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

September 2007

We thank colleagues at the University of Michigan for developing the Health and Retirement Survey, Surachai Khitatrakun for his advice, and seminar participants at the University of Chicago and the IRP Summer Research Workshop for helpful comments. We also are grateful to the NIA and NICHD, and the Center for the Demography of Health and Aging at UW-Madison for financial support that initiated our work. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA or any agency of the Federal Government.

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 the household at the 90th percentile of the lifetime earnings distribution relative to the earnings of the household at the 10th percentile (referred to as the 90-10 ratio) is 22.5. The 90-10 ratio for 1992 household net worth (including housing wealth) is 525. The coefficient of variation (the standard deviation divided by the mean) of lifetime income is 0.76. The coefficient of variation of net worth is 2.01. 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]

There is a large literature on life-cycle wealth accumulation. But surprisingly few studies examine the effects of children on consumption and wealth.[2] Children might be expected to affect wealth accumulation for at least three reasons. First, family size is correlated with lifetime earnings, so optimal asset accumulation will be correlated with children if wealth accumulation varies with a household's place in the income distribution.[3] Second, the number of children (and adults) in the household affects the utility of a given amount of (private) consumption, which in turn affects optimal consumption decisions. Third, with uncertain earnings (and uncertainty in health and lifespan), the timing of fertility can affect optimal consumption decisions.

This paper focuses on the effects that children have on life-cycle wealth accumulation. We start with a simple permanent income model with no uncertainty and complete markets to build intuition about the effects of children. But this framework does not come close to matching the

[1] A recent study documenting this fact is Dynan, Skinner, and Zeldes (2004).
[2] Browning (1992) is a notable exception, as is Attanasio and Browning (1995) and Browning and Ejrnæs (2002). We briefly discuss the latter two papers later.
[3] A common feature of many important papers on life-cycle wealth accumulation is to ask, given an earnings distribution, what is the implied distribution of wealth (see, for example, Modigliani and Brumberg, 1954; Deaton,

distribution of existing wealth. So we then look at the effects of children in the augmented life-cycle model discussed in Scholz, Seshadri, and Khitatrakun (2006). But both approaches take the arrival and timing of children as being exogenous: because fertility may be affected by wealth and earnings expectations, we also describe results from a model that incorporates endogenous fertility in the spirit of Barro and Becker (1988). Our conclusions about the importance of children in understanding wealth accumulation are consistent across modeling approaches.

We find that children have a large effect on household's net worth and consequently are an important factor in understanding the wealth distribution. We show, for example, that the effects of children are much larger than the effects of asset tests associated with cash and near-cash transfers, given earnings realizations and the social security system experienced by households in the HRS. This result is striking, given a conclusion of Hubbard, Skinner and Zeldes (1995) who write:

> "…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 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 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 time since they have

1991; Aiyagari, 1991; Hubbard, Skinner, and Zeldes, 1995; as well as more recent work, such as De Nardi, 2004).

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.

In the next section we describe our data and present descriptive statistics from the HRS about the number of children across income deciles, the timing of fertility across families, and the age-earnings profiles of households with different numbers of children. Section 2 briefly discusses children in a life-cycle model with no uncertainty. Since most expenses on children are borne by parents prior to retirement, families with children would be expected to have lower retirement wealth, all else being equal, than families without. But the life-cycle model with no uncertainty does not reflect the importance of precautionary saving and credit constraints on wealth accumulation. In sections 3 and 4 we present two additional models that more closely match features of the economy, we describe our policy experiments, and we present our results. Section 5 briefly discusses descriptive, reduced form regressions from the HRS motivated by our analytic work. The paper concludes with a discussion of other related considerations.

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

The HRS is a national panel study with an initial sample (in 1992) of 12,652 persons in 7,702 households. It oversamples blacks, Hispanics, and residents of Florida. The baseline 1992 study consisted of in-home, face-to-face interviews of the 1931-1941 birth cohort and their spouses, if they are married. Follow-up interviews were given by telephone in 1994, 1996, 1998, 2000, 2002, and 2004. For the analyses in this paper we exclude 379 married households where one spouse did not participate in the 1992 HRS, 93 households that failed to have at least one

---

Our paper also takes the earnings distribution as being exogenous.

year of full-time work, and 908 households where the highest earner began working full time prior to 1951.[4] Our resulting sample has 10,523 respondents in 6,322 households.

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.

I.1. Children in the HRS

There are strong correlations in the HRS between children, factors that likely influence wealth accumulation, and wealth itself. In Table 1 we summarize some characteristics of the HRS population by the number of children they have. Column 1 shows the modal number of children for the sample is two, but 31.8 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

---

[4] 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.

[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

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 two children. It falls monotonically with the number of children above 2. If we simply calculate the *percentage* of mean net worth given in the second-to-last column of Table 1 to lifetime earnings (the last column), it is larger (20.6 percent) for families with no children than it is for families with any positive number of children. For families with children the net-worth-to-lifetime-earnings percentage has a concave shape, starting at 16.5 percent for one-child families, peaking at 18.0 for three-child families, and falling to 13.3 percent for families with seven or more children. These figures provide suggestive evidence that net worth is not fully determined by lifetime earnings and 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 and marital status. The first two columns show median and mean net worth, the variable of central interest to this paper. It is clear 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. Similar patterns hold for single households (in 1992).[7] There is little systematic relationship between the age of completed fertility and lifetime income, despite the fact that the number of children is

---

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 and saving adequacy.

[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). Appendix 1 provides a bit more detail and the on-line appendix of Scholz, Seshadri, and Khitatrakun (2006) provides complete details of our approach.

[7] Single and married households are categorized based on their status in 1992.

negatively correlated with lifetime income. This suggests that higher income HRS households may be 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 HRS households.[8] There appears to be a small amount of spreading of the earnings trajectories, but in general, the slopes of the profiles look similar. Childless individuals and/or couples clearly have the lowest incomes over their lifetimes. Households with 2 and 3 children have the highest and most steeply sloped age-earnings profiles. The profiles flatten and are lower as the number of children increases beyond 3.

