University of Michigan Retirement Research Center

Working Paper

WP 2007-171

Trends in the Labor Force Participation of Married Women

Christopher House, John Laitner and Dmitriy Stolyarov

M R
R C

Project #: UM07-15

Trends in the Labor Force Participation of Married Women

Christopher House
University of Michigan

John Laitner
University of Michigan

Dmitriy Stolyarov
University of Michigan

October 2007

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
http://www.mrrc.isr.umich.edu/
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98362-5-04). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Trends in the Labor Force Participation of Married Women

Christopher House, John Laitner and Dmitriy Stolyarov

Abstract

This study seeks to quantify determinants, and costs, of the labor—force participation of married women. We use demographic and earnings data from the Health and Retirement Study. The earnings data constitute an unusually long panel but have the defect of lacking corresponding reports on work hours. By using a highly structured model and concentrating on the participation margin, we nevertheless feel that we can make substantial progress. Our preliminary regression results imply that married women's market work disrupts their household consumption slightly less than one half as much as men's work (relative to complete household retirement). We lay out a course of additional steps that can, we believe, clarify these results even more precisely in the near future.

Authors' Acknowledgements

This paper is preliminary and should not be cited without permission of the authors. This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98362-5). The opinions and conclusions are solely those of the authors and should not be considered as representing the opinions or policy of the Social Security Administration or any agency of the Federal Government.

Trends in the Labor Force Participation of Married Women

1. Introduction

The purpose of this study is to quantify the determinants of female labor–force participation. Our focus is married women. Our approach is to analyze jointly households' consumption–saving choices and the participation decisions of married women. We construct a model that can predict the time path of female labor–force participation given data on children, age, retirement, and earnings. The model allows us to utilize data from the Health and Retirement Study (HRS) (a) to estimate the relative importance of different factors that affect female labor supply, (b) to compare extensive–margin costs of participating in the labor market with intensive–margin costs of how many hours to work conditional on participating at all, and (c) to estimate the direct implicit costs associated with married women's labor supply. We believe that understanding alterations in female labor supply over the past fifty years, and predicting developments that may yet come, is of crucial importance both for anticipating consequences to large programs, such as Social Security, and for keeping economic modeling in touch with some of the most significant changes in the U.S. economy in recent history.

The HRS is a rich source of data on men and women's annual earnings as well as their retirement ages, family demographic variables, and retirement–age net worth. The HRS lacks, however, detailed data on hours worked at most ages. On the other hand, it can provide fairly comprehensive information on which years a woman worked in the labor force and which years she did not. By organizing our analysis around a very structured life–cycle model, this paper attempts to make extensive use of the data that is available.

Although this paper represents a preliminary report, we have non–linear least squares estimates showing that households offset married women's labor–force participation with annual increases of consumption expenditure of about 8 percent. In comparison, fully retired couples seem to decrease their consumption expenditures about 20 percent relative to ages when the male alone worked. Thus, according to one metric, married women's labor–force participation seems to be slightly less than half as disruptive to households as having the male participate (rather than be retired). Other parameter estimates in this paper imply an intertemporal rate of substitution in consumption of about 0.25. That is well within the conventional range, though our procedure is sufficiently nonstandard that our outcomes may constitute independent confirmation.

The organization of this paper is as follows. Section 2 presents our model. Section 3 outlines our estimation strategy. Section 4 describes our data. Section 5 presents our results. Section 6 discusses an extension to simultaneous consideration of male retirement and female labor–force participation at all ages. Section 7 concludes the paper.

2. Model

Our model analyzes married households' decisions of how to allocate women's time between home production and market work, how to allocate lifetime resources over consumption at different ages, and when to retire. At this point, our empirical work (see Sections 4-5) focuses exclusively on women's time allocations — specifically, on women's

year–by–year decisions of whether or not to enter the labor force. Section 6 notes several of the model's implications for retirement — though only from a theoretical standpoint.

We assume that households take prices, wages, and interest rates as given. Prices clear markets, so there is never quantity rationing. Time is continuous, and there is no uncertainty. Households derive utility from a final–consumption aggregate. The final–consumption aggregate depends on a combination of consumption of market goods, leisure, and home production. Let c_{is} denote household i's quantity of consumption of the market good at age s. Normalize the price of the market good to 1 at all times. Then we assume that the flow of utility that household i enjoys at age s depends on $\lambda_{is} \cdot c_{is}$, where the multiplicative factor λ_{is}, which can vary over time and across households, reflects the ability of household i to enjoy market consumption and depends upon its concurrent leisure and home production. Specifically, we assume that household i derives more utility from a given amount of market consumption when its husband and wife can spend more time in leisure and home production. Thus, the household enjoys a higher λ_{is} if the husband is retired and/or if the wife does not participate in the labor market. As a matter of definition, set $\lambda_{is} = 1$ when both the husband and wife participate in the labor force. Then we set $\lambda_{is} = \lambda > 1$ when the husband alone participates in the labor force; and, we set $\lambda_{is} = \Lambda > \lambda$ when neither the husband nor the wife participates in the labor force. For simplicity, we assume that an adult of either sex participates at any age in the labor force full time or not at all, that men participate until they retire (and subsequently not at all), and that wives stop participating when their husband retires — if not before.

Consider a household i that lives from age S_i to T_i. For simplicity, the household's age is its husband's age. Also for simplicity, the interest rate, r, is constant. The household includes a man, a woman, and, possibly, children. The size of the household at age s is N_{is} "equivalent adults" (see, for example, Tobin [1967]). The man and woman both retire when the household reaches age R_i, though, as stated, the woman may have stopped her market–work career earlier. At its inception, a household chooses its life–cycle expenditure profile on market consumption goods, c_{is}; its retirement age, R_i; and the woman's labor force participation profile $p_{is} \in \{0\,,\,1\}$. The notation is $p_{is} = 0$ at age s if the woman is not participating in the labor force, and $p_{is} = 1$ if she is. As stated, we assume $p_{is} = 0$ for all ages $s \geq R_i$. The household's state variables are its age and its asset level, a_{is}.

To optimize, household i solves

$$U(S_i\,,\,T_i) = \max_{c_{is}\,,\,p_{is}\,,\,R_i} \int_{S_i}^{R_i} e^{-\rho \cdot s} \cdot N_{is} \cdot u\big(\frac{\lambda_{is} \cdot c_{is}}{N_{is}}\big)\, ds + V(R_i\,,\,T_i\,,\,a_{iR_i})\,, \tag{1}$$

$$\text{subject to:} \quad \dot{a}_{is} = r \cdot a_{is} + y^M_{is} + p_{is} \cdot y^F_{is} - c_{is}\,,$$

$$a_{i0} = 0\,,$$

where

$$\lambda_{is} = \lambda(s\,,\,p_{is}) \equiv \begin{cases} \lambda\,, & \text{if } s \leq R_i\,,\, p_{is} = 0\,, \\ 1\,, & \text{if } s \leq R_i\,,\, p_{is} = 1\,, \\ \Lambda\,, & \text{if } s > R_i\,, \end{cases} \tag{2}$$

and

$$V(R\,,\,T_i\,,\,A) \equiv \max_{c_{is}} \int_R^{T_i} e^{-\rho\cdot s}\cdot u\big(\frac{\Lambda\cdot c_{is}}{N_{is}}\big)\,ds\,,$$

$$\text{subject to:}\quad \dot{a}_{is} = r\cdot a_{is} - c_{is}\,,$$

$$a_{iR} = A\,,$$

$$a_{i\,,\,T_i} \geq 0\,.$$

Expression (2) captures the change in the value of purchased consumption, c_{is}, with changes in household labor–market participation. We assume that the flow utility function is isoelastic, with intertemporal elasticity of substitution equal to $1/\gamma$. Thus,

$$u(x) \equiv \frac{x^{1-\gamma}}{1-\gamma}\,,\quad \gamma > 0\,.$$

The following proposition characterizes the optimal expenditure path for each household's consumption.

Proposition 1. *Let $(c_{is}\,,\,p_{is}\,,\,R_i)$ be the optimal solution to the household problem. Let Y_{is} be the present value of lifetime earnings of household i measured at age s, so that for any s, $Y_{is} = e^{r\cdot s}\cdot Y_{i0}$. Then*

$$c_{is} = \big[\frac{Y_{i0}}{\int_{S_i}^{T_i} e^{\sigma\cdot t}\cdot N_{it}\cdot[\lambda_{it}]^{\frac{1-\gamma}{\gamma}}\,dt}\big]\cdot[\lambda_{is}]^{\frac{1-\gamma}{\gamma}}\cdot e^{(\sigma+r)\cdot s}\,, \tag{3}$$

where $\sigma \equiv -r + \frac{r-\rho}{\gamma}$.

