Children and Household Wealth

John Karl Scholz
Department of Economics, the Institute for Research on Poverty, and NBER
University of Wisconsin – Madison
1180 Observatory Drive
Madison, Wisconsin 53706-1393
jkscholz@wisc.edu

Ananth Seshadri
Department of Economics
University of Wisconsin – Madison
1180 Observatory Drive
Madison, Wisconsin 53706-1393
aseshadr@ssc.wisc.edu

July 2009

Abstract:

This paper examines the effect that children have on wealth accumulation. The fact that there are few papers on this topic is puzzling, since children have implications for optimal retirement planning for nearly every American household. To examine this issue, we incorporate endogenous fertility choice in the spirit of Barro and Becker (1989) into an augmented life-cycle consumption model. We use the model and data from the Health and Retirement Study (HRS) to show four central results. First, the model matches closely joint patterns of wealth and fertility in the HRS. Second, children have a substantial effect on the level and dispersion of wealth and thus should be accounted for in typical retirement planning advice. Third, children affect the timing of credit constraints, particularly for households with low lifetime income. Fourth, once variation in family size is accounted for, means-tested transfer programs (and the asset tests associated with them) have relatively minor effects on wealth accumulation.

This paper examines the effect that children have on wealth accumulation. Despite a large literature on life-cycle wealth accumulation, few papers examine this issue, and none quantify, using a life-cycle framework, the effect of children on net worth. This omission is puzzling, since if children have a substantial effect on net worth, it has implications for optimal retirement planning for nearly every American household. Conventional financial planning advice, for example, relies heavily on a "replacement rate concept" that ignores children and hence is inappropriately naïve.

Children are also fundamental to understanding why the distribution of retirement wealth is much more dispersed than earnings. Using data from the Health and Retirement Study (HRS) and social security earnings records, the ratio of real lifetime earnings for a married household at the 90[th] percentile of the lifetime earnings distribution relative to the earnings of a married household at the 10[th] percentile (referred to as the 90-10 ratio) is 4.6. The 90-10 ratio for 1992 married household net worth (including housing wealth) is 43.9. Explaining the dispersion in wealth has been a longstanding challenge. A simple-minded framework that assumes earnings differences solely explain wealth differences across the rich and the poor is too simplistic.[1]

After presenting a set of simple stylized facts using data from the Health and Retirement Study, we use a lifecycle model to study the effects of children and wealth. Our approach integrates a model of endogenous fertility choice, along the lines of Barro and Becker (1989), with a model of life-cycle consumption, as pioneered by Modigliani and Brumberg (1954) and later extended by Deaton (1991), Hubbard, Skinner, and Zeldes (1995), and others. We make fertility endogenous because tax and transfer (and other) policies that are natural candidates for

---

[1] A recent study documenting this fact is Dynan, Skinner, and Zeldes (2004).

analysis with the lifecycle model may also affect fertility. [2] Moreover, as we show in this paper, the effects of children on consumption decisions can be substantial.

Households in our model derive utility from consumption and the quality and quantity of children. They face uncertainty over earnings, longevity, and late-in-life medical expenses. Given the parameters of the model and the stylized tax, transfer, and social security systems that households face, we calculate optimal decision rules for consumption and family size, household-by-household in the HRS. Then using data on fertility and earnings (provided by the Social Security Administration), we can compare *actual* fertility and wealth of HRS households to their *optimal* fertility and wealth, given our augmented life-cycle model. We then use the model to examine the effects of children on household wealth, study the interplay between children and credit constraints, and compare the effects of children relative to asset tests of cash- and near-cash transfers.

In Scholz, Seshadri, and Khitatrakun (2006) we showed that an augmented lifecycle model that incorporates uncertain earnings, longevity, and health closely matches the level and distribution of wealth for households in the 1992 wave of the HRS. This paper differs from our prior work in several ways. Among the technical modeling differences, the number of children that a family chooses to have is now a choice variable in the model and households can have a purposeful bequest motive. These modifications matter. The high fertility rates of the low earnings households, for example, play a crucial role in accounting for their very low asset holdings. Similarly, the low fertility rate of the higher lifetime earners along with their bequest motives help explain their high wealth holdings.

---

[2] Important examples of endogenous fertility models include Galor (2005), who studies the link between mortality and development, and Doepke (2004), who studies the effect of introducing skill-intensive technologies on fertility.

The applications examined in this paper, of course, also differ from our earlier work. We show our new model closely matches the joint distribution of fertility and wealth. Moreover children have a large effect on the level and dispersion of wealth. We provide perspective on the magnitude of the results in three different ways. First, within the model, we vary the number of children and timing of fertility to better understand the effect of children on wealth. Altering the *number* of children has a substantial effect on net worth across lifetime earnings deciles.

Second, we show the effects of children on wealth are much larger than the effects of asset tests associated with transfers, given earnings realizations and the social security system experienced by households in the HRS. This result is striking, given results of Hubbard, Skinner and Zeldes (1995) who show asset tests associated with transfer programs will cause lower lifetime income households to accumulate little wealth. Once children are accounted for in households, transfer programs have a much smaller effect on asset accumulation.

Third, we present reduced-form correlations from the HRS data on children and net worth that are consistent with our conclusion that children have a substantial effect on household wealth.

We also show that credit constraints are quantitatively important, and fertility and credit constraints interact in ways that significantly affect wealth accumulation. In particular, poorer households with more children are typically credit constrained for a longer time than their richer counterparts. Absent the systematic variation in family size with respect to income, the model implies that *richer* households would be credit constrained for longer, since they have steeper age-earnings profiles than poorer households. The wide dispersion in wealth holdings arises, in part, from the interaction between the earnings and fertility distributions in a world with uninsurable risks and borrowing constraints.

3

I. Facts about Children and Wealth for Households in the Health and Retirement Study

The HRS is a panel dataset initiated in 1992 with in-home, face-to-face interviews of 12,652 persons in 7,702 households who are members of (or spouses of someone in) the 1931-1941 birth cohort. Follow-up telephone interviews were given in 1994, 1996, 1998, 2000, 2002, and 2004. A large fraction of those in the original HRS cohort gave permission for researchers to gain access, under tightly restricted conditions, to their social security earnings records. Our analytic approach requires earnings realizations data, so we focus on the 1992 data. For our analysis sample we exclude 379 married households where one spouse did not participate in the 1992 HRS, 93 households that failed to have at least one year of full-time work, and 908 households where the highest earner began working full time prior to 1951.[3]

Single households in our data had an average of 2.8 children (recall all were born between 1931 and 1941), so the vast majority had previous marriages. We drop 2,121 single adult households in 1992 because we do not observe the earnings of former partners, and hence we cannot characterize accurately their lifetime resources and annual household earnings realizations.[4] Our final sample has 4,201 married households in 1992.

The survey covers a wide range of topics, including batteries of questions on health and cognitive conditions; retirement plans; subjective assessments of mortality probabilities and the quality of retirement preparation; family structure; employment status and job history; demographic characteristics; housing; income and net worth; and pension details.

---

[3] We drop the first group because we do not have information on spousal, and hence household, income. We drop the second group because we do not have information on transfer payments in years prior to the HRS survey and therefore we cannot model the lifetime budget constraint. We drop households where the highest earner started working before 1951 for computational reasons. Our procedures to impute missing and top-coded data are more complicated when initial values of the earnings process are missing.

I.1. Children in the HRS

In this subsection we examine patterns in the HRS data between children, factors that likely influence wealth accumulation, and wealth itself. In Table 1 we summarize some characteristics of the married HRS population by the number of children they have. Column 1 shows the modal number of children for the sample is two, but 39.5 percent of families have three or four children. Not surprisingly, as the number of children increases, the mean age of the primary earner when the last child is born increases. And the later fertility is completed, the smaller is the share of lifetime earnings received after the last child is born. As we discuss later, a substantial fraction of HRS households are credit constrained early in life. Since children increase household consumption requirements, the presence of children in the household and the timing of births may affect the length of the credit-constrained period.

The final three columns of Table 1 highlight patterns of net worth and lifetime income by the number of children in households.[5] We summarize the relationship in Figure 1. For each household we calculate the ratio of net worth (in 1992) to real (undiscounted) lifetime earnings and plot the median of these values for families, tabulated by the number of children they have.[6] The ratio of net worth (in 1992) to lifetime income is highest for families with no children. It falls monotonically with the number of children above two.[7] Figure 1 gives suggestive evidence

---

[4] For couples who have remarried, we implicitly assume that the current partners were always married and had their children together. This assumption would also hold if the earnings and fertility of prior spouses were identical to the current spouse. Our results are similar if we restrict the sample to couples who have never divorced.

[5] Net worth (private savings) is a comprehensive measure that includes housing assets less liabilities, business assets less liabilities, checking and saving accounts, stocks, bonds, mutual funds, retirement accounts including defined contribution pensions, certificates of deposit, the cash value of whole life insurance, and other assets, less credit card debt and other liabilities. It excludes defined benefit pension wealth, social security wealth, and future earnings. The concept of wealth is similar (and in many cases identical) to those used in other studies of wealth.

[6] In brief, our use of restricted access social security earnings records allows us to construct an unusually accurate measure of real lifetime earnings. We account for top-coding of social security earnings records, missing observations, and future earnings (making use of past earnings and individuals' expected retirement dates). A brief explanation is given in the Appendix.

[7] Similar patterns hold if we calculate the ratio of net worth to lifetime income using the numbers in Table 1.

that net worth is not fully determined by lifetime earnings and that children may have some effect on the dispersion of wealth.

Table 2 shows information similar to that presented in Table 1, but organized by lifetime earnings deciles. The first two columns show that the distribution of net worth is skewed rightward, as the means substantially exceed the medians. The mean number of children among married couples falls from 4.6 in the lowest lifetime income decile to 3.1 in the highest. There is little systematic relationship between the age of completed fertility and lifetime income, despite the fact that the number of children is negatively correlated with lifetime income. This suggests that higher income HRS households are delaying fertility relative to others. Lastly, there is a positive correlation between lifetime income and the fraction of lifetime earnings received after the last child was born. Given there is little systematic pattern in the ages at which the last child was born, this suggests that households with high lifetime incomes have more steeply shaped age-earnings profiles.