The descriptive data are consistent with at least three channels through which children may influence wealth. First, as is clear from Figure 2, family size is correlated with lifetime earnings.[9] 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

[8] Specifically, we plot a median log earnings using Stata's "graph twoway mbands" command.
[9] The same qualitative patterns hold for versions of Figure 2 that are restricted to married couples, to single households, or to households who have never changed marital status (or partners) given their 1992 status.

household earnings is $1,718,932 ($1,541,555).[10] Retirement consumption will be financed out of defined benefit pension wealth (mean is $106,041, median is $17,327);[11] social security wealth (mean is $107,577, median is $97,726);[12] and nonpension net worth (mean is $225,928, median is $102,600). The mean age of the household head is 55.7.[13]

II. Children and Wealth in a Life-Cycle Model with no Uncertainty

We briefly start providing intuition about the effect of children on household wealth using a simple Modigliani and Brumberg (1954) permanent income model, allowing family size to vary exogenously across the life-cycle. Assume the household solves

$$\max \sum_{j=0}^{T} \beta^{j} N_j U\left(c_j / N_j\right) \quad \text{subject to} \quad \sum_{i=0}^{T} \frac{c_j}{(1+r)^j} = \sum_{i=0}^{T} \frac{y_j}{(1+r)^j}$$

where c_j denotes consumption, y_j stands for earnings, β is the pure rate of time preference (generally thought to be less than one), r is the real interest rate, and N_j adjusts the utility value

[10]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 assume that the real interest rate is 2.21%, consistent with the 6.3 percent interest rates and 4 percent inflation assumed under the intermediate scenarios of the Pension Present Value Database.

[11] 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).

[12] We use 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).

[13]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.

of consumption for the number of children and adults in the household.[14] If preferences are CRRA with $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$, the Euler equation is given by $\left(\frac{c_j}{N_j}\right)^{-\gamma} = \left[\beta(1+r)\right] \left(\frac{c_{j+1}}{N_{j+1}}\right)^{-\gamma}$ and the marginal utility of household consumption (c_j) is equal across periods. The optimal solution is given by $c_j = \left(\frac{N_j}{\sum_{j=0}^{T} \frac{N_j \left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j}} \right) \left(\sum_{j=0}^{T} \frac{y_j}{(1+r)^j} \left[\beta(1+r)\right]^{j/\gamma} \right)$.

The first term (enclosed in parentheses) adjusts period *j* consumption for the number of adults and children in the household. The second term (enclosed in parentheses) simply denotes discounted lifetime earnings. When family size is large, the household consumes more, so, all else equal, a larger family size reduces the household's resources available for retirement.[15] Thus, in the life-cycle model with no uncertainty and perfect capital markets, 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

---

[14] We multiply utility by N_j so the marginal utility of consumption is equal across families of different sizes.

[15] The partial derivative of consumption with respect to family size is given by

$$\frac{\partial c_j}{\partial N_j} = \left(\sum_{j=0}^{T} \frac{y_j}{(1+r)^j} \left[\beta(1+r)\right]^{j/\gamma} \right) \left(\frac{1}{\sum_{j=0}^{T} \frac{N_j \left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j}} - \frac{N_j}{\left(\sum_{j=0}^{T} \frac{N_j \left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j} \right)^2} \frac{\left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j} \right),$$

which reduces to

$$\frac{\partial c_j}{\partial N_j} = \frac{\left(\sum_{j=0}^{T} \frac{y_j}{(1+r)^j} \left[\beta(1+r)\right]^{j/\gamma} \right)}{\sum_{j=0}^{T} \frac{N_j \left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j}} \left(1 - \frac{\frac{N_j \left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j}}{\left(\sum_{j=0}^{T} \frac{N_j \left[\beta(1+r)\right]^{j/\gamma}}{(1+r)^j} \right)} \right) > 0.$$

lifetime resources when children are present relative to families with fewer children (all else being equal).

If there is systematic variation between family size and lifetime earnings, Euler equations estimated from the life-cycle model that fail to account for family size will overstate the variation in discount factors needed to rationalize household's consumption choices. Indeed this is the basis for Lawrance (1991), who concludes that accounting for variation in family composition reduces the heterogeneity in discount factors estimated from a consumption Euler equation. Nevertheless, she finds that the remaining variation in discount factors is systematic – high earners are more patient.

Attanasio and Browning (1995) show that once one accounts for the variation in family size over the life-cycle, a flat age-consumption profile – consistent with the life-cycle model – 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. In the context of the simple framework described in the previous section, however, family size variation alone cannot explain the level and skewness of wealth.[16] Thus, we explore the interaction between precautionary motives and variation in family size to better understand the distribution of wealth.

The next section describes calculations from a life-cycle model with borrowing constraints and idiosyncratic shocks, where family size and the timing of births varies based on data from

---

[16] For example, in the life-cycle model above (with $\beta = 0.97$, $\gamma = 3$ and $r = 0.03$) and where households have their observed earnings realizations, married households in the bottom decile optimally choose to have zero assets when we observe them in the data (the average age is 56.5), while households in the top decile have \$66,382. This is accounted for by two key factors. First married households at the bottom decile have 4.6 kids while those in the top decile have 3.1 kids. Second, the ratio of resources available at retirement (social security wealth and defined benefit wealth) to lifetime earnings is about 25 percent for the bottom decile and only 10 percent for the top decile, thereby leading the richer households to want to transfer more resources towards retirement.

the HRS. We show how variation in the number and timing of children affect household wealth. As will become clear, a key mechanism is that since larger households have children attached with them for longer, on average, than their counterparts with fewer children, they will be borrowing constrained for a longer period of time. All else equal, this reduces the optimal wealth at retirement. Indeed, in what follows, we find the quantitative effect of this phenomenon is large.

III. A Model of Optimal Wealth Accumulation

We solve a simple life-cycle model, augmented to incorporate uncertain lifetimes, uninsurable earnings, uninsurable medical expenses, and borrowing constraints. A household derives utility $U(c)$ from period-by-period consumption in equivalent units, where $g(A_j, K_j)$ is a function that adjusts consumption for the number of adults A_j and children K_j in the household at age j.[17] Let c_j and a_j represent consumption and assets at age j. 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 $j-1 < R$. At age D, $p_D = 0$. The discount factor on future utilities is β. Expected lifetime utility is then

$$E\left[\sum_{j=S}^{D} \beta^{j-S} g(A_j, K_j)\, U\left(c_j / g(A_j, K_j)\right)\right].$$

The expectation operator E denotes the expectation over future earnings uncertainty, uncertainty in health expenditures, and uncertainty over life span.

[17]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 if a spouse dies. Similarly, single households in 1992 are modeled as making their lifecycle consumption decisions as if they were single throughout their working lives. They are assumed to remain single until death.

Consumption and assets are chosen to maximize expected utility subject to the constraints, [18]

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

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

$$c_j + a_{j+1} = y_j + a_j - \tau(e_j + ra_j), \; j \in \{S, ..., R\},$$

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

The first two equations define taxable income for working and for retired households.[19] The last two equations show the evolution of resources available for consumption. In these constraints e_j denotes labor earnings at age j. $SS(\cdot)$ are social security benefits, which are a function of 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. Medical expenditures are denoted by m_j and the interest rate is denoted by r.[20] The tax function $\tau(\cdot)$ depicts total tax payments as a function of earned and capital income

[18]The economic environment implies a borrowing constraint in the sense that asset balances are non-negative in every period.

