



Federal Reserve Bank of Chicago

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

R. Jason Faberman

March 2015

WP 2015-02

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

March 2015

R. Jason Faberman
Federal Reserve Bank of Chicago*

Abstract

This paper examines the role of home production in estimating life-cycle labor supply. I show that, consistent with previous studies, ignoring an individual's time spent on home production when estimating the Frisch elasticity of labor supply biases its estimate downwards. I also show, however, that ignoring other ways a household can satisfy the demand for home production biases its estimate upwards. Changes in this demand over the life-cycle have an income effect on labor supply, but the effect can be mitigated through purchases in the market and through the home production of other household members. When accounting all factors related to home production, I find that the "micro" Frisch elasticity is about 0.4 and the "macro" Frisch elasticity, which accounts for extensive margin adjustments, is about 0.9. If I only account for an individual's own home production effort, I find that the "macro" elasticity is about 1.6.

JEL Codes: D13, D91, J22

Keywords: home production, Frisch elasticity of labor supply, life-cycle time use

* 230 S. LaSalle St., Chicago, IL, 60604. jfaberman@frbchi.org. I thank Yongsung Chang, Eric French, Luojia Hu, Erik Hurst, Derek Neal, Joe Nichols, Peter Rupert, Randy Wright, and participants in several conferences and seminars for their thoughtful comments. I also thank Nathan Brownback, Rebecca Friedman, Jennifer Knudson, and Thomas Walstrum for their excellent research assistance. The views presented here are my own and do not necessarily represent the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1. Introduction

The intertemporal elasticity of substitution, or Frisch elasticity, of labor supply plays a critical role in business cycle theory. There remains a debate, however, on its magnitude. Early business cycle models, such as that of Kydland and Prescott (1982), required a relatively high Frisch elasticity to match the aggregate hours fluctuations observed in the data. Most studies, however, find that the Frisch elasticity is relatively small, typically between 0.0 and 0.5. These include earlier studies by MaCurdy (1981) and Altonji (1986) as well as more recent studies by French (2004, 2005), and Chetty et al. (2011).¹ Some recent studies classify these estimates as a “micro” Frisch elasticity because they only account for changes in the work hours of the employed. These studies argue that the “macro” Frisch elasticity, which accounts for extensive-margin changes into and out of employment over the life cycle, is more relevant for studying aggregate fluctuations. In these studies (Imai and Keane, 2004; Chang and Kim, 2006; Rogerson and Wallenius, 2009, 2011), “macro” Frisch elasticity estimates are much larger, between 1.2 and 3.8, despite a “micro” elasticity that is relatively small and comparable to other findings.²

Previous research has examined the role of home production in intertemporal labor supply. Several, such as Rupert, Rogerson, and Wright (1995, 2000) and Rogerson and Wallenius (2011) focus on the effect of home production on life-cycle labor supply. Others have examined how home production affects labor supply over the business cycle (Benhabib, Rogerson, and

¹ Ghez and Becker (1975), Smith (1977), and Browning, Deaton, and Irish (1995) also find relatively small Frisch elasticities of labor supply. These studies use aggregated data over the life cycle in one way or another, while Altonji, MacCurdy, and French use longitudinal micro data on individuals.

² Unlike the studies focused on the “micro” elasticity, these studies generally estimate the Frisch elasticity using a calibrated life-cycle model matched to moments of the aggregate data. Exceptions include Mulligan (1999), Fioritto and Zanella (2012), and Peterman (2012).

Wright, 1991; McGrattan, Rogerson, and Wright; 1997; Aguiar, Hurst, and Karabarbounis, 2013).

The Rupert, Rogerson, and Wright (2000) study is perhaps most relevant to my own. In it, they argue that previous estimates of the Frisch elasticity are biased downwards because they neglect the role of home production in the life-cycle labor supply decision. They show that accounting for time spent in home production over the life cycle can essentially double one's estimate of the Frisch elasticity.

In this paper, I revisit the role of home production in the labor supply decision over the life cycle. I use the data from the American Time Use Survey (ATUS), which includes detailed time diary data for individuals as well as the extensive questions on their labor force and demographic characteristics that are present in the Current Population Survey (CPS). The ATUS is much larger and more representative than most micro-level studies of labor supply and it is the largest by far of any study that examines the role of home production. I aggregate the data into synthetic cohorts and estimate the Frisch elasticity using panel data methods and an instrumental variables analysis. Earlier studies (e.g., MaCurdy, 1981; Altonji, 1986) used longitudinal data on individuals, such as the Panel Study of Income and Dynamics, but these data do not contain information on home production. The RRW study uses aggregate data pooled by age, which cannot account for cohort effects in life-cycle behavior. The synthetic cohort approach, which has been used in several life cycle studies of both labor supply and consumption, allows me to employ the panel data techniques on the ATUS despite its lack of longitudinal data.³ It also allows me to account for extensive-margin labor adjustments and

³ Browning, Deaton, and Irish (1985), Mulligan (1999), and Peterman (2012) all use the synthetic cohort approach to estimate the Frisch elasticity. Deaton (1985), Attanasio and Weber (1993), Attanasio and Browning (1995), and Attanasio et al. (1999) use a synthetic cohort analysis and methods similar to what I employ here to study life-cycle consumption.

estimate both a “micro” and a “macro” Frisch elasticity from the same data. Finally, I extend the standard model of life-cycle labor supply to include not only a role for one’s own home production, but also a role for the home production of a spouse (if present) and expenditures on goods and services that could potentially be produced at home. The synthetic cohort approach allows me to incorporate spousal home production using a “synthetic spouse” estimated from the time use data and expenditures on potentially home-produced goods and services using similarly aggregated data from the Consumer Expenditure Survey (CEX).

The extended life-model illustrates that the latter two measures are particularly important for estimating the Frisch elasticity. Changes in the demand for home production over the life cycle have an income effect on labor supply. For example, increases in the number of household children will induce an individual to work more (because of the income effect) but also induce the individual to spend more time on home production (because of the increase in its demand). The household can satisfy the increase in the demand for home production through other means, however, essentially mitigating the income effect on labor supply. Specifically, other household members could meet the increase in demand with their own production effort, or the household could purchase the required goods and services in the market. To the extent that ignoring an individual’s home production time will bias the Frisch elasticity of labor supply downward, ignoring these other margins will bias the Frisch elasticity upwards. To what degree it does so is an empirical question.

I use ATUS data from 2003 to 2012 and estimate the Frisch elasticity using different specifications within a synthetic cohort analysis. I begin by studying the wage and hours changes of the employed over the life cycle, which generates of what is often referred to as the “micro” Frisch elasticity. Ignoring a role for home production produces a Frisch elasticity estimate of

0.16, but that estimate rises to 0.44 when accounting for home production. Accounting for the home production of a spouse and for expenditures on household goods and services reduces the estimate but only slightly, to 0.39. These estimates are within the higher end of what the literature has found for the “micro” elasticity. An exact replication of the Rupert, Rogerson and Wright (2000, henceforth RRW) analysis produces essentially the same result, with slightly higher elasticity estimates than the original RRW study.

I estimate a “macro” Frisch elasticity by including both the employed and non-employed in the construction of the synthetic cohorts. I account for selection into non-employment by using the wage of individuals who were employed at the time of their last CPS interview but non-employed at the time of their ATUS interview, which occurs two to five months after the last CPS interview. Using this approach, I estimate a Frisch elasticity of 0.68 when excluding home production and 1.58 when including one’s own home production. This is on the low end of what previous studies have found using calibrated macro models, but still well above one. It also more than doubles the already-large “macro” elasticity of the case where home production is excluded. When accounting for household expenditures and the home production of the spouse, the estimate falls to 0.92, which is more in line with estimates found by Chetty et al. (2011). Estimates from first-differenced regressions and from separate regressions for men and women produce quantitatively similar patterns but lower overall elasticity estimates. In these cases, the fully-specified model produces “micro” elasticity estimates between 0.04 and 0.33 and “macro” elasticity estimates between 0.51 and 0.73.

My results confirm that, as has been argued previously, accounting for hours worked at home is important for estimating the Frisch elasticity of labor supply in a life-cycle setting. Ignoring home production produces a downward bias because it ignores the fact that hours

worked responds to both wage changes and changes in home production hours, which are substitutes for both work hours and leisure. At the same time, I show that accounting for all of the ways a household can respond to changes in the demand for home production is important as well. This is especially true when one considers extensive-margin labor adjustments. A household's ability to accommodate the change in demand through either expenditures or through the home production of other household members has a relatively large effect on the participation margin of labor supply. The model highlights the mechanisms driving this result. Changes in household composition, such as an increase in the number of children in the household, act like a preference shock to household utility. Individuals demand more household goods and services. I show that this acts like a reduction in their "full income," which makes the individual want to work more. The extent of the effect depends on the ratio of the wage to the market price of these goods and services, since this ratio determines how much of the demand will be satisfied by home production and how much will be purchased in the market. This occurs even if markets are complete, an important caveat since many life-cycle studies rely on the complete-markets assumption to avoid various empirical complications. The key insight of this paper is that the home production hours of a spouse and the ability to purchase home-produced goods in the market dampens one's own labor supply response to a change in the demand for home production. The large effects that I find for extensive-margin adjustments are not too surprising, given how well known it is that the presence of young children has a large effect on female labor supply. The model adds the insight that the female labor supply response will have an effect on male labor supply as well. Ultimately, I show that the net effect of all the ways a household can respond to changes in the demand for home production generates a "macro" Frisch elasticity estimate that is considerably larger than the "micro" elasticity estimate, but still smaller than what is required by most business cycle models to match the observed fluctuations

in aggregate hours. The results suggest that accounting for the complete household response to changes in the demand for home production is important for understanding for how labor supply responds to wages at the aggregate level.

2. Model

2.A. Labor Supply Theory

I motivate my analysis with a standard model of lifetime labor supply, augmented to incorporate a demand and technology for home production. The model appears in various forms in numerous studies of the intertemporal elasticity of labor supply.⁴ In the model, an individual i maximizes lifetime utility subject to an intertemporal budget constraint. Like the RRW study, I assume that consumers at time t receive utility from a composite consumption good, c_{it} , and household goods and services, z_{it} . The consumption good must be purchased in the market at a numeraire price, while household goods and services can either be purchased in the market in an amount x_{it} at a price p_t or produced at home. The amount of household goods required at any point in time depends on a preference shifter for the demand for household goods and services, η_{it} . Exogenous shocks to η_{it} over the life cycle will cause the demand for z_{it} to vary over time. One can think of changes in η_{it} over time as changes in household characteristics and composition, such as the number of children in the household. Changes in these characteristics over the life cycle are captured in a similar manner in a variety of life-cycle models of consumption and labor supply (e.g., Attanasio and Weber, 1993; Attanasio et al., 1999). The representative household member allocates her time between market work, n_{it} , which earns a

⁴ The seminal examples include Lucas and Rapping (1969), Ghez and Becker (1975), MaCurdy (1981), and Altonji (1986). Rupert, Rogerson, and Wright (2000) use a similar model and extend it to include home production.

real wage w_{it} , home production, h_{it} , and leisure. I allow for the presence of a spouse, who chooses his work hours to yield earnings of y_{it}^S and supplies h_{it}^S hours of home production for the household. These choices are exogenous to the utility maximization problem of the individual under study but factor into her optimization decision since they affect how she will satisfy her demand for z_{it} .⁵ The household member then solves

$$\max_{c_{it}, x_{it}, n_{it}, h_{it}} \sum_{t=1}^T \beta^t [u(c_{it}, z_{it}; \eta_{it}) - v(n_{it}, h_{it})]$$

subject to

$$\text{i.)} \quad \sum_{t=1}^T (1+r)^{-t} [c_{it} + p_t x_{it}] = A_0 + \sum_{t=1}^T (1+r)^{-t} [w_{it} n_{it} + y_{it}^S], \quad \text{and}$$

$$\text{ii.)} \quad 0 \leq n_{it} + h_{it} \leq H, \quad \text{with}$$

$$z_{it} = x_{it} + g(h_{it} + h_{it}^S).$$

Here, A_0 is the initial level of assets, β is the discount factor, r is the interest rate, and the total per-period time endowment of the individual is H . Household goods and services are produced at home by both the individual and her spouse using the production function $g(h)$, with $g' > 0$, $g'' < 0$ and $g(0) = 0$. The home production of the spouse and individual are perfect

⁵ Note that the model described is equivalent to a joint household maximization problem where a planner maximizes utility by choosing the consumption, work hours, and home production hours of each individual, as well as the amount of household goods and services purchased in the market, subject to a household budget constraint. The equivalence holds because the preference shocks, η_{it} , and the resulting demand for z_{it} that must be satisfied is at the household level, and the amount of purchased household goods, x_{it} , is the same in both cases. Within-household differences in η_{it} would break the equivalence. Empirically, any gender differences in the utility function (such as a comparative advantage in, or disutility from, home production) will be accounted for with either fixed effects for gender or separate regressions by gender.

substitutes. Similarly, the purchased household goods and services and the amount of these goods and services that come from home production are also perfect substitutes. 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

$$(1.1) \quad \beta^t u_{c,it} = (1+r)^{-t} \lambda,$$

$$(1.2) \quad \beta^t u_{z,it} = (1+r)^{-t} \lambda p_t,$$

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

$$(1.4) \quad g' u_{z,it} = v_{h,it}.$$

The Lagrange multiplier on the budget constraint is λ , which represents the marginal utility of wealth. Equation (1.3) usually forms the basis for estimating the Frisch elasticity. For example, assume that $v(\cdot) = \phi n_{it}^{\gamma_1} h_{it}^{\gamma_2}$, where ϕ is the disutility of (any type of) work and γ_1 and γ_2 allow for imperfect substitution between market work and home production. One can then express (1.3) in logs as

$$(2) \quad (\gamma_1 - 1) \ln n_{it} = \ln \lambda - \ln \phi \gamma_1 - t \ln \beta (1+r) + \ln w_{it} - \gamma_2 \ln h_{it}.$$

The Frisch elasticity in this case would be $1/(\gamma_1 - 1)$. One could theoretically estimate (2) using OLS, though for reasons I discuss below, one needs to use an instrumental variables approach in practice. As RRW argue, the key insight from (2) is that the omission of hours of home production from the estimating equation will introduce a (downward) omitted variable bias in the estimated Frisch elasticity.

2.B. Theoretical Considerations

Even if one correctly includes a measure of household work hours in the estimation of (2), there are theoretical issues that imply that an OLS estimation of (2) would yield inconsistent estimates of the Frisch elasticity. Most importantly, variations in other omitted variables will affect the (endogenous) behavior of h_{it} and therefore lead to an inconsistent estimate. This will be true even if markets are complete, an assumption that ensures that the marginal utility of wealth, λ , is constant over time. From (1.2) above, however, it is clear that changes in η_{it} over the life cycle will jointly affect the behavior of n_{it} and h_{it} .

To see this explicitly, assume that $u(c_{it}, z_{it}; \eta_{it}) = \ln c_{it} + \eta_{it} \ln z_{it}$. Then, (1.2) implies that $\eta_{it}/p_t z_{it} = [\beta(1+r)]^{-t} \lambda$. One can substitute this into (2) and combine it with the fact that $z_{it} = x_{it} + g(h_{it} + h_{it}^S)$ to yield

$$(3) \quad (\gamma_1 - 1) \ln n_{it} = \ln \eta_{it} - \ln(p_t \phi \gamma) + \ln w_{it} - \gamma_2 \ln h_{it} - \ln(x_{it} + g(h_{it} + h_{it}^S)).$$

The equation now includes all the endogenous variables relevant for hours choices, in addition to the preference shifter, η_{it} . Exogenous changes in η_{it} over the life cycle will be accommodated by changes in spousal home production and household goods and services purchased in the market, *in addition to* changes in one's own home production. To see this, consider the Euler equation implied by (1.2) and our assumption on $u(\cdot)$.

$$E_t \left(\frac{z_{i,t+1}}{z_{it}} \right) = \beta(1+r) E_t \left(\frac{\eta_{i,t+1}}{\eta_{it}} \frac{p_t}{p_{t+1}} \right)$$

If p is constant, then the demand for household goods and services will change proportionally with changes in η_{it} . Empirically, h_{it} , h_{it}^S , and x_{it} will matter for the estimation of the Frisch elasticity, but will depend on (unobservable) movements in η_{it} . So long as one has suitable

proxies for η_{it} , such as measures of household composition and characteristics, one can estimate the Frisch elasticity using an instrumental variables approach on (3).

