

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

November 2008

We thank colleagues at the University of Michigan for developing the Health and Retirement Survey, Surachai Khitatrakun for his advice, and seminar participants at Beijing University, Chicago, LSE, Penn, Texas, Tsinghua, University College London, and the IRP Summer Research Workshop as well as Bruce Meyer, Kevin Murphy, and Jon Skinner for helpful comments. We are also grateful to the CDHA at UW-Madison for financial support that initiated our work and for financial support from the Michigan Retirement Research Center, the NIA, and the NICHD. 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.

Abstract:

This paper examines the effect that children have on wealth accumulation. The small existing literature on this topic is puzzling, since children may 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 wealth accumulation, which has important implications for retirement planning. Third, variation in family size plays an important role in understanding the wide dispersion in wealth. Fourth, once variation in family size is 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 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 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. 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. We develop intuition using the simplest approach: a permanent income model with no uncertainty and complete markets. But this framework does not come close to matching the distribution of existing wealth. Moreover, it takes the arrival of children as being exogenous.

Because fertility may be affected by wealth, earnings expectations, and institutional features of the tax and transfer system, we develop our central results for an augmented life-cycle model

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

that incorporates endogenous fertility. Our approach combines two literatures. The first is models of life-cycle consumption pioneered by Modigliani and later extended by Deaton (1991), Hubbard, Skinner, and Zeldes (1995), Engen, Gale, and Uccello (1999), and others to include idiosyncratic shocks. The second is models of endogenous fertility choice pioneered by Barro and Becker (1989), who explicitly examine the trade-off between quantity and quality of children. By merging these two literatures, we present a very rich model capable of addressing a wide range of policy issues.

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, compare the effects of children relative to asset tests of cash- and near-cash transfers, and study the interplay between children and credit constraints.

We show that the model matches closely patterns of wealth and fertility in the HRS: our augmented life-cycle model with endogenous fertility does an excellent job explaining the joint distribution of fertility and wealth.² 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

² Scholz, Seshadri, and Khitatrakun (2006) show that a similar life-cycle model with exogenous fertility can closely match the net worth held by HRS households.

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 cash and near-cash 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 welfare programs (including old-age health insurance) will cause lower lifetime income households to accumulate little wealth. Once children are accounted for in the household, 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.

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 year of full-time work, and 908 households where the highest earner began working full time prior to 1951.³ We then drop 2,121 single adult households in 1992. These single households had an average of 2.8 children (recall all were born between 1931 and 1941), so the vast majority had previous marriages. Individuals in the lowest lifetime income decile had an average of 4.4 children. We do not observe the earnings of former partners, so we would do a much worse job characterizing the lifetime resources and annual earnings realizations of single individuals than we do for married couples. Hence, we drop these individuals, leaving a final sample of 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

1.1. 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 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.⁴ 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.⁵ The ratio of net worth (in 1992) to lifetime income is highest for families with no children. It

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.

⁴ 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 and saving adequacy.

⁵ 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. An on-line appendix of Scholz, Seshadri, and Khitatrakun (2006) provides complete details of our approach.

falls monotonically with the number of children above two.⁶ Figure 1 gives suggestive evidence 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 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. 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.⁷ 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.

⁶ Similar patterns hold if we calculate the ratio of net worth to lifetime income using the numbers in Table 1.

⁷ Specifically, we plot a median log earnings using Stata's "graph twoway mbands" command.

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).⁸ Retirement consumption will be financed out of defined benefit pension wealth (mean is \$133,241, median is \$45,248);⁹ social security wealth

⁸When 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.

⁹ 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).

(mean is \$135,338, median is \$133,547);¹⁰ and nonpension net worth (mean is \$208,549, median is \$142,885). The mean age of the household head is 55.8.¹¹

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

The Modigliani and Brumberg (1954) permanent income model, modified to allow family size to vary exogenously across the life-cycle, provides useful intuition about the effect of children on household wealth. Assume the household solves

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

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 of consumption for the number of children and adults in the household.¹² If preferences are

CRRA with $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$, the Euler equation is given by $\left(\frac{c_j}{N_j}\right)^{-\gamma} = [\beta(1+r)] \left(\frac{c_{j+1}}{N_{j+1}}\right)^{-\gamma}$ and the

marginal utility of household consumption (c_j) is equal across periods. The optimal solution is

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

¹⁰ 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).