[19]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.

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

for working households, and as a function of pension and capital income plus a portion of social security benefits for retired households.[21]

We simplify the problem by assuming households incur no out-of-pocket medical expenses prior to retirement and face no pre-retirement mortality risk. Therefore, the dynamic programming problem for working households 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.

III.1. Model Parameterization

We briefly discuss several key modeling decisions. Details for survival probabilities, the tax function, and medical expenses are given in Appendix 2. Further discussion and sensitivity analyses are given in Scholz, Seshadri, and Khitatrakun (2006).

We use constant relative risk-averse preferences, so $U(c)=\left\{\dfrac{c^{1-\gamma}}{1-\gamma},\right.$ when $\gamma \neq 1$. We set the discount factor as $\beta = 0.96$ and the coefficient of relative risk aversion (the reciprocal of the intertemporal elasticity of substitution) to $\gamma = 3$. We assume an annualized real rate of return of 4 percent.

Our equivalence scale comes from Citro and Michael (1995) and takes the form $g(A_j, K_j) = (A_j + 0.7K_j)^{0.7}$, where A_j indicates the number of adults (children) in the household and K_j indicates the number of children in the household. This scale implies that a two parent

[21]Specifically, taxable social security benefits for single taxpayers are calculated from the expression $\max(0, \min(0.5 * SS\ Benefits,\ Income - 0.5 * SS\ Benefits - 25{,}000))$. Taxable benefits for married couples are calculated similarly, but replacing 25,000 with 32,000. This approach approximates the law in effect in 1992.

family with 3 children consumes 66 percent more than a two parent family with no children. 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 corresponding numbers for these equivalence scales is 88 percent, 76 percent and 59 percent. Our scale lies in between these values.

One of the purposes of the paper is to contrast the effects of children on wealth with the effects of asset-tested transfer payments. To do this we model the benefits from public income transfer programs using a specification suggested by Hubbard, Skinner and Zeldes (1995). The transfer that a household receives while working is given by

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

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).[22] 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. We assume through this formulation that earnings, retirement income, and assets reduce public benefits dollar for dollar.

---

[22]The $\underline{c}$ in the model reflects the consumption floor that is the result of all transfers (including, for example, SSI). Moffitt (2002, http://www.econ.jhu.edu/People/Moffitt/DataSets.html) provides a consistent series for average benefits received by a family of four. To proxy for the effects of all transfer programs we use his “modified real benefit sum” variable, which roughly accounts for the cash value of food stamp, AFDC, and Medicaid guarantees. We weight state-level benefits by population to calculate an average national income floor. We use 1960 values for years prior to 1960 and use the equivalence scale described above to adjust benefits for families with different configurations of adults and children. We confirm that the equivalence scale adjustments closely match average

We aggregate individual earnings histories into household earnings histories. Earnings expectations are a central influence on life-cycle consumption decisions, both directly and through their effects on expected pension and social security benefits. 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 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 six groups according to marital status, education, and number of earners in the household, giving us six sets of household-group-specific parameters.[23] Estimates of the persistence parameters range from 0.58 for single 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 for married households with either one or two earners and in which the highest earner has at least a college degree, to 0.21 for single households without college degrees (Scholz, Seshadri, and Khitatrakun, 2006, give more details).

---

benefit patterns for families with different numbers of adults and children using data from the Green Book (1983, pp. 259–260, 301–302; 1988, pp. 410–412, 789).

[23]The six groups are (1) single without a college degree; (2) single with a college degree or more; (3) married, head without a college degree, one earner; (4) married, head without a college degree, two earners; (5) married, head with a college degree, one earner; and (6) married, head with a college degree, two earners. A respondent is an earner if his or her lifetime earnings are positive and contribute at least 20 percent of the lifetime earnings of the household.

III.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) 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 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: for example, birth years of children affect the scale economies of a household at any given age (as determined by the equivalence scale). Consequently, it is not sufficient to solve the life-cycle problem for just a few household types.

III.3. Policy Experiments and Results

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 wealth. Using the model and household data, we can incorporate the specific variation in both the number and timing of kids that we see in the HRS. It also allows us to conduct counterfactual policy experiments where we can alter features of the economic environment to better understand the effect that children have on wealth accumulation.

The baseline results presented in Table 3 are discussed in Scholz, Seshadri and Khitatrakun (2006). Here we discuss other features of the results. The model generates a distribution of optimal wealth that matches (in fact it slightly exceeds) the skewness of the actual wealth distribution (so, for example, we do not need to rely on bequest motives to replicate the distribution of wealth). The 90-25 ratio of unweighted net worth in the data is 20.7. In the simulated optimal wealth data it is 28.5 (we cannot compute the 90-10 ratio, since optimal wealth in the 10^{th} decile is $0). The coefficient of variation in the actual data is 2.1, in the simulated data it is 2.4. For ease of exposition when discussing our remaining results, we present data on median wealth that arises in the counterfactual environment with median optimal wealth in the baseline model. The qualitative results and conclusions are the same when using mean wealth levels as the benchmark (details are available on request).

The model also captures the gradient in median wealth by number of children – simulated optimal net worth increases as the number of children increases from 0 to 2 and declines monotonically thereafter – mirroring the pattern seen in the data. This pattern is partly a reflection of the earnings profiles shown in Figure 2, where lifetime earnings increase across households with 0, 1, and 2 children and then falls for households with more than 2 children. But as described below, the pattern is also a consequence of interactions between consumption, children, and wealth accumulation.

III.3.1.The Effect of the Number and Timing of Children

Our first experiment highlights the effects that heterogeneity in both the number and timing of children has on wealth. We assign each married couple the mean number of children (for all married couples), assuming they are born at the median age of married couples that have four children. Specifically, married couples are assumed to have 3.6 children, born at ages (of the head of household) of 23, 26, 29, and the 0.6 child at age 33. Similarly, all single households have 2.8 children, born at the ages of 23, 26, and the 0.8 child at age 29. 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.

As can be seen from Table 4, the effect of altering the timing and number of children is substantial. When the lowest income decile households have 3.6 children at the timing of the median 4-child household instead of 4.6 children at different times in the lifecycle, median optimal net worth increases from $1,350 to $16,403.[24] Children have two related effects. First, by having fewer children in the counterfactual simulations than they do in the data, child-oriented expenditures (and aggregate expenditures) are smaller and their retirement wealth is larger than it would be if they had more children. Second, children affect the length of time households will be credit constrained. The second and fourth columns of Table 4 report the ages at which the median household in each lifetime income decile is credit constrained in the baseline economy and in the counterfactual world where there is no variation in the number (and timing) of children. In the baseline economy, the median household in the lowest lifetime income decile is credit constrained until age 34. This figure drops to age 26 when there is no variation in the number of children. The timing and number of children has a substantial effect on when the household begins saving for retirement.