Intuitively, exogenous increases in the demand for household goods in services over the life cycle act like a negative income effect on labor supply. This is seen most clearly through a one-period, “full-income” version of the household’s budget constraint,

$$(4) \quad c + wl + wh - pg(h + h^s) = wH - pz + y^s + A,$$

where $l = H - n - h$ represents the implicit leisure hours after work hours and home production hours are chosen. The budget constraint’s “full income” measure is on the right-hand side of (4). It is assets ($y^s + A$) plus potential earnings (wH) less the value of the demand for home production (pz). Increasing η will increase z , making the household feel poorer, in the full-income sense. This has competing effects on labor supply. The income effect will induce the agent to work more to offset the decline in demand for market goods, c , and so that they can increase their household expenditures, x . Furthermore, the presence of a spouse can mitigate the negative income effect through either additional household income, which would allow an increase in purchased household goods and services, or through his own home production. Thus, equation (3) suggests that accounting for an agent’s own home production is likely not enough. One needs to also account for changes in the home production of other household members and expenditures on household goods and services to account for variations in η_{it} over the life cycle.

2.C. Empirical Considerations

Even if one has data on home production hours and expenditures on household goods and services, there are other issues that make it difficult to obtain consistent estimates of the Frisch elasticity from the data. First, there is measurement error in hours that can lead to

measurement error in wages when the latter are measured as earnings per period, i.e., the wage measure can suffer from division bias, as Altonji (1986) and others acknowledge. Division bias is particularly problematic with wage estimates derived from annual data like the Panel Study of Income Dynamics (PSID) or the National Longitudinal Survey of Youth (NLSY). The ATUS has earnings measured at either the hourly or weekly frequency, reducing the potential for division bias considerably, though it still exists for individuals who report their earnings weekly. The relatively small sample sizes of my synthetic cohort cells attenuate the division bias issue for these weekly reporters.

Second, there is an issue of endogeneity. Nearly every variable in the right-hand side of equation (3) is endogenous in the model it is based upon. The endogeneity of home production hours and household expenditures comes directly from the model. Wages are not necessarily endogenous but must be instrumented because of potential measurement error. Spousal home production hours are potentially endogenous as well because the spouse will likely also respond to changes in η_{it} over his life cycle. Thus, my analysis will need to address the endogeneity issue to a greater degree than previous studies.

Finally, even when focusing on life-cycle variation as the source for identifying the Frisch elasticity, there are several ways one can go about estimating it. One approach estimates the elasticity directly from longitudinal micro data, such as the PSID or NLSY (MaCurdy, 1981; Altonji, 1986; French, 2004). 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. It is more limited in examining extensive-margin labor adjustments. Another approach uses the micro data to calibrate a life cycle model of labor supply (Imai and Keane, 2004; Chang and Kim, 2006; Domeij and Floden, 2006; Rogerson and

Wallenius, 2009, 2011). These studies often focus on estimating the “aggregate” or “macro” Frisch elasticity (the elasticity that takes into account both intensive and extensive labor adjustments) given a calibrated value of the “micro” elasticity (i.e., the elasticity that only accounts for intensive-margin adjustments.) A third approach involves aggregating individuals based on their age or birth cohort and estimating the Frisch elasticity based on changes across observations as they age. The RRW study follows Lucas and Rapping (1969), Ghez and Becker (1975), and Smith (1977) in using individuals pooled by age as their unit of observation.⁶ This approach, however, suffers from the fact that one cannot control for cohort effects, which is a problem if the marginal utility of wealth or shocks to the demand for household goods and services have a cohort-specific component, perhaps because of business cycle or other aggregate effects on life-cycle labor supply. Aggregating repeated cross-sections into a time-series panel based on birth cohorts is known as the synthetic cohort approach and dates back to Deaton (1985). It has been used to study labor supply by Deaton, Browning, and Irish (1985), Mulligan (1999), and Peterman (2012). It has also been used extensively in studying life-cycle consumption behavior (Attanasio and Weber, 1993; Attanasio and Browning, 1995; Attanasio et al., 1999).

In this study, I use the synthetic cohort approach to estimate a specification based on equation (3), though I also create pooled age means as in RRW when replicating their results. The longitudinal micro data commonly used to study labor supply does not contain information on time spent on home production, and my ATUS sample consists of only repeated cross-sections of individuals. The synthetic cohort approach gives a panel dimension to the data that

⁶ Fiorito and Zanella (2012) use individuals pooled across years rather than age, and compare their results to the case where they instead use longitudinal individual data.

allows me to control for cohort effects and estimate the Frisch elasticity from the within-cohort variation over time. The approach also allows me to generate estimates of spousal home production from the ATUS data and merge in a measure of household expenditures from the Consumer Expenditure Survey (CEX). Finally, the approach allows me to estimate both a “micro” and “macro” Frisch elasticity based on whether I exclude or include the non-employed in the construction of the synthetic cohorts.

3. Data and Measurement

3.A. Data

My main source of data is the American Time Use Survey (ATUS) produced by the Bureau of Labor Statistics (BLS). The BLS draws the ATUS sample from the outgoing rotation groups of the Current Population Survey (CPS). The ATUS follows up with these respondents two to five months after their final CPS survey. The ATUS includes information on demographics, labor force status, and earnings for each respondent in a manner similar to the CPS.

Consequently, this information is available for individuals at the time of the ATUS survey and at the time of their last CPS interview. More importantly, the ATUS includes detailed time diary data for a single day. The time diary classifies respondents’ activities into a wide range of categories. The ATUS also has demographic, labor force, and earnings information for all other members of the household age 16 and older, which allows me to control for a variety of household characteristics despite the fact that I only have time diary data for a single individual within the household.

My sample includes all individuals with time diary data from repeated cross-sections of the ATUS between 2003 and 2012. The sample contains 101,253 individuals between the ages of

22 and 64. This is a large increase in the number of observations used in previous studies, in many cases by an order of magnitude.⁷ It allows for a synthetic cohort approach to estimation and allows me to account for a variety of demographic and household factors in a way that is not feasible with the data used in earlier studies.

I also use data from two other sources. The first is the Consumer Expenditure Survey (CEX), which provides a measure of expenditures on household goods and services (x_{it} in the model). The CEX contains quarterly data on expenditures across a broad range of detailed categories in addition to data on the income and demographics of the household. The survey allows me to directly match my expenditure estimates from the CEX to the ATUS data at the cohort-year level. My second source of additional data is the Merged Outgoing Rotation Group (MORG) samples of the CPS. These data contain much of the same information on earnings, labor force status, and demographics already contained in the ATUS sample, but have it for a much larger number of individuals. Consequently, estimates from the MORG are much less susceptible to sampling and measurement error. I use the MORG data to reduce measurement error in wage estimates derived from weekly earnings.

3.B. Sample Creation

My analysis uses either pooled means by respondent age or synthetic cohorts, grouped by birth year and followed throughout the sample period. In calculating pooled age means (the unit of observation in my replication of the RRW study), I pool all individuals by their age at their time of the interview across all sample years. Their mean hours and wage are calculated from

⁷ 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 data on a sample-weighted basis, with some adjustments made (described in the next section). In creating synthetic cohorts, I group individuals into 10 categories based on their birth year. My birth cohorts represent individuals grouped between four and seven years apart, with the oldest cohort born between 1940 and 1945 and the youngest cohort born between 1981 and 1987.⁸ I then calculate their (sample-weighted) mean estimates of hours and earnings separately for each year of the sample period. The median age is used for each cohort-year observation. Demographic statistics for each observation are calculated as either the sample-weighted mean (e.g., mean number of household children) or a sample-weighted fraction (e.g., percent female).

Measurement concerns arise when calculating the mean estimates of ATUS time diary data by sample cell. This is because the data are only reported for a single day and in general one would prefer weekly data. To see why, note that with daily data, employed individuals who report their time use would not be counted as employed unless they worked on their reporting day. Luckily, for the purpose of identifying employment status, I can use the usual weekly hours worked reported in the CPS-based portion of the survey. To get a direct, comparable measure of hours worked and home production hours, I divide individuals into their appropriate sample cells (by either age or cohort-year) and then aggregate their time diary data into synthetic weeks within each cell. This 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 for each day-of-week category, and sums them up to create the

⁸ The ATUS actually has data on both younger and older individuals, but I restrict my analysis to individuals within cohort-year observations where the median age is between 22 and 64. I use birth-year ranges that are larger for the youngest and oldest respondents to maintain comparable sample sizes across cohorts.

weekly average. My analysis uses the ATUS work hours as its dependent variable, but uses an adjusted measure of the usual hours worked when calculating a wage from weekly earnings data. This minimizes the effects of division bias.

3.C. Measurement

I count work hours as 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 measure household work hours as the total time spent in housework, child and adult care, pet care, vehicle care, shopping for goods, and purchasing services.

Wages are measured as average hourly earnings, where reported, or total weekly earnings divided by total usual work hours. I translate each wage into a real wage using the Consumer Price Index. The wage measure is a considerable improvement over those used in previous studies. First, the data are of high enough frequency to not be subject to possible short-term spells in joblessness, as is the case with the annual data in the NLSY and PSID.⁹ Second, the wages reported correspond to the individuals reporting the time-use data. For example, the RRW study used wages from an outside source, matching CPS wage data to the time-use data at the age-cohort level. Despite these advantages, I have to derive an hourly wage from weekly earnings and hours data for about one-third of my sample, which introduces the potential for division bias. I deal with this in two ways. First, I use usual weekly hours rather than the ATUS time-diary work hours when constructing my wage measure. Second, I use an adjusted measure of usual hours obtained from the predicted relationship between the usual hours

⁹ Note that this will tend to produce smaller estimates of the “micro” Frisch elasticity, since short-term spells of non-employment represent extensive-margin adjustments.

reported in the ATUS and the usual hours in the MORG data (estimated at the cohort-year or pooled age mean level). The wage used for individuals who report weekly is their (within-cell) mean weekly earnings divided by the adjusted usual hours measure, and the wage used for the sample cell is the weighted average between this wage measure and the wage for those who report their hourly earnings directly. I detail the full adjustment methodology in the appendix.

Table 1 reports the average wage and time spent by individuals aged 22 to 64 in various activities. It reports the estimates for all individuals, all workers, male workers, and female workers. Male workers have higher wages and considerably more hours of market work than female workers, 42.0 to 34.9, but women have more hours in home production, 26.4 to 17.2. Combined, women perform about two hours more total work and engage in 5.2 hours less leisure and socializing than male workers.

Figure 1 shows the behavior of market work and household work for the employed (top panel) and all individuals regardless of employment status (bottom panel) for men and women over the life cycle. The figure highlights the life-cycle differences between male and female labor supply. 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 exhibit greater variation in their home production hours over their life cycle, with a notable hump-shaped pattern that peaks during their thirties. Looking at all individuals shows the importance of considering the extensive margin of labor supply as both male and female hours exhibit greater variability over the life cycle in the bottom panel of Figure 1, particularly later in life. Notably, females engage in more home production than market work for most of their adult life. Finally, note that home production increases later in life for both males and females, a pattern

consistent with finding of substitution between money and time spent shopping among older workers by Aguiar and Hurst (2005).

My model suggests that I need a measure of the spouse's home production, but I only have time diary data for a single individual within each household. I therefore estimate spousal home production hours by creating a "synthetic spouse" for individuals that report having a spouse or partner present in the household, and assigning an estimate of mean home production hours based on the hours of the synthetic spouse and the demographic data for each cohort-year or pooled age cell. Specifically, I split each cohort-year or pooled age cell by gender, employment status, marital status, and the spouse's employment status (if a spouse is present). I then calculate the mean home production hours for each subgroup and assign these values to the complementary subgroup (e.g., married working males with a non-employed spouse are assigned the spousal home production hours of the appropriate married non-employed females with a working spouse). Mean spousal home production hours is the weighted mean of these assignments, where the weighting takes into account the fraction of the cell that is female, has a working spouse present, and is employed. The methodology is presented formally in the appendix.

Figure 2 presents the life-cycle behavior, by cohort-year, of individual and total home production hours (the sum of one's own hours and their synthetic spouse's home production hours). The top panel reports the estimates for employed respondents while the bottom panel reports the estimates for all respondents regardless of employment status. Spousal hours are included in the household totals regardless of the spouse's employment status in both cases. The pictures show that the hump-shaped pattern of home production over the life-cycle is much more pronounced for the household than it is for the individual. Two things are going on here.

First, individuals are getting married early in the life-cycle, increasing total household home production along the extensive margin by adding another household member. Second, individuals are able to pool their available hours when allocating time between work and home production, allowing individuals within the household to substitute home production for market work along the extensive margin. In other words, total household home production can reflect one household member exiting the labor force to care for children.

I obtain a measure of expenditures on household goods and services using data from the Consumer Expenditures Survey (CEX). I restrict the sample to include individuals who respond to the CEX for four straight quarters, and include either the household head (the primary survey respondent) or any spouse present. To be consistent with the ATUS sampling frame, I randomly include either the household head or the spouse to represent a household when both are present. I limit the measure to expenditures on goods and services that could plausibly be produced at home. Such a categorization is highly subjective, so I create three categories that increase in their scope of what could be produced at home. “Core” household goods and services include spending on food, household services (housekeeping, maintenance and repair), child care, and personal care services. “Related” expenditures include spending on transportation maintenance and gasoline, non-durable entertainment spending (including rentals related to leisure), and professional services. “Tangential” expenditures include spending on utilities, public transit, healthcare, books and educational supplies, and clothing. My measure aggregates the quarterly CEX data into cohort-year cells that are identical to the cells for the ATUS data. For each category, I generate the predicted relationship between the expenditure estimate and a vector of household and demographic characteristics, as well as cohort and year fixed effects. I interact the coefficients from this predicted relationship with the same vector of

characteristics calculated from the ATUS data. This generates a predicted estimate of household expenditures that accounts for sampling differences between the ATUS and CEX. The methodology is essentially the two-sample instrumental variables (TSIV) estimator derived by Angrist and Krueger (1992) and extended by Japelli, Pischke, and Soueles (1998). I provide more detail on it in the appendix.

Figure 3 reports the cohort-year observations of (log) real expenditures on “core” goods and services, “core” plus “related” goods and services, and all three expenditure categories combined for all individuals, regardless of employment status¹⁰. All three measures show the well-documented hump shape of consumption over the life cycle (e.g., see Aguiar and Hurst, 2013). The measure that includes all three categories exhibits less of a decline than the other two later in the life cycle, primarily because the “tangential” expenditures category includes expenditures on health care.

4. Evidence on the Labor Supply Elasticity

4.A. Replication of Previous Findings

I start with a replication of the analysis done by Rupert, Rogerson and Wright (2000) using the ATUS, since their study is the closest to my own in its empirical approach. It is not clear if using a broader, and larger data set covering a different time period will yield different results. RRW pool their time-use data and aggregate individuals into pooled age means. They then estimate a version of equation (3) on male workers aged 22 to 62. They use several alternative specifications for $v(n, h)$, which include $v(\cdot) = \phi L_t^\gamma$, $v(\cdot) = \phi(112 - L_t)^\gamma$, $v(\cdot) =$

¹⁰ The patterns for the subset of employed respondents is very similar.

$\phi(168 - s_t - L_t)^\gamma$, and $v(\cdot) = \phi \exp(\gamma L_t)$, where $L_t = n_t + h_t$, $H = 112$ is a measure of total hours that deducts 56 hours per week for sleep, and s_t is a measure of reported time spent on sleep or personal care. They estimate the Frisch elasticity using OLS regressions that are weighted by either cohort size or the variance of work hours. For the case where $v(\cdot) = \phi L_t^\gamma$, their estimating equation is

$$\ln(n_{it} + h_{it}) = \alpha + \psi t + \beta_0 \ln w_{it} + \varepsilon_{it},$$

with $\beta_0 = 1/(\gamma - 1)$ equal to the Frisch elasticity and ψ is included as a time trend.

Table 2 reports the results of replicating the RRW analysis using the ATUS data alongside the original estimated Frisch elasticities from their study. I report only results weighted by cohort size, since the variance-weighted results are nearly identical. In their study, RRW find that ignoring household work produces an estimated labor supply elasticity between 0.09 and 0.13, while accounting for it produces notably 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 0.24. When I include household production in the regressions, all four specifications produce higher elasticity estimates than the case where household work is excluded and nearly all produce higher elasticity estimates than those found by RRW. The one exception is again when sleep and personal care are included, which produces a nearly identical elasticity as the RRW study. The elasticity estimates when household work is included in the other specifications range between 0.40 and 0.55. 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 by 10 to 24 log points, while including household work with the ATUS data increases the point estimate by 17 to

32 log points. If anything, the replication of the RRW study reinforces their finding that accounting for household production is important for estimating the intertemporal elasticity of labor supply, and it produces estimates that are near the upper bound of those found in previous micro-level studies (e.g., Ghez and Becker, 1975; MaCurdy, 1981; and Altonji, 1986).