Appendix A supplies a proof.

Proposition 1 shows that optimal expenditure on consumption per equivalent adult is piecewise exponential, with discontinuities at times when λ_{is} and/or N_{is} change due to participation or retirement status changes or to changes in household composition.

Note that consumption c_{is} will change discontinuously when participation or family composition does. Let Δ_{is} denote the size of the consumption jump at age s. From (3),

$$\Delta_{is} = \frac{c_{i\,,\,s+0}/N_{i\,,\,s+0}}{c_{i\,,\,s-0}/N_{i\,,\,s-0}} = \big(\frac{\lambda_{i\,,\,s+0}}{\lambda_{i\,,\,s-0}}\big)^{\frac{1}{\gamma}-1}\,, \tag{4}$$

where the notation $s+0$ or $s-0$ denotes points in time arbitrarily after or before s. When λ_{is} rises from retirement or cessation of female labor–force participation, expenditure per adult equivalent decreases (increases) provided $\gamma > 1$ ($\gamma < 1$). The intuition is as follows. A household desires to smooth its flow of final–good aggregate,

$$\frac{\lambda_{it}\cdot c_{is}}{N_{is}}\,,$$

over its lifetime, yet consumption yields the highest payoff when λ_{is} is high. If $\gamma < 1$, the desire to smooth is relatively weak (indeed, it disappears entirely when $\gamma = 0$). Then complementarity with high λ_{is} predominates. In this case, a household chooses a high c_{is} for ages with high λ_{is}. On the other hand, if $\gamma > 1$, the desire to smooth is sufficiently powerful to predominate. In this case, a household intertemporally equalizes its marginal utility by choosing a low c_{is} at ages when λ_{is} is high.[1]

Most empirical studies of consumption behavior favor an intertemporal elasticity of substitution less than 1 — in other words, a value $\gamma > 1$. As we have just seen, that range implies that households will tend to reduce their consumption expenditure upon retirement — which accords well with a number of recent studies.[2] The same logic implies that when a woman joins the labor force, causing her household's λ_{is} to fall from λ to 1, her household's optimal market expenditure c_{is} should rise in compensation.

Labor–force participation increases a household's financial resources. That, in turn, reduces the household's overall marginal utility of spending on market consumption goods. And, during periods of labor–force participation, a household sacrifices efficiency in producing final–good aggregate. The following proposition derives the marginal condition for optimal participation that captures the tradeoff.

Proposition 2. *Let* $(c_{is}\,,\, p_{is})$ *be the optimal solution to the household problem given* R_i. *Then*

$$p_{is} = 1 \quad \textit{iff} \quad y^F_{is} - N_{is} \cdot \Big[\frac{Y_{i0}}{\int_{S_i}^{T} e^{\sigma \cdot t} \cdot N_{it} \cdot [\lambda_{it}]^{\frac{1-\gamma}{\gamma}}\, dt}\Big] \cdot e^{(\sigma+r)\cdot s} \cdot \frac{\gamma}{1-\gamma} \cdot [\lambda^{\frac{1-\gamma}{\gamma}} - 1] \geq 0\,, \quad (5)$$

where $\sigma \equiv -r + \frac{r-\rho}{\gamma}$.

Appendix A presents a proof.

Condition (5) shows that participation decisions depend upon a woman's potential market earnings, y^F_{is}; her household's lifetime resources, Y_{i0}; household size, N_{is}; and household age, s. Ceteris paribus, labor–force participation is more attractive for a woman with high potential market earnings, y^F_{is}, and less attractive for a woman whose household's lifetime earnings, Y_{i0}, are high. Female participation is less attractive at ages when a household is large relative to its average size. In other words, if a household has many children at home, the wife has strong incentives to remain at home as well. (If household size is unchanged over the life cycle, on the other hand, then N_{is} drops out of (5).)

Finally, suppose that the market interest rate exceeds the subjective discount rate (i.e., suppose that $r > \rho$). In this case, (5) shows that a woman's participation in the labor market is less likely for older households — provided earnings growth levels off beyond some age. Consumption growth reduces the marginal utility of more dollars to spend; however, gains from more time at home are proportionate, in our framework, to resources regardless

[1] See the discussion in Laitner and Silverman [2005, 2007].

[2] See, for instance, Hurd and Rohwedder [2003], Laitner and Silverman [2005, 2007], and Aguiar and Hurst [2005].

of the latter's level — and they are, in that sense, immune to diminishing returns. An analogous argument in Section 6 provides a story for household retirement.

3. Estimation Strategy

The model presented and analyzed in Section 2 can predict the labor–force participation of a woman in a given household at all ages. To generate such predictions from the model, we require parameter values for λ and γ, as well as household–level data on the present value of lifetime earnings, Y_{i0}; the time path of female earnings, y_{is}^F; and household size, N_{is}. We seek to obtain parameter values with the following estimation strategy. For a prospective vector of parameters $\vec{\theta}$, which is common to all households, we use the model and household–level data to predict the labor–force participation profile for each household. Then we compare the prediction with observed female behavior in the data. Finally, we adjust the prospective parameter vector $\vec{\theta}$ to minimize the difference between predicted and observed participation profiles.

Predicted Participation Profiles. Let $\mathcal{D}_i$ denote the data on household i. The data consists of two numbers, S_i and R_i, which specify a household's starting and retirement ages, and three functions of the household's age s, namely, N_{is}, y_{is}^M, and y_{is}^F, which specify, respectively, the household's size profile, male earnings, and female earnings. We calibrate values for ρ, r, and male and female lifespans, which are common to all households, and we treat $\vec{\theta}$ as the vector of parameters to be estimated. Given $\mathcal{D}_i$ and $\vec{\theta}$, the optimal labor–force participation condition (see (5)) generates a predicted participation profile $\widehat{p}_s(\mathcal{D}_i\,,\,\vec{\theta})$ for each household:

$$\widehat{p}_s(\mathcal{D}_i\,,\,\vec{\theta}) \equiv \chi\big(y_{is}^F - N_{is}\cdot\big[\frac{Y(p_{is}\,|\,\mathcal{D}_i\,,\,\vec{\theta})}{J(p_{is}\,|\,\mathcal{D}_i\,,\,\vec{\theta})}\big]\cdot e^{(\sigma+r)\cdot s}\cdot\frac{[\lambda^{\frac{1-\gamma}{\gamma}}-1]}{\frac{1-\gamma}{\gamma}}\geq 0\big)\,, \tag{6}$$

where

$$Y(p_{is}\,|\,\mathcal{D}_i\,,\,\vec{\theta}) \equiv \int_{S_i}^{R_i} e^{-r\cdot s}\cdot(y_{is}^M + p_{is}\cdot y_{is}^F)\,ds\,, \tag{7}$$

$$J(p_{is}\,|\,\mathcal{D}_i\,,\,\vec{\theta}) \equiv \int_{S_i}^{T_i} e^{\sigma\cdot s}\cdot N_{is}\cdot[\lambda(s\,,\,p_{is})]^{\frac{1-\gamma}{\gamma}}\,ds\,, \tag{8}$$

and $\chi(.)$ is an indicator function.

Calibrating Λ. We calibrate Λ as follows. Let ω_i be the sample weight for household i in our HRS data. Let

$$\Omega \equiv \sum_i \omega_i\,.$$

Then if $\overline{p_R}$ is the sample–average fraction of wives who work until their husband retires, we have

$$\overline{p_R} = \frac{1}{\Omega} \cdot \sum_i \omega_i \cdot p_{iR_i} \, .$$

In our model, if a husband and wife retire together, their household's consumption should jump (downward provided $\gamma > 1$) by

$$[\Lambda]^{\frac{1-\gamma}{\gamma}} \, .$$

If the wife retires earlier, then upon the husband's retirement the household's consumption jump is

$$[\frac{\Lambda}{\lambda}]^{\frac{1-\gamma}{\gamma}} \, .$$

Thus, the model predicts a sample–average change in household consumption at male retirement of

$$\Delta_R = [\Lambda]^{\frac{1-\gamma}{\gamma}} \cdot \overline{p_R} + [\frac{\Lambda}{\lambda}]^{\frac{1-\gamma}{\gamma}} \cdot (1 - \overline{p_R}) \, . \tag{9}$$

We set the left–hand side of (9) from an external source: Laitner and Silverman [2007] use Consumer Expenditure Survey data to measure the average drop in household consumption at retirement and find a fraction of about 0.20. Using, say, $\Delta_R = 0.20$, we can solve (9) for Λ:

$$\Lambda(\mathcal{D}_i \, , \vec{\theta} \, , \, \Delta_R) = \big(\frac{\Delta_R}{(1 - \overline{p_R}) \cdot \lambda^{\frac{\gamma-1}{\gamma}} + \overline{p_R}}\big)^{\frac{\gamma}{1-\gamma}} \, . \tag{10}$$

This provides our calibration of Λ.