The systematic variation of kids by lifetime income can be thought of as increasing the dispersion in earnings. Low lifetime income households have, on average, more children than do high lifetime income households. Therefore, the effective income available to the household after adjusting for family size (through the equivalence scale) falls by more for low-income households than it does for high-income households. Thus, fertility differences make the resources available for consumption even more dispersed than the distribution of earnings.

[24] Mean optimal asset holdings increase to $63,472 from $38,537.

Hence 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 2.4 in the baseline optimal net worth distribution to 1.7 when the variation in children is shut down.

III.3.2.The Effect of the Timing of Children

To study the effect of the timing of children, we allow each household to have the number of children that it actually has, but assume that all families with one child have the child at age 29 (the median age of birth for one-child families), all two-child families have their children at ages 26 and 30, and so on.[25]

Timing should matter for the following reason. Since children's consumption depends on their parent's consumption and since income increases with age, having children later on in life will mean more expenditures on children. Families that have children later in their life-cycle will, all else equal, have fewer resources at retirement. Thus, shutting down this variation, should lead to a smaller dispersion in wealth.

The results of eliminating variation in the timing of children are shown in the second column of Table 5. When there is no variation in the timing of births, wealth doubles in the lowest decile and increases 40 percent in the second lifetime income decile. While these percentage changes seem substantial, the dollar changes are much smaller than the combined effect of altering both the number and timing of children. Therefore, the bulk of the variation in wealth is caused by the variation in family size across households with different lifetime incomes.

[25] Ages for 3-child families are 24, 27, 31; 4-child families are 23, 26, 29, 33; 5-child families are 22, 25, 27, 30, 34; 6-child families are 22, 24, 26, 28, 31, 35; and 7-child families are 22, 24, 26, 28, 30, 32, 36.

III.3.3. The Effect of Heterogeneity in Earnings

It is difficult to assess exactly how large the effects of children are on wealth accumulation absent alternative counterfactual reference points. In this subsection we perform an experiment where we shut down household heterogeneity in earnings processes. Recall that we assume that earnings processes have a household specific component that governs the slope of the earnings profile. More educated households, for instance, have higher intercepts and steeper slopes of their expected age-earnings profiles than do less well educated households. Earnings expectations, of course, affect wealth accumulation.

Shutting down this source of heterogeneity would lead every household to draw its earnings shocks from a distribution with the same slope of the earnings profile. For instance, a college graduate would (incorrectly) assume that she would experience the same growth rate in earnings as would a high school graduate. This would lead the college graduate to accumulate less wealth (in all states of the world relative to the case in which she has a higher alpha) since the graduate would expect a lower future income. To be clear, a graduate who is now assigned a smaller slope coefficient is pleasantly surprised (on average) when she receives her earnings draws. However, she correctly recognizes that the persistence of the shock is high. Relative to the 'truth' wherein the slope is higher, her expectation of future income is lower. The lower expectation of future income leads her to accumulate less for retirement than in the case in which she is assigned a higher alpha to begin. Table 5 reports the results. Notice that the dispersion in wealth is smaller than in the baseline case. The noteworthy feature of the results is that the effect of earnings heterogeneity is about the same order of magnitude as that of heterogeneity in children.

III.3.4.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. Recall, $\underline{c}$ denotes the generosity of the transfer program. To study their effects we set $\underline{c}$ to zero, but assume that there exists a governmental program that insures individuals against out of pocket medical shocks. Thus our experiment effectively eliminates cash and near-cash transfers. The last column of Table 5 reports the results. Cash and near-cash transfer programs have very little effect on asset accumulation – the median net worth in the lowest decile increases from $1,350 to $1,483 when the consumption floor is set to zero.[26]

The structure, benefits, and receipt of transfers modeled in Hubbard, Skinner, and Zeldes and in our paper are very similar. 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.[27] In 1980, when the average HRS respondent was 44 years old, 25.3 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 (0.6 percent in our model, 2.3 percent in the PSID). A similar close correspondence holds across education groups for households in 1990.

The negligible effect of the transfer program arises due to two key differences in our work relative to Hubbard, Skinner and Zeldes. First, poorer households by virtue of their larger family

[26] In a recent careful study, Hurst and Ziliak (2006) find little effect on wealth accumulation from state-level changes in asset tests associated with the 1996 welfare reform.
[27] Our floor, of course, varies by year and by family composition.

size, optimally plan on having fewer resources for retirement, when their children will have left the household. Second, these households are credit constrained for a longer period of time and hence begin asset accumulation later on in life. This depresses wealth accumulation.

To summarize, the presence of children in the household (along with upward sloping age-earnings profiles) implies that low-income households are credit constrained while young, so they have little reason to save to smooth the discounted marginal utility of pre-retirement consumption. We also find that 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 the 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 very little disincentive effect of the transfer program on (already negligible) private asset accumulation. It is noteworthy that this holds despite the similarities in the way in which Hubbard, Skinner and Zeldes and we model the social security system.

While we focus on wealth in 1992 when the average household is 55.7 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 34, a substantially older age than for high-income households. Absent the 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 our view of the world, having 5 children (or the number of children observed in the HRS) alters optimal consumption choices sufficiently strongly to fully reconcile the low wealth holdings of the very poor with the data.

While we find small effects of the transfer program on wealth accumulation, our model implies a larger effect of transfer 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 the effect on asset accumulation is negligible. Simply put, our analysis implies that poor households have few assets in part 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, thereby leading to a smaller overall effect on consumption.

IV. A Model with Endogenous Fertility

To this point, we have not offered any theory of why families have children. Clearly, wealth and earnings expectations affect decisions about the number and timing of children, so endogenizing fertility is necessary to examine the robustness of the previous results. When writing a model of endogenous fertility we want to account for the joint distribution of wealth and fertility – a much more stringent test than simply matching wealth. To do this, we follow the pioneering work of Becker and Barro (1988) and assume that parents get utility from the quantity and the quality of their children. We do not model the timing of children, and instead assume that parents give birth to all their children at age B > S. Children are then in the household for 18 years. Parental preferences are given by

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

Specifically, parents care about the number of children, *f*, and utility per child, $U(c_j^k)$, while the child lives in the household. The function *b(f)* denotes the weight that parents place on quantity. Following Barro and Becker, we assume that *b(f)* is increasing and concave. We also assume that children entail a cost. The budget constraint during the period of time when the kids are attached to parents is given by

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

where

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

Notice that the only change from the previous model is that each child requires the fraction κ of the parent's earnings, over and above direct consumption needs. This captures the indirect time costs associated with bearing and rearing children. The presence of this fixed cost will imply that higher e_j households will have fewer children than their lower e_j counterparts.

