Are All Americans Saving 'Optimally' for Retirement?

John Karl Scholz
Institute for Research on Poverty and NBER
University of Wisconsin–Madison
and
Ananth Seshadri
University of Wisconsin–Madison

Prepared for the 10th Annual Joint Conference of the Retirement Research Consortium
"Determinants of Retirement Security"
August 7-8, 2008
Washington, D.C.

The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium (RRC). The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government or the RRC. We are grateful to our colleague Surachai Khitatrakun for very helpful advice, Jeff Kling for his help in acquiring the restricted HRS data, to the National Institute of Aging and the Social Security Administration for financial support, and to the professionals at the University of Michigan, RAND and the Federal Reserve Board of Governors who developed the HRS and SCF data that we rely on.

There is widespread concern expressed in newspapers, and in public policy and academic studies that a substantial fraction of Americans are preparing poorly for retirement. The headlines of newspaper articles – two recent examples are "Debt-Squeezed Gen X Saves Little" or "Retirement's Unraveling Safety Net" – suggest that individuals or the institutions that people rely on for retirement security are falling short.[1] Journalists likely take cues from the financial services industry and from writing by academics and other opinion leaders. An article in the 2007 *McKinsey Quarterly* (Court, Farrell, and Forsyth, 2007) states "One finding of our research was a segmentation indicating that only about a quarter of the boomers are financially prepared for their twilight years" (page 106). Munnell, Webb and Golub-Sass (2007) conclude "The National Retirement Risk Index has shown that even if households work to age 65 and annuitize all their financial assets, including the receipts from reverse mortgages on their homes, nearly 45 percent will be 'at risk' of being unable to maintain their standard of living in retirement."[3] A widely cited statistic from the National Income and Product Accounts (NIPA), the personal saving rate as a fraction of disposable personal income, has declined steadily since the early 1980s. The personal saving rate was negative in 2002, 2005 and 2006, the only years it has been negative since 1932. The personal saving rate was 13.2 percent in 1986. It would seem there are good reasons to worry about American's financial behavior.

But developing rigorous, systematic evidence on the degree to which people are preparing sensibly for retirement is difficult. In the first part of this paper we briefly summarize and interpret some of the evidence on the adequacy of retirement wealth accumulation. A key building block for many studies is the "replacement rate" concept. For reasons discussed below, we think replacement rates do not provide a sensible underpinning for assessing retirement financial preparedness. In the second part of the paper we present descriptive evidence on wealth holdings across U.S. birth cohorts and

---

[1] See "Debt-Squeezed Gen X Saves Little," Lynn O'Shaughnessy, USA Today, 5/20/2008, p. 5a.; or "Retirement's Unraveling Safety Net," Dale Russakoff, Washington Post, 5/15/2005, p. 1a.

[3] As additional examples, see "A Nation in Debt," Barbara Dafoe Whitehead," The American Interest Magazine, July/August 2008, III(6), and David Brooks, "The Great Seduction," New York Times, June 10, 2008.

subjective attitudes about their financial circumstances in retirement. This descriptive evidence does not seem consistent, in our view, with dire assessments of poor financial preparation.

In part three of this paper we implement the straightforward, but computationally complex approach used in Scholz, Seshadri, Khitatrakun (2006) to assess the adequacy of retirement wealth preparation of Americans born before 1954. Extensive data are required to implement this approach. We have recently gained access (under tightly restricted conditions), to the social security earnings histories for a much broader age range of Health and Retirement Study participants than we previously had available. We use these data to assess the degree to which all Americans born before 1954 have accumulated or are accumulating the wealth necessary to maintain pre-retirement living standards in retirement.

We cannot emphasize too strongly that the results of this paper are very preliminary. The steps needed to undo top-codes in the social security earnings data, estimate models for expectations about medical expenses and future earnings, calculate the value of defined contribution pension balances, and calculate optimal household decision rules given the economic model we write down are involved. While making use of the RAND-version of the HRS saved us considerable time, we nevertheless have done a great deal of complicated data manipulation, estimation and computation in a short period. We have been able to compare our results for the original HRS cohort with our previous, thoroughly checked and vetted results, which allows us to benchmark some of our work. But even though we appear to satisfactorily meet this specification check, other problems may have entered into the analysis and we have not given the underlying analysis as much scrutiny as we would like. Hence, the results of this paper are preliminary.