Figure 2 plots age-earnings profiles by family size for married HRS households.[8] There appear to be small differences in the earnings trajectories, but in general, the slopes of the profiles look similar. Couples with seven or more children have the lowest incomes over their lifetimes. Couples with 2 and 3 children have the highest age-earnings profiles, with one-child couples being only slightly lower. The profiles fall monotonically as the number of children increases beyond 3, though the differences are relatively small.

The descriptive data are consistent with at least three channels through which children may influence wealth. First, family size is correlated with lifetime earnings. Second, the number of children varies inversely with lifetime income. If children are costly, this alone will lead to

wealth differences (as a fraction of lifetime income) between high- and low-lifetime income households. Third, those with more children have children later in life so children are present in the household for a larger portion of adults' working years. Below, we systematically explore the implications of these facts in the context of the life-cycle model.

Appendix Table 1 provides means, standard deviations, and in some cases, medians for other variables important to this study. The mean (median) present discounted value of lifetime household earnings is $2,178,452 ($2,059,862).[9] Retirement consumption will be financed out of defined benefit pension wealth (mean is $133,241, median is $45,248);[10] social security wealth (mean is $135,338, median is $133,547);[11] and nonpension net worth (mean is $208,549, median is $142,885). The mean age of the household head is 55.8.[12]

II. A Model of Optimal Wealth Accumulation and Fertility Choice

An influential set of papers examine how children affect consumption profiles. Attanasio and Browning (1995) show that once one accounts for the variation in family size over the life-

---

[8] Specifically, we plot a median log earnings using Stata's "graph twoway mbands" command.

[9] When calculating present discounted values of earnings and social security wealth, we discount the constant-dollar sum of earnings (social security, or pensions) by a real interest rate measure (prior to 1992, we use the difference between the 3-month Treasury bill rate and the year-to-year change in the CPI-W; for 1992 and after we use 4 percent). For the defined benefit pension wealth, we use the intermediate scenarios (that assumes a real return of 2.21 percent) of the Pension Present Value Database available for the HRS data.

[10] The value of defined benefit pensions are calculated using the HRS "Pension Present Value Database" at http://hrsonline.isr.umich.edu/data/avail.html. The programs use detailed plan descriptions along with information on employee earnings. We use self-reported defined-benefit pension information for households not included in the database. The assumptions used in the program to calculate the value of defined contribution (DC) pensions – particularly the assumption that contributions were a constant fraction of income during years worked with a given employer – are likely inappropriate. Consequently, we follow others in the literature (for example, Engen et al., 1999, p. 159) and use self-reported information to calculate DC pension wealth.

Defined benefit pension expectations are formed on the basis of an empirical pension function that depends in a nonlinear way on union status, years of service in the pension-covered job, and expectations about earnings in the last year of work. We estimate the function with HRS data. Details are in Scholz, Seshadri, and Khitatrakun (2006).

[11] We develop a social security calculator to compute benefits based on the social security earnings histories (and for those who refused to release earnings, imputed earnings). Households in the model expect the social security rules in 1992 to prevail and develop expectations of social security benefits that are consistent with their earnings expectations. Details are in Scholz, Seshadri, and Khitatrakun (2006).

cycle, a flat age-consumption profile (consistent with the certainty life-cycle model where

$$\beta = \frac{1}{1+r}$$) obtains. Browning and Ejrnæs (2002) argue that precautionary motives may not play

an essential role in generating hump-shaped age-consumption profiles: taking proper account of

the ages and number of children may be sufficient.[13] Attanasio *et al.* (1999) demonstrate that a

life-cycle model with uncertainty and changing demographics can account for the life-cycle

profile of consumption.

Our paper differs from the above work in many ways. A critical difference, however, is that

we focus on *levels* of wealth and fertility, household by household. Explaining variations across

specific households in wealth and fertility requires us to solve a model of optimal consumption

with endogenous fertility and endogenous bequests. Matching the joint distribution of wealth

and fertility is a more ambitious goal than matching consumption profiles.

Banks *et al.* (1994, Figures 5-7) show in the life-cycle model with no uncertainty and

perfect capital markets, that larger families consume more of their income earlier in their life-

cycle and hence consume less in retirement. Put differently, larger families would appear to be

more impatient, consuming a greater share of lifetime resources when children are present,

relative to families with fewer children (all else being equal).

It is well known that accounting only for family size variation in the context of the

permanent income, life-cycle model is not enough to explain the level and skewness of wealth.[14]

The life-cycle model we solve, therefore, is augmented to incorporate uncertain lifetimes,

---

[12]The head of household is defined throughout the paper as the person in the household with the largest share of lifetime earnings. When we refer to the age or retirement date of the household, we are referring to the age or retirement date of the household head.

[13] Lawrance (1991) finds that accounting for variation in family composition reduces the heterogeneity in discount factors estimated from a consumption Euler equation.

uninsurable earnings, uninsurable medical expenses, and borrowing constraints. Wealth and

earnings expectations affect decisions about the number of children, so we make fertility an

endogenous variable. We follow the pioneering work of Barro and Becker (1989) and assume

that parents get utility from the quantity and the quality of their children.[15] We allow parents to

have purposeful bequest motives. We do not model the timing of children, and instead assume

that parents give birth to all their children at age $B > S$. In our calibration, we set B to 27.

Children are assumed to remain in the household for 18 years.

Parental preferences are given by

$$E\left[\sum_{j=S}^{D} \beta^{j-S} U(c_j) + \sum_{j=B}^{B+17} \beta^{j-S} d(f) U(c_j^k) + \beta^{j-S} d(f)\Phi(b) \right].$$

The operator E denotes the expectation over uncertain future earnings, health expenditures, and

life span. With probability p_j the household survives into the next period, so the household

survives until age j with probability $\prod_{k=S}^{j-1} p_k$, where $\prod_{k=S}^{j-1} p_k = 1$ if $j-1 < R$.[16] At age D,

$p_D = 0$. The discount factor on future utilities is β.

Parents care about the number of children, f, and utility per child, $U(c_j^k)$. The function $d(f)$

denotes the weight that parents place on the number of children, which we assume is increasing

and concave. We also assume that children are costly. The budget constraint during the period

of time when the kids are attached to parents is given by

---

[14] See, for example, Hubbard, Skinner, and Zeldes (1995) or DeNardi (2004).

[15] We assume that the household has a joy of giving motive – they value passing on bequests to their children. Solving the full blown dynastic model is prohibitively computationally-intensive since the parent would need to keep track of the child's state space (and the grandchild's) once the child (and the grandchild) forms his or her own household.

[16]We do not model marriage or divorce. Married households in 1992 are modeled as making their lifecycle consumption decisions jointly with their partner throughout their working lives. They become single only when a spouse dies.

$$c_j + fc_j^k + a_{j+1} = y_j + a_j - \tau\left(e_j + ra_j\right), j \in \{B,...,B+17\},$$

where

$$y_j = (1-\kappa f)e_j + ra_j + T(e_j, a_j, j, n_j), \ j \in \{S,...,R\}$$

and the borrowing constraint is

$$a_{j+1} \geq 0. \ [17]$$

Notice that each child, in addition to their own consumption needs, requires the fraction κ of the parent's earnings, where e_j denotes labor earnings at age j. This captures the indirect time costs associated with bearing and rearing children. This fixed cost implies that higher e_j households will tend to have fewer children than their lower e_j counterparts, though absent the fixed cost, they would have more children and consumption than their lower e_j counterparts.

The budget constraint during the retirement period is given by

$$c_j + a_{j+1} + m_j = y_j + a_j - \tau\left(SS\left(\sum_{j=S}^{R} e_j\right), DB(e_R) + ra_j\right), j \in \{R+1,...,D\} \text{ where}$$

$$y_j = SS\left(\sum_{j=S}^{R} e_j\right) + DB(e_R) + ra_j + T_R(e_R, \sum_{j=S}^{R} e_j, a_j, j, n_j), \ j \in \{R+1,...,D\}.$$

The variable y_j defines taxable income for working and for retired households.[18] Labor supply and the retirement date are exogenous. $SS(\cdot)$ are social security benefits, which are a function of

---

[17] The borrowing constraint is a natural consequence of the economic environment. The household faces uncertain earnings when working (where earnings draws can get arbitrarily close to zero) and health shocks and uncertain longevity when retired. Households cannot die with debt. Consequently, they maintain a nonnegative asset position at every age.

[18] To define a household's retirement date for those already retired, we use the actual retirement date for the head of the household. For those not retired, we use the expected retirement date of the person who is the head of the household.

aggregate lifetime earnings, and $DB(\cdot)$ are defined benefit receipts, which are a function of earnings received at the last working age. The functions $T(\cdot)$ and $T_R(\cdot)$ denote means-tested transfers for working and retired households. Transfers depend on earnings, social security benefits and defined benefit pensions, assets, the year, and the number of children and adults in the household, n. Out-of-pocket medical expenditures are denoted by m_j and the interest rate is denoted by r.[19] The tax function $\tau(\cdot)$ depicts total tax payments, which are a function of earned and capital income for working households, and a function of pension and capital income plus a portion of social security benefits for retired households.[20] The bequest that each child receives, b, equals fa_{j+1}. Parents' assets are distributed evenly across their children. As is standard, our specification for bequests incorporates both an accidental and an intentional component.

We simplify our computational problem by assuming households incur no out-of-pocket medical expenses prior to retirement and face no pre-retirement mortality risk. The dynamic programming problem for working households, therefore, has two fewer state variables than it does for retired households. During working years, the earnings draw for the next period comes from the distribution Φ conditional on the household's age and current earnings draw. We assume that each household begins life with zero assets.