Regression. Rather than searching for parameters to match each household's exact path of female labor–market participation, this paper focuses on the total length of each woman's market career. In other words, we choose parameters to match average female labor–force participation. Let

$$\widehat{\tau}(\mathcal{D}_i \, , \vec{\theta}) \equiv \int_{S_i}^{T} \widehat{p}_s(\mathcal{D}_i \, , \vec{\theta}) \, ds \, . \tag{11}$$

Then $\widehat{\tau}(\mathcal{D}_i \, , \vec{\theta})$ is the total predicted length of all labor–force participation spells for household i. Note that since our model is set in continuous time, $\widehat{\tau}(.)$ is a continuous variable that can take non–integer values. Let the actual total participation span for the wife of household i be

$$\tau_i \equiv \int_{S_i}^{T} p_{is} \, ds \, .$$

Then our regression model is

$$\tau_i = \widehat{\tau}(\mathcal{D}_i \, , \vec{\theta}) + \epsilon_i \, , \tag{12}$$

where ϵ_i is an *iid* random error. We estimate $\vec{\theta}$ from (12) using nonlinear least squares (NLLS).

We think of ϵ_i as arising in (12) from measurement error. Our participation data has observations at annual frequency, yet true participation spells may start or end in the middle of a year. We devise an imputation procedure for partial years of work below. The ϵ_i in (12) reflects measurement errors that remain.

One could think of identification for λ as coming from cross–sectional variation in female labor–force participation for households of different sizes, different total lifetime earnings, and different female earnings. One could think of identification for γ as coming from each household's changing participation with age.

This paper calculates $\widehat{\tau}(\mathcal{D}_i\,,\,\vec{\theta})$ for the right–hand side of (12) as follows. Suppose, as above, that ϵ_i reflects measurement errors in p_{is}, $s \in [S_i\,,\,R_i]$, stemming from limitations of our data set. Our observation of lifetime earnings, Y_{i0}, does not have this error — annual observations on male or female earnings automatically reflect actual months (and hours) of employment. If one is willing to overlook possible errors in $J(.)$ from using the data on p_{is}, one can compute $\widehat{\tau}_i$ given $\mathcal{D}_i$, Y_{i0}, and any prospective $\vec{\theta}$. Maintaining $\mathcal{D}_i$ and Y_{i0}, one can then adjust $\vec{\theta}$ to minimize the sum of squared errors in the regression. That is the technique that this paper uses. It leads to the results in Table 1 of Section 5.

Future drafts will recognize more fully the fact that λ_{is} depends on p_{is}. For given $(\mathcal{D}_i\,,\,Y_{i0}\,,\,\vec{\theta})$ and a prospective J_i (as in (8)), we can use (5) to generate p_{is} all s and, hence, a new J_i, say, J_i^*. This defines a mapping $\psi(.)$:

$$J_i^* = \psi(J_i\,,\,\mathcal{D}_i\,,\,Y_{i0}\,,\,\vec{\theta})\,.$$

We need to find a fixed point, say, $J_i^{FP} = J^{FP}(\mathcal{D}_i\,,\,Y_{i0}\,,\,\vec{\theta})$, of the new mapping:

$$J_i^{FP} = \psi(J_i^{FP}\,,\,\mathcal{D}_i\,,\,\vec{\theta})\,.$$

Then we should use this fixed point in (6) and (11) to generate

$$\widehat{\tau}\big(J^{FP}(\mathcal{D}_i\,,\,Y_{i0}\,,\,\vec{\theta})\,,\,\mathcal{D}_i\,,\,\vec{\theta}\big)\,.$$

This will be our next step.

4. Data

We use data from the Health and Retirement Study (HRS) to construct a comprehensive measure of household lifetime earnings and the earning profiles of women. To implement our empirical strategy, we also require data on age of male retirement and number of equivalent adults per household at each age. This section discusses our data set in detail, and Figures 1-2 present quantitative summary information.

We use the original survey cohort from the HRS, consisting of households in which the respondent is 51-61 in 1992 (Juster and Suzman [1995]). We restrict our attention to married couples that meet the following conditions: (1) each spouse has only been married once, (2) the couple has a complete earnings record (see below), (3) the birth dates for all

children in the household are available in the data set, and (4) both spouses have retired by 2002. Appendix B discusses other sample restrictions that we employ.

As a matter of definition, we associate the age of a household with that of the adult male. Set $S_i^M = \max\{18\,,\text{ years education} + 6\}$ for the male, and similarly define S_i^F for the female. Converting both to the male's age, the household's starting age is the minimum of the two. An individual is retired if he/she reports that status — though we assume that both spouses are retired when the male is. Our model abstracts from uncertainty about lifespan. For every household, we assume that the man lives to age 76 and the woman lives to age 80. These calibrations together with the spousal age difference imply a household–specific ending age T_i.

We aggregate family members into a single index of "equivalent adults." Before the household is married, the man and the woman live apart and, after age S_i^M the male contributes 1 to the household's equivalent adult total, and after age S_i^F the female contributes 1 as well. Once the couple marries, prior to its first child the number of household equivalent adults is 1.5. Each child, to a maximum of 2 in any given year, counts 0.5 in the household equivalent adult total. Upon the death of one spouse, the household's number of equivalent adults declines from 1.5 to 1.0 (assuming no remaining children at home).

Total lifetime resources, Y_{i0}, for household i is the present value at age 0 of male and female lifetime earnings. Future work will compute net–of–tax earnings (including federal income taxes and payroll taxes) and add Social Security and Medicare benefits.

We have annual Social Security earnings histories for both men and women in our sample. The HRS supplies survey data (and work hours) for 1992, 1994, 1996, 1998, 2000, and 2002. Prior to 1981, the Social Security histories are right censored at the OASDI maximum taxable amount; after 1981, earnings are censored at \$100,000, 250,000, and 500,000 (to protect respondent confidentiality). We replace right–censored amounts with imputations from an earnings dynamic equation (see House, Laitner, Stolyarov [2007]). From the same equation, we impute all missing male earnings figures. These steps implicitly assume that males participate full time in the labor market for all ages S_i^M to R_i.

We use a separate earnings dynamics equation to impute right–censored female earnings. We make an effort to impute non–FICA earnings from the equation for spans in which the woman reports uncovered jobs.[3] For all other years, we use reported female earnings, which can be 0.

Because many women have spells of non–participation in the labor force, we worry about years that bracket withdrawal from the market. In other words, we worry about setting $p_{is} = 1$ for all years with any female earnings. We make the following correction. If a bracketing age, say, s, is preceded or followed by two or more non–bracketed ages, we derive an extrapolated earnings figure, say, $\widehat{y}_{is}^F$, from the two nearest non–bracketed ages. If the extrapolation is larger than observed earnings, y_{is}^F, we set $p_{is} = y_{is}^F/\widehat{y}_{is}^F$.

Another complication is that our model requires estimates of y_{is}^F for ages s in which a woman does not participate in the labor market. We obtain estimates for such ages using

[3] Specifically, we use survey information on number of years with non–FICA earnings and the time interval for the latter. We assume the non–FICA earnings occur in an unbroken sequence. In our estimation below, we position the sequence to minimize the sum of squared errors in the regression. See Appendix B.