1. A brief (selective) overview of the existing literature

There are two major elements to data-based analyses that conclude Americans are saving too little for retirement. The first is the declining, sometimes negative NIPA personal saving rate figures. The second are studies that make use of the workhorse financial planning concept of replacement rates.

NIPA Saving Rates

There are difficult issues associated with the NIPA's effort to measure personal saving, which are capably discussed in Gale and Sabelhaus (1999). One important item is that accrued (and realized) capital gains are excluded from the saving measure. Thus, increases in consumption that might result from appreciating stock market and housing wealth will result in falling NIPA personal saving rates. Investment in consumer durables is also not treated as personal saving, and Gale and Sabelhaus note that "From the perspective of economic theory, the line between personal and corporate saving is thin and somewhat arbitrary." They write, after looking at both data from the NIPA and Flow of Funds Accounts, that "The official personal saving measures do not measure wealth accumulation in the form of capital gains…. They provide inconsistent treatment of durable goods, payments from corporations, inflation, and taxes. They are affected by demographic factors, and they provide no information on the distribution of saving across households." The conclude that both the NIPA and Flow of Funds measures substantially overstate the decline in personal saving in the period they study.

It seems clear to us from the discussion in Gale and Sabelhaus (1999) and elsewhere, that one should not make inferences about the saving behavior of individual households based on the aggregate NIPA or Flow of Funds personal saving rate. Thus, they are badly flawed indictors of the degree to which Americans are adequately preparing for retirement.

Replacement Rates

The replacement rate – the amount of income in retirement needed to maintain pre-retirement living standards – is a simple, intuitively appealing concept. Typical financial planning advice suggests that replacement rates should be 70 to 85 percent of pre-retirement income.[4] Target replacement rates are less than 100 percent for three reasons. First, upon retirement, households typically will face lower taxes than they face during their working years, if for no other reason than social security is more lightly taxed than wages and salaries. Second, households typically also save less in retirement. Prior to retirement, households save to augment employer-provided and government-

[4] Applications will use different measures of pre-retirement income, such as income in the year immediately prior to retirement, average income during the working life, or income in the "n years" immediately prior to retirement.

provided pensions. During retirement, households typically decumulate resources, so the need to save is no longer a claim on available income. Third, work related expenses fall in retirement.

Low income household are thought to need higher replacement rates than higher-income households because, prior to retirement, they have lower tax burdens and saving rates than their more affluent counterparts.

Many studies use replacement rates as their standard for assessing the adequacy of wealth accumulation. Court, Farrell, and Forsyth (2007), for example, write that "Our analysis also indicates that 60 percent of boomers will need to work (following formal retirement) just to maintain 80 percent of their current consumption."

An ambitious study of financial preparedness for retirement is Munnell, Webb and Golub-Sass (2007), which finds that under a best case scenario, 43 percent of American households will be at risk of being unable to maintain their standard of living in retirement. One needs to make strong assumptions to go from the Survey of Consumer Finances (SCF) – a high quality *cross-sectional* dataset on household wealth – to assess the adequacy of wealth accumulation at retirement. We mention three of these.[5] First, they take the components of net worth observed for each household in the SCF and extrapolate these to age 65 (and then annuitize net worth). They assume that each household's place in the net worth distribution remains the same over time. Second, they need to estimate lifetime income. All men and women in the SCF are assumed to have the median earnings profile drawn from restricted social security earnings records from the Health and Retirement Study. There is, of course, a great deal of actual variation in earnings realizations across households. In the third step, if the household's replacement rate (after annuitizing all wealth sources) is below the designated target for a household with the given income, the household is "at risk." Forty three percent of households are found to be more than 10 percent below the target. Households who are younger than those in the original HRS cohort (born between 1931 and 1941) are particularly likely to be at risk.

[5] For more details on their approach, see Munnell, Webb, and Delorme (2006). Munnell, Webb, and Golub-Sass (2007) provide evidence for Late Boomers (born 1955-64) and GenXers (1965-72). Our data do not cover these cohorts. Forecasting problems likely increase with the length of the forecast and optimal wealth accumulation for younger households tends to be very low, both because of children and upward sloping age-earnings profiles.