The decision problem faced by households includes three more choice variables than the standard life-cycle consumption problem: the fertility rate, f; consumption per child, c_j^k; and bequests, b. The first order conditions with respect to c_j^k, f, and b are given by

---

[19]Medical expenses are drawn from the Markov processes $\Omega_{jm}(m_{j+1} \mid m_j)$ for married households. Medical expenses after the death of a spouse are assumed to be half of those drawn from the distribution for married couples.

$$c_j^k : U'(c_j) = d(f)U'(c_j^k),$$

$$f : U'(c_B)\left[c_B^k + \kappa e_B - \frac{\partial T}{\partial f} \right] = \frac{\partial d(f)}{\partial f} EV_{B+1}(\bullet) + d(f)E\frac{\partial V_{B+1}(\bullet)}{\partial f},$$

$$b : f\frac{\partial U(c_j)}{\partial b} = d(f)E\frac{\partial V_{B+1}(\bullet)}{\partial b}$$

In the above equations, $V_{B+1}(\bullet)$ stands for the value function at age $B+1$. Consider the first order condition for fertility choice: the left hand side represents the marginal cost of an additional child, which is increasing in earnings. There are two reasons for this. First, families with higher earnings will spend more on their children's consumption. Second, each child costs the fraction κ of parent's earnings. The right hand side represents the marginal benefit of an additional child: the two terms stand for the marginal effect on quantity, holding quality fixed, and the marginal effect on quality, holding quantity fixed.

Instead of the Barro-Becker motivation for children, readers might wonder whether parents, particularly with low lifetime incomes, have children as an investment with the expectation that children might support them in old age. There is little evidence of children playing this role in the United States. Gale and Scholz (1994) review older evidence, which shows only small transfer flows from children to parents in data from the Survey of Consumer Finances (SCFs). More recent data from the 2004 SCF yields a similar conclusion. Roughly 20 percent of U.S. households report having received a "substantial" gift or inheritance. Fewer than 1 percent of these gifts were from children.[21] Sixteen percent of the sample reports having made a gift in the

---

[20]Specifically, taxable social security benefits for married couples are calculated from the expression $\max(0, \min(0.5 * SS\ Benefits,\ Income - 0.5 * SS\ Benefits - 32,000))$. This approach approximates the law in effect in 1992.

[21] Overwhelmingly gifts were from older relatives, primarily parents, grandparents, and aunts and uncles.

previous year. About 18 percent of these gifts in 2003 went to individuals in an older generation (parents or grandparents), so there is a striking difference in reports of transfers given and transfers received, which is consistent with parents understating transfers received from children when asked in surveys. Nevertheless, the relatively infrequent occurrence of transfer receipt from a younger generation (5 million out of 112 million households), and relatively small conditional amounts of transfers ($2,000), makes the investment model of children a less appealing explanation of U.S. fertility choices than the Barro-Becker framework.

II.1. Model Parameterization[22]

We assume households have constant relative risk-averse preferences, so $U(c) = \left\{ \dfrac{c^{1-\gamma}}{1-\gamma} \right.$.

We follow our earlier work and a number of prior papers in assuming $\gamma = 3$, the discount factor $\beta = 0.96$, and the annualized real rate of return is 4 percent. Following Jones and Schoonbroodt (2009) the altruism function is given by $d(f) = b_0 f^{b_1}$, $b_1 < 0$. The original Barro-Becker formulation relied on utility being positive and hence $\gamma < 1$. Our current formulation allows us to use the same parameter values for the utility function that is standard in life-cycle models. We assume that the bequest function is given by $\Phi(b) = \dfrac{b^{1-\phi}}{1-\phi}$. This isoelastic functional form is used elsewhere in the literature (see for instance, DeNardi, 2004).

Earnings expectations are a central influence on life-cycle consumption and fertility decisions, both directly and through their effects on expected pension and social security benefits. We aggregate individual earnings histories into household earnings histories (a brief

---

[22] A brief description of our approach to addressing top-coded social security earnings records, the tax function, and expectations about out-of-pocket medical expenses is given in the appendix.

discussion is in the appendix). The household model of log earnings (and earnings expectations) is

$$\log e_j = \alpha^i + \beta_1 AGE_j + \beta_2 AGE_j^2 + u_j,$$

$$u_j = \rho u_{j-1} + \varepsilon_j,$$

where, as mentioned above, e_j is the observed earnings of the household i at age j in 1992-dollars, α^i is a household specific constant, AGE_j is age of the head of the household, u_j is an AR(1) error term of the earnings equation, and ε_j is a zero-mean i.i.d., normally distributed error term. The estimated parameters are α^i, β_1, β_2, ρ, and σ_ε.

We divide households into four groups according to education and the number of earners in the household, resulting in four sets of household-group-specific parameters.[23] Estimates of the persistence parameters range from 0.61 for married one-earner households without college degrees to 0.76 for married households with two earners, in which the highest earner has at least a college degree. The variance of earnings shocks ranges from 0.08 to 0.10.

One purpose of this paper is to assess the importance of key factors affecting household wealth. Hubbard, Skinner and Zeldes (1995) conclude that the asset tests associated with means-tested transfers is one such factor, particularly for low-income households. To examine their results, we model public income transfer programs using their specification. Specifically, the transfer that a household receives while working is given by

$$T = \max\left\{0, \underline{c} - \left[e + (1+r)a\right]\right\},$$

---

[23]The groups are (1) married, head without a college degree, one earner; (2) married, head without a college degree, two earners; (3) married, head with a college degree, one earner; and (4) married, head with a college degree, two

14

whereas the transfer that the household will receive upon retiring is

$$T_R = \max\left\{0, \underline{c} - \left[SS(E_R) + DB(e_R) + (1+r)a\right]\right\}.$$

This transfer function guarantees a pre-tax income of $\underline{c}$, which we set based on parameters drawn from Moffitt (2002).[24] Subsistence benefits ($\underline{c}$) for a one-parent family with two children increased sharply, from $5,992 in 1968 to $9,887 in 1974 (all in 1992 dollars). Benefits have trended down from their 1974 peak—in 1992 the consumption floor was $8,159 for the one-parent, two-child family. Following Hubbard, Skinner, and Zeldes, this formulation implies that earnings, retirement income, and assets reduce public benefits dollar for dollar.

Our model introduces four new parameters relative to a standard life-cycle model with exogenous fertility: b_0, b_1, ϕ and κ. The parameter κ measures the time cost of children. According to Haveman and Wolfe (1995) the cost per child computed as the reduction in the mother's time spent in the paid labor force valued at the market wage is about 9.5 percent of parent's earnings. Consequently we set κ at 0.095. This leaves us with three parameters we need to set: b_0, b_1 and ϕ.

Given the functional forms, the first order condition for the optimal choice of consumption is given by $c_j^k = b_0^{1/\gamma} f^{b_1/\gamma} c_j$. Hence total family consumption is given by

$fc_j^k + c_j = \left(1 + b_0^{1/\gamma} f^{1+b_1/\gamma}\right)c_j$. To calibrate the remaining parameters, we make use of the

equivalence scale from Citro and Michael (1995), $g(A_j, K_j) = (A_j + 0.7K_j)^{0.7}$, where A_j

indicates the number of adults in the household and K_j indicates the number of children in the

household. There are other equivalence scales, including ones from the OECD (1982),

Department of Health and Human Services (Federal Register, 1991) and Lazear and Michael

(1980). The scale we use implies that a two parent family with 3 children consumes 66 percent

more than a two parent family with no children. The corresponding numbers for the other

equivalence scales of 88 percent, 76 percent and 59 percent bracket our scale.

We use the equivalence scale to match the consumption needs of a household with 2

children relative to those with one, as well as the consumption requirements of a household with

4 children relative to those with one. These two moments, together with the requirement that we

match the sample mean fertility rate for a "representative" household with mean earnings

realizations pins down the remaining parameters, which are $b_0 = 0.44$, $b_1 = -2.56$, and $\phi = 2.02$.

Notice that the utility function exhibits greater curvature in consumption than in bequests,

implying that bequests are a luxury good as discussed in Kopczuk (2009). We then use the same

parameters for all households to examine the model's predictions.

II.2. Model Solution

We solve the dynamic programming problem by linear interpolation on the value function.

For each household in our sample we compute optimal decision rules for consumption (and

hence asset accumulation) and fertility from the oldest possible age (D) to the beginning of

working life (S) for any feasible realizations of the random variables: earnings, health shocks,

and mortality. These decision rules differ for each household, since each faces stochastic draws

16

from different earnings distributions (recall that α_i is household specific). Household-specific

earnings expectations also directly influence expectations about social security and pension

benefits. Other characteristics also differ across households. Consequently, it is not sufficient to

solve the life-cycle problem for just a few household types.

II.3. Results and Policy Experiments

A key feature of our analysis is that we compute optimal decision rules for *each* household

in the HRS. Using the optimal rules, households' actual earnings draws and the rate of return

assumption, we obtain household-level predictions for fertility and wealth. The model then

allows us to conduct counterfactual policy experiments where we can alter features of the

economic environment to enhance understanding of the effect that children have on wealth

accumulation.

We summarize the joint density of fertility and wealth in Table 3, where we compare our

predictions to their actual values, classifying the results by deciles of lifetime income. The first

and third columns of the Table repeat (for convenience) data on median and mean net worth

shown in Table 2. The second and fourth columns show the median and mean optimal net worth

targets for married HRS households. These targets include resources that could be accumulated

in real and financial assets, the current value of defined contribution pensions, including 401(k)s,

and housing net worth, given households' expected defined benefit pension and social security

benefits.

The optimal wealth target for the median households in the lowest decile of the lifetime

earnings distribution is $16,684. The mean target for bottom decile households is $69,405.

These low targets are largely a consequence of three factors. First, lifetime earnings are low for

bottom decile households and social security is mildly progressive. Second, the number of

17

children in this cohort is inversely related to lifetime earnings (married couples in the bottom decile of lifetime earnings had 4.6 children, couples in the highest decile had 3.1 children). The variation in fertility across lifetime income deciles has an effect similar to imposing a higher discount rate for low-income households. Both will reduce optimal wealth accumulation, all else being equal. Third, the average age of households is 55.8, so the average household will work (and accumulate wealth) for many additional years before retiring.

The optimal targets are increasing in lifetime income. However, observed median and mean net worth of married HRS households exceed the optimal targets in each lifetime income decile. As discussed in Scholz, Seshadri, and Khitatrakun (2006), there is no evidence that these households are failing to accumulate sufficient resources to maintain living standards in retirement.