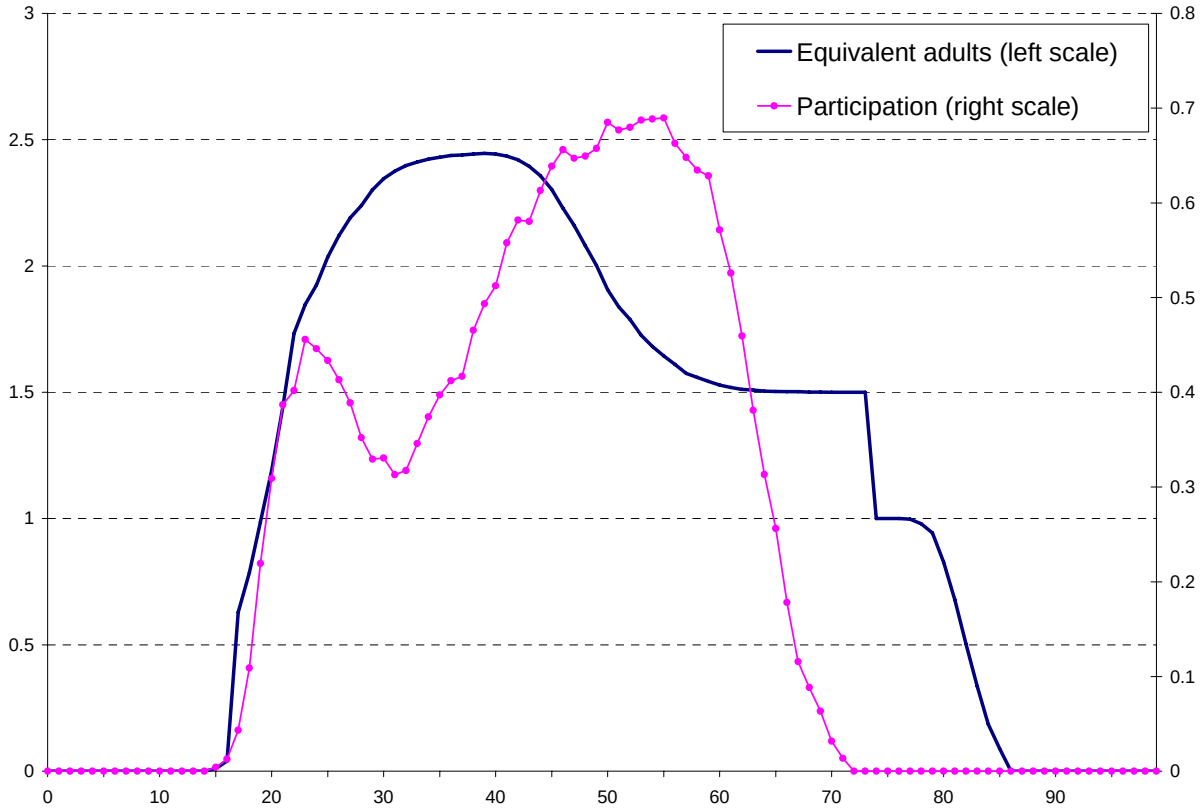

Figure 1. Household size and participation profile.

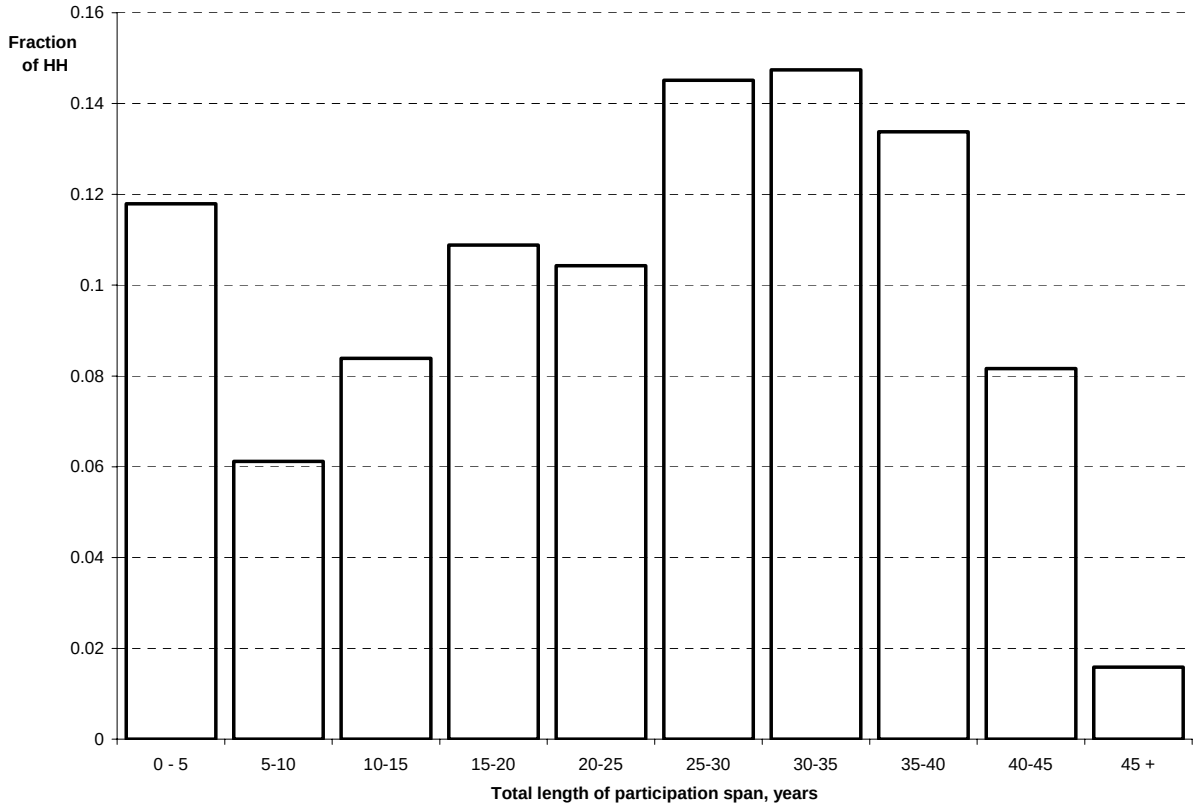

Figure 2. Distribution of participation spans

cubic interpolations (or extrapolations) from the nearest ages with participation.

Finally, we set the time preference parameter $\rho = 0.02$, and we assume a constant net–of–tax interest rate of $r = 0.0442$ per year (e.g., Laitner and Silverman [2007]).

Summary Statistics. We present summary information on our sample in figures that follow.

Figure 1 shows the sample average number of equivalent adults, N_{is}, and the sample average female labor–force participation, p_{is}. Almost all households in the sample have at least 2 children living with them at some point. The smooth hump–shaped profile of the average number of adult equivalents reflects heterogeneity in the timing of births and the timing of children's departure from the household. The sudden drop in household size at 76 reflects male mortality. There is no equivalently sharp drop at age 80 when women die because the graph is plotted against household age. Household age, by definition, equals the age of the adult male; thus, household age for female mortality varies with the age difference of spouses.

The life–cycle female labor–force participation pattern exhibits two peaks. The early peak (at male age 23) corresponds to women finishing their education and working prior to bearing children. The second, larger peak corresponds to mothers of grown children returning to work. Female labor–force participation drops steeply at retirement age.

Figure 2 depicts the distribution of households by the length of female lifetime labor–force participation. Almost 7 percent of women in the sample never worked in the market at all — hence the peak at participation spans of less than 5 years. Another slightly larger peak is at 30-35 total years in the labor force, which corresponds to participating almost without interruption.

5. Results

We estimate regression equation (12) with non–linear least squares for three different values for the drop in household consumption at retirement: a 10 percent drop, a 15 percent drop, and a 20 percent drop. In each case, we calculate a corresponding value of Λ from (10), as explained in Section 3. (The drop in consumption at retirement is Δ_R in the notation of Section 3.) Table 1 presents results.

In Table 1, the point estimates of λ center around 1.115 and are not sensitive to the chosen value of Δ_R. The implied value of Λ, on the other hand, does (not surprisingly) co–vary strongly with the size of the drop in consumption at retirement. The point estimate of γ is roughly 4.0, which is similar to a number of other estimates in the literature (see, for example, the discussion in Auerbach and Kotlikoff [1987]), but which is higher than the value estimated in Laitner and Silverman [2007]. The standard errors are somewhat sensitive to the method that we use to construct potential female earnings, although the point estimates are not.

Section 3 describes the exact computations leading to the results in Table 1. It also indicates more sophisticated steps that we intend to use for additional estimates in the future.

