intertemporal Substitution," *Journal of Monetary Economics* 46(3): 557-79.

Smith, James P. (1977). “Family Labor Supply Over the Life Cycle,” *Explorations in Economic Research* (National Bureau of Economic Research) 4(2): 205-76.

Table 1. Wages and Time Use Patterns

	All Individuals	All Workers	All Male Workers	All Female Workers	Married Male Workers
Real hourly wage (2003 $)	18.57	18.72	20.44	16.82	22.52
Usual work hours	41.98	41.85	45.00	38.37	46.05
Total reported work & work-related time	30.52	38.60	41.81	35.06	42.77
Core work	30.27	38.38	41.55	34.88	42.51
Work-related time	0.25	0.22	0.26	0.18	0.26
Commuting	2.54	3.27	3.74	2.75	3.94
Job search	0.22	0.06	0.06	0.05	0.04
Total household work	25.30	21.71	17.40	26.47	19.14
Home work (incl. pet and vehicle care)	13.17	11.05	8.90	13.42	9.53
Child and adult care	6.59	5.43	4.26	6.73	5.33
Shopping & purchasing services	5.58	5.23	4.25	6.32	4.29
Leisure & socializing	40.31	37.20	39.77	34.37	38.26
Personal care	6.09	6.11	5.00	7.35	4.78
Sleep	58.72	57.44	56.89	58.05	55.78
N	52,454	35,621	16,787	18,834	10,920

Notes: Estimates are for the listed individuals aged 22-62 in the 2003-2007 surveys of the ATUS. Wages are in 2003 dollars. Time use is hours per week.

Table 2. Estimated Intertemporal Labor Supply Elasticities

Specification	Without Household Work		Including Household Work	
	RRW	ATUS	RRW	ATUS
$y = \ln h$	0.117	0.232	0.274	0.412
	(0.119)	(0.059)	(0.072)	(0.045)
$y = \ln (112 - h)$	0.092	0.125	0.280	0.417
	(0.054)	(0.033)	(0.071)	(0.048)
$y = \ln (168 - s - h)$	0.099	0.027	0.338	0.336
	(0.054)	(0.037)	(0.084)	(0.052)
$y = h$	0.126	0.217	0.222	0.390
	(0.085)	(0.057)	(0.051)	(0.044)

Note: The table reports the regression coefficients from regressing the listed dependent variable on ln w_{it} and age. Regressions without household work use $h_t = h_{it}^m$, while those with household work use $h_t = h_{it}^m + h_{it}^H$. The columns labeled "RRW" report the results from Rupert, Rogerson, and Wright (2000), and the columns labeled "ATUS" report results when using age cohorts created from pooled ATUS data for 2003-07 for working males aged 22-62. All regressions are weighted by the size of the age cohorts. Standard errors are in parentheses.

Table 3. Estimates of the Intertemporal Labor Supply Elasticity, Various Worker Groups