A common criticism of the life-cycle model is that it is unable to match the skewness of the observed wealth distribution: our model does so. The 90-25 ratio of weighted net worth in the data is 10.2. In the simulated optimal wealth data it is 9.8. The coefficient of variation in the actual data is 1.8, in the simulated data it is 1.9. In the data the top one percent of households hold 8.6 times the mean wealth holdings. In the model simulations, the number is 8.4. Absent fertility choice and bequests, this ratio drops to 5.4, suggesting that the combination of the variation in fertility rates and bequests helps explain the right tail of the wealth distribution.

The model also does a good job matching the distribution of children across lifetime earnings deciles. Recall that we impose a time cost parameter, κ, for children that we set at 9.5 percent. Thus, it is impossible in our framework for any family in the model to "optimally" have more than 10 children. Yet, about 2 percent of HRS households have more than 10. Many

factors not captured in the model – fecundity and the timing of marriage being two obvious ones – will affect fertility patterns. So it perhaps is not surprising that we fail to match observed behavior precisely.

Figure 4 shows a box and whisker plot of the actual number of children in households (on the horizontal axis) and optimal number of children (on the vertical axis). The boxes indicate the 25[th] and 75[th] percentiles of the distribution of the optimal number of children. The lines in the middle of the box indicate the median number of children.[25] The mean optimal number in the endogenous fertility model fairly closely matches the data as the actual number of children, shown on the horizontal axis, varies from 0 to 8 or more. There is considerable dispersion in our calculations of the optimal number of children, but the correlation between the optimal number and actual number of children is 0.74. The endogenous fertility model, even with its parsimonious specification of the fertility process, does a good job matching the level and dispersion of the number of children in married HRS households.

II.3.1. *The Effect of Children on the Cross-Sectional Distribution of Wealth*

To assess the quantitative significance of the systematic variation in birth rates by lifetime earnings, we assign each married couple the mean number of children (for all married couples). Specifically, married couples are assumed to have 3.7 children, born at age 27. Allowing households to have "fractional" children ensures that the aggregate number of children in the simulated economy matches the number of children born to HRS households. This consistency is essential if children, in fact, are shown to have an important effect on wealth. Assuming that households are *endowed* with the mean number of children, we then examine the predictions for wealth. Children are, of course, still assumed to be costly: there is a resource cost governed by

19

κ and there is a consumption cost, which, when we fix the number of children, is exogenous and governed by the equivalence scales.

As can be seen from Table 4, the effect of altering the number of children is substantial. When the lowest income decile households have 3.7 children instead of 4.6 children, median optimal net worth increases to $26,452 from $16,684. Mean optimal net worth increases to $103,369 from $69,405. In the highest income decile the mean number of children increases to 3.7 from 3.1. Median optimal net worth falls to $266,379 from $350,986. Mean optimal net worth falls to $471,857 from $585,261. There is very little change in the middle lifetime income deciles because these households, on average, have close to the mean number of children. The wealth changes shown in Table 4 are driven by two distinct factors. First, for households that have fewer (more) children in the counterfactual simulations than they do in the data, child-oriented and aggregate expenditures are lower (higher) than they otherwise would be, which increases (decreases) their retirement wealth. Second, children affect the length of time households will be credit constrained.

The interactions between children and credit constraints are also highlighted in Table 4. In the baseline economy, column 2 shows the median household in the lowest lifetime income decile is credit constrained until age 37. As shown in column 4, when there is no variation in the number of children the median household in the lowest lifetime income decile begins saving for retirement at age 26.

Low lifetime income households have, on average, more children than do high lifetime income households, which means the effective income available to these household after

---

[25] Recall that, given our approach to parameterizing the number of children, households can have a non-integer number of children.

adjusting for family size (through the equivalence scale) falls by more due to children than it does for high-income households. Thus, fertility differences make the resources available for consumption even more dispersed than the distribution of earnings, which implies that asset variation decreases when we shut down the variation in the number and timing of kids. Indeed the coefficient of variation of optimal net worth drops from 1.8 in the baseline optimal net worth distribution to 1.5 when the variation in children is shut down.

II.3.2. The Effect of the Timing of Children

The previous experiment altered the number of children a household has. We do not offer a theory for the timing of births, but timing might matter. Households with children born later in life will spend a greater share of lifetime resources on children than other households, due to the interaction of upward-sloping age-earnings profiles and credit constraints. Therefore, all else equal, families who have children later would be expected to have lower net worth when we observe them in the HRS than otherwise equivalent households.

In fact, we show that the quantitative magnitude of the effect of timing on wealth is small. To show this, we compare the previous results where married couples have 3.7 children born at age 27 to a specification where married couples still have 3.7 children, but they are born at the modal ages at which 4-child families have children: namely, at ages 23, 26, 29, and the 0.7 child at age 33.

The results of our analysis on the timing of fertility are shown in the last column of Table 4. In the lowest lifetime income deciles, the effects on optimal wealth of changing the timing of births is small relative to changing both the number of children and the timing of births. Put differently, the change in optimal wealth accumulation when comparing columns 5 and 3 is tiny compared to the change in optimal wealth accumulation when comparing columns 1 and 3. We

21

conclude that variation across households in the timing of birth has relatively little effect on the patterns of wealth accumulation. Therefore, our failure to offer a theory for the timing of children is a quantitatively minor shortcoming of our analysis.

II.3.3. *The Effect of Transfer Programs*

Hubbard, Skinner and Zeldes (1995) argue that households with low earnings have little wealth (as a percentage of lifetime income) because asset tests associated with means-tested transfer programs discourage saving. To study the effect of transfer programs we set $\underline{c}$ to zero (recall that $\underline{c}$ denotes the generosity of transfer programs).

The structure, benefits, and receipt of transfers modeled here are very similar to Hubbard, Skinner, and Zeldes (1995). They model a consumption floor of $7,000 in 1984 dollars. Our floor in 1984 (based on data provided by Moffitt) is roughly $6,300 dollars.[26] In 1980, when the average HRS respondent was 44 years old, 24.4 percent of households with less than a high school degree received transfers in our model. Hubbard, Skinner and Zeldes report that 23.7 percent of households age 40 to 49 without a high school degree received transfers in the 1984 PSID. A small percentage of college graduates receive transfers in these years (1.2 percent in our model, 2.3 percent in the PSID). A similar close correspondence holds across education groups for households in 1990.[27]

Columns 1 and 2 in Table 5 repeat the optimal median net worth targets and optimal number of children in the endogenous fertility model. The effects of eliminating cash and near-cash transfers in the endogenous fertility model are shown by comparing the first two columns with

---

[26] Our floor, of course, varies by year and by family composition.
[27] The calculations in this paragraph are based on a sample that also includes single households in 1992, since the full population is needed to compare our results to population averages.

columns 3 and 4. The effects of transfer programs are modest: the fertility rate of the poorest households decreases slightly and the simulated median optimal net worth increases by about two-thirds the amount that it did when we changed the number of children. The wealth increase results from two factors. First, holding fertility fixed, the household wants to increase its wealth to provide insurance, which it previously received through transfer programs. Second, eliminating transfer programs reduces the fertility rate of program recipients. The reduction in fertility leads the household to cut back on children's consumption as well. Since the change in the fertility rate is small, this also has a small effect on wealth accumulation. Hence, the overall effect is also small.

Hubbard, Skinner, and Zeldes (1995) discuss a life-cycle model with no children and find a large effect on wealth of asset tests associated with transfer programs. We find a small effect. To show the difference is due to incorporating children in our model, rather than some other modeling decisions, we replicate the Hubbard, Skinner, and Zeldes results using our model. In particular, we *eliminate* all children from households: the sample is then just composed of married, childless couples. Optimal wealth by lifetime income decile for married couples is given in column 5 in Table 5. Then in column 6 we show the effect of eliminating transfers, setting $\underline{c} = 0$. Doing so leads to a very large increase in optimal wealth accumulation, as households take a greater responsibility to self-insure, which is similar to the Hubbard, Skinner, and Zeldes (1996) results.

Children are the major difference explaining our results and those of Hubbard, Skinner and Zeldes. Poorer households by virtue of their larger family size, optimally plan on having fewer resources for retirement. To see why, consider the following example that focuses on two married couples, one with five children, and the other with no children but identical in every

23

other way. The husband and wife in the five-child family becomes accustomed to a lower standard of living prior to retirement than the childless couple, since a significant fraction of the five-child family's resources are used to support their children. Put differently, considerably less wealth is needed to maintain the accustomed living standards of the husband and wife in the five-child family, after their children have left the house, than is needed to maintain the living standards of the otherwise identical childless couple.

This general point is even more important for low-income households. A substantial portion of households, even in the bottom decile, have social security benefits exceeding the consumption floor and thereby assign a very low probability of using safety net programs in the future. The fact that social security benefits cannot be borrowed against and that replacement rates for the poor are (almost) sufficient to cover their reduced consumption requirements in retirement (given that their household size is now much smaller) implies that there is little disincentive effect of transfer programs on their already negligible private asset accumulation.

While we focus on wealth in 1992 when the average household is 55.8 years of age, the model also implies low wealth levels for this cohort earlier in their life-cycle. Indeed a striking aspect of the simulations is that the average household in the bottom decile is borrowing constrained until age 36, while households in the 7[th] lifetime income decile, for example, begin saving for retirement at age 28. Absent demographic variation, the exact opposite holds: richer households, by virtue of their steeper earnings profiles, will be borrowing constrained for a longer period of time. Thus the addition of children into the analysis leads to the prediction that poorer households, despite their flatter earnings profile, will choose not to save for a substantial part of their life cycle, even when there is no disincentive effects of transfer programs. Indeed in

24

our view of the world, having 5 children (or the number of children observed in the HRS) alters optimal consumption choices sufficiently strongly to largely reconcile the low wealth holdings of the very poor with the data: the asset tests associated with transfer programs play a minor role in wealth accumulation.[28]

While we find small effects of transfers on wealth accumulation, our model implies a larger effect of these programs on consumption (and hence welfare) than implied by the Hubbard, Skinner and Zeldes analysis. If transfer programs have a substantial (negative) effect on asset accumulation, then their effect on consumption is smaller than in a world in which their effect on asset accumulation is negligible. Simply put, our analysis implies that poor households have few assets partly due to commitments to their children. The presence of a transfer program increases consumption by a large magnitude, since, in the absence of the transfer program, they would have few resources to support consumption. In contrast, had we assumed that there was no variation in family size, cutting back on the transfer program would have increased asset accumulation, leading to a smaller overall effect on consumption.