¹¹ 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.

¹² We multiply utility by N_j so the marginal utility of consumption is equal across families of different sizes.

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.¹³ 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 lifetime resources when children are present, relative to families with fewer children (all else being equal).

The magnitude of the effect of children on consumption for the equivalence scale of Citro and Michael (1995), $N_j = (A_j + 0.7K_j)^{0.7}$, is shown in Figure 3, which is similar to Figures 5-7 in Banks *et al.* (1994). Total expenditure at each age in the Figure is normalized relative to expenditure when no children are in the household.¹⁴ The calculations show that expenditure when 5 children are in the household is more than twice the level it is when no children are present. Between the ages 21 and 58 (the average age of the HRS sample), the 5-child family will have consumed 50 percent more than the 0-child (using the equivalence scale). The discussion above focuses on household expenditure, which we will treat as being synonymous with consumption. Aguiar and Hurst (2005) discuss the differences between consumption and expenditure and show that life-cycle model implications appear to hold more closely for

¹³ It is straightforward to show the partial derivative of consumption with respect to family size,

$$\frac{\partial c_j}{\partial N_j} > 0.$$

¹⁴ A_j is the number of adults and K_j is the number of children in the equivalence scale. The other parameters used to generate Figure 3 are $r = .03$, $\beta = 0.97$, $\gamma = 3$ and children are born when the household head is 22, 25, 28, 29 and 31. All children leave the household at age 18.

consumption than they do for expenditure. Aguiar and Hurst (2007) demonstrate that effective prices paid vary over the life-cycle and this can help explain life-cycle consumption. Children presumably play an important role in understanding life-cycle consumption.

In work looking directly at children and consumption profiles, 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 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.¹⁵ 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, here we focus on two. First, family size in our framework is endogenous. This innovation is important, particularly when we examine the implications of policy changes. Second, previous work shows the life-cycle model can successfully match profiles of consumption. We take this work another step forward, focusing on *levels* of wealth and fertility, household by household. Explaining variations across specific households in the level of wealth and fertility is more demanding than just matching profiles.

¹⁵ Lawrance (1991) finds that accounting for variation in family composition reduces the heterogeneity in discount factors estimated from a consumption Euler equation.

Accounting for family size variation alone in the context of the simple permanent income, life-cycle model is not enough to explain the level and skewness of wealth.¹⁶ Thus, in the next section we describe calculations from a life-cycle model with borrowing constraints and idiosyncratic shocks, where family size is endogenous. Variation in earnings generates cross-household differences in the number of children and household wealth. A key mechanism is that consumption in households when children are present exceeds the consumption of otherwise equivalent but childless households. All else equal, this reduces the optimal wealth at retirement for households with children. Indeed, in what follows, we find that the quantitative effect of this phenomenon is large. But households with children will also be borrowing constrained for a longer period of time than would households without children, and this mitigates the effect of children on wealth.

III. A Model of Optimal Wealth Accumulation and Fertility Choice

We solve a life-cycle model, augmented to incorporate uncertain lifetimes, 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. Our goal in writing a model of endogenous fertility is to account for the joint distribution of wealth and fertility – a considerably more stringent test than simply matching wealth, as done in Scholz, Seshadri, and Khitatrakun (2006) and elsewhere. To accomplish this, we follow the pioneering work of Barro and Becker (1989) and assume that parents get utility

¹⁶ For example, in the life-cycle model above, 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, while households in the top decile have \$67,437. 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 mean of the ratio of resources available at retirement (social security wealth and defined benefit wealth) to lifetime earnings is about 23 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.

from the quantity and the quality of their children.¹⁷ 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

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

The expectation operator E denotes the expectation over future earnings uncertainty, uncertainty in health expenditures, and uncertainty over 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 $j-1 < R$.¹⁸ At age D , $p_D = 0$. The discount factor on future utilities is β .

As shown above parents care about the number of children, f , and utility per child, $U(c_j^k)$.