	Men		Women		All Workers		Married Men	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
t	-0.0017	-0.0030	0.0003	-0.0013	-0.0012	-0.0029	0.0006	-0.0026
	(0.0011)	(0.0014)	(0.0008)	(0.0007)	(0.0007)	(0.0007)	(0.0013)	(0.0017)
$\ln w_t$	0.232	0.353	0.015	0.247	0.165	0.380	-0.049	0.098
	(0.058)	(0.097)	(0.051)	(0.061)	(0.040)	(0.066)	(0.083)	(0.096)
$\ln h_t^H$	---	-0.140	---	-0.330	---	-0.267	---	-0.253
		(0.090)		(0.067)		(0.069)		(0.097)
$\bar{R}^2$	0.39	0.41	-0.04	0.36	0.32	0.50	-0.04	0.09
	White Men		**All White**		**All Black**		**All Hispanic**	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
t	-0.0016	-0.0044	-0.0011	-0.0028	0.0018	-0.0001	0.0002	0.0008
	(0.0011)	(0.0015)	(0.0007)	(0.0007)	(0.0018)	(0.0017)	(0.0018)	(0.0015)
$\ln w_t$	0.232	0.467	0.146	0.383	0.110	0.173	0.210	0.305
	(0.060)	(0.097)	(0.042)	(0.071)	(0.110)	(0.096)	(0.115)	(0.099)
$\ln h_t^H$	---	-0.246	---	-0.272	---	-0.389	---	-0.356
		(0.090)		(0.067)		(0.105)		(0.086)
$\bar{R}^2$	0.35	0.45	0.22	0.43	0.06	0.29	0.12	0.38
	High School or Less		**Some College**		**College Degree**		**Postgraduate**	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
t	-0.0013	-0.0017	-0.0005	-0.0015	0.0003	-0.0001	0.0001	-0.0015
	(0.0013)	(0.0012)	(0.0016)	(0.0018)	(0.0010)	(0.0018)	(0.0019)	(0.0023)
$\ln w_t$	0.158	0.294	0.193	0.320	-0.037	0.189	0.198	0.293
	(0.123)	(0.097)	(0.094)	(0.135)	(0.054)	(0.081)	(0.121)	(0.141)
$\ln h_t^H$	---	-0.351	---	-0.131	---	-0.185	---	-0.143
		(0.121)		(0.101)		(0.053)		(0.110)
$\bar{R}^2$	-0.01	0.16	0.18	0.20	-0.04	0.19	0.07	0.08

Note: The table reports the regression coefficients from regressing $\ln h_{it}^m$ on the listed dependent variables as specified in equation (3) in the text. All regressions use age cohorts created from pooled ATUS data for 2003-07 for the listed groups of workers aged 22-62. All regressions are weighted by the size of the age cohorts. Standard errors are in parentheses.

Table 4. Synthetic Cohort Panel Regression Estimates

	Age		Age × Gender × Marital Status × Education		Age × Gender × Marital Status × Race	
	(1)	(2)	(3)	(4)	(5)	(6)
t	-0.0010	-0.0026	0.0004	-0.0005	-0.0012	-0.0029
	(0.0007)	(0.0007)	(0.0006)	(0.0006)	(0.0007)	(0.0007)
$\ln w_t$	0.115	0.327	0.060	0.127	0.165	0.380
	(0.062)	(0.071)	(0.038)	(0.036)	(0.040)	(0.066)
$\ln h_t^H$	---	-0.283	---	-0.232	---	-0.267
		(0.067)		(0.028)		(0.069)
Fixed Effects?	No	No	Yes	Yes	Yes	Yes
R^2	0.09	0.40	0.46	0.54	0.43	0.50
N	38	38	456	456	456	456

Note: The table reports the regression coefficients from regressing $\ln h_{it}^m$ on the listed dependent variables as specified in equation (3) in the text. All regressions use the noted disaggregation of cohorts created from pooled ATUS data for 2003-07 for the listed groups of workers aged 25-62. Fixed effects regressions use dummy variables for either gender, marital status, and education group, or gender, marital status, and race, as appropriate. All regressions are weighted by the size of the age cohorts. Standard errors are in parentheses.

Table 5. Synthetic Cohort Panel Regression Estimates, Controlling for Household Characteristics

	Age		Age × Gender × Partner Present × Education			
	(1)	(2)	(3)	(4)	(5)	(6)
t	-0.0037	-0.0017	0.0002	-0.0005	0.0009	0.0015
	(0.0019)	(0.0016)	(0.0005)	(0.0005)	(0.0007)	(0.0006)
$\ln w_t$	0.195	0.223	0.060	0.120	0.034	0.060
	(0.121)	(0.101)	(0.036)	(0.035)	(0.039)	(0.035)
$\ln h_t^H$	---	-0.327	---	-0.215	---	-0.253
		(0.084)		(0.026)		(0.026)
$\ln child_{it}$	-0.030	0.008	---	---	0.010	0.025
	(0.018)	(0.018)			(0.005)	(0.005)
$\ln w_t^s$	0.046	0.046	---	---	-0.045	-0.031
	(0.068)	(0.056)			(0.033)	(0.030)
$\ln h_t^{m,s}$	-0.009	0.006	---	---	-0.067	-0.055
	(0.111)	(0.092)			(0.050)	(0.045)
Fixed Effects?	No	No	Yes	Yes	Yes	Yes
R^2	0.17	0.45	0.48	0.55	0.50	0.59
N	38	38	456	456	456	456