Table 1. Parameter Estimates
Health and Retirement Study Data 1992-2002[a]

Parameter	Assumed Consumption Drop at Retirement:[b] $\Delta_R =$		
	10 percent	15 percent	20 percent
γ	4.4511	4.0552	3.8487
(S.E.)	(0.0522)	(0.0527)	(0.0728)
λ	1.1156	1.1140	1.1155
(S.E.)	(0.0041)	(0.0040)	(0.0039)
Summary Statistics			
R^2	0.6682	0.6688	0.6685
Observations	441	441	441
Addendum: Derived Parameter Values			
Λ	1.2148	1.3216	1.4411
$\lambda^{\frac{\gamma-1}{\gamma}}$	1.0885	1.0847	1.0842
Λ/λ	1.0889	1.1864	1.2912

a. See text.
b. See text.

Interpretation. The point estimates and equation (3) imply that household expenditure per adult equivalent should rise by a factor $\lambda^{\frac{\gamma-1}{\gamma}}$ when a woman joins the labor force. Her household compensates for its loss of services (from her home production and leisure) with higher expenditure. Table 1 implies the size of this compensation is roughly 8 percent of household spending on consumption. The ratio Λ/λ in Table 1 measures the additional value of consumption services that the household enjoys from leisure in retirement.

Figure 3 plots the sample average of household earnings and the sample average consumption–expenditure profile implied by our parameter estimates when the drop of household consumption at retirement is 20 percent. Again, the smoothness of both series is a consequence of aggregation over the heterogeneous households in our sample. The hump in the household consumption profile corresponds to the presence of minor children in the household. The average consumption profile is not very sensitive to λ — perhaps because most changes in female labor–force participation occur after children leave home so that the impact of a jump in λ on total household consumption is often mitigated by small household size at the time of the change.

As in Figure 1, Figure 3 reveals a sharp drop in consumption expenditure at age 76. This follows from our assumption that men die at this age. Although adding uncertainty about lifespan would complicate the picture, mortality would still cause steep declines at the right–hand side of Figure 3.

6. Extension: Optimal Retirement

Although analysis so far treats retirement age as exogenous to families, our model also has implications for optimal retirement — recall Section 2. There are at least two reasons to consider optimal retirement decisions. First, one would expect households simultaneously to choose female participation, life–cycle consumption trajectories, and retirement age. Thus, modeling optimal retirement age should yield additional insights on other life–cycle choices. Second, if our analysis yields more first–order conditions for optimal behavior, we should have more conditions for estimating parameters. This section briefly considers how one could incorporate optimal retirement age into our framework and expand the scope of our regression analysis to include estimation of Λ.

Returning to optimization problem (1) and analyzing the first–order condition for optimal R, we have

Proposition 3. *If household i chooses its optimal retirement age $R = R_i$, then R satisfies*

$$y_{iR} = N_{iR} \cdot \Big[\frac{Y_{i0}}{\int_{S_i}^{T_i} e^{\sigma \cdot s} \cdot N_{is} \cdot [\lambda_{is}]^{\frac{1-\gamma}{\gamma}}\, ds}\Big] \cdot e^{(\sigma+r)\cdot R} \cdot \frac{\gamma}{1-\gamma} \cdot \Big([\Lambda]^{\frac{1-\gamma}{\gamma}} - [\lambda_{i\,,\,R-0}]^{\frac{1-\gamma}{\gamma}}\Big)\,, \quad (13)$$

where $\sigma \equiv -r + \frac{r-\rho}{\gamma}$ and $y_{is} \equiv y_{is}^M + y_{is}^F$.

Appendix A presents a proof.

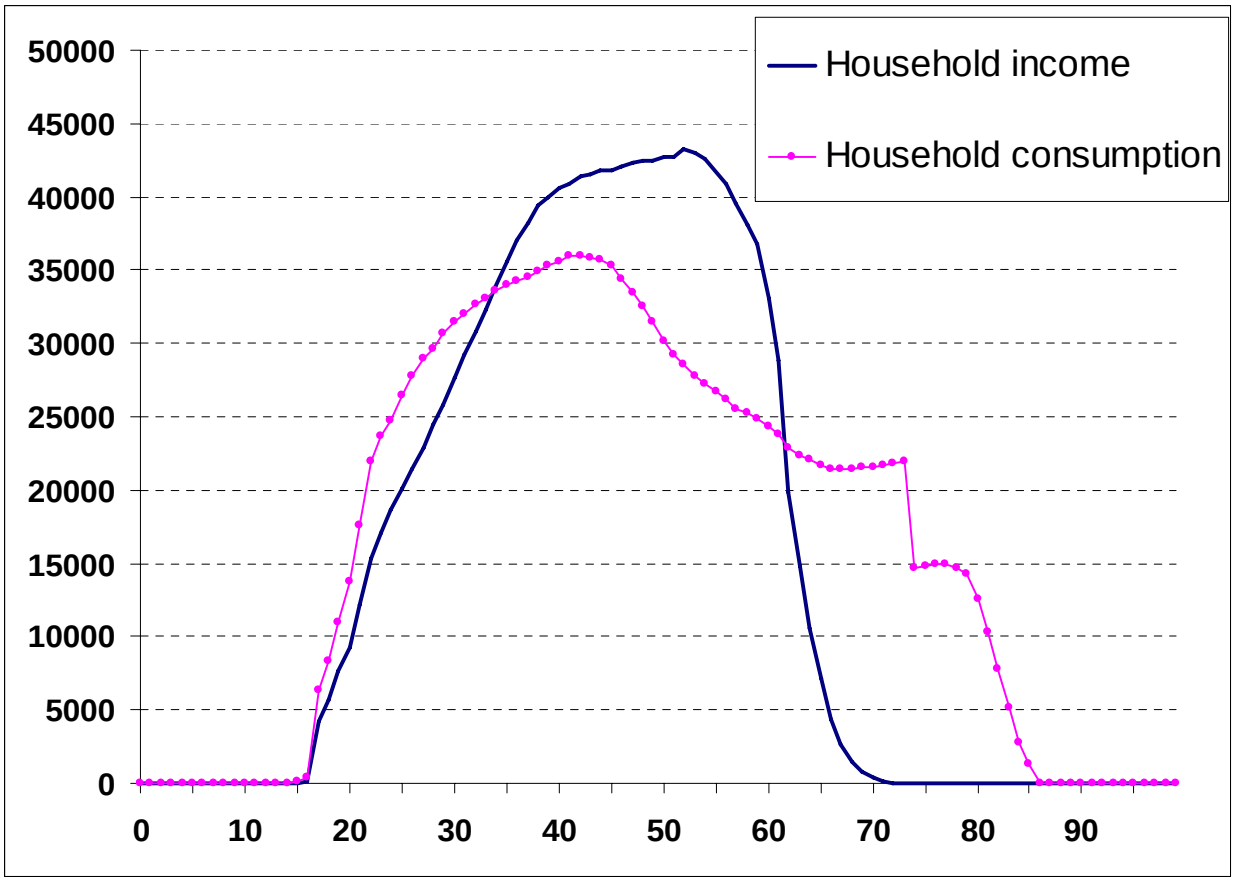

Figure 3. Sample average household income and consumption profiles, 1984 dollars.

Equation (13) balances forgone earnings against utility gains from retirement. An intuitive argument is as follows. Suppose that both spouses work prior to R and then retire together. Drop the subscript i for expositional simplicity. Let

$$c \equiv \frac{Y_0}{\int_S^T e^{\sigma \cdot s} \cdot N_s \cdot [\lambda_s]^{\frac{1-\gamma}{\gamma}} \, ds} \, .$$

So, c is consumption just prior to retirement — in other words, $c_{R-0} = c$. Retirement at age R entails loss of earnings flow y_R. Marginal utility per equivalent adult is

$$[N_R]^{\gamma} \cdot [c_{R-0}]^{-\gamma} \, .$$

The earnings flow loss per equivalent adult is y_R/N_R. Thus, the sacrifice at retirement of utility flow is

$$\frac{y_R}{N_R} \cdot [N_R]^{\gamma} \cdot c^{-\gamma} \, . \tag{14}$$

The gain in utility flow from retirement has two components. Consumption just after retirement is

$$c \cdot \Lambda^{\frac{1-\gamma}{\gamma}} \cdot \Lambda \, .$$

Hence, the household gain in utility flow after retirement is

$$\frac{[N_R]^{\gamma}}{1-\gamma} \cdot \left[[c \cdot \Lambda^{\frac{1-\gamma}{\gamma}} \cdot \Lambda]^{1-\gamma} - c^{1-\gamma} \right] . \tag{15}$$

The second component of gain emerges because, as we have noted above, consumption after retirement will be smaller (greater) than c if $\gamma > 1$ ($\gamma < 1$). The dollar value of the flow saving for the household is

$$N_R \cdot c \cdot (1 - \Lambda^{\frac{1-\gamma}{\gamma}}) \, .$$

The dollar flow per equivalent adult is the same divided by N_R. In utility terms, the second component then is

$$\frac{N_R \cdot c}{N_R} \cdot (1 - \Lambda^{\frac{\gamma-1}{\gamma}}) \cdot [N_R]^{\gamma} \cdot c^{-\gamma} \, . \tag{16}$$

At the optimal retirement age R, (14) should equal (15) plus (16). Noting that in this case $\lambda_{iR} = 1$, such equality yields (13).