Following Barro and Becker (1989), the function $b(f)$ 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

$$c_j + f c_j^k + a_{j+1} = y_j + a_j - \tau(e_j + r a_j), j \in \{B, \dots, B+17\},$$

where

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

¹⁷ Unlike Becker and Barro (1988), we assume that the household does not care about children once they have left the household. 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.

Notice 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 implies that higher e_j households, where e_j denotes labor earnings at age j , will have fewer children than their lower e_j counterparts.

The budget constraint during the retirement period is given by

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

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

The variable y_j defines taxable income for working and for retired households.¹⁹ Labor supply and the retirement date are exogenous. $SS(\cdot)$ 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

¹⁸We 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.

¹⁹To 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.

²⁰Medical expenses are drawn from the Markov processes $\Omega_{jm}(m_{j+1} | m_j)$ for married households. Medical expenses drawn from the distribution for single households (after the death of a spouse) are assumed to be half of those drawn from the distribution for married couples.

income 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 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 two more choice variables than the standard life-cycle consumption problem: 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

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

In the above equation, $V_{B+1}(\bullet)$ stands for the value function at age $B+1$. 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

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

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.²² Sixteen percent of the sample reports having made a gift in the 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.

*III.1. Model Parameterization*²³

We assume households have constant relative risk-averse preferences, so

$$U(c) = \begin{cases} c^{1-\gamma} \\ 1-\gamma \end{cases}, \quad 0 < \gamma < 1. \quad \text{The restriction that } \gamma \text{ lies between 0 and 1 is designed to ensure that}$$

²² Overwhelmingly gifts were from older relatives, primarily parents, grandparents, and aunts and uncles.

²³ 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. More complete details on these and the

utility is always a positive number.²⁴ We assume an annualized real rate of return of 4 percent and the altruism function is given by $b(f) = b_0 f^{b_1}$, $0 < b_1 < 1$. We set the discount factor to $\beta = 0.96$.

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 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 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.²⁵ Estimates of the persistence parameters range from 0.58 for single households without college degrees to

defined benefit pension function, and social security calculations are given in the published and on-line appendices to Scholz, Seshadri, and Khitatrakun (2006).

²⁴ 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.

²⁵The 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.

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).

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. To examine their results in the context of our framework, which (among other things) accounts for the variation in fertility across families, we model the benefits from public income transfer programs using the Hubbard, Skinner, and Zeldes transfer specification. Specifically, the transfer that a household receives while working is given by

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

whereas the transfer that the household will receive upon retiring is

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

This transfer function guarantees a pre-tax income of \underline{c} , which we set based on parameters drawn from Moffitt (2002).²⁶ 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-

²⁶The \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 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).

parent, two-child family. We assume through this formulation 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 , 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

$f c_j^k + c_j = (1 + b_0^{1/\gamma} f^{1+b_1/\gamma}) c_j$. To calibrate the remaining parameters, we make use of the Citro and Michael (1995) equivalence scale mentioned earlier: $g(A_j, K_j) = (A_j + 0.7K_j)^{0.7}$, where A_j indicates the number of adults (children) in the household and K_j indicates the number of children in the household. 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 are 88 percent, 76 percent and 59 percent. Our scale lies between these values.

Imposing the equivalence scale and using expression for total family consumption, we have the condition

$$(1 + b_0^{1/\gamma} f^{1+b_1/\gamma}) \equiv 2^{0.7\gamma} (1 + 0.7f)^{0.7\gamma}.$$

This condition 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 $b_0 = 0.48$, $b_1 = 0.53$ and $\gamma = 0.64$. These parameters lie within the range of values in the fertility literature (see for instance Doepke, 2004). We then use the same parameters for all households to examine the model’s predictions.

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) 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 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.

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 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.

The optimal wealth target for the median households in the lowest decile of the lifetime earnings distribution is \$16,186 (including housing wealth). The mean target for bottom decile households is \$68,007. 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 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 would 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. But as emphasized in Scholz, Seshadri, and Khitatrakun (2006), observed median and mean net worth of married HRS households exceed the optimal targets in each lifetime income decile. Thus, from these

descriptive data, there is no evidence that these households are failing to prepare sufficient resources to maintain living standards in retirement.