Given the often substantial idiosyncratic and aggregate shocks that households receive in middle and older ages, it is unclear how accurately the wealth extrapolations will mimic the actual wealth holdings of SCF households upon retirement. Hence, the resources available to households in retirement may be misstated. Perhaps more importantly, there will be substantial forecast errors (both positive and negative) in estimates of lifetime earnings, particularly when anchored by the earnings reports found in a single cross-section of data. But unlike many forecasting exercises, this is not a case where upside errors cancel out downside errors, leaving an arguably plausible average estimate. Those whose lifetime earnings are overstated are more likely than they should be to be classified as being at risk. To see why, consider a household that, by assumption, has accumulated exactly the retirement resources needed to maintain living standards. If the forecast of lifetime earnings is overstated, which likely will occur in roughly half the cases, the household will appear to have insufficient wealth, not because wealth accumulation is too low but because the estimate of lifetime earnings (and hence, pre-retirement living standards) is overstated. If earnings forecast errors are substantial and symmetric, it is perhaps not surprising that upwards of 50 percent of the population is found to be at risk.

The replacement rate concept is also flawed. We elaborate on the following in Scholz and Seshadri (2007), but consider the following example. One married couple has five children. A second is identical in every way, except they have no children. Because the family with five children consumes more prior to retirement (when the children are around), they should optimally accumulate less retirement wealth than the otherwise identical zero-child family. The reason is that 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 family's resources are used to support their children. Social security replacement rates of, say, 50 percent may fully meet the retirement consumption needs of the parents, once the children are out of the house. Financial planning rules of thumb, and specifically replacement rates, ignore the role that children play in optimal life-cycle wealth decisions. Replacement rates also do not account for differences in timing of income and wealth shocks that occur over household members' working lives.

Some Papers Suggest Most Americans are Preparing Well for Retirement

The life-cycle model, augmented for uncertain lifespan, uncertain incomes, and old-age health shocks is a natural starting point for assessing the adequacy of wealth accumulation for retirement. Even if one does not think people behave in a manner consistent with the model – it is the natural *normative* benchmark for assessing adequacy. The lifecycle model's prescription for optimal behavior is both straightforward and intuitive: households will maximize wellbeing when they equate the discounted marginal utility of consumption across periods, consuming lifetime resources by the time they die. Put differently, given lifetime resources and preferences, lifecycle model consumption choices will maximize appropriately discounted lifetime well-being.

Hubbard, Skinner, and Zeldes (1995) and Engen, Gale, and Uccello (1999) use life-cycle models to simulate the expected distribution of wealth for representative household types. Hubbard, Skinner and Zeldes note that when realistic features of the tax and transfer system are modeled, the distribution of optimal wealth that results from the life-cycle model matches the distribution observed in data. Engen, Gale and Uccello conclude, using their best judgment regarding model and data, that "households are largely saving adequately, but other interpretations are possible."

Scholz, Seshadri, and Khitatrakun (2006) go a step further, and examine the household-specific implications of an augmented lifecycle model. We found that fewer than 20 percent of households born between 1931 and 1941, members of the original Health and Retirement Study (HRS) cohort, had less wealth than would be suggested by an optimal household-specific target. These targets were computed from an augmented life-cycle model with uncertain earnings, lifespan, end-of-life health shocks, supplemented with Social Security earnings records and other economic and demographic data from the HRS. The wealth deficit of those who were undersaving was generally small. A critical unresolved issue, however, is the degree to which these results hold for other cohorts, particularly those born after 1941.

Love, Palumbo and Smith (2008) study the wealth trajectories of households in retirement, showing they do not decumulate wealth as quickly as one would expect from

the no-uncertainty life-cycle model. They show that while wealth in levels falls with age for elderly households in the HRS, "annuitized" wealth does not. Annuitized wealth reflects both the steady flow of annual income that could be drawn from a given level of wealth and the fact that as years go by, remaining lifespan will (generally) fall. Rising annuitized wealth as households move through their retired years is not the pattern one would expect to see in the data if people systematically were saving too little for retirement.