The decision problem now entails two more choice variables – the fertility rate, f, and consumption per child, c_j^k. The first order conditions with respect to c_j^k and f are given by

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

and

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

In the above equation, $V_{B+1}(\bullet)$ stands for the value function at age *B+1*. We continue to assume that the utility function is of the CRRA variety but assume that $U(c) = \left\{ \frac{c^{1-\gamma}}{1-\gamma},\ 0<\gamma<1. \right.$ The restriction that γ lies between 0 and 1 is designed to ensure that utility is always a positive number.[28] We assume that the discount factor function is given by $b(f) = b_0 f^{b_1}, 0 < b_1 < 1.$

Our model introduces four new parameters: b_0, b_1, γ 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 γ.

Given the functional forms, the first order conditions 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 make sure that the structure of preferences is similar to the structure that we used in the exogenous fertility version of the model, the condition

$$\left(1 + b_0^{1/\gamma} f^{1+b_1/\gamma}\right) = 2^{0.7\gamma}(1 + 0.7f)^{0.7\gamma}$$

must be satisfied. This condition ensures that the equivalence scale is (approximately) an equilibrium implication of the endogenous fertility model. This condition together with the requirement that we match the fertility rate for the median family pins down the remaining parameters.[29] The resulting parameter values are $b_0 = 0.66$, $b_1 = 0.57$ and $\gamma = 0.61$. These parameters lie within the range of values in the fertility literature (see for instance Doepke, 2004).

With the parameters specified, we are now in a position to examine the implications of the endogenous fertility model. As in the baseline case, we solve the model household-by-household and report the predictions by deciles of lifetime income in Table 6.[30] The model matches the wide variation in fertility rates strikingly well. The predictions for wealth (in levels) are also shown in Table 6 – again the model is reasonably successful in accounting for the wide variation in wealth.

We are now in a position to analyze the effects of transfer programs on fertility and wealth. To do this we again shut down cash and near-cash transfer programs. In Table 6 we see that the fertility rate of the poorest households decrease slightly and the simulated optimal net worth increases slightly. What happens is that as transfer incomes are eliminated, the household cuts down on the number of children. The reduction in income resulting from eliminating transfer programs leads the household to cut back on consumption, children's consumption and the number of children. The corresponding increase in wealth is due to two reasons. First, holding fertility fixed, the household wants to increase its wealth to provide insurance, which it previously received through the transfer program. This effect arises in the exogenous fertility model, but it is small. Second, eliminating the transfer program reduces the fertility rate of program recipients. 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.

[28] An alternative is to keep the value of γ at 3 but to add a constant B to the utility function. The constant needs to be high enough to ensure that utility is always positive. This approach yields very similar quantitative results.

We draw two conclusions from the endogenous fertility model. First, the endogenous fertility model is able to capture the joint distribution of fertility and wealth strikingly well. Second, the transfer program does not have a large effect on wealth accumulation even after accounting for an economic explanation for why families have children.

V. 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. Table 7 presents one reduced form median regression specification showing the correlation of children and net worth.[31] The sample is restricted to married couples and excludes the self-employed and includes the combined annual earnings of both partners between the ages 22 and the age of the household head in 1991. It is clear, relative to the baseline childless household, that families with children have less net worth. The patterns by parity are uneven: like Figure 1, the gradient is generally declining with the number of children for families with 2 to 7 children, though the differences across adjacent parities are generally not statistically significant. The sharpest difference appears for one-child families.

Instead of the Becker-Barro 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 for that in the United States. Gale and Scholz (1994) review the older evidence showing relatively minor transfer flows from

---

[29] By construction, we match the fertility rate of only the median household.

[30] In Table 6 we restrict the sample to only married couples.

[31] The OLS estimates are -$68,624 for one-child families, -$74,259 for 2-child families, -$93,105 for 3-child families, -$72,553 for 4-child families, -$99,810 for 5-child families, -$103,418 for 6-child families, and -$111,705 for families with 7 or more children.

children to parents in data from the Survey of Consumer Finances. The HRS data corroborate these results.

VI. Conclusions

A large number of potential explanations for wealth dispersion have been proposed in the literature. Some argue that the life-cycle model with uncertainty must be augmented with a bequest motive to match the observed skewness of the wealth distribution. Others suggest that the poor have higher discount rates than richer households. While there has been a lot of attention paid to variation in discount factors and bequest motives, we find it surprising that very little attention has been paid to examining the effect that children have on wealth accumulation.

We study the effect of children in the context of a life-cycle model with uninsurable income risks and borrowing constraints, as well as a variant of the same model with endogenous fertility. Our study yields two conclusions – first, the variation in family size plays a very important role in understanding the wide dispersion in wealth. Second, once variations in family size are accounted for, means-tested cash and near-cash transfer programs have very little effect on wealth accumulation.

References

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

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. 1992. "Children and Household Economic Behavior." *Journal of Economic Literature* 30(3) (September): 1434–1475.

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.

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.

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. Background Material and Data on Programs Within the Jurisdiction of the Committee on Ways and Means. Committee on Ways and Means, U.S. House of Representatives, Washington D.C.

Green Book. 1988. Background Material and Data on Programs Within the Jurisdiction of the Committee on Ways and Means. 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. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy* 40: 59–125.

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.

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.

McGrattan, Ellen R. and Edward C. Prescott. 2003. "Average Debt and Equity Returns: Puzzling?" *American Economic Review Papers and Proceedings*. 93(2) (May): 392–397.

Modigliani, Franco and R. Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross Section Data." In Post-Keynesian Economics, 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.