The size of the ATUS sample allows me to replicate the analysis for different demographic groups. I do so using the specification where $v(\cdot) = \phi n_{it}^{\gamma_1} h_{it}^{\gamma_2}$, which allows for different elasticities for home production and market work, but leave all other aspects of the original RRW analysis the same. The results are in Table 3. I estimate the Frisch elasticity separately for all workers, male workers, female workers, married male workers, married female workers, all white workers, and all nonwhite workers. In every case, the addition of home production increases the estimated Frisch elasticity. Depending on the group, the increase is between 17 and 34 log points, and the estimated elasticity when home production is included is between 0.23 and 0.50. The elasticity is about the same for men and women, but the increase when home production is included is considerably higher for women. Compared to nonwhite workers, whites have higher elasticity estimates and a larger increase when home production is included.

4.B. Synthetic Cohort Analysis

Next, I estimate the Frisch elasticity using the synthetic cohort analysis first introduced by Deaton (1985) and used extensively in studies of consumption behavior over the life-cycle (Attanasio and Weber, 1993; Attanasio and Browning, 1995; Attanasio et al., 1999), and in several studies of life cycle labor supply (Browning, Deaton, and Irish, 1985; Mulligan, 1999; Peterman, 2012). By aggregating individuals based on their birth year, I can follow their cohorts over time and exploit the within-cohort time-series variation to estimate the Frisch labor supply

elasticity over the life cycle. Compared to the pooled age means approach, the synthetic cohort analysis uses substantially more observations in its estimation.

Table 4 reports basic demographic statistics for the cohorts used in the analysis, including the average cell size of each cohort-year observation and the years used in estimation. I restrict my analysis to cohort-year observations of employed individuals (both male and female) with a median cell age between 22 and 64. I estimate the following equation across i cohorts and t years,

$$(5) \quad \ln n_{it} = \alpha_i + \psi t + \beta_0 \ln w_{it} + \beta_1 \ln h_{it} + \beta_2 \ln h_{it}^S + \beta_3 \ln x_{it} + X_{i(t)} \delta + \varepsilon_{it},$$

where α_i is a cohort fixed effect, ψ is a time trend, and $X_{i(t)}$ is a vector of fixed observable characteristics (the fraction female and the fraction white, non-Hispanic). The notation for the latter has the time subscript in parentheses because these fractions are fixed only in theory. Sampling error across years for the same cohort will cause differences in these characteristics over time. The β_0 coefficient is the Frisch elasticity estimate. For household expenditures, I include both “core” and “related” household goods and services as my preferred measure, though I obtain very similar estimates if I include only “core” household expenditures or if I add “tangential” expenditures.¹¹ I again weight all regressions by cohort size. For the fully-specified model, I correct the standard errors for the fact that the expenditure measure is derived from two samples and uses the TSIV estimator of Angrist and Krueger (1992) and Jappeli, Pischke, and Soueles (1998). When accounting for the endogeneity of expenditures, my approach amounts to

¹¹ Results using the alternate expenditure measures are in the appendix.

what I refer to as a two-stage, two-sample instrumental variables (2S-TSIV) estimation. I detail the estimator and its econometric properties in the appendix.

For a variety of reasons discussed in Section 2, OLS estimates of the Frisch elasticity are likely biased. To deal with this, I estimate equation (5) using a two-stage least squares (2SLS) approach that treats wages, own home production, spousal home production, and household expenditures as endogenous. In the first stage, I instrument these variables with a variety of household and demographic characteristics. The choice of instruments is driven by those used in previous studies of life-cycle labor supply (such as MacCurdy, 1981) and life-cycle consumption (such as Attanasio and Weber, 1993). Specifically, I control for education with the fractions of each cohort-year with some college, a college degree, or a graduate degree. The education instruments are commonly used in previous labor supply studies to instrument for wages. I also include instruments that capture characteristics of the household. I draw these instruments from the literature on life-cycle consumption. These studies, which include Attanasio and Weber (1993), and Attanasio et al. (1999), use these measures to capture changes in household composition and characteristics that drive shocks to preferences over the life cycle, i.e., shocks analogous to shocks to η_{it} in the model here. These instruments include the fraction of the cohort-year cell with a spouse or partner present, its mean household size, the fraction with any children present in the household, the mean number of children aged 0 to 6, the mean number of children aged 7 to 18, the fraction with a working spouse present, and the log of the spouse's mean real wage (conditional on working). The instruments for spousal employment status and wage are particularly important in my study, since I not only care about how much spousal income contributes to expenditures, but how the spousal wage will affect his work and home production hours decision.

Table 5 presents the results. When I include the wage as the only endogenous variable on the right-hand side, the estimated Frisch elasticity is 0.16. Including home production increases the elasticity to 0.44, consistent with the findings of the previous section and the RRW study. When the full model is estimated, the estimate falls slightly to 0.39, but the decline is not statistically (or economically) significant. Thus, at least for the “micro” Frisch elasticity, there is little difference between estimates that only account for one’s own home production and estimates that take all factors related to home production into account.

4.C. Accounting for Extensive Labor Adjustment

Next, I estimate a “macro” Frisch elasticity that accounts for home production. Extensive-margin labor adjustments (i.e., movements between employment and non-employment) account for about two-thirds of the volatility in aggregate hours, leading to a focus on these adjustments as the relevant margin for macroeconomic models of the business cycle (see Hansen, 1985, and Rogerson, 1988). Studies that estimate a “macro” Frisch elasticity, which accounts for these adjustments, tend to find that the elasticity is large, with most estimates falling between 1.2 and 3.8 (see Imai and Keane, 2004; Chang and Kim, 2006; Rogerson and Wallenius, 2009). Figures 1 and 2 show that both work hours and home production hours exhibit greater life-cycle variation for all individuals than for the employed alone. This is especially true for male work hours and female home production hours. Therefore, one might expect that accounting for home production will have an even larger effect on the estimated “macro” elasticity.

I see if this is the case by re-estimating the Frisch elasticity including all individuals, not just the employed, in my cohort-year observations. Including these individuals, however, creates a selection issue when it comes to the wage. The non-employed have zero work hours by

definition, but I do not observe a wage for them. These individuals are arguably non-employed because they face a different outside option, on average, than those who are employed. Therefore, an estimate of the “macro” elasticity that accounts for changes in the employment status of these individuals but only uses the wages of the employed will be biased. To deal with this, I appeal to the fact that a fraction of those non-employed at the time of their ATUS interview were employed during their final CPS interview. The final CPS interviews occur two to five months prior to the ATUS interview and approximately 15 percent of the non-employed in my sample report a wage during the CPS interview. I use these data to calculate a wage that I then impute to all non-employed. I deal with potential division bias in CPS weekly earnings in the same manner as I did for the ATUS weekly earnings (i.e., using a predicted measure of usual weekly hours based on MORG data). I also Winsorize the distribution of reported earnings because the sample of non-employed with a CPS wage is small and sensitive to outliers. Figure 5 reports the real wage over the life cycle measured for individuals employed in the ATUS and individuals not employed at the time of their ATUS survey but employed in their final CPS interview. The CPS wage is a noisier measure, but the figure shows that there is a clear selection effect reflected in the wage. The CPS wage is 17 percent lower than the ATUS wage for individuals aged 25-29 and is 24 percent lower for individuals aged 44-49. The wage used in the estimation of (5) is the weighted average of the ATUS wage and the CPS wage, where the weight is the fraction of the cell that is employed.

The results of this analysis are in Table 6. When including only wages in the regression, the Frisch elasticity rises from the “micro” estimate of 0.16 to a “macro” estimate of 0.68. When I add one’s own home production, the estimate increases even further to 1.58, which is within the range of estimates obtained by the macro studies mentioned earlier. When I estimate the

elasticity using the fully-specified model, however, the estimate falls to 0.92, which is higher but not statistically different from the 0.68 estimate when including only wages in the regression.

The decline in the estimate when adding spousal home production and household expenditures is consistent with the intuition of the model in Section 2. These factors mitigate the income effect changes in household composition (and the related change in the demand for household goods and services) have on labor supply. When one only accounts for an individual's home production, and not the home production effort of other household members, they capture the income effect caused by changes in household composition but over-estimate the labor supply response because they ignore the household's ability to have other individuals meet the demand for home production or purchase the demanded household goods and services in the market. My estimates suggest that when all factors are accounted for, the "macro" Frisch elasticity estimate is 64 log points lower than when one only controls for an individual's home production, and only somewhat higher than the estimate obtained when ignoring the role of home production completely. Furthermore, the estimate of 0.92 lies below the range of estimates for the "macro" elasticity obtained from calibrated macro models of the life cycle. Instead, my estimates are very similar to those obtained in the meta-analysis of Chetty et al. (2011), who do not account for home production but estimate a "micro" Frisch elasticity of 0.54 (compared to 0.39 here) and a "macro" Frisch elasticity of 0.82 (compared to 0.92 here). The results suggest that estimates that account for home production must do so in a way that accounts for all the ways a household can respond to changes in the demand for household goods and services. Otherwise, one will over-estimate the labor supply response to wages. The results suggest that the effects of changes in home production are particularly strong over the life cycle for extensive-margin adjustments. It is unclear how much estimates of a macro

calibration exercise would be affected by extending the analysis to include the home production behavior of the household. Recent work by Rogerson and Wallenius (2011) incorporates a role for home production in such an exercise, though it focuses on the labor supply decision at retirement rather than over the full life cycle. Recent work by Dyrda, Kaplan, and Rios-Rull (2012) suggests that the implications for the cyclical behavior of labor supply could be large. They show that accounting for cyclical changes in household composition, such as the young moving back in with their parents, can have a large effect on the labor supply elasticity.

4.D. Alternative Estimation Specifications

Finally, I examine the effect of using different specifications for my analysis on the elasticity estimates. Specifically, I estimate the elasticity using the model in first differences rather than levels, and I estimate the elasticity separately for men and women using the synthetic cohort approach. I also estimate the model including household expenditures throughout. Since they appeal to longitudinal micro data, many of the previous studies on labor supply use a first-difference specification. In addition, many previous studies have focused only on the labor supply of men. Therefore, these exercises are meant to provide added comparability to previous research rather than act as a robustness check on my main analysis. The final exercise addresses the fact that changes in expenditures over the life cycle might capture changes in the marginal utility of wealth (assumed to be constant in the model), which would overstate the effect of including it on the Frisch elasticity estimate.

First-differencing the regression in equation (5) yields,

$$(6) \quad d \ln n_{it} = \psi + \beta_0 d \ln w_{it} + \beta_1 d \ln h_{it} + \beta_2 d \ln h_{it}^s + \beta_3 d \ln x_{it} + dX_{i(t)}\delta + u_{it}.$$

Note that the fixed demographic characteristics (percent female and percent white, non-Hispanic) remain in the equation because of sampling differences over time within cohorts. I

estimate equation (6) using 2SLS as before. The variables used as instruments are also the same but are similarly first differenced. As before, $\ln x_{it}$ uses an estimate predicted from coefficients obtained from the CEX and household and demographic characteristics from the ATUS, and standard errors for the fully-specified model are corrected using the 2S-TSIV approach described in the appendix.

Estimates of the “micro” Frisch elasticity (using only employed individuals) and “macro” Frisch elasticity (using both employed and non-employed individuals) are in Table 7. The “micro” estimates are generally similar to those estimated in levels (Table 5), but are less precise and their instruments are considerably weaker in first differences. The elasticity estimate is higher when only the wage is included in the estimation (0.28 compared to 0.16), and rises to 0.35 when one’s own home production is included. It falls only slightly to 0.33 when including spousal home production and household expenditures (compared to 0.39 in the levels specification), but the estimate is no longer statistically significant.

The “macro” Frisch elasticity estimate is 0.54 when only wages are included (compared to 0.68 in the levels specification). Using first differences, adding in home production only increases the estimate to 0.74 (rather than 1.58 in the levels specification), and the fully-specified model yields an elasticity estimate of only 0.53, which again is statistically insignificant given the relatively large standard errors of the first-differenced model.

Table 8 reports the results of replicating the analysis using the levels specification from equation (5) for separate cohort-year observations of men and women. The “micro” estimates are for either employed men or employed women, while the “macro” estimates are for either all men or all women. I only report the coefficient on wages for each specification, though the full

results are included in the appendix. Each column represents the specification from the same column of Table 5 (employed individuals) or Table 6 (all individuals).

The “micro” Frisch elasticity for men is 0.19, which is very similar to the 0.16 estimated for all employed individuals. Including home production increases their elasticity only slightly, however, to 0.24. When the full model is specified, the male “micro” elasticity for men is essentially zero. For women, including only wages in the specification yields a “micro” elasticity of only 0.08, but adding in their own home production increases the estimate to 0.26. There is only a slight decrease in the estimated elasticity, to 0.21, when using the full model.

For men, the “macro” elasticity is unchanged at 0.73 when moving from a specification that includes only wages to the full specification. It only rises to 0.91 when only male home production hours are included. Doing the same exercise for women increases their “macro” elasticity estimate from 0.38 to 0.51, with an estimate of 0.88 when only female home production is included. Thus, when all factors relating to home production are taken into account and the extensive-margin changes in labor supply are allowed for, men have a somewhat higher labor supply elasticity (though the difference with women is not significant), but have a smaller response of that elasticity to the inclusion of home production. This is not surprising given the evidence in Figure 1, which shows that women do more home production, have home production profiles that vary more over the life cycle, and are more likely adjust their labor hours along the extensive margin.

Finally, Table 9 reports the results from including expenditures on household goods and services throughout. I report the results of adding (log) household expenditures to the specifications from the first two columns of Table 5 (for the employed) and Table 6 (for all individuals) to see how it affects the elasticity estimates prior to including expenditures in the

full specification. Note that adding expenditures to the third specification in each table just reproduces the results for the full specification. Controlling for expenditures decreases the elasticity estimates in all cases, suggesting that expenditures are capturing some movements in the marginal utility of wealth, though the decreases are relatively small, and the qualitative results of the main analysis still hold. When excluding home production, the “micro” Frisch elasticity estimate falls from 0.16 to 0.08 and the “macro” Frisch elasticity falls from 0.68 to 0.30. When only accounting for one’s home production the estimates fall from 0.44 to 0.36 and from 1.58 to 1.26, respectively. Note that the estimates from the full specification (Tables 5 and 6) yielded estimates of 0.39 and 0.92, respectively.

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 changes in home production over the life cycle will lead to a downward bias when estimating this elasticity. I develop a life-cycle model that allows for substitution along other margins in response to changes in the demand for home production, such as changes in the number of household children. The model allows individuals to purchase required household goods in the market and allows for the presence of a spouse, who may also either work or engage in home production. The key insight from the model is that these additional margins allow an individual to mitigate the negative income effect caused by an increase in the demand for home production. I show that ignoring these additional margins in an analysis of home production and labor supply will produce estimates that overstate the Frisch elasticity.

I take this extended the model to the data using a synthetic cohort analysis. Synthetic cohorts provide my data with a panel dimension that allows a more robust treatment of potential changes in the marginal utility of wealth over time and also allows for more sophisticated econometric methods. The elasticities that I estimate using wages alone are comparable to the elasticity estimates obtained from earlier studies that use micro data, while the estimates that include home production more than double these estimates, just as in RRW. Estimating a fully-specified model based on my extended life-cycle model reduces the estimates. The reduction is small for the “micro” Frisch elasticity, which is around 0.4 when changes in the home production of one’s spouse and expenditures on household goods and services are accounted for.

Estimates of the “macro” elasticity, however are much more responsive to how one accounts for home production. Accounting for one’s own home production hours yields a significant increase in the “macro” elasticity estimate, from about 0.7 to about 1.6 when compared to the case where one ignores home production entirely. When one moves to the fully-specified model, however, the elasticity falls to 0.9, which is still higher than the 0.7 estimated when ignoring home production entirely, but not significantly different from it. Most models of the business cycle require a considerably larger elasticity to match the observed fluctuations in aggregate hours. I obtain even lower estimates of the “macro” elasticity when estimating the model in first differences and estimating it separately for men and women.