2. Descriptive evidence on the adequacy of saving

There are three pieces of descriptive evidence that provide additional perspective on the degree to which Americans are preparing well for retirement. The first shows the net worth held by the typical member of broadly specified birth cohorts, at comparable ages. If some birth cohorts are preparing well, while others are doing less well, one might expect to see evidence of this when comparing cohort patterns of wealth accumulation over time.

The second simply compares the wealth, lifetime income, and wealth-to-income ratios of HRS cohorts. As the HRS has matured, new cohorts have been added. The 2004 version of the data, which we rely on for this paper, includes households from the AHEAD cohort, born before 1924; Children of Depression Age (CODA) cohort, born between 1924 and 1930; the original HRS cohort, born between 1931 and 1941; the War Baby cohort, born between 1942 and 1947; and the Early Boomer cohort, born between 1948 and 1953. Again, if there are substantial differences in behavior between HRS birth cohorts, one might expect to see clues in the descriptive data.

The third makes use of two subjective questions posed in the HRS to retired households: (a) how satisfied are you with retirement, and (b) how are the retirement years compared to before? A comparison of responses to these questions, and how they relate to net worth may be revealing about the degree to which people have prepared well for retirement.

Cohort Patterns of Wealth Accumulation

Figure 1 shows data from the Survey of Consumer Finances (from 1962, 1983, and every three years between 1989 and 2004) for two population cohorts: households who are age 25 to 39 when we begin to follow them, and households who are age 40 to

54. We plot the evolution of median net worth for 25 to 39 year olds in 1962, in 1983 (there is no SCF-like survey conducted in the 1970s), and in 1992. We also plot the evolution of median net worth for the three older cohorts: those who were 40 to 54 in 1962, 1983, and 1992.

Figure 1: Median Net Worth of Cohorts, Full Population (2004 dollars)