The argument is almost the same if the household's wife has stopped working prior to R. Then $\lambda_{iR} = \lambda$. Let c be as above, and let

$$C \equiv c \cdot \lambda^{\frac{1-\gamma}{\gamma}} \, .$$

The analogue of (14) is

$$\frac{y_R}{N_R} \cdot [N_R]^{\gamma} \cdot [C \cdot \lambda]^{-\gamma} \cdot \lambda = \frac{y_R}{N_R} \cdot [N_R]^{\gamma} \cdot \lambda^{1-\gamma} \cdot C^{-\gamma} = \frac{y_R}{N_R} \cdot [N_R]^{\gamma} \cdot c^{-\gamma} \,. \tag{17}$$

The analogue of (15) is

$$\frac{[N_R]^{\gamma}}{1-\gamma} \cdot \left[[c \cdot \Lambda^{\frac{1-\gamma}{\gamma}} \cdot \Lambda]^{1-\gamma} - [c \cdot \lambda^{\frac{1-\gamma}{\gamma}} \cdot \lambda]^{1-\gamma} \right] . \tag{18}$$

And, the replacement for (16) is (using the marginal utility formula from (17))

$$\frac{N_R \cdot c}{N_R} \cdot (\lambda^{\frac{1-\gamma}{\gamma}} - \Lambda^{\frac{1-\gamma}{\gamma}}) \cdot [N_R]^{\gamma} \cdot c^{-\gamma} \,. \tag{19}$$

Equation (13) emerges if we equate (17) with the sum of (18) and (19).

Future work will solve (13) as we solved (11) in Section 3, yielding

$$\widehat{R}(\mathcal{D}_i \,,\, \vec{\theta})$$

with

$$\vec{\theta} \equiv (\lambda \,,\, \gamma \,,\, \Lambda) \,.$$

We will then estimate (12) and

$$R_i = \widehat{R}(\mathcal{D}_i \,,\, \vec{\theta}) + \eta_i \tag{20}$$

as a system.

7. Conclusion and Future Work

This paper's model depicts the labor–force participation trajectory for married women over their life cycle. The model assumes that a household's productivity in utilizing consumption drops when the husband and/or wife have fewer hours per week at home. Earnings compensate for lost time at home. If her market compensation is too low, a woman will choose not to work. The model predicts that labor–force participation will be greatest for women with high wages and/or relatively few children. When a household retires, the household's productivity in using consumption rises still further.

We use the model, together with micro–level data available from the HRS on household earnings and demographic characteristics, to estimate key parameters of the model. Our estimates predict that a household's consumption should rise roughly 8-10 percent to compensate for lost home production when women enter the labor market. This is 40-50 percent of the decline other studies predict when a household retires completely. In a sense, therefore, we find that married women's labor–force participation disrupts a household's efficiency at home production about 40-50 percent as much as having a retired household return to work would.

We have outlined how we expect to extend the analysis to consider women's labor–force participation and household retirement jointly. Another extension that we expect to

pursue will make a woman's productivity at home dependent on her number of children and their ages. We expect the model then to complement existing work on the value of women's home production (see, for example, House et al. [2007], Benhabib et al. [1991], and Rupert et al. [2000]) and the cost of raising children.

References

[1] Aguiar, Mark; and Erik Hurst "Consumption vs. Expenditure," *Journal of Political Economy*, October 2005, 113(5), pp. 919-948.

[2] Auerbach A. and Kotlikoff L., *Dynamic Fiscal Policy*, (Cambridge: Cambridge University Press, 1987).

[3] Benhabib, Jess, Rogerson, Richard, and Wright, Randall, "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy* 99(6) (December 1991), 1166–1187.

[4] French, Eric, "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behavior," *Review of Economic Studies* 72(2) (April 2005), 395-427.

[5] Guner, Nezih, and Knowles, John, "Marital Instability and the Distribution of Wealth," Working Paper, University of Pennsylvania, July 2004.

[6] House, Christopher, Laitner, John P., Stolyarov, Dmitriy, "Valuing Lost home Production for Dual Earner Couples" forthcoming, *International Economic Review*, 2007.

[7] Hurd, Michael; Rohwedder, Susann "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement" NBER working paper 9586 (2003).

[8] Juster, F.T., and Suzman R., "The Health and Retirement Study: An Overview," *Journal of Human Resources*. Supplement JHR30-S, S7-S56, 1995.

[9] Laitner, John P., and Silverman, Daniel, "Estimating Life–Cycle Parameters from Consumption Behavior at Retirement," NBER working paper 11163 (2005).

[10] Laitner, John P., and Silverman, Daniel, "Consumption, Retirement and Social Security: Evaluating the Efficiency of Reform that Encourages Longer Careers," University of Michigan working paper, Dec 2006.

[11] Rupert, Peter, Rogerson, Richard, Wright, Randall, "Homework in labor economics: Household production and intertemporal substitution." *Journal of Monetary Economics* 46(3) (December 2000), 557-579.

[12] Social Security Administration, *Annual Statistical Supplement, 2001, to the Social Security Bulletin*. (Washington: U.S. Government Printing Office, 2001).

[13] Tobin, James, "Life Cycle Saving and Balanced Growth", in W. Fellner (ed.), *Ten economic Studies in the Tradition of Irving Fisher*. (New York: Wiley, 1967).

Appendix A: Proofs of the Propositions

This appendix contains the proofs of Propositions 1 and 2.

Proposition 1 *Let (c_{is}, p_{is}, R_i) be the optimal solution to the household problem. Let Y_{is} be the present value of lifetime earnings of household i measured at age s, so that for any s, $Y_{is} = e^{rs}Y_{i0}$. Then*

$$c_{is} = \left[\frac{Y_{i0}}{\int_{S_i}^{T} e^{\sigma s} N_{is} \lambda_{is}^{\frac{1-\gamma}{\gamma}} ds}\right] N_{is} \cdot \lambda_{is}^{\frac{1-\gamma}{\gamma}} \cdot e^{\frac{r-\rho}{\gamma}s}, \tag{10}$$

where $\sigma = -r + \frac{r-\rho}{\gamma}$.