Figure 4 gives a scatterplot of optimal net worth against actual net worth, for married HRS households with optimal and actual wealth between \$0 and \$1,000,000. The curved line gives a cubic spline of the median values of observed and optimal net worth.²⁷ If household net worth was exactly the same as optimal net worth, the ordered pairs of actual and optimal net worth for the HRS sample would map out the 45° line. The scatterplot gives striking visual evidence that most HRS households have saved very near or above their optimal retirement targets.

A second striking aspect of Figure 4 is that it illustrates how a well-specified life-cycle model can closely account for variation in cross-sectional household wealth accumulation. A linear regression of actual net worth against predicted net worth and a constant shows the model explains 74 percent of the cross-household variation in wealth (that is, the R^2 is 74 percent).

A common criticism of the life-cycle model is that it is unable to match the skewness of the observed wealth distribution without augmenting the model with purposeful bequest motives. The studies that develop this observation, however, miss two critical features that we incorporate. First, they lack detailed data on lifetime earnings realizations. Second, family size does not vary across households in these models. With these features, our model generates a distribution of optimal wealth that matches the skewness of the actual wealth distribution. The 90-25 ratio of weighted net worth in the data is 10.2. In the simulated optimal wealth data it is 9.9. The coefficient of variation in the actual data is 1.8, in the simulated data it is 2.0. Our model does not match the wealth holdings of extremely high lifetime income households: to

²⁷The median band is smoothed by dividing households into 30 groups based on observed net worth. We use Stata's "connect(s) bands(30)" option for the figure.

understand the wealth holdings of the top 1 percent or so of households, a bequest motive may be important (see, for example, De Nardi, 2004).

The model also does a fairly 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.

We show the precise comparison between model and data in Figure 5, which shows a box and whisker plot of the actual number of children in households (on the horizontal axis) and optimal number of children in the house (on the vertical axis). The boxes indicate the 25th and 75th percentiles of the distribution of the optimal number of children. The lines in the middle of the box indicate the median number of children.²⁸ The mean optimal number in the endogenous fertility model fairly closely matches the data as the actual number of children 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 0.64. Given the parsimonious specification of the fertility process, we conclude that the endogenous fertility model does a credible job matching the level and dispersion of the number of children in married HRS households.

²⁸ Recall that, given our approach to parameterizing the number of children, households can have a non-integer number of children.

III.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 3.6 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 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 3.6 children instead of 4.6 children, median optimal net worth increases to \$26,232 from \$16,186. Mean optimal net worth increases to \$99,462 from \$68,007. In the highest income decile the mean number of children increases to 3.6 from 3.1. Median optimal net worth falls to \$255,442 from \$305,605. Mean optimal net worth falls to \$462,183 from \$552,056. 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 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 36. This figure drops to age 25 when there is no variation in the number of children. It is clear from the results that the timing and number of children has a substantial impact on when the household begins saving for retirement.

The systematic variation of kids by lifetime income increases the dispersion in earnings. 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 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 2.2 in the baseline optimal net worth distribution to 1.6 when the variation in children is shut down.

III.3.2. The Effect of the Timing of Children

The previous experiment changes two things: the number of children and when they were born. We do not offer a theory for the timing of births, but timing might matter. Since children's consumption depends on their parent's consumption, income increases with age, and some households are credit constrained, households with children born later in life will spend a greater share of lifetime resources on children than other households. 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 3.6 children born at age 27 to a specification where married couples still have 3.6 children, but they are born at the modal ages at which 4-child families have children: namely, at ages 23, 26, 29, and the 0.6 child at age 33.

The results of our analysis of the effects that the timing of fertility has on net worth are shown in the last column of Table 4. In the lowest lifetime income deciles, the effects of changing the timing of births has only a minor effect on optimal wealth, relative to changing both the number and 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 conclude that variation across households in the timing of birth has relatively little effect on the patterns of wealth accumulation compared to the total number of children households have. Therefore, our failure to offer a theory for the timing of children is a quantitatively minor shortcoming of our analysis.

III.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). We nevertheless assume that a governmental program (like Medicaid) exists that insures individuals against out of pocket

medical shocks. Thus our experiment eliminates cash and near-cash transfers, but not health insurance for indigent elderly.