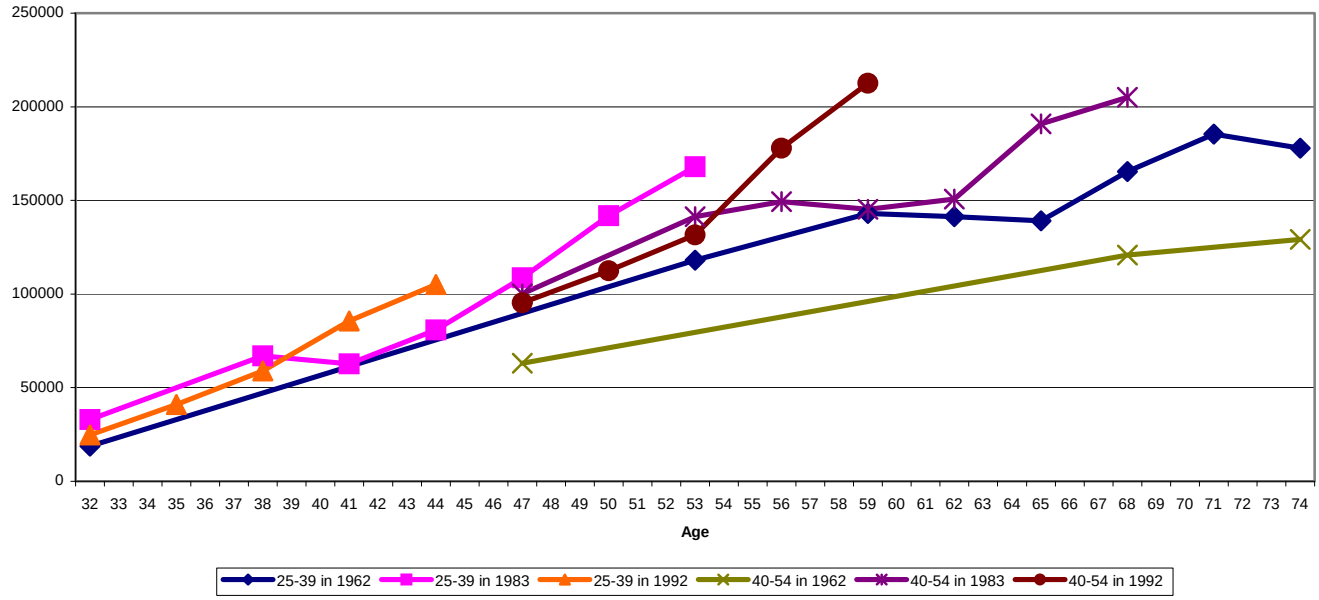

Each symbol in the Figure plots the median net worth at the middle age in the given age band (for example, households age 40 to 54 are plotted as if they were 47 years old). The figures show the evolution of median net worth for the *same sets of households over time*, since (aside from mortality, immigration and emigration) we know households who are 25 to 39 in 1962 (as defined by the head's age) will be 46 to 60 in 1983, 52 to 66 in 1989, and so on until their final observation as 67 to 81 year olds in 2004.

There are three noteworthy aspects of Figure 1. First, the cohort defined as 40 to 54 in 1962 (the line marked by "x" in the lower right portion of the figure) has significantly lower net worth than the other cohorts. Individuals in this cohort were children or young adults during the Depression and were young adults during World War II. Opportunities for human capital acquisition and wealth accumulation were more limited for this cohort than they were for subsequent cohorts. Second, median net worth grows steadily for each cohort. The patterns shown here are difficult to reconcile with

assertions that living standards for typical Americans are declining. Third, each successive cohort ends up with somewhat more wealth after the last two periods of observation (in 2001 and in 2004) than the cohort before it. To see this, at any given age (fixing age on the horizontal axis), the most recently born of the given age group has greater net worth (read straight down, which holds age constant). This shows that net worth (in levels) is growing across cohorts, even through the period of weak economic and stock market performance between 2001 and 2004.

Wealth and Income Across HRS Cohorts

Table 1 provides data from the HRS on wealth accumulation and lifetime income across cohorts. The table provides useful magnitudes for interpreting the model-based simulation results shown later. Slightly less than half of the sample is from the original HRS cohort born between 1931 and 1941.[6] This is the group intensively studied in Scholz, Seshadri, Khitatrakun (2006). It is useful to have these households in the sample, since they allow us to benchmark our preliminary treatment of the new HRS sample by comparing new results for the original HRS cohort with the results from our earlier work. The remaining portion of the sample is roughly evenly split between the remaining four cohorts.

[6] There are, in fact, over 12,000 households in the 2004 HRS. We restrict the sample in this preliminary draft, however, to those who allow researchers access to their social security earnings records (under tightly restricted agreements). Our earlier work (Scholz, Seshadri, and Khitatrakun, 2006) and Haider and Solon (2000) suggest the observed characteristics of those in the original HRS cohort who agreed to release their data are very similar to those who do not. As our work develops we will give this selection concern greater scrutiny. But based on preliminary work, we think this sample is likely representative of American households over 51 in 2004.

Table 1: Household Wealth, Income, and Wealth-To-Income Ratios of the HRS Cohorts, 2004 dollars

	Number of Observations	Age of Head	Median Net Worth	Median Lifetime Income	Median Wealth-to-Income Ratio
AHEAD	1,380	85	$140,000	$512,273	0.290
CODA	1,035	77	208,300	779,527	0.303
HRS	4,174	68	249,700	1,151,943	0.226
War Babies	940	59	164,349	1,155,174	0.153
Early Boomers	965	53	64,500	831,355	0.079
Total	8,495	69	182,000	946,791	0.209

Note: One household in our sample does not fit into any cohort. Authors' calculations based on the RAND version of the HRS, the regular HRS waves combined with the restricted access Social Security earnings data.

The age of head column reminds us of the age ranges that are used when drawing samples for the new HRS cohorts. AHEAD households will clearly be affected by survivorship bias, namely, some members of this cohort will have died, and mortality is likely correlated with household resources. Hence, the AHEAD sample will likely be composed disproportionately of wealthier, higher-income members of the cohort.