Note: The table reports the regression coefficients from regressing $\ln h_{it}^m$ on the listed dependent variables as specified in equation (3) in the text. All regressions use the noted disaggregation of cohorts created from pooled ATUS data for 2003-07 for the listed groups of workers aged 25-62. Fixed effects regressions use dummy variables for either gender, presence of partner, and education group. All regressions are weighted by the size of the age cohorts. Standard errors are in parentheses.

Table 6. Synthetic Cohort Panel Regression Estimates, Instrumental Variables Estimation

	Age				Age × Gender × Partner Present × Education			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
t	-0.0064	-0.0148	0.0051	0.0002	0.0001	0.0012	0.0004	0.0017
	(0.0061)	(0.0070)	(0.0085)	(0.0075)	(0.0010)	(0.0009)	(0.0011)	(0.0011)
$\ln \hat{w}_t^{IV}$	0.571	1.338	-0.538	0.044	0.060	-0.069	0.091	0.040
	(0.546)	(0.628)	(0.681)	(0.616)	(0.091)	(0.089)	(0.095)	(0.035)
$\ln h_t^H$	---	-0.162	---	-0.315	---	-0.198	---	-0.250
		(0.072)		(0.093)		(0.026)		(0.026)
$\ln child_{it}$	---	---	0.024	0.023	---	---	0.011	0.028
			(0.027)	(0.024)			(0.005)	(0.005)
$\ln w_t^s$	---	---	0.104	0.077	---	---	-0.039	-0.039
			(0.076)	(0.067)			(0.033)	(0.033)
$\ln h_t^{m,s}$	---	---	-0.030	0.016	---	---	-0.072	-0.072
			(0.119)	(0.105)			(0.050)	(0.050)
Fixed Effects?	No	No	No	No	Yes	Yes	Yes	Yes
R^2	0.03	0.15	0.13	0.36	0.48	0.54	0.50	0.58
N	38	38	38	38	456	456	456	456

Note: The table reports the second stage regression coefficients from the instrumental variables estimation of regressing $\ln h_{it}^m$ on the listed dependent variables as specified in equation (3) in the text. The wage is aggregated from the predicted values obtained from an individual-level first-stage regression that uses state fixed effects and a dummy for living in an urban area as instruments. All regressions use the noted disaggregation of cohorts created from pooled ATUS data for 2003-07 for the listed groups of workers aged 25-62. Fixed effects regressions use dummy variables for either gender, presence of partner, and education group. All regressions are weighted by the size of the age cohorts. Standard errors are in parentheses.