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.²⁹ 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 1990.³⁰

Columns 1 and 2 in Table 5 repeat the optimal 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 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 half 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 the transfer program. Second, eliminating the transfer program 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

²⁹ Our floor, of course, varies by year and by family composition.

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) find a large effect on wealth of asset tests associated with transfer programs. We find a small effect. But we have written down a different model, so perhaps it is not surprising that we get different results. 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 4 in Table 5. Then in column 5 we show the effect of eliminating cash and near-cash transfers. Doing so leads to a very large increase in optimal wealth accumulation, as households take a greater responsibility to self-insure. Indeed, the magnitude of the effect shown by comparison columns 4 and 5 in Table 5 are similar to that found by Hubbard, Skinner, and Zeldes.³¹

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, recall the example from Figure 3 with two married couples, one with five children, and the other with no children but identical in every 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

³⁰ 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.

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 very little disincentive effect of the transfer program on (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, an older age than for high-income households. 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 our view of the world, having 5 children (or the number of children observed in the HRS) alters optimal consumption choices sufficiently

³¹ The net worth to permanent income ratio for the 20th percentile household in their model was 0.35. It is 0.37 in our model with no children when we include both single individuals and married couples.

strongly to largely reconcile the low wealth holdings of the very poor with the data: the asset tests associated with cash and near-cash transfers play a minor role in wealth accumulation.³²

While we find small effects of cash and near-cash 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 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, leading to a smaller overall effect on consumption.

III.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.

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

³² In a regression-based analysis, Hurst and Ziliak (2006) find little effect on wealth accumulation from state-level

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

changes in asset tests associated with the 1996 welfare reform.

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. We interpret these estimates as being consistent with children having a substantial effect on the net worth of households approaching retirement.

IV. 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 risks, borrowing constraints, and endogenous fertility. Our study yields three main conclusions.

First, an amalgam of a standard Modigliani-style life-cycle model and a Barro-Becker type fertility choice model provides a remarkably good description of the joint distribution of fertility and wealth. To the best of our knowledge, such a framework has not been used to study the variation in fertility across households. Indeed, some speculate that one may need to appeal to differences in tastes across households in order to explain the microeconomic variation in fertility choice – our study accomplishes this without resorting to such differences.

Second, children have a substantial effect on wealth accumulation, which has important implications for retirement planning. Financial planning rules of thumb ignore the role that children play in optimal life-cycle wealth decisions. The “replacement rate” – the percentage of pre-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. Sensible financial planning advice should reflect that fact.

Third, variation in family size plays an important role in understanding the wide dispersion in wealth. Several 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 do richer households. We show that that a sensibly specified augmented life-cycle model with endogenous fertility can match the wealth distribution, except, perhaps, at the very top. Two features of our modeling approach are central to this result: first, it is critical to observe earnings realizations when examining the implications of the life-cycle model. Second, 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.

Finally, 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).

Hubbard, Skinner, and Zeldes emphasize the asset tests associated with the Aid to Families with Dependent Children program at the beginning of their paper. Moreover, we are able to replicate their results using our model. But 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. Nevertheless, it is clearly the case that eliminating the safety net for medical shocks would sharply increase wealth accumulation (and lower lifetime well-being) throughout the income distribution. Exploring the role of health shocks and health-related institutions is a worthwhile topic for future work.

References

- Aguiar, Mark and Erik Hurst. 2005. "Consumption vs. Expenditure." *Journal of Political Economy* 113(5): 919–948.
- Aguiar, Mark and Erik Hurst. 2007. "Lifecycle Prices and Production", *American Economic Review*, 97(5).
- 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.
- 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. 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.
- 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 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 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	\$16,186	\$109,006	\$68,007	4.6	4.5
2	65,400	41,040	159,462	105,013	4.2	4.2
3	87,500	54,898	165,335	130,327	3.8	4.0
4	105,286	70,172	211,178	145,433	3.6	3.9
Middle	126,200	84,805	201,292	157,711	3.7	3.8
6	135,000	101,106	233,372	190,998	3.5	3.7
7	175,456	118,750	290,330	232,028	3.3	3.6
8	203,852	170,929	317,147	257,595	3.3	3.5
9	261,000	209,324	438,408	342,850	3.3	3.3
Highest	433,800	305,605	680,671	552,056	3.1	3.2
All Married Couples	142,885	108,551	280,549	218,141	3.7	3.8