Median net worth is a comprehensive wealth measure, reflecting the value of stocks, bonds, mutual funds as well as other financial instruments, the value of houses and real estate (less the associated debt), and defined contribution pension fund balances. The patterns of net worth are not conclusive. Early boomers have much less net worth than their counterparts in the HRS and War Babies cohorts, which could be consistent with the idea that this group of households is failing to prepare appropriately for retirement. Of course, we expect there to be a natural lifecycle pattern of wealth accumulation. Households in the early baby boom cohort in 2004 typically will have more than another decade in the paid labor market. Moreover, many will have children who have recently left the household. As emphasized by Scholz and Seshadri (2007), children have a substantial, negative effect on wealth accumulation. Hence, we expect there to be a great deal of wealth accumulated in the high-earning years between the time children leave the house and retirement. Moreover, we expect retired households to decumulate wealth. Hence, the patterns of net worth in Table 1 may be precisely what we would expect to see for life-cycle households.

The median lifetime income column is the sum of real household earnings up to 2004.[7] Most households in the oldest three cohorts have retired: for these cohorts, income is higher, the younger the households. Members of the War Babies and Early Boomers will typically work more years in the paid labor market. The final column of Table 1 reports the median of the ratio of net worth to cumulative earnings to date. As with the net worth (in levels) figure, the ratios may be consistent with problems in wealth accumulation, or may reflect precisely the pattern we would expect to see if the lifecycle model capably summarizes behavior. In descriptive regressions where wealth-to-income is the dependent variable, and conditioning variables include age and indicator variables for educational attainment, defined benefit pensions for the husband and wife, and indicators for cohort, the cohort indicators are insignificant, individually and jointly. The final section of the paper takes a more rigorous look at patterns of wealth accumulation across cohorts and by individual.

Subjective Views of Financial Satisfaction in Retirement

HRS respondents were asked two questions about their subjective views of retirement. The responses are summarized in Table 2. It is critical to understand that few households in the younger cohorts – the War Babies and Early Boomers – are actually retired. Those who are retired in these cohorts likely incurred some health or employment shock that led to unexpected negative changes in economic circumstances. For this reason, the samples are quite small in these cohorts. Also, there many fewer responses to the retirement comparison question than to the retirement satisfaction question.

Over the entire population, only 9 percent of households find retirement not at all satisfying. Nineteen percent of households find their living standards worse in retirement than they were prior to retirement. Responses to these subjective questions, while far from definitive, are consistent with the idea that households in the HRS are on track to achieving financially secure retirements, particularly over the portions of the sample (the AHEAD, CODA, and HRS cohorts) where there are substantial numbers of retirees.

[7] We have imputed earnings for households whose reported earnings are capped by the social security earnings limit. Our approach is described in section 3.

Table 2: Subjective Views of Retirement Financial Well-being, HRS Cohorts, 2004

	How satisfying is retirement?			Retirement years compared to before?		
	Very	**Moderately**	**Not at all**	**Better**	**Same**	**Worse**
AHEAD	55.8 percent	40.1 percent	4.1 percent	31.7 percent	52.4 percent	15.9 percent
CODA	56.1	36.4	7.5	40.3	36.1	23.5
HRS	53.1	37.3	9.7	48.4	35.9	15.7
War Babies	45.7	35.3	19.1	44.4	30.9	24.7
Early Boomers	19.2	41.1	39.7	30.4	29.0	40.6
Overall	53.3	37.8	8.9	44.7	36.4	18.9

Note: 4,156 households answered the "how satisfying" question. 1,047 answered the other.

3. Model-Based Calculations of the Adequacy of Wealth Accumulation

To avoid confusion about the specific model we have in mind, we start this section by describing our baseline model that incorporates uncertain lifetimes, uninsurable earnings, uninsurable medical expenses, and borrowing constraints.

We assume 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.[8] 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 uncertain future earnings, health expenditures, and life span.

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

[8] Married households in 2004 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 2004 are modeled as making their lifecycle consumption decisions as if they were single throughout their working lives.

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

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

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

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

The first two equations define taxable income for working and for retired households.[9] 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.[10] The tax function $\tau(\cdot)$ depicts total tax payments as a function of earned and capital income for working households, and as a function of pension and capital income plus a portion of social security benefits for retired households.

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.

[9]In the baseline model, we define a household's retirement date for those already retired as 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. The head is defined as being the person with the highest lifetime earnings.

[10]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.