II.4. Children and Wealth in the HRS Data

The models we analyze suggest that children are a significant determinant of wealth accumulation. A natural question to ask is whether these patterns are observable in the HRS data. To a certain extent we have already answered the question, as Figure 1 shows net worth as a percentage of lifetime earnings is declining with children once a family has two children. But many papers estimate reduced form models examining correlates of net worth. Here we briefly discuss the correlation between children and net worth in a common regression context.

---

[28] In a regression-based analysis, Hurst and Ziliak (2006) find little effect on wealth accumulation from state-level changes in asset tests associated with the 1996 welfare reform.

Table 6 summarizes the results of a set of reduced form regression models of non-DB-pension, non-social-security net worth. Each regression includes 47 covariates for the household's earnings between ages 18 and 65, as well as the variables shown in the Table: three education indicators, age, and DB pension, race and ethnicity indicator variables. The central variables of interest are the number of children in the first two empirical models, and the number of children interacted with indicator variables for lifetime income quintile in the final two empirical models. The samples for the regressions are the same as those used elsewhere in the paper: namely married couples from the 1992 wave of the HRS. We report the results for OLS regression and for median regressions. Median regression is commonly used in studies examining factors correlated with wealth, due to the skewed distribution of wealth.

Mean net worth in the unweighted sample is $254,025. The coefficient of "number of children" in the mean regression is -$6,384. Varying the number of children from 0 to its mean value of 3.77 would reduce net worth by $24,068, or 9.5 percent of the mean. The magnitude of the child coefficient is somewhat smaller in the median regression, the coefficient of -$2,601 would imply a 6.3 percent reduction from the median net worth of $124,000 when moving from 0 to the median number of 3 children.

These mean and median estimates mask what appear to be much larger effects of children on wealth in the middle lifetime income quintiles. These estimates are shown in the last two columns of Table 6. Following common practice in the literature, we focus on the median regression estimates (the results for the mean regression are shown in column 3). Estimates for children in the bottom income quintile are insignificantly different from zero (and the point estimate is incorrectly signed). Median net worth in the second lifetime income quintile is

$76,670. The estimate of -$3,876 suggests a change from 0 to 3 children would lower net worth $11,628, or 15.2 percent. Similar calculations for the middle quintile suggest a 20.2 percent reduction, from a median net worth of $111,722; and for the fourth quintile, a 14.8 percent reduction from a median of $167,466. The regression estimates are consistent with our prior conclusion that children have a substantial effect on the net worth of households approaching retirement.[29]

III. Conclusions

There are many life-cycle consumption papers but, despite the ubiquity of children, relatively few papers in the literature study the relationship between children and wealth. We examine this relationship in the context of a life-cycle model with uninsurable income, health, and longevity risks, borrowing constraints, endogenous fertility, and purposeful bequest motives. Our study yields four main conclusions.

We are the first to develop an amalgam of an augmented life-cycle model and a Barro-Becker model of fertility choice. We show the model closely matches the joint distribution of wealth and fertility, household-by-household. Children play a critical role in matching wealth dispersion. Low-income families have more children than high-income families. This, coupled with the fact that social security replaces a greater share of lifetime earnings for low-income families than for high-income families goes a long way toward reconciling the observed distribution of wealth.

Second, financial planning rules of thumb should reflect the important role that children play in optimal life-cycle wealth decisions. The "replacement rate" – the percentage of pre-

---

[29] Results for median regression specifications using simulated optimal net worth rather than actual net worth yield very large, highly significant, negative child coefficients in the lowest three income quintiles in the mean regression and the lowest two income quintiles in the median regression.

retirement income that is covered by retirement income – is a workhorse concept in retirement planning. AON Consulting and Georgia State University (2004) issue a widely quoted measure, suggesting that target replacement rates should range between 75 and 86 percent, depending on family income. The Society of Actuaries recommends a replacement rate of 90 percent (LIMRA and the Society of Actuaries, 2003). But, the resources needed to equate the discounted marginal utility of consumption for a husband and wife after four children have left the house are less than what is needed for an otherwise identical childless couple.

Third, children have a large effect on the timing of credit constraints for households with low lifetime income. Without children, the typical household in the lowest (second) lifetime earnings decile would be credit constrained until age 26 (28). With their endogenous fertility choices (which closely match their actual choices), they are credit constrained until 37 (35).

Fourth, once variations in family size are accounted for, means-tested cash and near-cash transfer programs (and the asset tests associated with them) have relatively minor effects on wealth accumulation. This result is striking, given the conclusion of the widely cited Hubbard, Skinner and Zeldes (1995) paper:

> "…the presence of asset-based means testing of welfare program can imply that a significant fraction of the group with lower lifetime income will not accumulate wealth. *The reason* [emphasis added] is that saving and wealth are subject to an implicit tax rate of 100 percent in the event of an earnings downturn or medical expense large enough to cause the household to seek welfare support. This effect is much weaker for those with higher lifetime income…" (p. 393).

We show that children, not asset tests associated with cash and near-cash transfers, largely account for the low levels of wealth accumulation by households with low lifetime income.

28

References

Aon Consulting and Georgia State University. 2004. "The Aon Consulting/Georgia State University 2004 Retirement Income Replacement Ratio Study."

Attanasio, Orazio P. and Martin Browning. 1995. "Consumption Over the Life Cycle and Over the Business Cycle." *The American Economic Review* 85(5) (December): 1118–1137.

Attanasio, Orazio, P., James Banks, Costas Meghir and Guglielmo Weber. 1999. "Humps and Bumps in Lifetime Consumption," Journal of Business and Economic Statistics 17(1): 22-35.

Banks, James, Richard Blundell and Ian Preston, 1994. "Life-cycle expenditure allocations and the consumption costs of children." *European Economic Review*, 1994, vol. 38, pp. 1391–1410.

Barro, Robert and Gary S. Becker, 1989. "Fertility Choice in a Model of Economic Growth," Econometrica 57 (2): 481—501.

Becker, Gary S. and Robert J. Barro. 1988. "A Reformulation of the Economic Theory of Fertility," *Quarterly Journal of Economics*, CII(1), 1-25.

Browning, Martin and Mette Ejrnæs. 2002. "Consumption and Children." http://www.econ.ku.dk/browning/Pdf%20papers/2002/Consumption_and_Children.pdf. August

Citro, Constance F. and Robert T. Michael. 1995. *Measuring Poverty: A New Approach*. Washington, D.C.: National Academy Press.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*, 59 (5): 1221–48.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743-768.

Doepke, Matthias. 2004. "Accounting for Fertility Decline during the Transition to Growth," *Journal of Economic Growth*, 9(3), 347-383, September 2004.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy* 112(2) (April): 397–444.

Engen, Eric M., William G. Gale and Cori R. Uccello. 1999. "The Adequacy of Retirement Saving." *Brookings Papers on Economic Activity* 2: 65–165.

Federal Register (1991), 56, n 34.

Galor, Oded. "From Stagnation to Growth: Unified Growth Theory," *Handbook of Economic Growth*, P. Aghion and S. Durlauf, eds. (Amsterdam: North Holland, 2005).

Gale, William G. and John Karl Scholz. 1994. "Intergenerational Transfers and the Accumulations of Wealth." *Journal of Economic Perspectives* 8(4), Fall, 145-160

Gouveia, Miguel and Robert R. Strauss. 1994. "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis." *National Tax Journal* 47(2) (June): 317–39.

Gouveia, Miguel and Robert R. Strauss. 1999. "Effective Tax Functions for the U.S. Individual Income Tax: 1966–89." *Proceedings of the 92nd Annual Conference on Taxation, National Tax Association:* 155–65.

Green Book. 1983. <u>Background Material and Data on Programs Within the Jurisdiction of the Committee on Ways and Means</u>. Committee on Ways and Means, U.S. House of Representatives, Washington D.C.

Green Book. 1988. <u>Background Material and Data on Programs Within the Jurisdiction of the Committee on Ways and Means</u>. Committee on Ways and Means, U.S. House of Representatives, Washington D.C.

Haveman, Robert, and Barbara Wolfe. 1995. "The Determinants of Children's Attainments: A Review of Methods and Findings," *Journal of Economic Literature*, 33, 1829-1878.

Hubbard, Glen R., Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy*, 103(2) (April): 360–99.

Hurst, Erik and James P. Ziliak. 2006. "Do Welfare Asset Limits Affect Household Saving? Evidence from Welfare Reform." *Journal of Human Resources* 41(1): 46-71.

Jones, Larry and Alice Schoonbroodt. 2009 "Complements Versus Substitutes and Trends in Fertility Choice in Dynastic Models." Mimeo, University of Minnesota.

Kopczuk, Wojciech. 2009. "Economics of Estate Taxation: Review of Theory and Evidence." Mimeo, Columbia University.

Lawrance, Emily C. 1991. "Poverty and the Rate of Time Preference." *Journal of Political Economy* 99(1) (February): 54–77.

Lazear, Edward and Michael, Robert T. 1980. "Family Size and the Distribution of Real Per Capita Income." *American Economic Review* 70, 91-107.

LIMRA International, Inc. and the Society of Actuaries. 2003. "Retirement Planning Software."

Modigliani, Franco and R. Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross Section Data." In <u>Post-Keynesian Economics</u>, K.K. Kurihara (ed.), pp. 388–436. New Brunswick, NJ: Rutgers University Press.

Moffitt, Robert. 2002. "Documentation for Moffitt Welfare Benefits File." Manuscript, February 22, http://www.econ.jhu.edu/People/Moffitt/DataSets.html.

Organization for Economic Cooperation and Development. 1982. The OECD List of Social Indicators.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy,* 114(4), August: 607-643.