My results suggest that changes in the demand for home production can have a powerful effect on labor supply, but that those effects are partially mitigated by the pooling of the time devoted to market work and home work between spouses, including the option for one spouse to drop out of the labor force entirely and focus on home production. My evidence

shows women are considerably more likely to respond in this manner. Consequently, their labor supply is more responsive to home production demands. These demands can also be mitigated by expenditures on household goods and services, which substitute for home production. The presence of a second wage-earner in the household will make such expenditures more affordable. Economists have known for some time that changes in the demand for home production can have an effect on labor supply. What was less apparent was that the home production of other household members could have large but countervailing effects on the Frisch elasticity. My results show that these effects are most important for extensive-margin labor adjustments, which in turn are the adjustments that are most important for aggregate hours fluctuations. Implicitly, this suggests that aggregate changes in household composition, such as a rising incidence of single-parent households (where spousal home production is non-existent) or cyclical changes in household composition of the kind studied by Dyrda, Kaplan, and Rios-Rull (2012), can have significant macroeconomic implications through their effects on aggregate labor supply.

References

- Aguiar, Mark, and Hurst, Erik (2005) "Consumption versus Expenditure," *Journal of Political Economy*, 113(5): 919-48.
- (2013). "Deconstructing Life-Cycle Expenditure," *Journal of Political Economy*, 121(3): 437-92.
- Aguiar, Mark, Hurst, Erik, and Karabarbounis, Loukas (2013) "Time Use During Recessions," *American Economic Review*, 103(5): 1664-96.
- Altonji, Joseph G. (1986). "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy* 94(3, Part 2): S176-S215.
- Angrist, Joshua D., and Krueger, Alan B. (1992). "The Effect of Age at School Entry on Attainment: An application of Instrumental Variables with Moments from Two Samples," *Journal of the American Statistical Association*, 87: 328-36.
- Attanasio, Orazio P., and Browning, Martin (1995). "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, 85(5): 1118-37.
- Attanasio, Orazio P., Banks, James, Meghir, Costas, Weber, Guglielmo (1999). "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics* 17(1): 22-35.
- Attanasio, Orazio P., and Weber, Guglielmo (1993). "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103(6): 1121-57.
- Benhabib, Jess, Rogerson, Richard, and Wright, Randall (1991). "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy* 99(6): 1166-87.
- Browning, Martin, Deaton, Angus, and Irish, Margaret (1985). "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica* 53(3): 503-43.
- 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.
- Chetty, Raj, Guren, Adam, Manoli, Day, and Weber, Andrea (2011) "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review, Papers and Proceedings*, 101(3): 471-75.
- Deaton, Angus (1985). "Panel Data from Time Series of Cross Sections," *Journal of Econometrics*, 30(1-2): 109-26.
- Domeij, David and Floden, Martin (2006). "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates are Biased," *Review of Economic Dynamics*, 9(2): 242-62.

Dyrda, Sebastian, Kaplan, Greg, and Ríos-Rull, José-Víctor (2012). "Business Cycles and Household Formation: The Micro vs the Macro Labor Elasticity," NBER Working Paper No. 17880

Fiorito, Riccardo, and Zanella, Giuilo (2008). "Labor Supply Elasticities: Can Micro be Misleading for Macro?" Italian Ministry of Economy and Finance Working Paper 2008-04.

French, Eric (2004). "The Labor Supply Response to (Mismeasured but) Predictable Wage Changes," *Review of Economics and Statistics*, 86(2): 602-13.

----- (2005). "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behaviour," *Review of Economic Studies*, 72(3): 395-427.

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).

Hansen, Gary (1985). "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics* 16(2): 309-27.

Imai, Susumu, and Keane, Michael P. (2004). "Intertemporal Labor Supply and Human Capital Accumulation," *International Economic Review* 45(2): 601-41.

Jappelli, Tullio, Pischke, Jörn-Steffen, and Soueles, Nicholas S. (1998). "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources." *Review of Economics and Statistics*, 80(2): 251-62.

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.

McGrattan, Ellen R., Rogerson, Richard, and Wright, Randall (1997). "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy," *International Economic Review*, 38(2): 267-90.

Mulligan, Casey B. (1999) "Substitution over Time: Another Look at Life-Cycle Labor Supply," *NBER Macroeconomics Annual*, 13: 75-134.

Peterman, William B. (2012). "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity: A Sensitivity Analysis," Federal Reserve Board of Governors, Finance and Economics Discussion Series No. 2012-75.

Rogerson, Richard (1988). "Indivisible Labor, Lotteries, and Equilibrium," *Journal of Monetary Economics* 21(1): 3-16.

Rogerson, Richard, and Wallenius, Johanna (2009). "Micro and Macro Elasticities in a Life Cycle Model with Taxes," *Journal of Economic Theory* 144(6): 2277-92.

-----, (2011). "Retirement, Home Production, and Labor Supply Elasticities," Stockholm School of Economics, mimeo.

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 Employed	Employed Males	Employed Females
Real hourly wage (2003 \$)	---	\$ 18.03	\$ 19.58	\$ 16.23
Usual work hours (CPS measure)	---	41.45	44.34	38.30
Reported work time (time diary)	29.30	38.26	41.31	34.92
Total household work	24.62	21.06	17.02	25.49
Leisure & socializing	43.01	39.00	41.12	36.69
Personal care & sleep	66.41	64.80	63.15	66.60
<i>N</i>	101,253	66,388	31,441	34,947

Notes: Estimates are for the listed individuals aged 22-64 in the 2003-2012 surveys of the ATUS. Wages are in 2003 dollars. Time use is hours per week.

Table 2. Estimated Intertemporal Elasticities of Labor Supply

Specification	Without Home Production		Including Home Production	
	RRW	ATUS	RRW	ATUS
$y = \ln L$	0.117 (0.119)	0.242 (0.044)	0.274 (0.072)	0.415 (0.034)
$y = \ln (112 - L)$	0.092 (0.054)	0.131 (0.025)	0.280 (0.071)	0.407 (0.030)
$y = \ln (168 - s - L)$	0.099 (0.054)	0.037 (0.025)	0.338 (0.084)	0.331 (0.034)
$y = L$	0.126 (0.085)	0.228 (0.042)	0.222 (0.051)	0.550 (0.042)

Note: The table reports the regression coefficients from regressing the listed dependent variable on $\ln w_{it}$ and age. Regressions without household work use $L = h$, while those with household work use $L = n + h$. 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 2003-12 for working males aged 22-62. All regressions are weighted by the size of the age cohorts. Standard errors are in parentheses.

Table 3. Estimated Intertemporal Elasticities of Labor Supply

Sample	Without Home Production	Including Home Production	
	$\ln w_t$	$\ln w_t$	$\ln h_t$
All Employed	0.167 (0.042)	0.568 (0.096)	-0.447 (0.100)
Male Employed	0.229 (0.047)	0.392 (0.096)	-0.193 (0.100)
Female Employed	0.031 (0.056)	0.386 (0.086)	-0.385 (0.079)
Married Male Employed	0.044 (0.053)	0.165 (0.059)	-0.303 (0.089)
Married Female Employed	0.027 (0.065)	0.235 (0.077)	-0.365 (0.093)
White Employed	0.122 (0.044)	0.488 (0.100)	-0.387 (0.098)
Nonwhite Employed	0.218 (0.066)	0.383 (0.086)	-0.249 (0.092)

Note: The table reports the regression coefficients from regressing $\ln n_t$ on a constant, time trend, and either $\ln w_t$ alone (first column) or $\ln w_t$ and $\ln h_t$ (second and third columns). All regressions use pooled age means, created from pooled ATUS data for 2003-12, as their observations. The means are for the listed groups of workers aged 22-64 (regression $N = 43$). All regressions are weighted by the size of the pooled age mean cell. Standard errors are in parentheses.

Table 4. Synthetic Cohort Summary Statistics

Birth Cohort	Median Age, 2003	Mean Cell Size	Percent Female	Percent White	Percent Married	Percent College Ed.	Years Used
1940-45	60	303	53.1	80.3	68.2	31.0	2003-07
1946-50	55	521	50.6	77.6	70.4	35.5	All
1951-55	50	721	50.4	77.2	70.2	34.9	All
1956-59	45	710	49.2	74.9	70.4	32.2	All
1960-63	42	778	49.4	72.0	66.8	34.2	All
1964-67	38	809	46.3	68.2	67.4	35.3	All
1968-71	33	812	45.5	67.2	67.2	39.2	All
1972-75	29	730	45.9	63.2	64.4	39.7	All
1976-80	25	778	45.0	64.7	50.0	38.7	All
1981-87	20	720	48.5	64.2	22.4	25.0	2006-12

Note: The table reports summary statistics for each group of birth cohorts generated from the ATUS. Statistics represent means across survey years (2003-12), unless otherwise noted.

Table 5. Synthetic Cohort Panel Regression Estimates, Employed Men and Women

	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.161** (0.076)	0.438** (0.130)	0.522** (0.166)	0.391* (0.236)
log(own home production), $\ln h_t$		-0.359** (0.137)	-0.237 (0.198)	-0.329 (0.226)
log(spousal home production), $\ln h_t^S$			-0.087 (0.096)	-0.022 (0.127)
log(household expenditures), $\ln x_t$				0.135 (0.179)
Adj. R-squared	0.220	0.262	0.229	0.251
F-test of Overidentification [<i>p</i> -value]	2.28 [0.027]	1.49 [0.177]	1.37 [0.232]	1.68 [0.139]
First-stage <i>F</i> -test of excluded instruments [<i>p</i> -value]. ¹	21.86 [0.000]	11.02 [0.000]	84.61 [0.000]	7.15 [0.000]

Note: The table reports the regression coefficients from regressing log(reported work hours), $\ln n_t$, on the listed dependent variables, using the specification from equation (6) in the text. The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. All regressions also include cohort fixed effects, a year trend, and controls for the share of the cohort-year cell that is female and that is white, non-Hispanic. All regressions are weighted by cohort-year cell size. Instrumental variables regressions are estimated using two-stage least squares. The instruments include fraction of the cohort-year cell that has a spouse or partner present, the fraction with any children present, mean household size, the mean number of children ages 0 to 6, the mean number of children ages 7 to 18, the fraction of individuals with a working spouse, the mean spousal real wage, and the fractions of the cell with some college, a college degree, or a postgraduate degree. Standard errors are in parentheses.

1. Row lists the *F*-value for the test of excluded instruments for the endogenous regressor added in each column: log(real wage) in column (1), log(own home production) in column (2), log(spousal home production) in column (3), and log(household expenditures) in column (4).

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Table 6. Synthetic Cohort Estimates, Accounting for Extensive Margin Adjustments

	IV, All Individuals			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.679** (0.154)	1.581** (0.257)	1.501** (0.350)	0.920* (0.489)
log(own home production), $\ln h_t$		-1.489** (0.310)	-1.651** (0.574)	-1.942** (0.556)
log(spousal home production), $\ln h_t^s$			0.114 (0.340)	0.343 (0.343)
log(household expenditures), $\ln x_t$				0.639* (0.387)
Adj. R-squared	0.786	0.752	0.753	0.804
F-test of Overidentification [<i>p</i> -value]	7.57 [0.000]	0.83 [0.581]	0.95 [0.477]	0.95 [0.468]
First-stage <i>F</i> -test of excluded instruments [<i>p</i> -value]. ¹	11.25 [0.000]	17.47 [0.000]	81.04 [0.000]	8.16 [0.000]

Note: The table reports the 2SLS estimates from regressing log(reported work hours), $\ln n_t$, on the listed dependent variables, using the specification from equation (6) in the text. The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include either all individuals or all employed individuals plus non-employed with a wage in their outgoing CPS interview, as indicated in the table. The wage used is either the wage of employed individuals in the ATUS or a weighted average of the ATUS-reported wage and the wage reported by the non-employed in the CPS outgoing interview.

1. Row lists the *F*-value for the test of excluded instruments for the endogenous regressor added in each column: log(real wage) in column (1), log(own home production) in column (2), log(spousal home production) in column (3), and log(household expenditures) in column (4).

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Table 7. Synthetic Cohort Panel Regression Estimates, Estimation in Log-First Differences

	IV, Employed Only			
	(1)	(2)	(3)	(4)
log(real wage), $d \ln w_t$	0.281 (0.180)	0.346** (0.175)	0.478** (0.178)	0.334 (0.251)
log(own home production), $d \ln h_t$		-0.290* (0.176)	-0.071 (0.393)	-0.278 (0.241)
log(spousal home production), $d \ln h_t^s$			-0.189 (0.289)	-0.045 (0.105)
log(household expenditures), $d \ln x_t$				0.139 (0.244)
Adj. R-squared	0.051	0.094	0.072	0.065
F-test of Overidentification [<i>p</i> -value]	2.39 [0.020]	2.76 [0.011]	2.70 [0.016]	3.12 [0.009]
First-stage F-test of excluded instruments [<i>p</i> -value]. ¹	3.18 [0.002]	1.83 [0.071]	26.83 [0.000]	5.15 [0.000]
	IV, All Individuals			
	(1)	(2)	(3)	(4)
log(real wage), $d \ln w_t$	0.539** (0.243)	0.743** (0.265)	0.624** (0.292)	0.533 (0.328)
log(own home production), $d \ln h_t$		-0.805** (0.358)	-1.012** (0.419)	-1.031** (0.404)
log(spousal home production), $d \ln h_t^s$			0.153 (0.174)	0.141 (0.168)
log(household expenditures), $d \ln x_t$				0.212 (0.394)
Adj. R-squared	0.057	0.093	0.099	0.104
F-test of Overidentification [<i>p</i> -value]	1.47 [0.178]	0.81 [0.600]	0.88 [0.528]	1.09 [0.378]
First-stage F-test of excluded instruments [<i>p</i> -value]. ¹	1.93 [0.055]	2.03 [0.042]	12.07 [0.000]	1.58 [0.132]

Note: The table reports the 2SLS estimates from regressing log(reported work hours), $\ln n_t$, on the listed dependent variables, using the specification from equation (7) in the text. The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include either employed individuals or all individuals regardless of employment status, as indicated in the table. The non-employed use a wage derived from their CPS outgoing interview.

1. Row lists the *F*-value for the test of excluded instruments for the endogenous regressor added in each column: log(real wage) in column (1), log(own home production) in column (2), log(spousal home production) in column (3), and log(household expenditures) in column (4).

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Table 8. Frisch Elasticity Estimates by Gender

	(1)	(2)	(3)	(4)
Employed Males	0.186* (0.098)	0.240* (0.175)	0.231* (0.128)	0.038 (0.188)
All Males	0.731** (0.153)	0.905** (0.168)	0.796** (0.196)	0.731** (0.246)
Employed Females	0.076 (0.095)	0.259** (0.122)	0.235* (0.133)	0.211 (0.135)
All Females	0.379** (0.178)	0.883** (0.225)	0.843** (0.240)	0.514* (0.278)

Note: The table reports the Frisch elasticity estimate (coefficient on log wages) from the 2SLS regressions of $\log(\text{reported work hours})$, $\ln n_t$, on the model specified in equation (6) in the text. Specifications in each column are identical to those in columns (1) through (4) of Table 5 (employed only) or Table 6 (all individuals). The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include the demographic listed for each row. When included, the non-employed use a wage derived from their CPS outgoing interview.