Table 1: Variation in Age of Last Birth, Earnings, and Net Worth by Number of Children, Weighted 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	9.2	Not Applicable	Not Applicable	$87,471	$213,720	$1,038,604
1	8.8	28.9	83.8	95,500	206,911	1,255,927
2	24.9	30.5	82.6	144,500	272,842	1,528,333
3	22.5	31.7	79.4	127,000	265,723	1,472,372
4	9.3	33.1	76.0	98,000	207,173	1,335,529
5	9.0	33.9	73.4	93,269	199,081	1,272,693
6	6.4	35.0	70.6	62,000	161,280	1,172,262
7 or more	9.9	37.1	67.6	51,923	143,543	1,083,206
Full Population	100.0	32.4	77.7	102,600	225,928	1,338,754

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, Weighted

Married Couples

Lifetime Earnings Decile /1	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	$35,450	$111,991	4.6	35.3	69.1
2	65,600	166,974	4.1	33.4	74.2
3	90,962	171,847	3.9	32.7	77.3
4	114,000	199,800	3.5	32.5	77.9
Middle	124,348	238,961	3.7	32.3	78.2
6	136,672	214,699	3.6	32.4	78.3
7	184,000	286,538	3.3	32.1	79.0
8	206,253	330,984	3.3	32.7	79.0
9	266,800	451,280	3.3	32.4	80.3
Highest	433,326	687,277	3.1	33.3	82.1
All Married Couples	142,885	280,549	3.7	32.9	77.4

Singles

Lifetime Earnings Decile /1	Median 1992 Net Worth	Mean 1992 Net Worth	Mean Number of Children	Mean Age of Head When Last Child is Born /2	Mean %age of Earnings After Last Child is Born
Lowest	$1,000	$83,556	4.4	31.5	65.6
2	4,942	42,252	3.5	31.2	71.4
3	9,600	46,481	3.5	31.4	75.0
4	14,423	63,616	3.1	30.1	83.8
Middle	32,020	84,142	3.0	30.3	82.2
6	35,000	81,134	2.7	30.2	82.1
7	59,950	135,709	2.5	30.7	81.0
8	91,347	164,248	2.3	30.8	82.3
9	86,500	188,393	2.0	31.8	79.5
Highest	129,808	315,067	2.1	32.7	78.3
All Singles	39,000	121,682	2.8	31.0	78.4

Notes

/1 Earnings deciles are defined separately for married couples and singles

Table 3: Actual and Optimal Median and Mean Net Worth, Unweighted HRS Data

Lifetime Earnings Decile	Median 1992 Net Worth	Median Optimal Net Worth	Mean 1992 Net Worth	Mean Optimal Net Worth
Lowest	$2,885	$1,350	$44,872	$38,537
2	21,050	10,749	73,767	56,447
3	37,750	24,281	97,291	67,629
4	61,565	36,539	142,990	103,014
Middle	80,938	45,733	162,037	110,753
6	99,300	63,639	179,298	120,961
7	118,462	74,250	212,584	143,187
8	157,000	93,618	254,822	172,105
9	213,000	127,082	329,334	220,750
Highest	353,500	221,434	639,505	433,869
Full Population	88,200	52,889	204,109	139,071

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				
	Baseline		**No Variation in kids**	
Decile of Lifetime Earnings Distribution	**Median Net Worth**	**Credit Constrained Until Age…**	**Median Net Worth**	**Credit Constrained Until Age…**
Lowest	$1,350	34	$16,403	26
2	10,749	32	27,584	27
3	24,281	31	31,475	27
4	36,539	29	38,576	28
5	45,733	28	45,638	28
6	63,639	27	64,372	29
7	74,250	27	67,463	30
8	93,618	29	87,394	31
9	127,082	30	115,394	31
Highest	221,434	32	180,463	34

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

Table 5: Effect of Altering the Timing of Children, Earnings, and The Transfer System on Median Optimal Net Worth, HRS Data

	Median Optimal Net Worth			
Decile of Lifetime Earnings Distribution	**Baseline Model**	**No variation in Timing**	**No variation In alpha**	**No Means tested transfer**
Lowest	$1,350	$2,674	$14,356	$1,483
2	10,749	14,563	25,674	11,302
3	24,281	27,946	32,564	25,056
4	36,539	37,956	39,561	36,897
5	45,733	46,475	45,637	46,088
6	63,639	62,197	64,573	63,858
7	74,250	72,183	66,674	74,382
8	93,618	91,364	80,675	93,656
9	127,082	122,362	110,263	127,131
Highest	221,434	210,573	170,483	221,437

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

Table 6: Results on Fertility and Net Worth for a Model with Endogenous Fertility, Married Couples, HRS Data

Decile of Lifetime Earnings Distribution	Optimal Net Worth (Exogenous Fertility)	Birth Rate (HRS Data)	Predicted Birthrate (Endogenous Fertility Model)	Net Worth, No Transfers, Endogenous Fertility	Birthrate, No Transfers, Endogenous Fertility
Lowest	$20,714	4.6	4.5	$26,221	4.3
2	38,254	4.1	4.2	41,573	4.1
3	53,894	3.9	4.0	54,903	4.0
4	71,996	3.5	3.7	72,035	3.7
5	74,718	3.7	3.5	74,734	3.5
6	79,159	3.6	3.4	79,163	3.4
7	111,280	3.3	3.3	111,282	3.3
8	134,092	3.3	3.3	134,092	3.3
9	153,326	3.3	3.3	153,326	3.1
Highest	270,442	3.1	3.2	270,442	3.0

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

Table 7: Median Regression Estimates of Household Net Worth: The Sample Is Restricted to Married, Non-Self-Employed Couples			
	Coefficients	**Standard Errors**	**T-Statistics**
High School Graduate	23,934.3	4,987.9	4.80
College Graduate	74,515.0	12,142.2	6.14
Post-College Education	103,614.4	22,899.2	4.52
Age	2,359.3	727.5	3.24
Do You Hold a DB Pension	5,169.8	3,610.7	1.43
1 Child	-62,871.0	13,090.1	-4.80
2 Children	-36,207.0	13,087.3	-2.77
3 Children	-48,026.9	14,135.1	-3.40
4 Children	-51,591.0	19,202.3	-2.69
5 Children	-63,353.1	14,718.2	-4.30
6 Children	-56,888.6	26,360.6	-2.16
7 or More Children	-68,721.7	12,311.9	-5.58
Constant	-85,025.4	48,722.8	-1.75

Notes: The regressions include annual earnings for households between ages 22 to 65.