Proof: Let μ_t be the co-state variable on the budget constraint. Fix an arbitrary participation profile p_{is}. This pins down λ_{is} from (2). The first-order condition on c_{is} reads

$$e^{-\rho t} N_{is} \frac{\lambda_{is}}{N_{is}} u'\left(\frac{\lambda_{is} \cdot c_{is}}{N_{is}}\right) = \mu_{is} = \mu_{i0} e^{-rs}, \text{ all } s \leq T \Longleftrightarrow$$

Let Y_{is} be the present value of lifetime earnings of household i measured at age s, so that for any s, $Y_{is} = e^{rs}Y_{i0}$. Since $a_{S_i} = a_T = 0$,

$$\left[\frac{1}{\mu_{i0}}\right]^{\frac{1}{\gamma}} \cdot \int_{S_i}^{T} e^{\sigma s} N_{is} \lambda_{is}^{\frac{1-\gamma}{\gamma}} ds = Y_{i0} \tag{11}$$

Then

$$c_{is} = \left[\frac{Y_{i0}}{\int_{S_i}^{T} e^{\sigma s} N_{is} \lambda_{is}^{\frac{1-\gamma}{\gamma}} ds}\right] N_{is} \cdot \lambda_{is}^{\frac{1-\gamma}{\gamma}} \cdot e^{\frac{r-\rho}{\gamma}s} = c_0\left(S_i, R_i, N_{is}, p_{is}, Y_{i0}; \vec{\theta}\right) \cdot N_{is} \cdot \lambda_{is}^{\frac{1-\gamma}{\gamma}} \cdot e^{\frac{r-\rho}{\gamma}s}$$

$$\mu_{is} = \left[\frac{\int_{S_i}^{T} e^{\sigma s} N_{is} \lambda_{is}^{\frac{1-\gamma}{\gamma}} ds}{Y_{i0}}\right]^{\gamma} e^{-rs} = \mu_0\left(S_i, R_i, N_{is}, p_{is}, Y_{i0}; \vec{\theta}\right) \cdot e^{-rs}$$

■

Proposition 2. *Let (c_{is}, p_{is}) be the optimal solution to the household problem given R_i. Then*

$$p_{is} = 1 \text{ iff } y_{is}^F - N_{is}\left[\frac{Y_{i0}}{\int_{S_i}^{R_i} e^{\sigma s} N_{is} \lambda_{is}^{\frac{1-\gamma}{\gamma}}}\right] e^{(\sigma+r)s} \frac{\left(\lambda^{\frac{1-\gamma}{\gamma}} - 1\right)}{\frac{1-\gamma}{\gamma}} \geq 0, \tag{12}$$

where $\sigma = -r + \frac{r-\rho}{\gamma}$.

Proof: Suppose the optimal value function $V(t,a)$ is known and $t < R$. Represent the household problem as a discrete choice problem.

$$V(t,a) = \max\left\{V_1(t,a), V_0(t,a)\right\}$$

where

$$\begin{aligned} V_1(t,a) &= \max_{c_t}\left(e^{-\rho t} N_t u\left(\frac{c_t}{N_t}\right) dt + V(t+dt, a+da_1)\right) \\ V_0(t,a) &= \max_{c_t}\left(e^{-\rho t} N_t u\left(\frac{\lambda c_t}{N_t}\right) dt + V(t+dt, a+da_0)\right) \end{aligned}$$

$$da_1 = \left(ra + y_t^M + y_t^F - c_t\right) dt,$$
$$da_0 = \left(ra + y_t^M + y_t^F - c_t\right) dt$$

are the value functions and the laws of motion for assets that correspond to $p_t = 1$ and $p_t = 0$ respectively. To derive the condition for participation, compare $V_1(t,a)$ and $V_0(t,a)$:

$$V_1(t,a) = V(t,a) + \frac{\partial V}{\partial t}(t,a)\,dt + \frac{\partial V}{\partial a}(t,a)\left(ra + y_t^M\right)dt +$$
$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right)dt + \frac{\partial V}{\partial a}(t,a)\cdot y_t^F dt, \tag{13}$$

$$V_0(t,a) = V(t,a) + \frac{\partial V}{\partial t}(t,a)\,dt + \frac{\partial V}{\partial a}(t,a)\left(ra + y_t^M\right)dt +$$
$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\lambda c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right)dt, \tag{14}$$

Solve the maximization problems in (13) and (14) to express c_t through $\frac{\partial V}{\partial a}$:

$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right) = \frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\left[\frac{\partial V}{\partial a}\right]^{1-\frac{1}{\gamma}}$$

$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\lambda c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right) = \frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\cdot\lambda^{\frac{1-\gamma}{\gamma}}\left[\frac{\partial V}{\partial a}\right]^{1-\frac{1}{\gamma}}$$

$$V_1(t,a) - V_0(t,a) = \left(\frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\left[\frac{\partial V}{\partial a}\right]^{1-\frac{1}{\gamma}}\left(1-\lambda^{\frac{1-\gamma}{\gamma}}\right) + \frac{\partial V}{\partial a}y_t^F\right)dt$$

Therefore,

$$V_1(t,a) - V_0(t,a) \geq 0$$

is equivalent to

$$\frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\left(1-\lambda^{\frac{1-\gamma}{\gamma}}\right) + \left[\frac{\partial V}{\partial a}\right]^{\frac{1}{\gamma}}y_t^F \geq 0 \Longleftrightarrow$$
$$\left[\frac{\partial V}{\partial a}\right]^{\frac{1}{\gamma}}y_t^F \geq \frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\left(\lambda^{\frac{1-\gamma}{\gamma}} - 1\right) \tag{15}$$

Kamien and Scwartz (p. 137) show that

$$\frac{\partial V}{\partial a}(t,a) = \mu_t = \mu_0 e^{-rt}.$$

Then (15) simplifies to

$$y_t^F \geq \frac{\gamma}{1-\gamma}N_t\left[\frac{1}{\mu_0}\right]^{\frac{1}{\gamma}}e^{\frac{r-\rho}{\gamma}t}\left(\lambda^{\frac{1-\gamma}{\gamma}} - 1\right)$$

■

Proposition 3 *If the retirement age R is chosen optimally, then R satisfies*

$$y_{R-0} = \frac{\gamma}{1-\gamma}\left[\frac{Y_{i0}}{\int_{S_i}^{T}e^{\sigma s}N_{is}\lambda_{is}^{\frac{1-\gamma}{\gamma}}ds}\right]e^{\frac{r-\rho}{\gamma}R}\left(N_{R+0}\Lambda^{\frac{1-\gamma}{\gamma}} - N_{R-0}\lambda_{R-0}^{\frac{1-\gamma}{\gamma}}\right) \tag{16}$$

Proof Let

$$V(t,a) = \max\left\{V_W(t,a), V_R(t,a)\right\}$$

where

$$V_W(t,a) = \max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\lambda_t c_t}{N_t}\right)dt + V(t+dt, a+da_W)\right)$$
$$V_R(t,a) = \max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\Lambda c_t}{N_t}\right)dt + V(t+dt, a+da_R)\right)$$

$$da_W = (ra + y_t - c_t)\,dt,$$

$$da_R = (ra - c_t)\,dt$$

are the value functions and the laws of motion for assets that correspond to work and retirement, respectively. To derive the condition for optimal retirement, compare $V_W(t,a)$ and $V_R(t,a)$:

$$V_W(t,a) = V(t,a) + \frac{\partial V}{\partial t}(t,a)\,dt + \frac{\partial V}{\partial a}(t,a)\,radt +$$

$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\lambda_t c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right)dt + \frac{\partial V}{\partial a}(t,a)\cdot y_t dt, \tag{17}$$

$$V_0(t,a) = V(t,a) + \frac{\partial V}{\partial t}(t,a)\,dt + \frac{\partial V}{\partial a}(t,a)\,radt +$$

$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\Lambda c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right)dt, \tag{18}$$

Solve the maximization problems in (17) and (18) to express c_t through $\frac{\partial V}{\partial a}$:

$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right) = \lambda_t^{\frac{1-\gamma}{\gamma}}\frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\left[\frac{\partial V}{\partial a}\right]^{1-\frac{1}{\gamma}}$$

$$\max_{c_t}\left(e^{-\rho t}N_t u\left(\frac{\lambda c_t}{N_t}\right) - \frac{\partial V}{\partial a}(t,a)\cdot c_t\right) = \Lambda^{\frac{1-\gamma}{\gamma}}\frac{\gamma}{1-\gamma}N_t\cdot e^{-\frac{\rho t}{\gamma}}\left[\frac{\partial V}{\partial a}\right]^{1-\frac{1}{\gamma}}$$

If R is the optimal retirement age, then at the moment of retirement, the household must be indifferent between retiring and continuing to work

$$\begin{aligned} 0 &= V_W(R-0,a) - V_R(R+0,a) = \\ &= \left(\frac{\gamma}{1-\gamma}e^{-\frac{\rho R}{\gamma}}\left[\frac{\partial V}{\partial a}\right]^{1-\frac{1}{\gamma}}\left(N_{R-0}\lambda_{R-0}^{\frac{1-\gamma}{\gamma}} - N_{R+0}\Lambda^{\frac{1-\gamma}{\gamma}}\right) + \frac{\partial V}{\partial a}\cdot y_{R-0}\right)dt. \end{aligned}$$

Substituting μ_t for $\frac{\partial V}{\partial a}$ in the above expression and using (11) gives (8). ■

Appendix B: Data

This section presents a detailed description and construction details for our data inputs.