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Table 9. Synthetic Cohort Estimates, Controlling for Expenditures

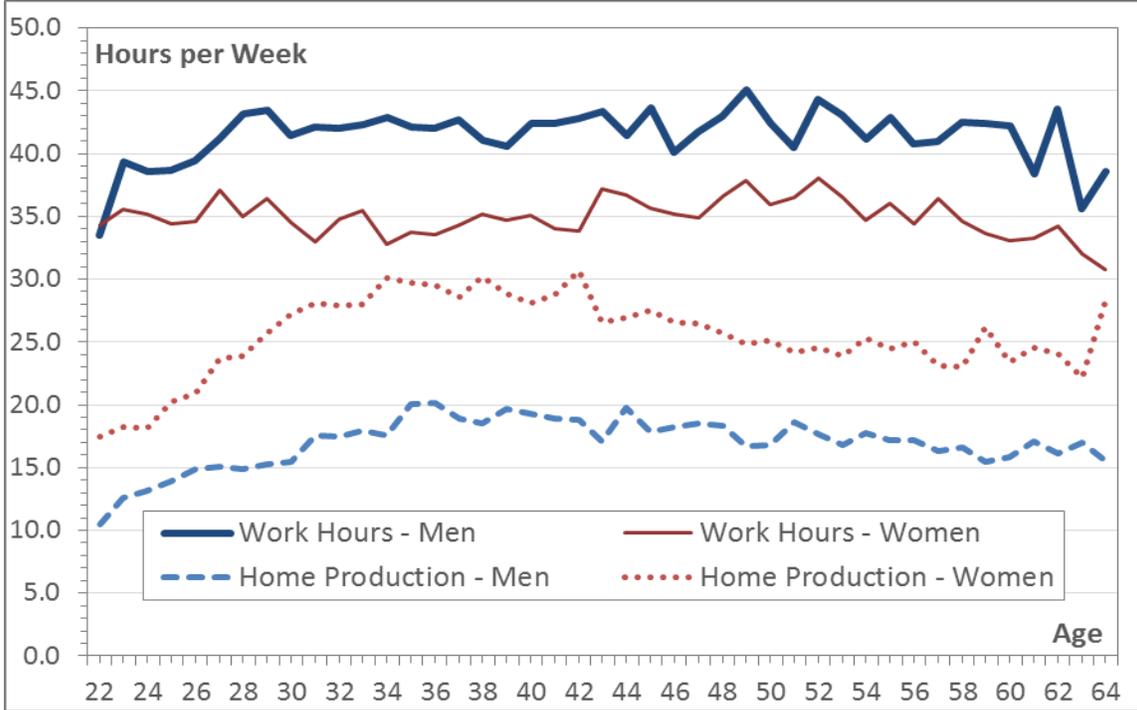
	IV, Employed Only		IV, All Individuals	
	(1)	(2)	(1)	(2)
$\log(\text{real wage}), \ln w_t$	0.082 (0.103)	0.359** (0.148)	0.300 (0.265)	1.264** (0.315)
$\log(\text{own home production}), \ln h_t$		-0.360** (0.139)		-1.461** (0.270)
$\log(\text{spousal home production}), \ln h_t^S$				
$\log(\text{household expenditures}), \ln x_t$	0.154 (0.132)	0.156 (0.130)	0.611* (0.335)	0.482 (0.333)
Adj. R-squared	0.214	0.266	0.802	0.804
F-test of Overidentification	2.25	1.48	8.73	1.00
[p -value]	[0.034]	[0.190]	[0.000]	[0.441]

Note: The table reports the 2SLS estimates from regressing $\log(\text{reported work hours})$, $\ln n_t$, on the listed dependent variables, using the specification from equation (6) in the text. The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include either all individuals or all employed individuals plus non-employed with a wage in their outgoing CPS interview, as indicated in the table. The wage used is either the wage of employed individuals in the ATUS or a weighted average of the ATUS-reported wage and the wage reported by the non-employed in the CPS outgoing interview.

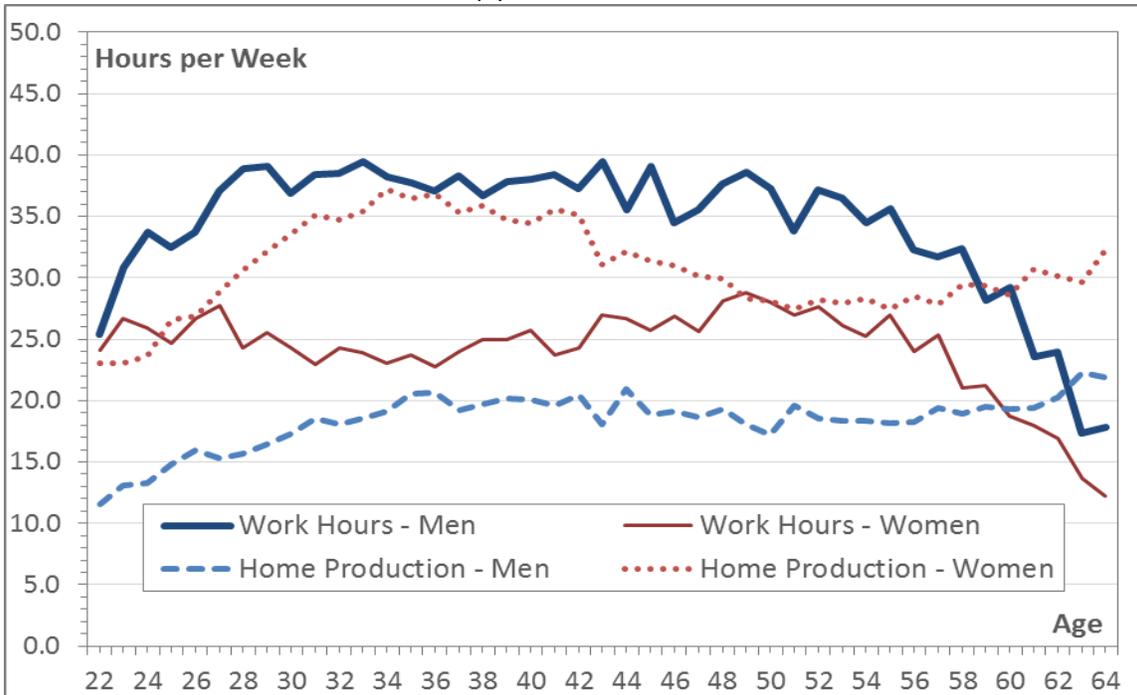
** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Figure 1. Hours Worked and Home Production over the Life Cycle
 (a) All Employed

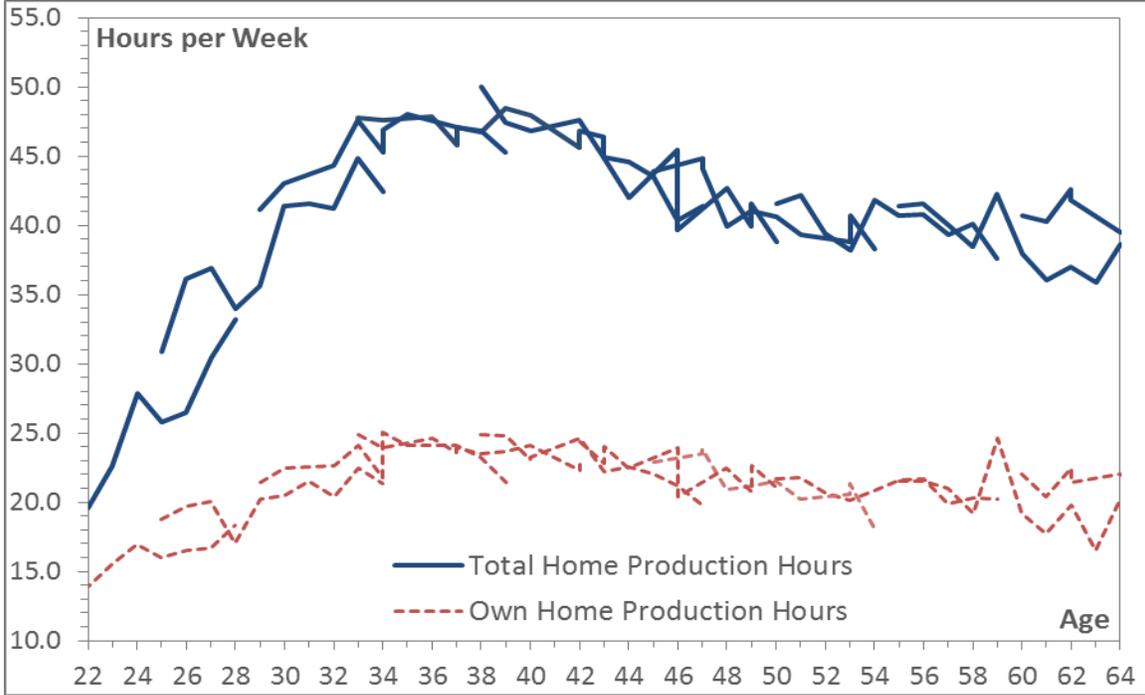


(b) All Individuals

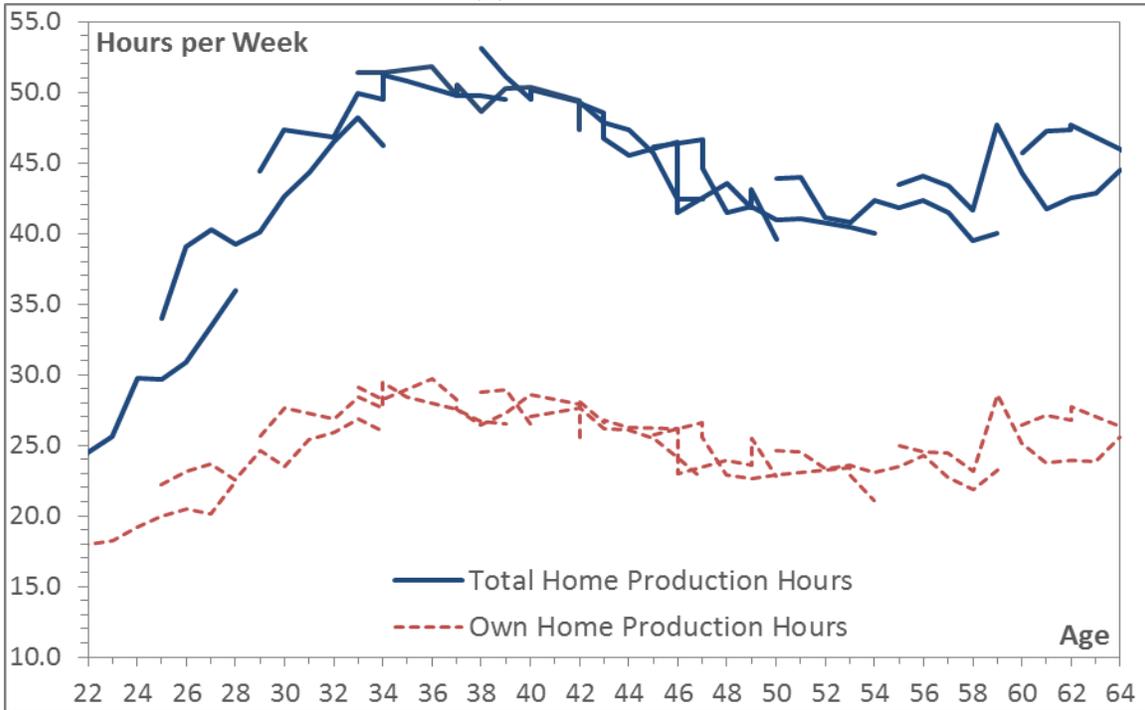


Notes: Each panel illustrates the total hours spent working and total hours spent on household work for 22 to 64 year olds. The upper panel reports the pooled means across years for all individuals and the bottom panel reports the pooled means for those reporting positive wages and work hours in ATUS data between 2003 and 2012.

Figure 2 Individual and Total Household Home Production over the Life Cycle
 (a) All Employed

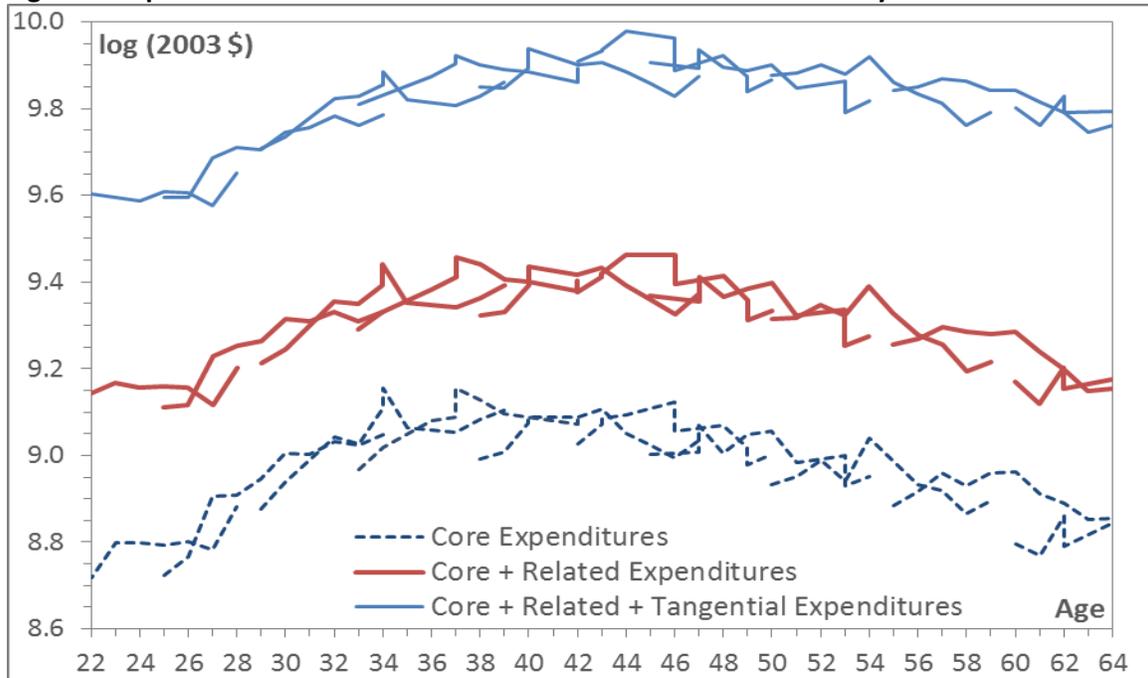


(b) All Individuals



Notes: Figure illustrates the mean hours spent on household work for employed 22 to 64 year olds by cohort and age calculated from ATUS data for 2003-12. Each continuous line represents one cohort. Dashed lines represent hours per week performed by the respondent, while solid lines represent total hours per week performed by both the respondent and spouse or partner, if any is present.

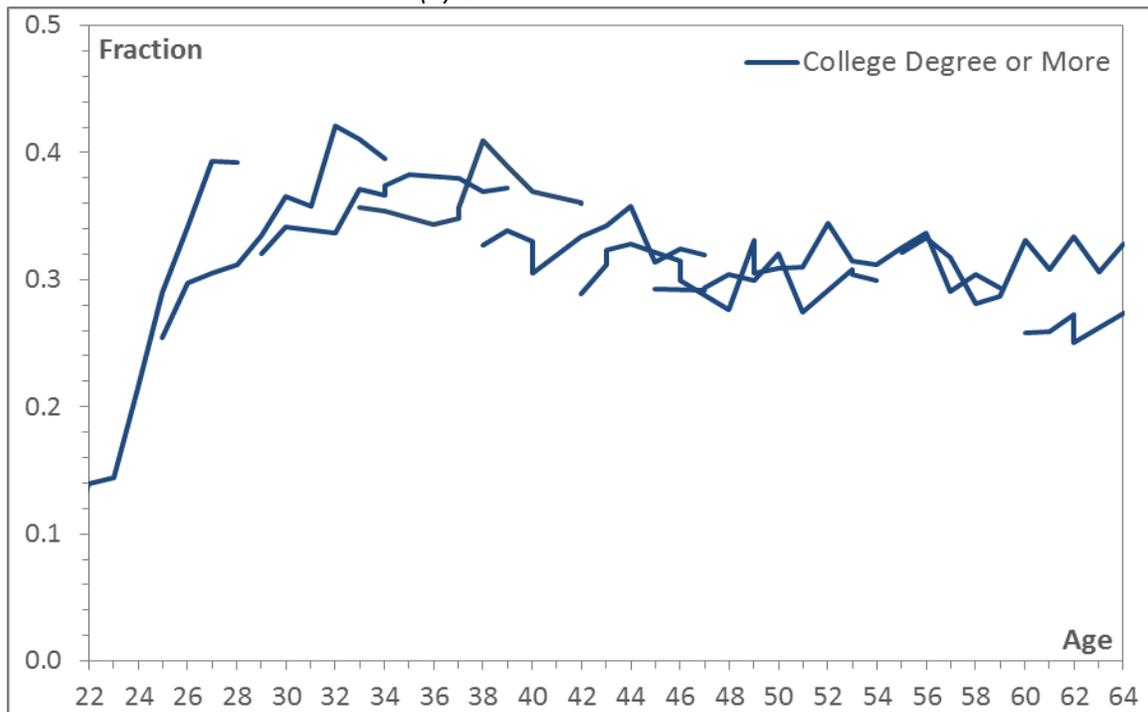
Figure 3. Expenditures on Household Goods and Services over the Life-Cycle



Notes: Figure reports log real expenditures on household goods and services for all individuals aged 22-64 by cohort and age calculated using ATUS and CEX data for 2003-12. Each continuous line represents one cohort. See text for details of what is included in core, related and tangential expenditures on household goods and services.