Scholz John Karl, Ananth Seshadri. 2008. "Are All Americans Saving 'Optimally' for Retirement?" mimeo prepared for the 10[th] Annual Joint Conference of the Retirement Research Consortium, Washington D.C., August 7-8

Figure 1: Net Worth in 1992 as a Percentage of Summed, Real Lifetime Earnings, By Family Size, Married Couples, HRS Data
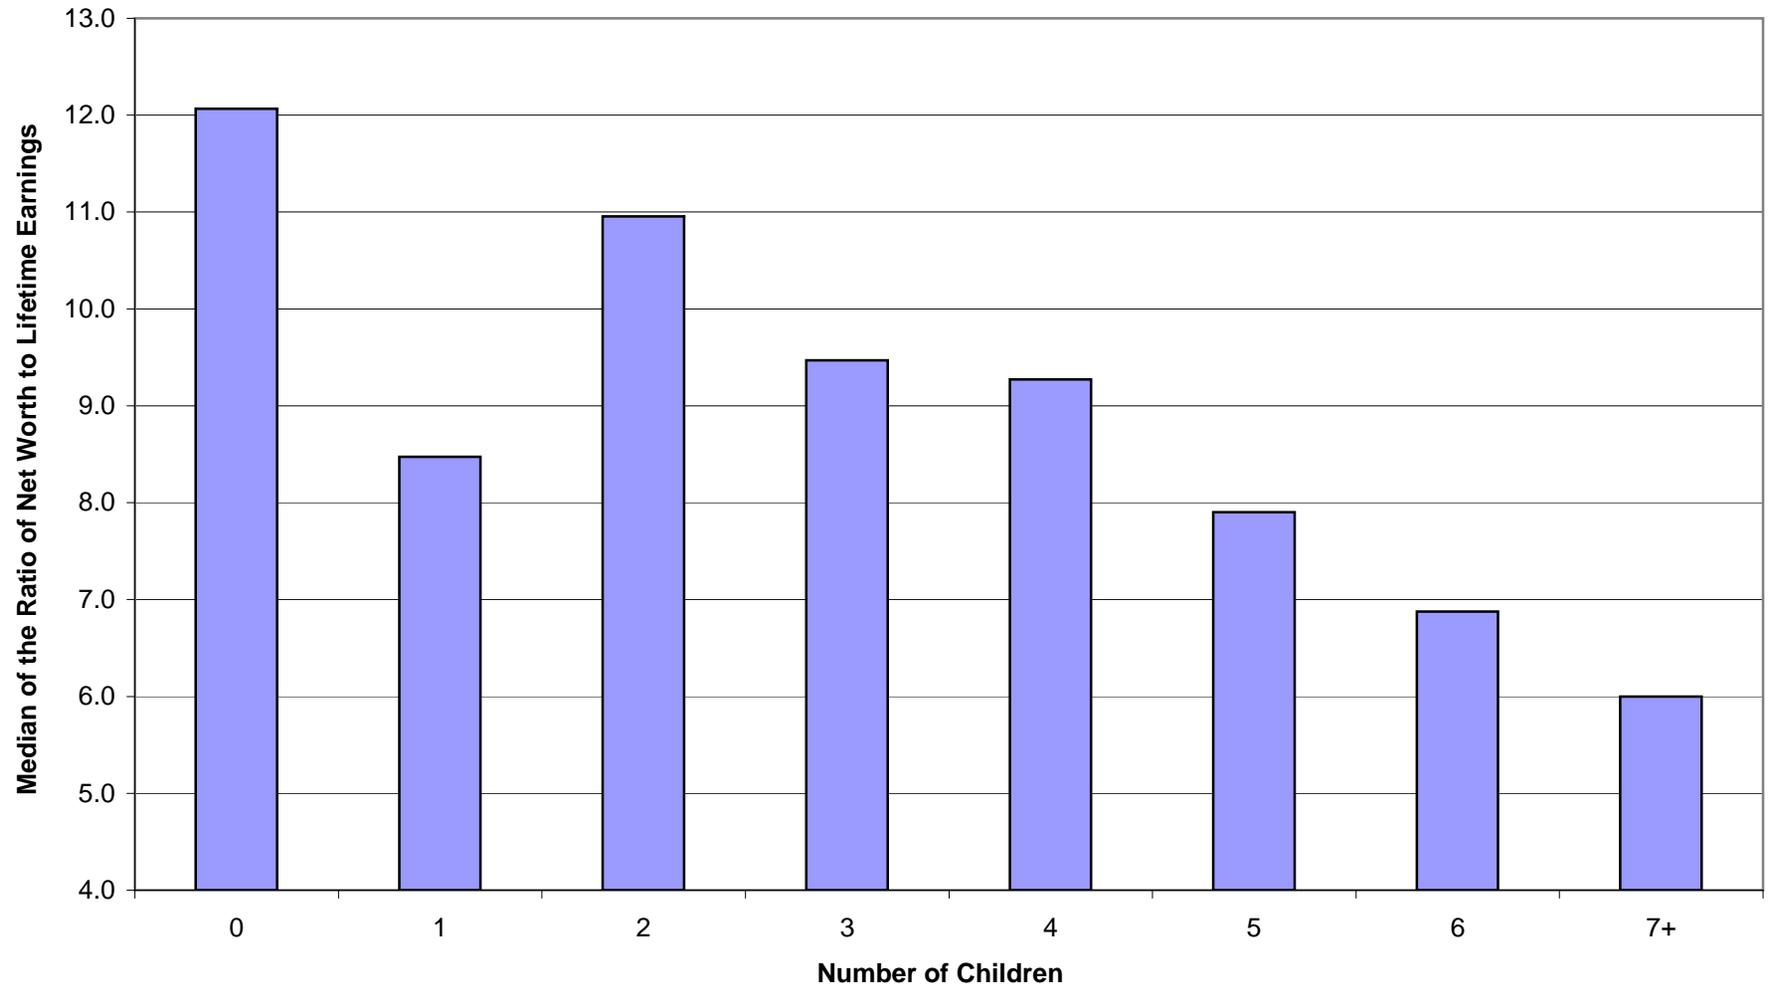
Figure 2: Median Age-Log Earnings Profiles by Family Size, Married Couples
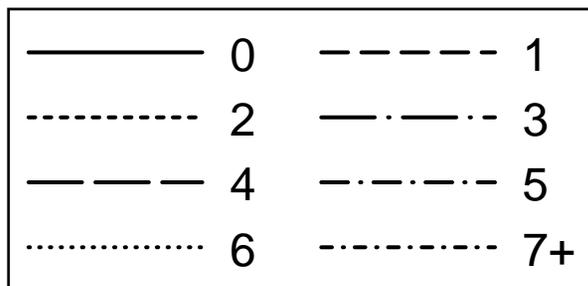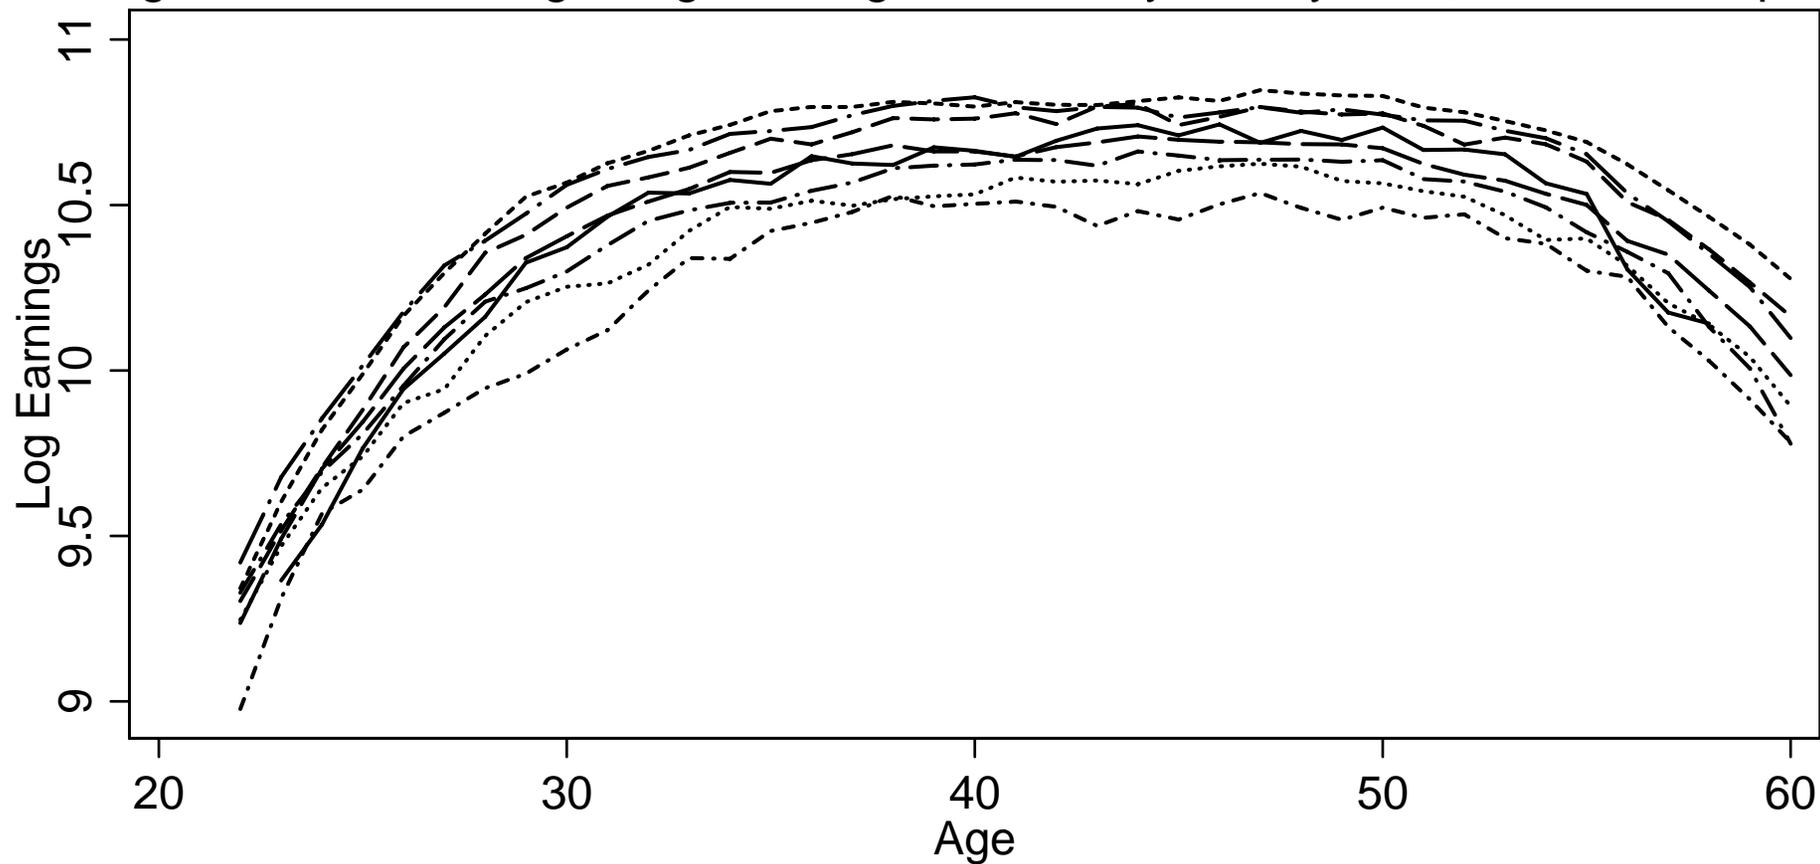
Figure 3: Optimal by Actual Children, Married Couples
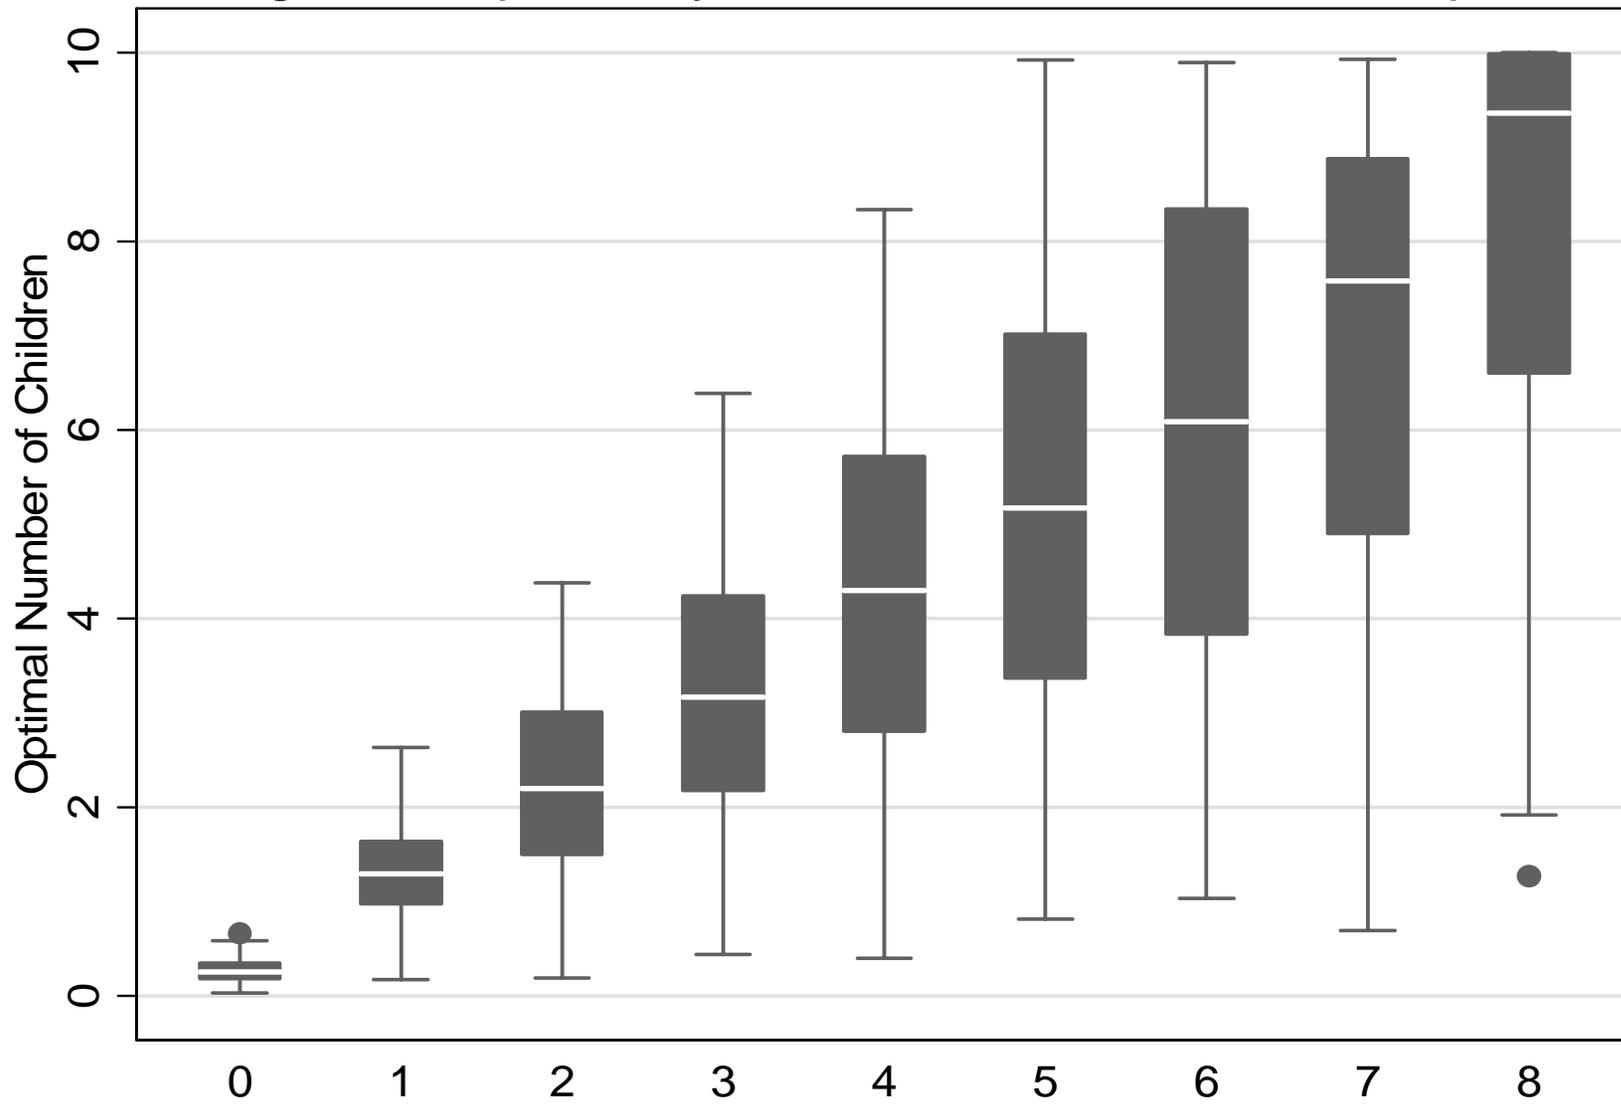
Table 1: Variation in Age of Last Birth, Earnings, and Net Worth by Number of Children for Married Couples, Weighted 1992 HRS Data