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

Baseline			No Variation in Kids		
Lifetime Earnings Decile	Median Net Worth	Credit Constrained Until Age...	Median Net Worth	Credit Constrained Until Age...	Median Optimal Net Worth, No Variation in Timing
Lowest	\$16,186	36	\$26,232	25	\$27,394
2	41,040	34	63,853	27	64,283
3	54,898	32	72,573	28	71,295
4	70,172	31	79,290	29	80,305
Middle	84,805	30	87,183	30	88,394
6	101,106	29	101,453	31	103,291
7	118,750	28	111,593	32	111,440
8	170,929	30	150,783	32	154,203
9	209,324	31	181,401	33	185,221
Highest	305,605	33	255,442	34	251,115
All Married Couples	108,551	31	104,204	30	105,394

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	\$16,186	4.5	\$21,204	4.3	\$37,302	\$87,293
2	41,040	4.2	45,320	4.0	60,551	110,395
3	54,898	4.0	57,679	3.9	74,584	135,893
4	70,172	3.9	73,231	3.8	95,782	143,295
Middle	84,805	3.8	86,446	3.7	102,492	152,930
6	101,106	3.7	102,782	3.6	136,902	163,993
7	118,750	3.6	118,892	3.6	159,205	189,724
8	170,929	3.5	170,991	3.5	210,893	233,125
9	209,324	3.3	209,378	3.3	248,639	261,463
Highest	305,605	3.2	305,690	3.2	365,337	378,147
All Married Couples	108,551	3.8	110,395	3.7	140,226	177,496

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 (91,741)***	-182,395 (37,893)***	-378,306 (96,233)***	-202,598 (38,021)***
High School Graduate	40,234 (17,891)**	13,073 (4,805)***	47,544 (17,938)***	14,350 (6,042)**
College Degree	58,987 (26,668)**	51,811 (11,840)***	57,090 (26,681)**	52,002 (13,521)***
Post-College Degree	129,288 (29,732)***	81,925 (16,800)***	127,863 (29,735)***	73,548 (19,304)***
Age	7,244 (1,518)***	3,590 (665)***	7,862 (1,539)***	3,663 (579)***
Has a Defined Benefit Pension	-161,840 (14,733)***	-29,466 (5,638)***	-153,614 (14,847)***	-24,108 (8,784)***
African-American	-86,256 (21,259)***	-37,732 (5,815)***	-92,098 (21,258)***	-35,270 (7,677)***
Hispanic	-71,204 (25,429)***	-15,204 (7,468)**	-86,735 (25,660)***	-17,534 (7,874)**
Number of Children	-6,384 (2,902)**	-2,601 (828)***		
Number of Children * Bottom Lifetime Income Quintile			3,753 (4,009)	246 (808)
Number of Children * 2nd Lifetime Income Quintile			-5,157 (4,376)	-3,876 (1,006)***
Number of Children * Middle Lifetime Income Quintile			-13,670 (4,292)***	-7,514 (1,499)***
Number of Children * 4th Lifetime Income Quintile			-20,705 (5,064)***	-8,267 (3,099)***
Number of Children * Highest Lifetime Income Quintile			-12,902 (6,105)**	-4,936 (3,232)
<u>All regressions include annual earnings between ages 18 and 65</u>				
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%			

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: Consumption Over the Life-Cycle, Married Couple with Five Children



$$n_j = (A_j + 0.7K_j)^{0.7}; \quad r = .03; \quad \beta = 0.97; \quad \gamma = 3$$

Figure 4: Scatterplot of Optimal and Actual Net Worth, Married Couples, 1994 HRS





Figure 5: Optimal by Actual Children, Married Couples



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

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

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 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

1989.³³ We use the 1966 parameters for years before 1966 and the 1989 parameters for 1990 and 1991.

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

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

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

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 106.³⁴ 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 other.

³³Estimated 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 $a_1 \rightarrow 0$, 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.

³⁴Older 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.