We briefly discuss several key modeling decisions. Further discussion is given in Scholz, Seshadri, and Khitatrakun (2006). We use constant relative risk-averse preferences, so $U(c) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma}, \text{ when } \gamma \neq 1. \end{cases}$ In our baseline parameterization, 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 in the household and K_j indicates the number of children in the household. This scale implies that a two parent 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 Organization for Economic Cooperation and Development (1982), Department of Health and Human Services (Federal Register, Volume 56, Number 34, February 20, 1991) and Lazear and Michael (1980). The corresponding numbers for these equivalence scales in this example are 88 percent, 76 percent, and 59 percent. Our scale lies in between these values.

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 $T = \max\{0, \underline{c} - [e + (1+r)a]\}$, 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).[11] We assume through this formulation that earnings, retirement income, and assets reduce public benefits dollar for dollar.

We aggregate individual earnings histories into household earnings histories. Earnings expectations are a central influence on life-cycle consumption decisions, both

[11]The $\underline{c}$ in the model reflects the consumption floor that is the result of all transfers (including, for example, SSI). Moffitt (2002) provides a consistent series for average benefits received by a family of four from 1960 to 1998. We assume that the parameters for years prior to 1960 and after 1998 are the same as the closest year for which we have data. We adjust (and verify) amounts for different family sizes using equivalence scales.

directly and through their effects on expected pension and social security benefits. The household model of log earnings (and earnings expectations) is

$\log e_j = \alpha^i + \beta_1 AGE_j + \beta_2 AGE_j^2 + u_j$, where $u_j = \rho u_{j-1} + \varepsilon_j$ and

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. They are available on request (but due to space constraints, they are not included here).

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.[12] Estimates of the persistence parameters range from 0.69 for one-earner married couples without college degrees to 0.74 for married households with two earners, in which the highest earner has at least a college degree.

The specification for out of pocket medical expenses 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 eight waves of the HRS.

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

[12]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.

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 the earning expectation parameter, 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.

Steps Needed to Develop the Analysis Sample

As mentioned above, we start with the Rand HRS Data, which pulls HRS data for respondents and spouses across waves into a single analysis file with consistent variable definitions across waves. We add a good deal of information to the Rand data and put it on a household basis. This includes adding information on child ages, defined contribution pension benefits from past and current jobs, defined benefit pension coverage from past and current jobs, and we add in the restricted access social security earnings data.

Earnings data from 1951 through 1977 are potentially censored, that is, the earnings report is not allowed to exceed the social security taxable earnings cap. Beginning in 1978, we have access to uncapped W-2 earnings reports. Among those with positive earnings, 22.5 percent of households have earnings capped in 1971, while 3.2 percent were capped in 1951. We impute earnings above the taxable earnings limit using Tobit regressions where earnings are the dependent variable and covariates include indicator variables for marital status, census regions, race and ethnicity, birth year, gender, and education group. To add a dynamic element to the earnings imputations, we include variables for the household's position in the aggregate earnings distribution in each of the preceding 4 years. We replace capped earnings in cases where the predicted earnings from the regression exceed capped earnings. The predictions typically exceed the capped amounts for more than 80 percent of the capped observations.

Optimal Wealth Accumulation Across HRS Cohorts

Table 3 presents information on mean and median optimal wealth targets, the percentage of households in each HRS cohort that has accumulated less than their optimal target, and the median net worth shortfall, conditional on failing to meet the optimal target. The targets represent the amount of non-DB pension, non-social security net worth that the household should have accumulated, at the time we observe them in the 2004 HRS, to be on track to equate the discounted marginal utility of consumption over their remaining life. In addition to assuming the preference parameters discussed above, we assume that households continue working until their expected retirement date (future wages are drawn from the fixed-effect earnings expectations function, estimated by household type), the social security system that households anticipate when making annual consumption decisions is the one in effect in 2004, and the health shocks households experience in old age are the ones we estimate based on eight waves of out-of-pocket medical expenses from the HRS cohorts (these shocks are correlated through an AR(1) error term). Presenting the optimal wealth targets as we have done assumes, implicitly, that housing wealth is fungible and can be used to support consumption in old age.