Figure 1: Net Worth in 1992 as a Percentage of Summed, Real Lifetime Earnings, By Family Size, HRS Data

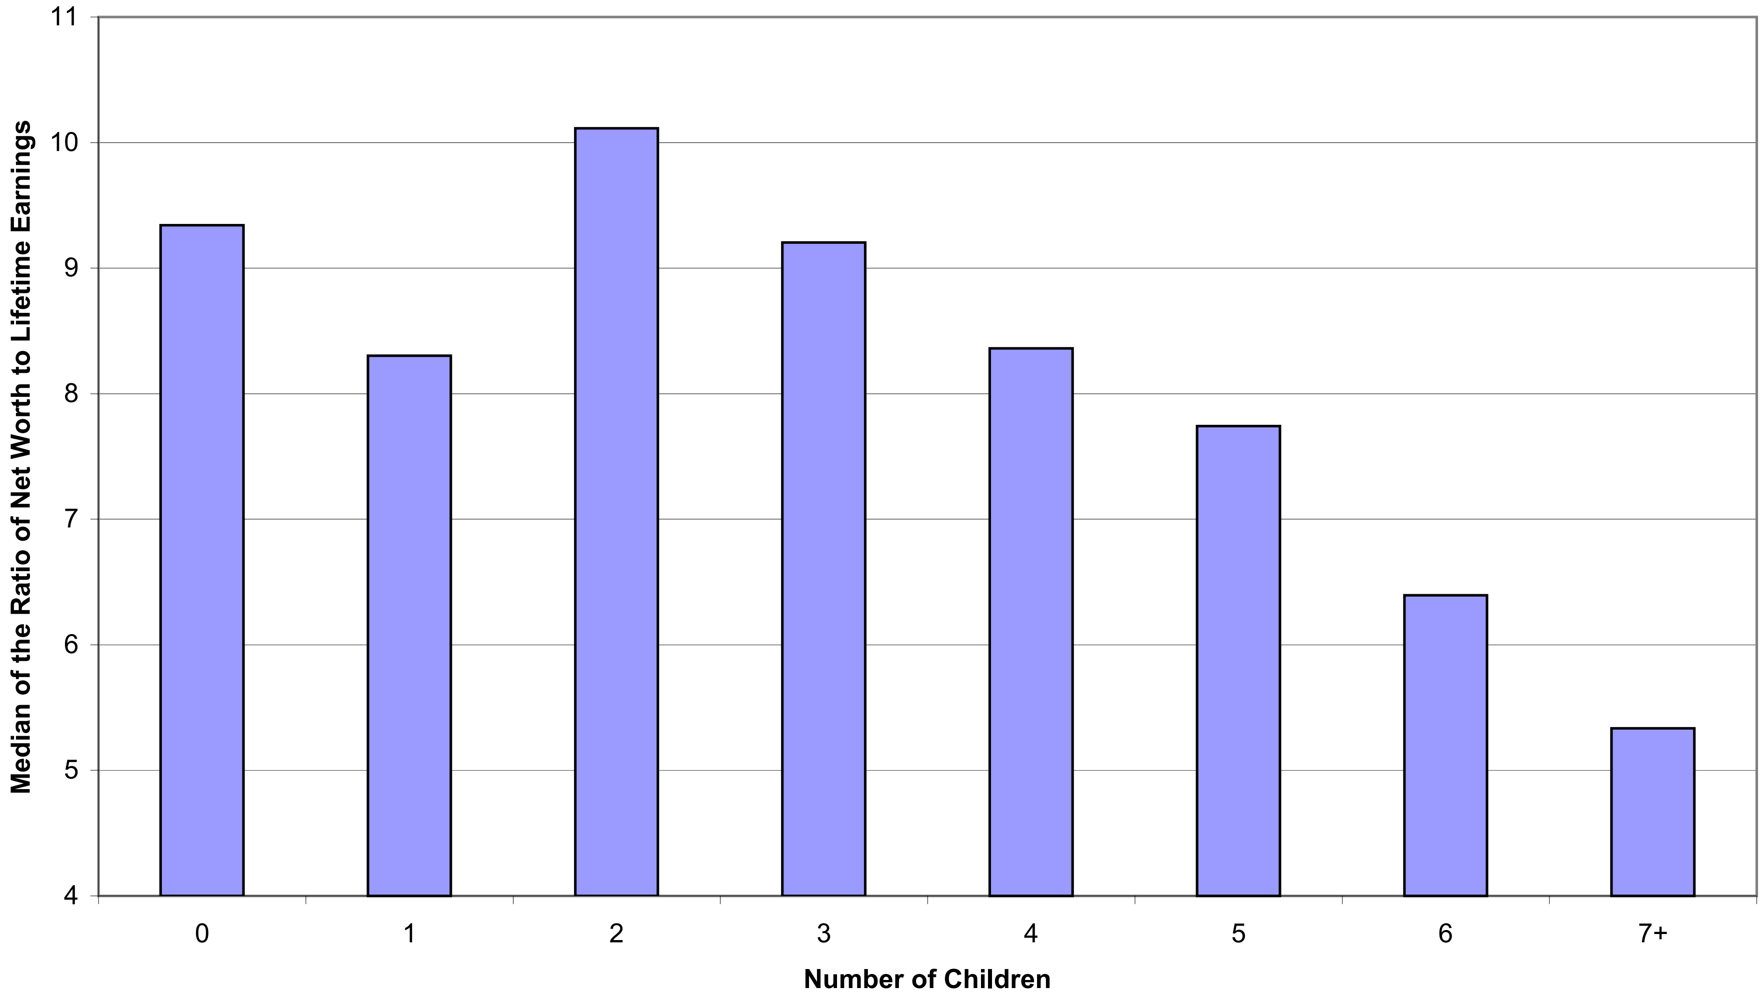

Figure 2: Median Age-Log Earnings Profiles by Family Size

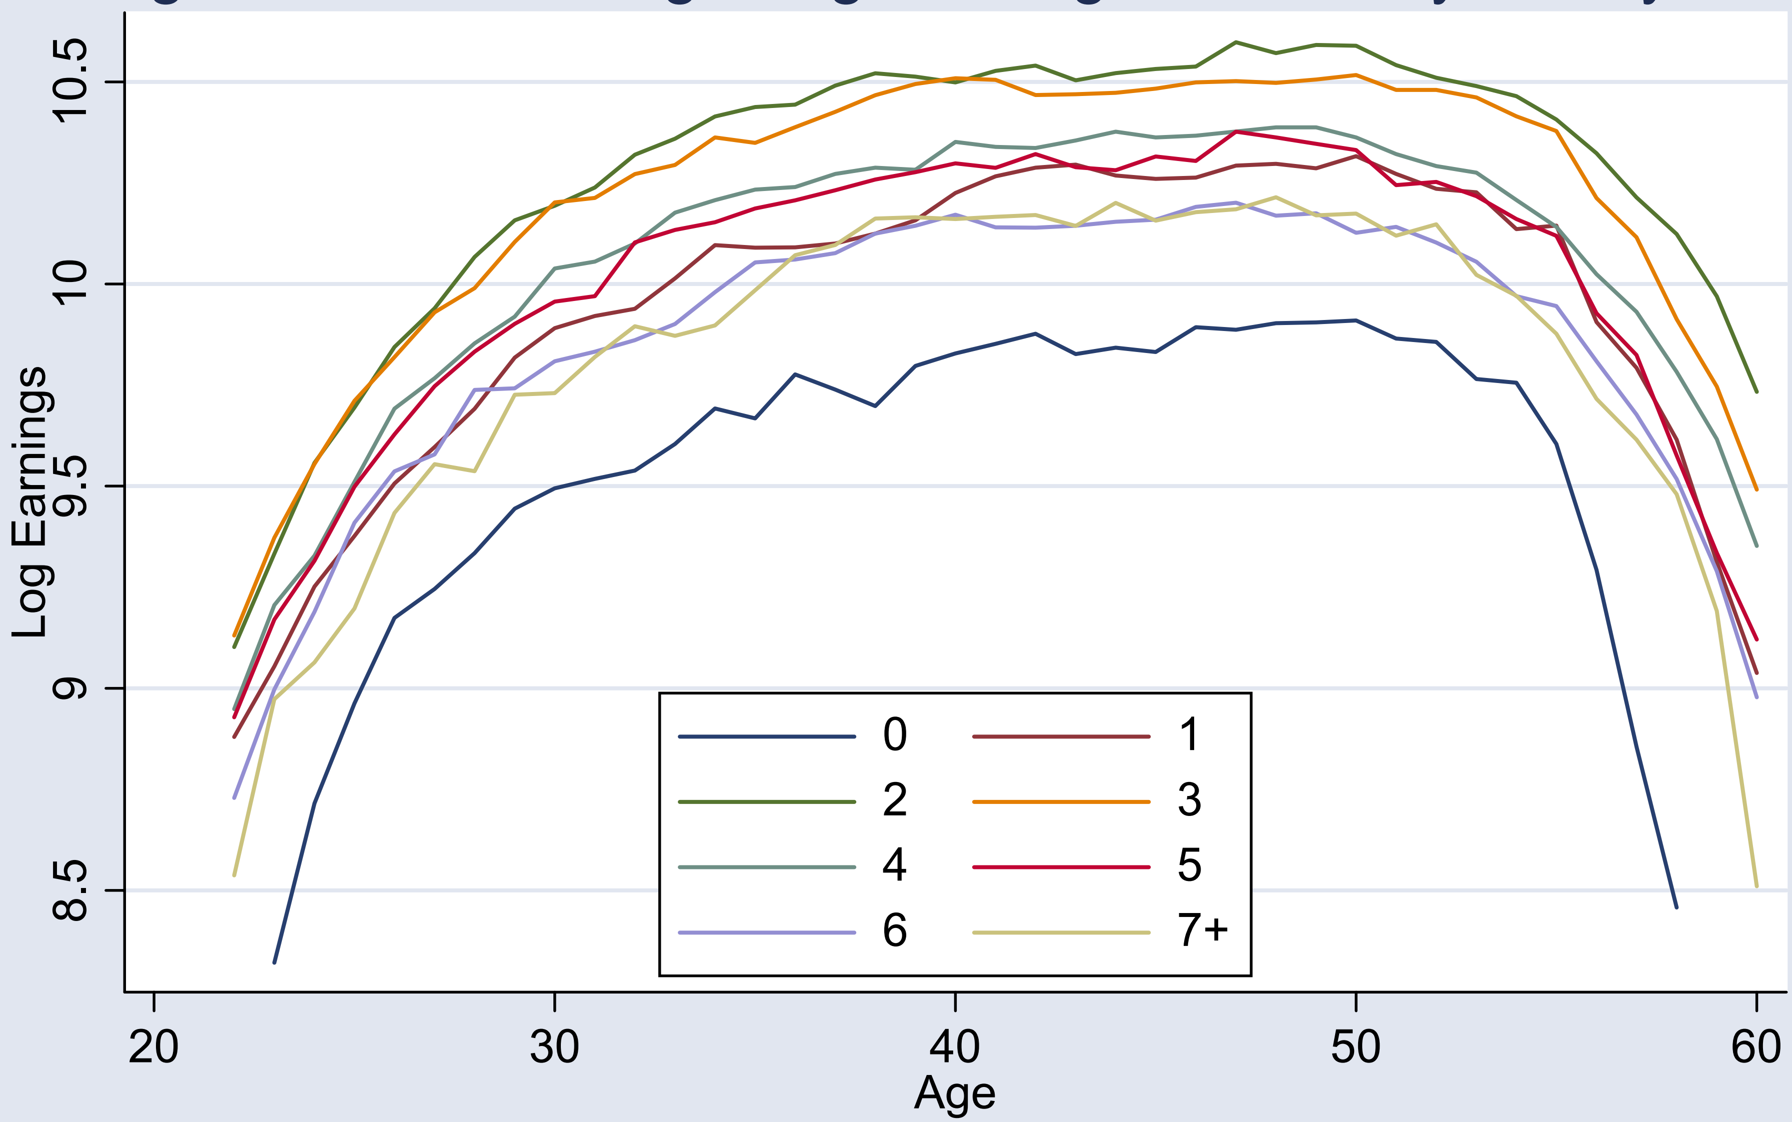

Source: HRS and restricted access social security earnings data and authors' calculations.

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

Variable	Mean	Median	Standard Deviation
Present Discounted Value of Lifetime Earnings	$1,718,932	$1,541,555	$1,207,561
Defined Benefit Pension Wealth	$106,041	$17,327	$191,407
Social Security Wealth	$107,577	$97,726	$65,397
Net Worth	$225,928	$102,600	$464,314
Mean Age (years)	55.7		4.7
Mean Education (years)	12.7		3.4
Fraction Male	0.70		0.46
Fraction Black	0.11		0.31
Fraction Hispanic	0.06		0.25
Fraction Couple	0.66		0.48
No High School Diploma	0.22		0.41
High School Diploma	0.55		0.50
College Graduate	0.12		0.33
Post-College Education	0.10		0.30
Fraction Self-Employed	0.15		0.35
Fraction Partially or Fully Retired	0.29		0.45

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

Appendix 1: 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.

Appendix 2: 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.[32] We use the 1966 parameters for years before 1966 and the 1989 parameters for 1990 and 1991.

[32]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 four groups of households: (1) single without a college degree, (2) single with a college degree, (3) married without a college degree, and (4) married with a college degree, using the 1998 and 2000 waves of the HRS, which provide medical expense information on households age 27 to 106.[33] We use the age and education of the head of household in the empirical model. Results are given in the third section of the Appendix. 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 others.

[33]Older cohorts from the AHEAD and two new cohorts were added to the HRS in 1998, which gives us a broader range of ages to estimate medical expense profiles after retirement. These new cohorts were not matched to their social security earnings records, so they cannot be used for our baseline analysis.