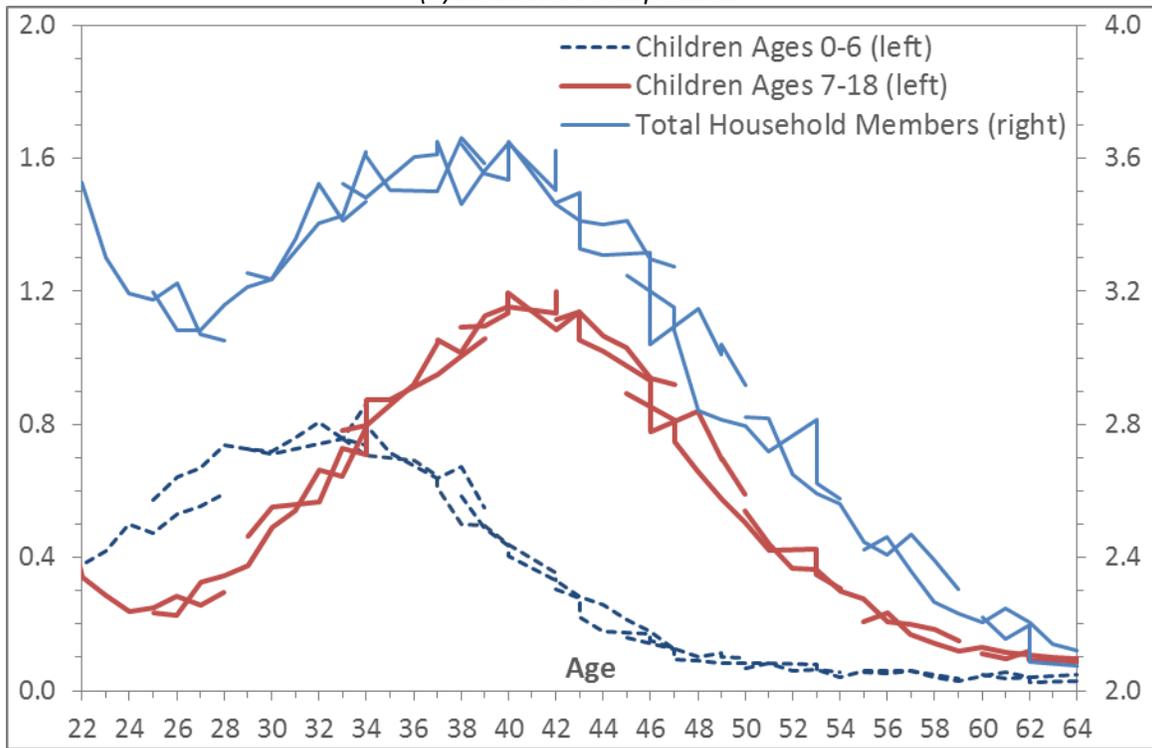
Figure 4. Demographic and Household Characteristics over the Life Cycle

(a) Educational Attainment

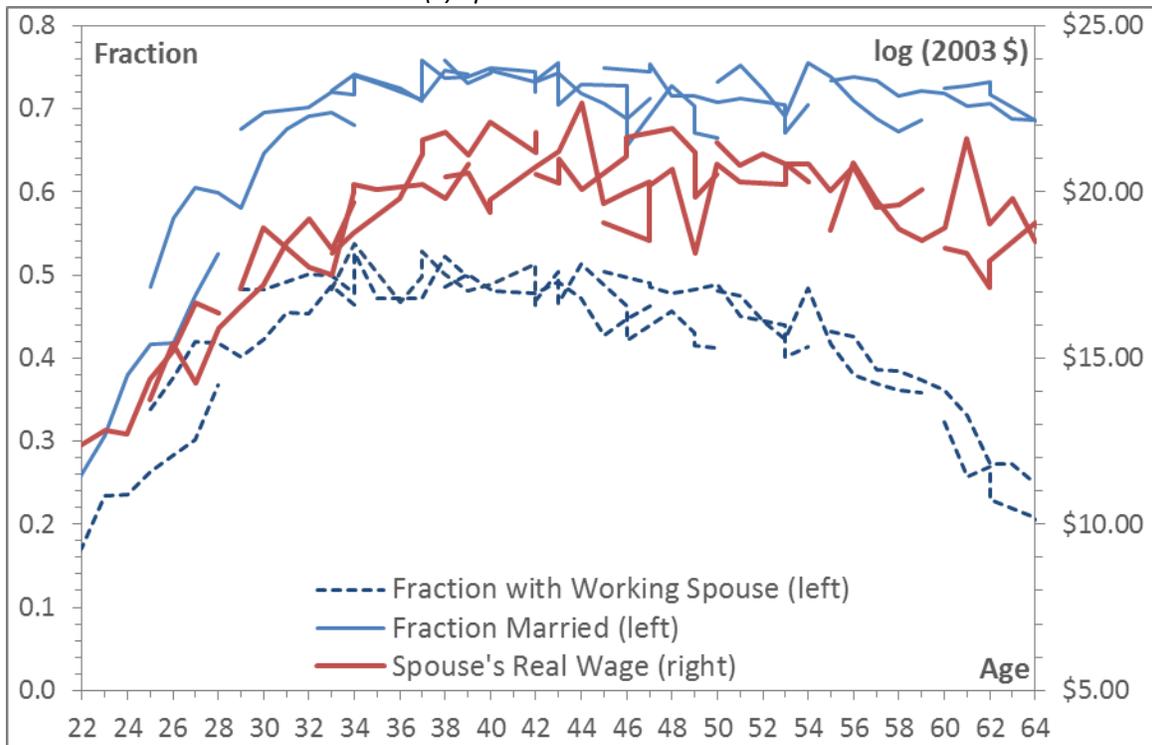


(continued on next page)

(b) Household Composition

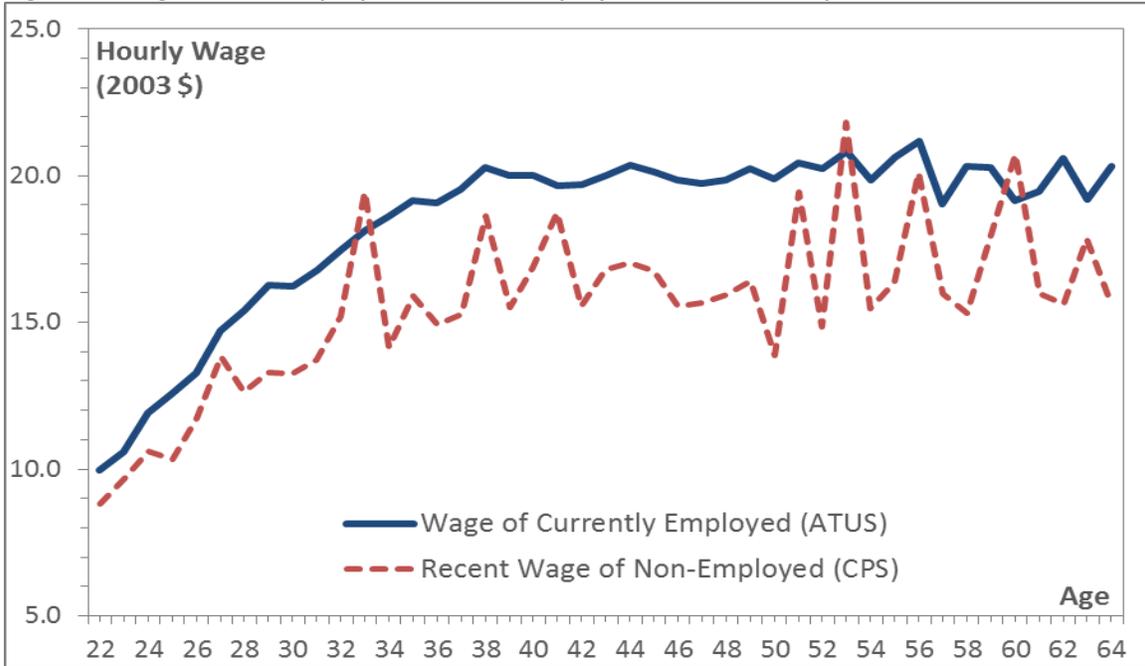


(c) Spousal Characteristics



Notes: Each panel of figure reports the listed household or demographic means for all individuals aged 22-64 by cohort and age calculated using ATUS data for 2003-12. Each continuous line represents one cohort. The spouse's real wage is conditional on having a working spouse present.

Figure 5. Wages of the Employed and Non-Employed over the Life Cycle



Notes: Figure reports real hourly wages of 22-64 year olds by age calculated using ATUS data for 2003-12. The solid line depicts the mean wage of those employed at the time of the ATUS survey, while the dashed line represents the mean previous wage of those non-employed at the time of their ATUS interview but employed at the time of their last CPS interview.

Appendix

A. Measurement

Wages

I adjust the wage measure to minimize the incidence of measurement error due to the use of reported usual hours in the calculation of an hourly wage measure from weekly earnings data. In doing so, I use data from the Merged Outgoing Rotation Groups (MORG) supplement of the Current Population Survey. The MORG data are ideal for this adjustment because they represent a much larger sample than what is in the ATUS and it is also the ATUS sample frame.

There are two issues with using weekly earnings to derive a wage measure. The first is division bias. Namely, work hours appear in the denominator of the wage measure and as the dependent variable in my regressions, causing a downward bias in the coefficient on wages. I can address this with the fact that my dependent variable is a measure of work hours obtained directly from time-diary data while my measure of “usual” weekly work hours is reported as part of the CPS-reinterview portion of the ATUS survey. The second issue is sampling error that arises from the relatively small sample sizes of my synthetic cohort-year cells. I use the usual hours measure from the MORG to deal with the second issue.

In both the CPS and ATUS, respondents can report their earnings as either hourly or weekly. The following adjustment only pertains to those that report their hours weekly. One can express the mean wage for a sample cohort cell i in year t as

$$w_{it} = \theta_{it}^{hr} w_{it}^{hr} + (1 - \theta_{it}^{hr}) w_{it}^{wk} \quad (\text{A.1})$$

where w_{it}^{hr} is the (sample-weighted) mean hourly earnings within the cell for respondents who report their earnings hourly, w_{it}^{wk} is the (sample-weighted) mean hourly earnings for respondents who report their wages weekly, and θ_{it}^{hr} is the (weighted) fraction of cell respondents that report their earnings

hourly. Furthermore, $w_{it}^{wk} = E_{it}^{wk} / n_{it}^{wk}$, where E_{it}^{wk} is the mean weekly earnings reported by those who report their earnings weekly, and n_{it}^{wk} is the usual weekly hours worked reported by those who report their earnings weekly.

My adjustment uses an estimate of usual weekly hours, estimated at the cohort-year level in my main analysis, that comes from the predicted relationship between n_{it}^{wk} measured from the ATUS data, and n_{it}^{wk} measured from the MORG data. Specifically, I divide the MORG sample into the same cohort-years cells I use in my analysis of ATUS data and generate an estimate of n_{it}^{wk} for each cohort-year observation. I then regress the ATUS estimate of n_{it}^{wk} on the MORG estimate of n_{it}^{wk} to obtain a predicted estimate of usual hours, \hat{n}_{it}^{wk} . I then use the predicted hours estimate to get a predicted estimate of average hourly earnings, $\hat{w}_{it}^{wk} = E_{it}^{wk} / \hat{n}_{it}^{wk}$. Finally, I replace w_{it}^{wk} with \hat{w}_{it}^{wk} in equation (A.1) above to obtain an adjusted measure of average hourly earnings.

I use the same adjustment method throughout my analysis. Thus, when I estimate the Frisch elasticity using pooled age means, I pool both the ATUS and MORG data into age cells, calculate the above measures across pooled age categories, and perform the same adjustment. When estimating the Frisch elasticity while including the non-employed, I use either the same wage estimate as above, or a wage estimate that includes a weighted average of the above estimate and an estimate obtained from the wages reported in the respondents' CPS outgoing interview. In the latter case, the CPS-derived wage is adjusted in a similar manner but with an additional adjustment to the weekly earnings measure to account for the fact that the number of respondents within each cohort-year cell that are non-employed but reported an earned wage in their outgoing CPS interview is a very sparse sample.

Specifically, to generate an estimate of the mean CPS wage within a cohort-year cell, I follow the estimation approach outlined above, but calculate a mean CPS wage in accordance with equation (A.1) with the following adjustments. First, all measures, θ_{it}^{hr} , w_{it}^{hr} , E_{it}^{wk} , and n_{it}^{wk} come from the outgoing

CPS interview. The sample, and therefore the fraction reporting an hourly wage, θ_{it}^{hr} , is based on only those individuals who were non-employed during the ATUS interview and report positive earnings and positive usual hours in the outgoing CPS interview. With the prevalence of sparse cells, outliers in the reporting of weekly earnings can lead to large fluctuations in the adjusted CPS wage over time within cohorts. Therefore, I Winsorize the distribution of reported weekly earnings across all individuals who only report their CPS earnings weekly, adjusting the top and bottom 5 percent of the distribution when doing so. The usual hours worked during the outgoing CPS interview are again based on a predicted relationship with the same mean hours estimated from the MORG sample as before. The wage estimate for those reporting weekly is now $\hat{w}_{it}^{wk} = \tilde{E}_{it}^{wk} / \hat{n}_{it}^{wk}$, where \tilde{E}_{it}^{wk} represents the mean weekly earnings estimated from the Winsorized distribution. All other aspects of the adjustment are the same as described earlier, and the adjusted wage used in the estimation of the Frisch elasticity is the sample-weighted average of the wage reported for those currently employed during the ATUS interview and the CPS wage reported for those non-employed but employed at the time of their outgoing CPS interview.

Spousal Home Production Hours

I account for spousal home production by creating mean data for “synthetic spouses” from the ATUS data for each cohort-year observation. The fact that I only have time diary data for the respondent prevents a direct measurement of spousal home production hours. The synthetic spouse is created using relevant information on gender, marital status, employment status, and spousal employment status within each cohort cell to generate the mean home production hours of the cell’s representative synthetic spouse.

Specifically, I split each cohort-year or pooled age cell into males and females, then further split each group into those with and without a spouse or partner present, and into those with and without a working spouse, and calculate the sample-weighted weekly mean hours spent in home production from

their time diary data. Since the ATUS surveys an individual from the household randomly (thus keeping the share female at roughly 50 percent for all cells), the average ages of respondents within each cell and their spouse/partner tend to be roughly the same. Nevertheless, I account for differences in spousal age. I use the estimated shares of those with a spouse or partner present, the employment statuses of the individuals and their spouse/partners, the median age of their spouses, and the gender distribution within each cell to assign a mean spousal home production value, properly weighted, to each cohort-year or pooled age cell.

Formally, suppress the cohort and year subscripts (or age subscripts, in the case of the pooled age mean regressions) and consider individuals within a single cohort-year cell with sex $j = \{m, f\}$ (male, female), employment status $k = \{e, ne\}$, and spousal employment status $l = \{se, sn\}$, with marital status defined as having a spouse or partner present or not. This categorization creates 16 possible age-conditional spousal home production values within each cohort-year cell, though only eight are relevant in practice since spousal home production is zero for the non-married by definition.

I start with the subsample of individuals to those that have a spouse or partner present across all cohort-years to allow for potential matching across cohorts. I then calculate, for each age regardless of cohort or survey year, mean home production hours by gender, employment status, and spouse's employment status. In doing so, I aggregate the time diary data into synthetic weeks just as I do with the other hours measures from the ATUS. Next, for each $\{j, k\}$ category within each cell, I calculate the fraction of individuals with a spouse or partner present, the fraction with a working spouse (conditional on having one present), and the median age of the spouse, where the latter is calculated separately for employed and non-employed spouses.