Number of Children	Percentage of Total Population	Mean Age When Last Child is Born	Mean %age Earnings After Last Child is Born	Median Net Worth	Mean Net Worth	Mean Undiscounted Lifetime Earnings (1992 dollars)
0	3.1	Not Applicable	Not Applicable	$192,000	$316,952	$1,775,255
1	6.5	29.6	84.2	134,200	282,528	1,728,810
2	24.9	31.0	82.7	182,000	327,299	1,854,470
3	22.4	32.1	79.5	163,155	322,252	1,816,224
4	17.1	33.5	75.5	132,000	259,855	1,644,518
5	9.7	34.7	73.0	118,800	239,207	1,560,737
6	6.3	35.3	71.1	98,580	205,403	1,498,221
7 or more	10.1	37.4	66.7	73,100	182,037	1,343,040
All Married Couples	100.0	32.9	77.4	142,885	280,549	1,696,928

Source: Authors' calculations from Health and Retirement Study data, as described in the text.

Table 2: Variation in Net Worth, Fertility and Earnings by Lifetime Earnings Deciles, Married Couples, 1992 HRS Data, Weighted

			Married Couples		
Lifetime Earnings Decile	**Median 1992 Net Worth**	**Mean 1992 Net Worth**	**Mean Number of Children**	**Mean Age of Head When Last Child is Born**	**Mean %age of Earnings After Last Child is Born**
Lowest	$33,000	$109,006	4.6	35.5	68.6
2	65,400	159,462	4.2	33.7	73.6
3	87,500	165,335	3.8	32.4	77.8
4	105,286	211,178	3.6	32.7	77.5
Middle	126,200	201,292	3.7	32.4	78.0
6	135,000	233,372	3.5	32.5	77.8
7	175,456	290,330	3.3	31.8	79.6
8	203,852	317,147	3.3	32.8	78.6
9	261,000	438,408	3.3	32.5	80.3
Highest	433,800	680,671	3.1	33.1	82.1
All Married Couples	142,885	280,549	3.7	32.9	77.4

Notes: Authors' calculations from 1992 HRS data

Table 3: Actual and Optimal Net Worth and Children, Married Couples, weighted 1992 HRS Data

Lifetime Earnings Decile	Median 1992 Net Worth	Median Optimal Net Worth	Mean 1992 Net Worth	Mean Optimal Net Worth	Mean Number of Children	Mean Optimal Number of Children
Lowest	$33,000	$16,684	$109,006	$69,405	4.6	4.5
2	65,400	43,405	159,462	107,878	4.2	4.2
3	87,500	58,167	165,335	133,239	3.8	4.0
4	105,286	72,621	211,178	150,685	3.6	3.8
Middle	126,200	93,276	201,292	162,949	3.7	3.8
6	135,000	107,122	233,372	197,664	3.5	3.7
7	175,456	127,192	290,330	243,222	3.3	3.5
8	203,852	183,075	317,147	269,594	3.3	3.4
9	261,000	218,950	438,408	361,473	3.3	3.3
Highest	433,800	350,986	680,671	585,261	3.1	3.2
All Married Couples	142,885	115,148	280,549	228,073	3.7	3.7

Source: Authors' calculations from Health and Retirement Study data, as described in the text.