Figure 1. Wages and Hours Worked Over the Life Cycle

(a) All Workers

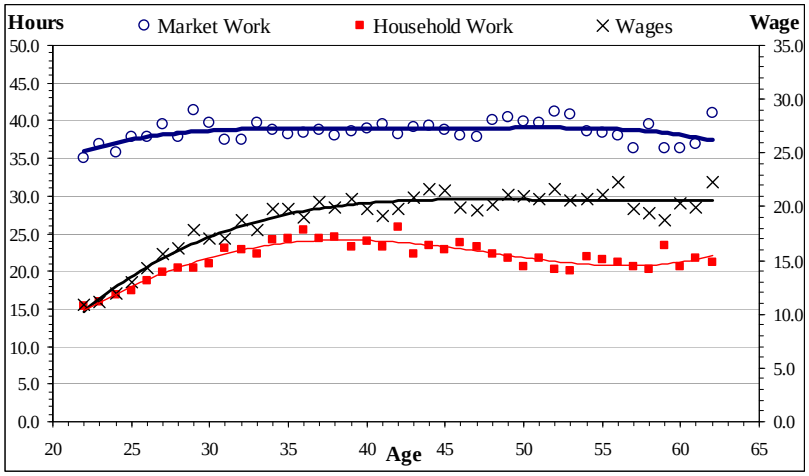

(b) Married Men

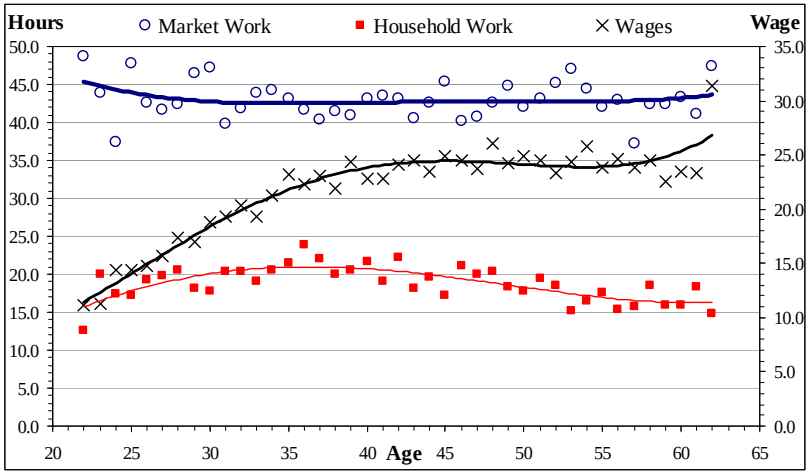

(c) Men

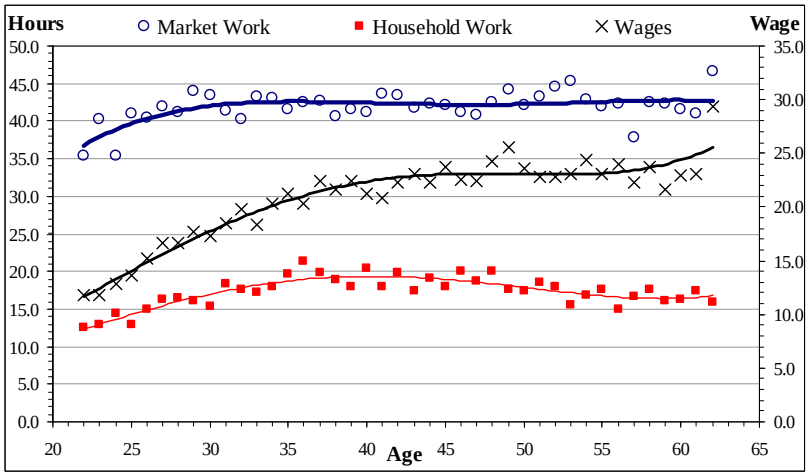

(d) Women

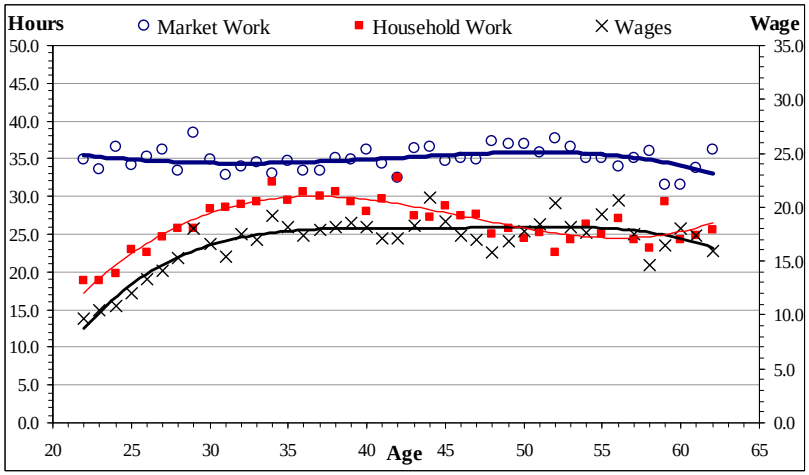

Notes: Each panel illustrates the real wage, total hours spend working and total hours spent on household work for 22 to 62 year olds reporting positive wages and work hours in ATUS data pooled over 2003 to 2007. Trends for each are estimated using a fourth-order polynomial.