The calculation of spousal home production hours rests on the assumption that the calculated home production means can be matched to members of the opposite sex with a spouse or partner present and complementary employment statuses for themselves and their spouses. That is,

$$(A.1) \quad h^{jkn}(a) = h^{j'nk}(a^S),$$

where j' denotes the opposite sex and a^S is the median age of the spouse. Given the matching assumption and the group-specific mean home production hours, deriving mean spousal home production for the cohort-year cell is simply a matter of aggregation. First, note that home production hours are matched on age by matching the mean hours for a given age to the median spousal age within each group. For each gender and employment status (of respondent), let μ^{jk} equal the fraction married and λ^{jk} equal the fraction with a working spouse, conditional on being married. Then mean spouse home production hours for the $\{j, k\}$ cell is

$$(A.2) \quad H^{jk} = \mu^{jk} [\lambda^{jk} h^{jk,se}(a^S) + (1 - \lambda^{jk}) h^{jk,sn}(a^S)].$$

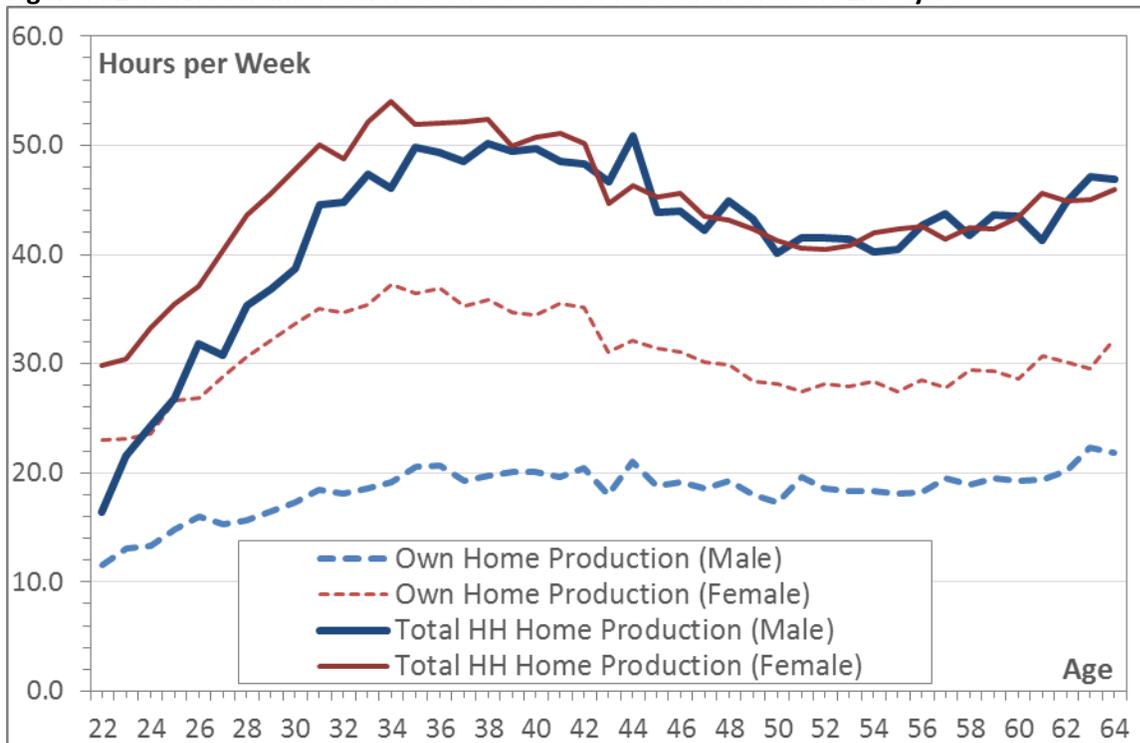
Remember that spousal home production is zero for the $1 - \mu^{jk}$ fraction of individuals in the group.

Letting the θ^f be the fraction of the cell that is female, the full aggregation is then

$$(A.3) \quad h^S = \theta^f \sum_k H^{fk} + (1 - \theta^f) \sum_k H^{mk}.$$

When I split the samples in my analysis by gender, the gender-specific spousal home production estimate is simply $\sum_k H^{fk}$. Figure A.1 below presents the respondent's home production and total household home production (the sum of respondent and spousal home production) for all individuals by gender. It is the analog to Figure 2 (bottom panel) in the main text.

Figure A.1 Individual and Total Household Home Production over the Life Cycle



Notes: Figure illustrates the mean hours spent on household work for employed 22 to 64 year olds by cohort and age calculated from ATUS data for 2003-12. Each continuous line represents one cohort. Dashed lines represent hours per week performed by the respondent, while solid lines represent total hours per week performed by both the respondent and spouse or partner, if any is present.

Household Goods and Services Expenditures

I measure expenditures on household goods and services using data from the Consumer Expenditure Survey (CEX). In doing so, I include expenditure categories that best resemble expenditures that could be substituted with time. I group these goods into three categories, since which expenditures fit this definition is debatable. “Core” household goods and services include spending on food, household services (housekeeping, maintenance and repair), child care, and personal care services. These represent the goods and services that are easiest to substitute with time. “Related” expenditures include spending on transportation maintenance and gasoline, non-durable entertainment spending (including rentals related to leisure), and professional services. “Tangential” expenditures include spending on utilities, public transit, healthcare, books and educational supplies, and clothing. My main

expenditure measure includes expenditures in the “core” and “related” categories, but I also experiment with including only “core” expenditures, and including “tangential” expenditures. Estimates of equation (5) in the main text using these different measures are reported in Table C.1.

I restrict the CEX sample to include individuals who respond to the survey for four straight quarters, and include both household heads (the primary survey respondent) and any spouse present. I then identify all household members and their spouses within the sample age range (i.e., born between 1940 and 1987). To avoid double-counting households, I randomly draw either the head of household or the spouse (when a spouse is present) for inclusion in the sample. This generates a sampling procedure that mimics the ATUS sampling methodology, which draws a single individual aged 16 or older from each household regardless of head-of-household status. There will remain small differences among the young in my samples because my methodology will fail to capture individuals living with their parents (since they are neither heads of household nor a spouse). My method for generating a predicted measure of household expenditures (akin to a two-stage least squares estimator, described in the next section) addresses this sampling difference. I group all individuals into cohort-year observations that are identical to the cohort-year observations created from the ATUS. When studying the extensive margin, I include the non-employed (with the same sampling approach applied to heads-of-households and spouses) and similarly aggregate the cohort-year observations. I then estimate the (sample-weighted) mean statistics for the demographic and household characteristics of each cell. These statistics include the fraction female, the fraction white and non-Hispanic, the fraction with a spouse or partner present, the fraction with a working spouse, the fractions with either some college, a college degree, or a graduate degree, the fraction with a child present in the household, the mean number of children aged 0 to 6, the mean number of children aged 7 to 18, and mean household size. I then regress the log of mean household expenditures on these variables in addition to fixed effects for cohort and survey year.

All regressions are weighted by the size of the cohort-year cell, and regressions that include the non-employed include an additional regressor for the fraction employed.

The regression provides a predicted relationship between the observable characteristics and household expenditures. I can calculate the same explanatory variables from the ATUS data, similarly aggregated for each cohort-year observation. Interacting the regression coefficients from the CEX data with the cohort-year data from the ATUS provides me with an estimate of household expenditures for each cohort-year cell that is free of error due to sampling differences between the ATUS and CEX within cohort-year cells. The approach is identical to the two-sample instrumental variables approach developed by Angrist and Krueger (1992) and extended to a similar setting as the one here by Japelli, Pischke, and Soueles (1998). In the following section, I show that my IV estimates using this predicted measure are consistent, and require a fairly straightforward adjustment to the standard errors to be consistent as well.

B. Corrected Standard Errors for Household Expenditures

The estimates of the Frisch elasticity in my main specification (equation (5) in the text) use a measure of household expenditures that derived using predictions from CEX data interacted with household and demographic characteristics from the ATUS. The approach is similar to the two-sample instrumental variables approach developed by Angrist and Krueger (1992) and extended by Jappeli, Pischke, and Soueles (1998). As Japelli et al. show, the estimator that uses the predicted values of household expenditures using coefficients from the first data set (the CEX) and data from the second data set (the ATUS) will be consistent, but an adjustment needs to be made to the standard errors for it to be asymptotically efficient. I make such an adjustment to my standard errors, but the approach is complicated by the fact that my predicted variable is itself endogenous and is instrumented using two-stage least squares (2SLS). To deal with this, I develop what I call a two-stage, two-sample IV estimator

(2S-TSIV). I show that the estimator is consistent and derive the correction needed to generate efficient standard errors.

This exercise is much easier if one simplifies the problem by aggregating variables into familiar matrices that are easier to work with. Begin with the full 2SLS specification of equation (5) in the text.

$$(B.1a) \quad \ln n_{it} = \alpha_i + \psi t + X_{i(t)}\tilde{\delta} + \beta_{01} \ln w_{it} + \beta_{02} \ln h_{it} + \beta_{03} \ln h_{it}^s + \beta_1 \widehat{\ln x_{it}} + \epsilon_{it},$$

$$(B.21b) \quad \begin{bmatrix} \ln w_{it} \\ \ln h_{it} \\ \ln h_{it}^s \\ \widehat{\ln x_{it}} \end{bmatrix} = \lambda_i + \xi t + X_{i(t)}\tilde{\eta} + V_{it}\mu + \epsilon_{it},$$

$$(B.1c) \quad \ln x_{it} = Z_{it}\pi + u_{it}.$$

I have changed the notation slightly to account for the fact that household expenditures is a predicted value and to accommodate the simplified notation that follows. The first three terms on the right-hand side of (B.1b) represent the included exogenous variables from the first stage of the 2SLS regression, while the vector V_{it} represents the instruments. The vector Z_{it} represents the variables used to predict household expenditures.¹

Suppress the time and cohort subscripts and let data from the CEX use a 1 subscript and data from the ATUS use a 2 subscript. Furthermore, let A_j represent the vector of included exogenous variables in data set j and δ represent the coefficients on these estimates. Let W_j represent the vector of endogenous variables other than household expenditures and β_0 their coefficients. Finally, simplify the notation to let $y_j \equiv \ln n_j$ and $L_j \equiv \ln x_j$. Then, we can rewrite the estimation of (B.1) as

$$(B.2a) \quad y_2 = A_2\delta + W_2\beta_0 + \hat{L}_{21}\beta_1 + \epsilon_2,$$

$$(B.2b) \quad \begin{bmatrix} W_2 \\ \hat{L}_{21} \end{bmatrix} = A_2\eta + V_2\mu + \epsilon_2,$$

¹ Note that in my implementation of the 2S-TSIV estimator, some variables appear in both V_{it} and Z_{it} . This will not matter so long as the variables are exogenous.

$$(B.2c) \quad L_1 = Z_1\pi + u_1.$$

Where $\hat{L}_{21} = Z_2\hat{\pi}$ is the predicted household expenditures derived from interacting the household and demographic characteristics from the ATUS with the coefficients derived from regressing actual expenditures on these characteristics using the CEX data. Keep in mind that that the two-sample strategy arises because I have data on L in the CEX but not the ATUS. In deriving the covariance matrix, I also have to deal with the fact that I have data on W in the ATUS but not the CEX.

The notation for proving consistency and deriving the efficient estimator becomes easier if we aggregate some matrices further. Let $\hat{X}_{21} = [W_2 \quad \hat{L}_{21}]$, $\beta = [\beta_0 \quad \beta_1]$, $C_2 = [A_2 \quad V_2]$, $\rho = [\eta \quad \mu]$, $\hat{D}_{21} = [A_2 \quad \hat{X}_{21}]$, and $\theta = [\delta \quad \beta]$. Then the equations in (B.2) become

$$(B.3a) \quad y_2 = \hat{D}_{21}\theta + \varepsilon_2,$$

$$(B.3b) \quad \hat{X}_{21} = C_2\rho + \epsilon_2,$$

$$(B.3c) \quad L_1 = Z_1\pi + u_1.$$

Asymptotic consistency and efficiency rely on the following assumptions:

1. $\text{plim } \hat{\pi} \rightarrow \pi$. This asserts the consistency of $\hat{\pi}$. Since $\hat{\pi} = (Z_1'Z_1)^{-1}Z_1'L_1$, this relies on the orthogonality assumption that $E(Z_1'u_1) = 0$.
2. $E(C_2'\varepsilon_2) = 0$. This is the orthogonality condition for the instruments used in the 2SLS estimation of \hat{X}_{21} . This simply assumes that the instruments are valid.
3. $\text{plim } (D_2'\Phi_{CC}D_2) = \text{plim } (D_1'\Phi_{CC}D_1) = \text{plim } (D_{12}'\Phi_{CC}D_{12}) \rightarrow \Sigma_D'\Phi_{CC}\Sigma_D$, where $D_j = [A_j \quad W_j \quad L_j]$, $D_{12} = [A_2 \quad W_2 \quad L_1]$ (the matrix of right-hand side variables from the ATUS save for (log) household expenditures, which is instead represented by actual (log) expenditures from the CEX), and $\Phi_{CC} \equiv C_2(C_2'C_2)^{-1}C_2'$ is the standard symmetric idempotent weighting matrix of a 2SLS regression.² This assumption asserts that sample moments derived from each sample converge to the population moment.
4. The data $\{y_2, A_2, W_2, Z_2, V_2\}$ and $\{L_1, Z_1\}$ are jointly independent.

² The matrix D_{12} can only exist when the number of observations in the ATUS and CEX are equal. This condition is satisfied because I estimate the model in (B.3) using the identical cohort-years in each data set.

With these assumptions, it is straightforward to show that $\hat{\theta}_{IV} = (\hat{D}_{21}'\Phi_{CC}\hat{D}_{21})^{-1}\hat{D}_{21}'\Phi_{CC}y_2$ is a consistent estimator of θ . Consistency is based on the first three assumptions. Since $\text{plim } \hat{\pi} \rightarrow \pi$, $\text{plim } \hat{D}_{21}'\Phi_{CC}\hat{D}_{21} = \text{plim } D_2'\Phi_{CC}D_2 \rightarrow \Sigma_D'\Phi_{CC}\Sigma_D$. In the limit, $y_2 = D_2\theta + \varepsilon_2$, so we have $\hat{\theta}_{IV} = (\hat{D}_{21}'\Phi_{CC}\hat{D}_{21})^{-1}\hat{D}_{21}'\Phi_{CC}D_2\theta + (\hat{D}_{21}'\Phi_{CC}\hat{D}_{21})^{-1}\hat{D}_{21}'\Phi_{CC}\varepsilon_2$, which converges in probability to θ since $E(C_2'\varepsilon_2) = 0$.

To obtain the covariance matrix, define $g(\theta) \equiv \hat{D}_{21}'\Phi_{CC}y_2 + \hat{D}_{12}'\Phi_{CC}D_{12}\theta$, where $\hat{D}_{12} = [A_2 \quad W_2 \quad \hat{L}_{11}]$ and $\hat{L}_{11} = Z_1\hat{\pi}$, i.e., the predicted value from the regression that generates $\hat{\pi}$. Premultiply the first term of $g(\theta)$ by $(\hat{D}_{21}'\Phi_{CC}\hat{D}_{21})^{-1}$ and the second term by $(\hat{D}_{12}'\Phi_{CC}\hat{D}_{12})^{-1}$. Appealing to the same consistency check as before and the assumption that the two samples are jointly independent, we have that, asymptotically, $(\hat{\theta} - \theta) \sim N(0, \text{cov}(\hat{\theta}_{IV}) + \text{cov}(\hat{\omega}))$, where $\hat{\omega} \equiv (\hat{D}_{12}'\Phi_{CC}\hat{D}_{12})^{-1}\hat{D}_{12}'\Phi_{CC}D_{12}\theta$. Therefore, the correct covariance matrix is the sum of the covariance matrix from the 2SLS regression that uses \hat{L}_{21} and the covariance matrix of $\hat{\omega}$. One can obtain the latter by using the same 2SLS estimation in the regression of $D_{12}\theta$ on \hat{D}_{12} , where $D_{12}\theta$ is itself the predicted value from the 2SLS regression of y_2 on D_{12} . Thus, one can obtain the correct standard errors for the estimator using standard regression software, just as in the cases of Angrist and Krueger (1992) and Jappeli, Pischke, and Soueles (1998).

C. Additional Results

This section reports additional results of the analysis. Table C.1 reports the Frisch elasticity estimates and the coefficients on (log) household expenditures for the fully-specified model using alternative measures of household expenditures that either only include “core” household spending, or instead include “core,” “related,” and “tangential” measures of spending. The first column reports the results from the main analysis using “core” and “related” expenditures (Tables 5 and 6, column 4), and the last column reports the results of using “core” and “related” expenditures without accounting for

differences in demographic and household characteristics between the ATUS and CEX, and using the predicted expenditures from the CEX data alone. All but the last column have their standard errors corrected as described in the previous section. As one can see, the alternative measures of household expenditures have little effect on the estimates of the Frisch elasticity. For the “micro” Frisch elasticity, the estimate from the main analysis is 0.39 while the other estimates all range between 0.29 and 0.34. For the “macro” Frisch elasticity, the estimate from the main analysis is 0.92 while the other estimates are between 0.91 and 0.93.

Tables C.2 and C.3 report the full regression results for the Frisch elasticities reported for men and women, respectively in Table 8 of the main text. The results are discussed in more detail there.