Table 4: The Effects of Eliminating Variation in the Number and Timing of Children, Married Couples, weighted 1992 HRS data

| Lifetime Earnings Decile | Baseline | | No Variation in Kids | | Median Optimal Net Worth, No Variation in Timing |
	Median Net Worth	Credit Constrained Until Age…	Median Net Worth	Credit Constrained Until Age…	
Lowest	$16,684	37	$26,452	26	$27,485
2	43,405	35	64,859	28	65,283
3	58,167	33	74,506	29	73,049
4	72,621	32	80,563	29	81,032
Middle	93,276	30	86,384	30	87,364
6	107,122	29	103,625	31	104,950
7	127,192	28	112,849	31	113,049
8	183,075	30	155,967	32	156,049
9	218,950	31	187,489	33	188,263
Highest	350,986	33	266,379	34	265,830
All Married Couples	115,148	31	109,624	30	108,094

Source: Authors' calculations from Health and Retirement Study data, as described in the text.

Table 5: Effects of the Transfer System on Optimal Median Net Worth and Children, Married Couples, weighted 1992 HRS data

Lifetime Earnings Decile	Optimal Net Worth, Endogenous Fertility	Optimal Fertility	Optimal Net Worth, Endogenous Fertility, No Transfers	Optimal Fertility, No Transfers	Optimal Net Worth, No Children	Optimal Net Worth, No Children, No Transfers
Lowest	$16,684	4.5	$23,125	4.2	$39,586	$89,208
2	43,405	4.2	55,579	4	65,600	118,504
3	58,167	4.0	67,467	3.9	79,372	141,207
4	72,621	3.8	78,476	3.8	96,375	144,589
Middle	93,276	3.8	97,296	3.7	103,282	154,567
6	107,122	3.7	110,048	3.6	138,399	162,354
7	127,192	3.5	128,882	3.5	170,577	198,761
8	183,075	3.4	183,864	3.4	220,391	245,463
9	218,950	3.3	219,463	3.3	265,084	280,945
Highest	350,986	3.2	351,193	3.2	396,706	409,384
All Married Couples	115,148	3.7	119,362	3.7	151,029	183,049

Source: Authors' calculations from Health and Retirement Study data, as described in the text.

Table 6: Correlates of Household Net Worth, Married Couples, 1992 HRS Data

Dependent Variable: Net Worth	OLS	Median Regression	OLS	Median Regression
Constant	-306,180	-182,395	-378,306	-202,598
	(91,741)***	(37,893)***	(96,233)***	(38,021)***
High School Graduate	40,234	13,073	47,544	14,350
	(17,891)**	(4,805)***	(17,938)***	(6,042)**
College Degree	58,987	51,811	57,090	52,002
	(26,668)**	(11,840)***	(26,681)**	(13,521)***
Post-College Degree	129,288	81,925	127,863	73,548
	(29,732)***	(16,800)***	(29,735)***	(19,304)***
Age	7,244	3,590	7,862	3,663
	(1,518)***	(665)***	(1,539)***	(579)***
Has a Defined Benefit Pension	-161,840	-29,466	-153,614	-24,108
	(14,733)***	(5,638)***	(14,847)***	(8,784)***
African-American	-86,256	-37,732	-92,098	-35,270
	(21,259)***	(5,815)***	(21,258)***	(7,677)***
Hispanic	-71,204	-15,204	-86,735	-17,534
	(25,429)***	(7,468)**	(25,660)***	(7,874)**
Number of Children	-6,384	-2,601		
	(2,902)**	(828)***		
Number of Children * Bottom Lifetime Income Quintile			3,753	246
			(4,009)	(808)
Number of Children * 2nd Lifetime Income Quintile			-5,157	-3,876
			(4,376)	(1,006)***
Number of Children * Middle Lifetime Income Quintile			-13,670	-7,514
			(4,292)***	(1,499)***
Number of Children * 4th Lifetime Income Quintile			-20,705	-8,267
			(5,064)***	(3,099)***
Number of Children * Highest Lifetime Income Quintile			-12,902	-4,936
			(6,105)**	(3,232)

All regressions include annual earnings between ages 18 and 65

Observations	4,201	4,201	4,201	4,201
R-squared	0.24		0.24	
Standard errors in parentheses,		* significant at 10%; ** significant at 5%; *** significant at 1%		

Appendix Table 1: Descriptive Statistics for the Married Couple Health and Retirement Study Sample (dollar amounts in 1992 dollars)

Variable	Mean	Median	Standard Deviation
Present Discounted Value of Lifetime Earnings	$2,178,452	$2,059,862	$1,157,896
Defined Benefit Pension Wealth	$133,241	$45,248	$213,095
Social Security Wealth	$135,338	$133,547	$61,818
Net Worth	$208,549	$142,885	$511,460
Mean Age (years)	55.8		5.3
Mean Education (years)	12.9		3.5
Male Has the Highest Income	0.87		0.34
Fraction Black	0.07		0.25
Fraction Hispanic	0.06		0.24
No High School Diploma	0.20		0.40
High School Diploma	0.55		0.50
College Graduate	0.14		0.35
Post-College Education	0.11		0.31
Fraction Self-Employed	0.17		0.38
Fraction Partially or Fully Retired	0.27		0.44

Source: Authors' calculations from the 1992 HRS. The table is weighted by the 1992 HRS household analysis weights.

Appendix

Earnings

   Two issues arise in using earnings information. First, social security earnings records are not available for 22.8 percent of the respondents included in the analysis. Second, the social security earnings records are top-coded (households earn more than the social security taxable wage caps) for 16 percent of earnings observations between 1951 and 1979. From 1980 through 1991 censoring is much less of an issue, because we have access to W-2 earnings records, which are very rarely censored.

   We impute earnings histories for those individuals with missing or top-coded earnings records assuming the individual log-earnings process

$$
\begin{aligned}
y_{i,0}^* &= x_{i,0}'\beta_0 + \varepsilon_{i,0} \\
y_{i,t}^* &= \rho y_{i,t-1}^* + x_{i,t}'\beta + \varepsilon_{i,t}, \ \ t \in \{1,2,...,T\} \\
\varepsilon_{i,t} &= \alpha_i + u_{i,t}
\end{aligned}
\tag{1}
$$

where $y_{i,t}^*$ is the log of latent earnings of the individual i at time t in 1992 dollars, $x_{i,t}$ is the vector of i's characteristics at time t, and the error term $\varepsilon_{i,t}$ includes an individual-specific component α_i, which is constant over time, and an unanticipated white noise component, $u_{i,t}$. We employ random-effect assumptions with homoskedastic errors to estimate equation (1).

   We estimate the model separately for four groups: men without a college degree, men with a college degree, women without a college degree, and women with a college degree. In Scholz, Seshadri, and Khitatrakun (2006) we give details of the empirical earnings model, coefficient estimates from that model, and describe our Gibbs sampling procedure that we use to impute earnings for individuals who refuse to release or who have top-coded social security earnings histories. Our approach is appealing in that it uses information from the entire sequence of individual earnings, including are uncensored W-2 data from 1980-1991, to impute missing and top-coded earnings.

Additional Model Parameters

   Survival Probabilities: These are based on the 1992 life tables of the Centers for Disease Control and Prevention, U.S. Department of Health and Human Services (http://www.cdc.gov/nchs/data/lifetables/life92_2.pdf).

   Taxes: We model an exogenous, time-varying, progressive income tax that takes the form

$$
\tau(y) = a_0\left(y - \left(y^{-a_1} + a_2 \right)^{-1/a_1} \right),
$$

where y is in thousands of dollars. Parameters are estimated by Gouveia and Strauss (1994, 1999), and characterize U.S. effective, average household income taxes between 1966 and 1989.[30] We use the 1966 parameters for years before 1966 and the 1989 parameters for 1990 and 1991.

---

[30]Estimated parameters, for example, in 1989 are $a_0 = 0.258$, $a_1 = 0.768$ and $a_2 = 0.031$. In the framework, $a_1 = -1$ corresponds to a lump sum tax with $\tau(y) = -a_0 a_2$, while when $a_1 \to 0$, the tax system converges to a proportional tax system with $\tau(y) = a_0 y$. For $a_1 > 0$ we have a progressive tax system.

Out of Pocket Medical Expenses: The specification for household medical expense profiles for retired households is given by

$$\log m_t = \beta_0 + \beta_1 AGE_t + \beta_2 AGE_t^2 + u_t,$$

$$u_t = \rho u_{t-1} + \varepsilon_t, \ \varepsilon_t \sim N(0, \sigma_\varepsilon^2),$$

where m_t is the household's out-of-pocket medical expenses at time t (the medical expenses are assumed to be \$1 if the self-report is zero or if the household has not yet retired), AGE_t is age of the household head at time t, u_t is an AR(1) error term and ε_t is white-noise. The parameters to be estimated are β_0, β_1, β_2, ρ, and σ_ε.

We estimate the medical-expense specification for married households without a college degree, and married households with a college degree, using the 1998 and 2000 waves of the HRS, which provide medical expense information on households age 27 to 106.[31] We use the age and education of the head of household in the empirical model. The persistence parameters for medical shocks cluster tightly between 0.84 and 0.86 across groups. The variance of shocks is lower for households with greater education within a given household type (married or single), presumably reflecting higher rates of insurance coverage for households with college degrees relative to other.

---

[31]Older cohorts from the AHEAD and two new cohorts were added to the HRS in 1998, which gives us a broad range of ages that we use to estimate medical expense profiles after retirement. Far fewer members of these cohorts were matched to their social security earnings records, so these observations are less well-suited than the data we use for our primary analysis.