Sample Criteria. We use the original survey cohort from the HRS, consisting of households in which the respondent is age 51-61 in 1992 (Juster and Suzman 1995). Our analysis focuses on married couples. The survey waves 1992, 1994, 1996, 1998, 2000, and 2002 have 4663 married couples.[2] Because wealth at retirement may be significantly affected by marriage history (e.g. Guner and Knowles 2004), we limit attention to couples in which each spouse has been married only once. This reduces the sample to 3046 households. The number of couples with single–marriage spouses for which we can construct lifetime earnings for both spouses is 1582 (see below). We require birth dates for all children, which further limits the sample to 1581 households. We include in our analysis only those household observations for which both the man and the woman are retired by 2002 (see below how the retirement age for the household is constructed), so their full earnings history is realized and observed.

We further restrict the sample as follows. As a protection against coding errors, we exclude any household observation with negative HRS net worth or comprehensive net worth above $5 million. We exclude males who are disabled when they retire, males who retire but later return to work, males who retire before age 56 or after age 68, couples with age difference exceeding 6 years, and males or females less than 4 years short of the mean age of death (74 and 80, respectively). After restricting the sample to households in which both spouses are retired and making the other adjustments, our final sample has 441 households.

Household starting age and retirement age We define the age of the household to be the age of the adult male. The household begins when either the man or the woman becomes independent. We define independence as the maximum of age 18 or the individual's years of education plus 6. Prior to marriage, the man and the woman live apart and thus each contribute 1 to the household's equivalent adult total. Once they marry, prior to their first child, the number of household equivalents is 1.5. Formally, let S_i^M be the maximum of 18 and the adult male's years of education plus 6, let S_i^F be the same for the adult female, and let D_i be the female's birth date minus the male's. The household begins at age S_i with

$$S_i = \min\left\{S_i^M\,,\; S_i^F + D_i\right\}.$$

(Note that D_i serves only to restate the age of the household in terms of the man's age.)

Our definition of "retirement" is as follows. The HRS asks individuals whether they are retired; it separately asks whether their retirement status is fully retired, partly retired, or not retired. In our analysis, an individual is retired if he (or she) answers yes to the first question, lists his/her retirement status as fully retired, or works less than 500 hours per year and does not list his/her retirement status as not retired. We exclude from the sample males who never worked. Our procedure classifies a household as retired when the male stops work.

Household size Our sample includes data from households with differing numbers of children and differing birth dates for children, age of marriage, etc. To incorporate this heterogeneity, we aggregate family members into a single index of "equivalent adults." Specifically, for household i at (male) age s, let the number of equivalent adults N_{is} be

$$N_{is} = n_{is}^A + n_{is}^S \cdot \alpha^S + n_{is}^C \cdot \alpha^C\,, \tag{19}$$

where is $n_{is}^A = \{0, 1, 2\}$ is the number of non-spouse adults in the household, $n_{is}^S = \{0, 1\}$ is the number of spouses and $n_{is}^C = \{0, 1, 2\}$ is the number of children in the household. The number of non-spouse adults is number non-spouse adults n_{is}^A is 0 until male or female begins work; 1 when male or female begins work; 2 when both are working but they have not yet married; 1 after they marry; 1 until both die - then 0. The number of spouses n_{is}^S is 1 between the age of marriage and the death of the first spouse, and is 0 otherwise. n_{is}^C is number of children age 0-20 in household at current male age, subject to a maximum of 2.

Male earnings For both men and women, the HRS provides annual earnings (as well as market hours) in each survey wave. If an HRS participant signs a permission waiver, we also have his or her Social Security Administration (SSA) annual–earnings history for the years 1951-91.[3]

[2] Unless otherwise noted, all HRS data is in public–use files — see http://hrsonline.isr.umich.edu/.

[3] Access to this data is more restricted than the HRS survey data. See http://hrsonline.isr.umich.edu/.

The SSA earning histories are a unique resource. They are not, however, without disadvantages that complicate analysis. (i) The Social Security System does not cover all jobs — the linked data provide no information on non-FICA employment. (ii) SSA earnings data are right–censored in some cases because they only track earnings up to the year's statutory maximum income level subject to the Social Security tax.[4] (iii) Social Security records do not include work hours or wage rates.

The earnings of individuals who have non-FICA jobs or who do have FICA jobs, but earn above the statutory maximum, need to be imputed. The imputation methodology is described in House, Laitner and Stolyarov (2007).

Female earnings As with male earnings, female earnings from the non-FICA jobs and those above the Social Security cap need to be imputed. In addition, our analysis requires female earnings (actual or potential) at every age, regardless of her participation decision. The HRS data only has actual earnings during participation years, so the potential earnings during non-participation spans need to be imputed as well. The imputation methodology is similar to that for males (see House, Laitner and Stolyarov, 2007), except in the case of female earnings we control for the fact that some earnings observations may come from part-time work.

Female participation profile The female participation profile is constructed from the participation data for FICA and non-FICA jobs that have different level of detail on participation. For FICA jobs, participation $p_{is}^{FICA} = \{0, 1\}$ is observed in each year. However, we suspect that years that bracket participation spells might correspond to fractional paticipation $0 < p_{is} < 1$, and these fractions need to be imputed to construct the female participation profile. The imputation is done as follows. Let I_{il} denote the participartion spell number $l \geq 1$ for household i,defined as a set of consecutive years where $p_{is} = 1$. The start of participation spell is $t_{il}^S = \min I_{il}$, and the end is $t_{il}^E = \max I_{il}$. Earnings in years t_{il}^S and t_{il}^E may be well below the earnings for other years in the participation spell, which we think indicate fractional participation. Consequently, we use the earnings data for years inside the participation spell to extrapolate the earnings in the start and end years:

$$\hat{y}_{i,t_{il}^S}^{FICA} = f\left(\left\{y_{is}^{FICA}\right\}_{t_{il}^S < s < t_{il}^E}\right),$$

where $f(\cdot)$ can vary with extrapolation method. If the extrapolated earnings are above actual, we interpret this as fractional participation. Formally,

$$\hat{p}_{is} = \max\left\{1, \frac{y_{is}^{FICA}}{\hat{y}_{is}^{FICA}}\right\},\ s = \left\{t_{il}^S, t_{il}^E\right\},\ \text{all } l \geq 1.$$

Thus we construct the fractional participation profile from FICA jobs as

$$p'_{is} = \begin{cases} \hat{p}_{is}, & s = \left\{t_{il}^S, t_{il}^E\right\},\ \text{all } l \geq 1 \\ p_{is}^{FICA}, & \text{all other } s \end{cases}.$$

For non-FICA jobs, the only information in the survey is the number of years τ^{NF} worked in a non-FICA job and the interval of years I^{NF} when this employment occurred. We assume that all τ^{NF} years of employment occurred in one spell and "locate" this spell in the interval I^{NF} in such a way as to minimize the distance $|\tau_i - \hat{\tau}_i|$ (note that both τ_i and $\hat{\tau}_i$ depend on where the non-FICA participation spell occurs). The outcome of this minimization is a non-FICA participation profile p''_{is}.

Finally, we set

$$p_{is} = \max\left\{p'_{is}, p''_{is}\right\},\ \text{all } s$$

and the potential female earnings

$$\tilde{y}_{is}^F = \begin{cases} y_{is}^{FICA}, & \text{if } p'_{is} > 0 \\ y_{is}^{NON-FICA} & \text{if } p'_{is} = 0 \text{ and } p''_{is} > 0 \\ y_{is}^F & \text{if } p'_{is} = 0 \text{ and } p''_{is} = 0 \end{cases}. \tag{20}$$

[4]The SSA also provides linked W2 tax reports annually for1980-91. Although the W2 records are right–censored for confidentiality, the upper limit is substantially higher than the Social Security earnings cap — $125,000 for earnings under $250,000; $250,000 for earnings under $500,000, and $500,000 for earnings above that amount. In practice, we assume right–censoring at $125,000 for all W2 amounts at or above $125,000. The W2 amounts include non–FICA earnings — and separately identify the latter. They omit some tax deferred pension amounts. Although they also omit self–employment earnings, they identify Social Security measures of the latter. In practice, an individual may have multiple jobs, and we add the corresponding W2 amounts.

This treatment implicitly assumes that in a course of one year a woman can hold either a FICA job or a non-FICA job, but not both. If a woman does not participate, her potential earnings are taken from the earnings dynamics equation.