Table C.1. Synthetic Cohort Wage Elasticity Estimates, Alternative Expenditure Measures

	IV, Employed Only			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.391* (0.236)	0.291 (0.258)	0.342 (0.260)	0.312 (0.276)
log(household expenditures), $\ln x_t$	0.135 (0.179)	0.270 (0.231)	0.270 (0.231)	0.170 (0.184)
Expenditure Measure	ATUS, Core & Related	ATUS, Core Only	ATUS, All Household	CEX, Core & Related
	IV, All Individuals			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.920* (0.489)	0.917* (0.486)	0.928* (0.520)	0.914* (0.540)
log(household expenditures), $\ln x_t$	0.639* (0.387)	0.799* (0.444)	0.742* (0.483)	0.480 (0.370)
Expenditure Measure	ATUS, Core & Related	ATUS, Core Only	ATUS, All Household	CEX, Core & Related

Note: The table reports the 2SLS estimates from regressing log(reported work hours), $\ln n_t$, using the specification from equation (3) in the text and alternate measures of household expenditures ($\ln x_t$). The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include either all employed individuals or all individuals regardless of employment status, as indicated in the table. The non-employed use a wage derived from their CPS outgoing interview.

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Table C.2. Synthetic Cohort Panel Regression Estimates, Males Only

	IV, Employed Males			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.186* (0.098)	0.240* (0.175)	0.231* (0.128)	0.038 (0.188)
log(own home production), $\ln h_t$		-0.101 (0.195)	-0.441 (0.325)	-0.649* (0.390)
log(spousal home production), $\ln h_t^S$			0.119 (0.095)	0.390** (0.172)
log(household expenditures), $\ln x_t$				0.112 (0.355)
Adj. R-squared	0.095	0.112	0.154	0.112
F-test of Overidentification [p-value]	0.93 [0.501]	1.18 [0.325]	1.39 [0.224]	0.78 [0.592]
First-stage F-test of excluded instruments [p-value]. ¹	13.86 [0.000]	2.26 [0.023]	51.15 [0.000]	9.31 [0.000]
	IV, All Males			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.731** (0.153)	0.905** (0.168)	0.796** (0.196)	0.731** (0.237)
log(own home production), $\ln h_t$		-0.617** (0.286)	-0.953** (0.427)	-0.924** (0.414)
log(spousal home production), $\ln h_t^S$			0.160 (0.153)	0.141 (0.152)
log(household expenditures), $\ln x_t$				0.154 (0.341)
Adj. R-squared	0.765	0.780	0.784	0.799
F-test of Overidentification [p-value]	1.28 [0.261]	0.92 [0.505]	0.92 [0.496]	1.16 [0.338]
First-stage F-test of excluded instruments [p-value]. ¹	8.83 [0.000]	2.96 [0.004]	47.69 [0.000]	11.66 [0.000]

Note: The table reports the 2SLS estimates from regressing log(reported work hours), $\ln n_t$, on the listed dependent variables, using the specification from equation (3) in the text. The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include either all employed males or all males regardless of employment status, as indicated in the table. The non-employed use a wage derived from their CPS outgoing interview.

1. Row lists the F-value for the test of excluded instruments for the endogenous regressor added in each column: log(real wage) in column (1), log(own home production) in column (2), log(spousal home production) in column (3), and log(household expenditures) in column (4).

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Table C.3. Synthetic Cohort Panel Regression Estimates, Females Only

	IV, Employed Females			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.076 (0.095)	0.259** (0.122)	0.235* (0.133)	0.211 (0.135)
log(own home production), $\ln h_t$		-0.283** (0.134)	-0.339* (0.180)	-0.506** (0.206)
log(spousal home production), $\ln h_t^S$			0.038 (0.082)	0.088 (0.088)
log(household expenditures), $\ln x_t$				0.240* (0.139)
Adj. R-squared	0.137	0.208	0.203	0.210
F-test of Overidentification [<i>p</i> -value]	1.01 [0.438]	0.83 [0.578]	0.93 [0.488]	0.54 [0.773]
First-stage F-test of excluded instruments [<i>p</i> -value]. ¹	14.86 [0.000]	8.72 [0.000]	27.97 [0.000]	5.89 [0.000]
	IV, All Females			
	(1)	(2)	(3)	(4)
log(real wage), $\ln w_t$	0.379** (0.178)	0.883** (0.225)	0.843** (0.240)	0.514** (0.246)
log(own home production), $\ln h_t$		-0.839** (0.234)	-0.955** (0.391)	-1.600** (0.408)
log(spousal home production), $\ln h_t^S$			0.146 (0.292)	0.284 (0.280)
log(household expenditures), $\ln x_t$				1.100** (0.310)
Adj. R-squared	0.679	0.699	0.695	0.728
F-test of Overidentification [<i>p</i> -value]	3.92 [0.001]	1.93 [0.068]	2.16 [0.049]	0.41 [0.872]
First-stage F-test of excluded instruments [<i>p</i> -value]. ¹	7.03 [0.000]	22.10 [0.000]	18.20 [0.000]	8.28 [0.132]

Note: The table reports the 2SLS estimates from regressing log(reported work hours), $\ln n_t$, on the listed dependent variables, using the specification from equation (3) in the text. The regressions are run on cohort-year observations for cells with a median age between 22 and 64 years old ($N = 92$) during the 2003-12 period. Cohort-year cells include either all employed females or all females regardless of employment status, as indicated in the table. The non-employed use a wage derived from their CPS outgoing interview.

1. Row lists the *F*-value for the test of excluded instruments for the endogenous regressor added in each column: log(real wage) in column (1), log(own home production) in column (2), log(spousal home production) in column (3), and log(household expenditures) in column (4).

** Denotes significance at the 5% level.

* Denotes significance at the 10% level.

Working Paper Series

A series of research studies on regional economic issues relating to the Seventh Federal Reserve District, and on financial and economic topics.

Examining Macroeconomic Models through the Lens of Asset Pricing <i>Jaroslav Borovička and Lars Peter Hansen</i>	WP-12-01
The Chicago Fed DSGE Model <i>Scott A. Brave, Jeffrey R. Campbell, Jonas D.M. Fisher, and Alejandro Justiniano</i>	WP-12-02
Macroeconomic Effects of Federal Reserve Forward Guidance <i>Jeffrey R. Campbell, Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano</i>	WP-12-03
Modeling Credit Contagion via the Updating of Fragile Beliefs <i>Luca Benzoni, Pierre Collin-Dufresne, Robert S. Goldstein, and Jean Helwege</i>	WP-12-04
Signaling Effects of Monetary Policy <i>Leonardo Melosi</i>	WP-12-05
Empirical Research on Sovereign Debt and Default <i>Michael Tomz and Mark L. J. Wright</i>	WP-12-06
Credit Risk and Disaster Risk <i>François Gourio</i>	WP-12-07
From the Horse's Mouth: How do Investor Expectations of Risk and Return Vary with Economic Conditions? <i>Gene Amromin and Steven A. Sharpe</i>	WP-12-08
Using Vehicle Taxes To Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden <i>Thomas Klier and Joshua Linn</i>	WP-12-09
Spending Responses to State Sales Tax Holidays <i>Sumit Agarwal and Leslie McGranahan</i>	WP-12-10
Micro Data and Macro Technology <i>Ezra Oberfield and Devesh Raval</i>	WP-12-11
The Effect of Disability Insurance Receipt on Labor Supply: A Dynamic Analysis <i>Eric French and Jae Song</i>	WP-12-12
Medicaid Insurance in Old Age <i>Mariacristina De Nardi, Eric French, and John Bailey Jones</i>	WP-12-13
Fetal Origins and Parental Responses <i>Douglas Almond and Bhashkar Mazumder</i>	WP-12-14

Working Paper Series *(continued)*

Repos, Fire Sales, and Bankruptcy Policy <i>Gaetano Antinolfi, Francesca Carapella, Charles Kahn, Antoine Martin, David Mills, and Ed Nosal</i>	WP-12-15
Speculative Runs on Interest Rate Pegs The Frictionless Case <i>Marco Bassetto and Christopher Phelan</i>	WP-12-16
Institutions, the Cost of Capital, and Long-Run Economic Growth: Evidence from the 19th Century Capital Market <i>Ron Alquist and Ben Chabot</i>	WP-12-17
Emerging Economies, Trade Policy, and Macroeconomic Shocks <i>Chad P. Bown and Meredith A. Crowley</i>	WP-12-18
The Urban Density Premium across Establishments <i>R. Jason Faberman and Matthew Freedman</i>	WP-13-01
Why Do Borrowers Make Mortgage Refinancing Mistakes? <i>Sumit Agarwal, Richard J. Rosen, and Vincent Yao</i>	WP-13-02
Bank Panics, Government Guarantees, and the Long-Run Size of the Financial Sector: Evidence from Free-Banking America <i>Benjamin Chabot and Charles C. Moul</i>	WP-13-03
Fiscal Consequences of Paying Interest on Reserves <i>Marco Bassetto and Todd Messer</i>	WP-13-04
Properties of the Vacancy Statistic in the Discrete Circle Covering Problem <i>Gadi Barlevy and H. N. Nagaraja</i>	WP-13-05
Credit Crunches and Credit Allocation in a Model of Entrepreneurship <i>Marco Bassetto, Marco Cagetti, and Mariacristina De Nardi</i>	WP-13-06
Financial Incentives and Educational Investment: The Impact of Performance-Based Scholarships on Student Time Use <i>Lisa Barrow and Cecilia Elena Rouse</i>	WP-13-07
The Global Welfare Impact of China: Trade Integration and Technological Change <i>Julian di Giovanni, Andrei A. Levchenko, and Jing Zhang</i>	WP-13-08
Structural Change in an Open Economy <i>Timothy Uy, Kei-Mu Yi, and Jing Zhang</i>	WP-13-09
The Global Labor Market Impact of Emerging Giants: a Quantitative Assessment <i>Andrei A. Levchenko and Jing Zhang</i>	WP-13-10

Working Paper Series *(continued)*

Size-Dependent Regulations, Firm Size Distribution, and Reallocation <i>François Gourio and Nicolas Roys</i>	WP-13-11
Modeling the Evolution of Expectations and Uncertainty in General Equilibrium <i>Francesco Bianchi and Leonardo Melosi</i>	WP-13-12
Rushing into American Dream? House Prices, Timing of Homeownership, and Adjustment of Consumer Credit <i>Sumit Agarwal, LuoJia Hu, and Xing Huang</i>	WP-13-13
The Earned Income Tax Credit and Food Consumption Patterns <i>Leslie McGranahan and Diane W. Schanzenbach</i>	WP-13-14
Agglomeration in the European automobile supplier industry <i>Thomas Klier and Dan McMillen</i>	WP-13-15
Human Capital and Long-Run Labor Income Risk <i>Luca Benzoni and Olena Chyruk</i>	WP-13-16
The Effects of the Saving and Banking Glut on the U.S. Economy <i>Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti</i>	WP-13-17
A Portfolio-Balance Approach to the Nominal Term Structure <i>Thomas B. King</i>	WP-13-18
Gross Migration, Housing and Urban Population Dynamics <i>Morris A. Davis, Jonas D.M. Fisher, and Marcelo Veracierto</i>	WP-13-19
Very Simple Markov-Perfect Industry Dynamics <i>Jaap H. Abbring, Jeffrey R. Campbell, Jan Tilly, and Nan Yang</i>	WP-13-20
Bubbles and Leverage: A Simple and Unified Approach <i>Robert Barsky and Theodore Bogusz</i>	WP-13-21
The scarcity value of Treasury collateral: Repo market effects of security-specific supply and demand factors <i>Stefania D'Amico, Roger Fan, and Yuriy Kitsul</i>	WP-13-22
Gambling for Dollars: Strategic Hedge Fund Manager Investment <i>Dan Bernhardt and Ed Nosal</i>	WP-13-23
Cash-in-the-Market Pricing in a Model with Money and Over-the-Counter Financial Markets <i>Fabrizio Mattesini and Ed Nosal</i>	WP-13-24
An Interview with Neil Wallace <i>David Altig and Ed Nosal</i>	WP-13-25

Working Paper Series *(continued)*

Firm Dynamics and the Minimum Wage: A Putty-Clay Approach <i>Daniel Aaronson, Eric French, and Isaac Sorkin</i>	WP-13-26
Policy Intervention in Debt Renegotiation: Evidence from the Home Affordable Modification Program <i>Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru</i>	WP-13-27
The Effects of the Massachusetts Health Reform on Financial Distress <i>Bhashkar Mazumder and Sarah Miller</i>	WP-14-01
Can Intangible Capital Explain Cyclical Movements in the Labor Wedge? <i>François Gourio and Leena Rudanko</i>	WP-14-02
Early Public Banks <i>William Roberds and François R. Velde</i>	WP-14-03
Mandatory Disclosure and Financial Contagion <i>Fernando Alvarez and Gadi Barlevy</i>	WP-14-04
The Stock of External Sovereign Debt: Can We Take the Data at ‘Face Value’? <i>Daniel A. Dias, Christine Richmond, and Mark L. J. Wright</i>	WP-14-05
Interpreting the <i>Pari Passu</i> Clause in Sovereign Bond Contracts: It’s All Hebrew (and Aramaic) to Me <i>Mark L. J. Wright</i>	WP-14-06
AIG in Hindsight <i>Robert McDonald and Anna Paulson</i>	WP-14-07
On the Structural Interpretation of the Smets-Wouters “Risk Premium” Shock <i>Jonas D.M. Fisher</i>	WP-14-08
Human Capital Risk, Contract Enforcement, and the Macroeconomy <i>Tom Krebs, Moritz Kuhn, and Mark L. J. Wright</i>	WP-14-09
Adverse Selection, Risk Sharing and Business Cycles <i>Marcelo Veracierto</i>	WP-14-10
Core and ‘Crust’: Consumer Prices and the Term Structure of Interest Rates <i>Andrea Ajello, Luca Benzoni, and Olena Chyruk</i>	WP-14-11
The Evolution of Comparative Advantage: Measurement and Implications <i>Andrei A. Levchenko and Jing Zhang</i>	WP-14-12

Working Paper Series *(continued)*

Saving Europe?: The Unpleasant Arithmetic of Fiscal Austerity in Integrated Economies <i>Enrique G. Mendoza, Linda L. Tesar, and Jing Zhang</i>	WP-14-13
Liquidity Traps and Monetary Policy: Managing a Credit Crunch <i>Francisco Buera and Juan Pablo Nicolini</i>	WP-14-14
Quantitative Easing in Joseph's Egypt with Keynesian Producers <i>Jeffrey R. Campbell</i>	WP-14-15
Constrained Discretion and Central Bank Transparency <i>Francesco Bianchi and Leonardo Melosi</i>	WP-14-16
Escaping the Great Recession <i>Francesco Bianchi and Leonardo Melosi</i>	WP-14-17
More on Middlemen: Equilibrium Entry and Efficiency in Intermediated Markets <i>Ed Nosal, Yuet-Yee Wong, and Randall Wright</i>	WP-14-18
Preventing Bank Runs <i>David Andolfatto, Ed Nosal, and Bruno Sultanum</i>	WP-14-19
The Impact of Chicago's Small High School Initiative <i>Lisa Barrow, Diane Whitmore Schanzenbach, and Amy Claessens</i>	WP-14-20
Credit Supply and the Housing Boom <i>Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti</i>	WP-14-21
The Effect of Vehicle Fuel Economy Standards on Technology Adoption <i>Thomas Klier and Joshua Linn</i>	WP-14-22
What Drives Bank Funding Spreads? <i>Thomas B. King and Kurt F. Lewis</i>	WP-14-23
Inflation Uncertainty and Disagreement in Bond Risk Premia <i>Stefania D'Amico and Athanasios Orphanides</i>	WP-14-24
Access to Refinancing and Mortgage Interest Rates: HARPing on the Importance of Competition <i>Gene Amromin and Caitlin Kearns</i>	WP-14-25
Private Takings <i>Alessandro Marchesiani and Ed Nosal</i>	WP-14-26
Momentum Trading, Return Chasing, and Predictable Crashes <i>Benjamin Chabot, Eric Ghysels, and Ravi Jagannathan</i>	WP-14-27
Early Life Environment and Racial Inequality in Education and Earnings in the United States <i>Kenneth Y. Chay, Jonathan Guryan, and Bhashkar Mazumder</i>	WP-14-28

Working Paper Series *(continued)*

Poor (Wo)man's Bootstrap
Bo E. Honoré and Luojia Hu

WP-15-01

Revisiting the Role of Home Production in Life-Cycle Labor Supply
R. Jason Faberman

WP-15-02