Table 3: Optimal Net Worth (excluding social security and DB pensions) and Percentage Failing to Meet Their Optimal Targets

Cohort	Median Optimal Wealth Target	Mean Optimal Wealth Target	Percentage Below Optimal Target	Median Conditional Deficit
AHEAD	$55,359	$172,394	1.9	$3,705
CODA	76,785	241,384	2.0	1,714
HRS	88,190	289,714	2.6	4,760
War Babies	57,211	219,112	5.2	12,337
Early Boomers	26,831	127,864	10.2	16,306
Full Sample	67,131	238,031	3.6	7,927

The first column of Table 3 shows the median level of optimal wealth for households in each HRS cohort. Optimal median wealth targets are low for Early Baby Boomer households – this presumably is due to the consumption needs of children, who for the typical early boomer, will have just left (or be just leaving) the household.

Optimal wealth targets are higher for households in the War Baby cohort, as they continue to work in the paid labor market and accumulate wealth for retirement. The optimal amounts peak for the HRS cohort, and then it falls as households age, as households presumably decumulate wealth in retirement. The qualitative patterns shown in column 1 of Table 3 mirror the patterns of actual net worth shown in Table 1.

The second column shows the mean optimal net worth targets across cohorts. The fact that mean targets far exceed the median targets shows the wealth distribution is highly skewed. But the same qualitative pattern across cohorts is apparent. Mean optimal targets are lowest for the Early Boomers, rise sharply for the War Babies, and peak in the HRS cohort, whose average age in the 2004 HRS is 68. The mean targets then decline as households move through retirement.

Columns 3 and 4 provide the first formal estimates of the degree to which households outside the original HRS cohort are preparing well for retirement. Only 1.9 percent of households in the AHEAD cohort have net worth that is below their optimal targets. Conditional on not meeting the target, the magnitude of the deficit is $3,705. Consequently, it appears that "undersaving" is of little practical consequence for households in the AHEAD cohort, though we qualify this assessment with three observations. First, as mentioned earlier in the paper, the empirical results in the paper are preliminary. Second, there will be differential mortality by economic status in the AHEAD cohort. It may be the case that individuals in households that were below their optimal targets were also more likely to die. Third, the fact that people are at or above their optimal targets means simply that they are in position, given their social security and defined benefit entitlements, to maintain the discounted marginal utility of consumption over time. If a household had, for example, a living standard below the poverty line during their working years, they would still likely have a below poverty income during retirement, even when they have met their optimal targets. Thus, "optimal" does not necessarily imply socially desirable: it simply suggests that given available resources, people are not consuming more than they should if they wish to maximize lifetime utility.

Almost all members of the CODA and HRS cohorts also have accumulated wealth that equals or exceeds their optimal wealth targets. Of those who have not, their

median conditional deficits are small. The evidence for the HRS cohort is fully consistent with the results of Scholz, Seshadri, and Khitatrakun (2006). There we found that 15.6 percent of the HRS cohort had accumulated less than their optimal targets in 1992. But the median conditional deficit in 1992 was $5,260 and the stock and housing markets performed very strongly in the 1990s. Consequently, we would expect considerably fewer households in the HRS cohort in 2004 to have wealth below their optimal targets than we found for the same households in 1992.

A larger percentage of households in the War Baby and Early Boomer cohorts are below their optimal targets. Just over 5 percent of households in the War Babies cohort are below their optimal targets: the conditional deficit is $12,337. Just over 10 percent of households in the Early Boomer cohort are below their optimal targets: the conditional deficit is $16,306. Our preliminary evidence suggests that few households, in general, are failing to accumulate sufficient wealth to maintain their living standards in retirement. But to the extent there is an issue, it appears that under-saving is somewhat more likely for the younger cohorts of the HRS relative to the older cohorts.

We examine the correlates of undersaving by estimating descriptive probit regressions, where the dependent variable takes the value one if the household's net worth is less than the optimal target (because of space constraints, we do not report the specific regression estimates). Significant negative correlates include income decile (households in the sixth decile and above are significantly less likely to have net worth below their optimal targets); being in the HRS cohort relative to the AHEAD cohort; being older; being married; and having a college degree or more. Being non-Caucasian white is significantly, positively correlated with failing to meet the optimal target. The probability also increases with the body mass index of the household head. The results here differ from our earlier work. There, no covariate besides marital status was significantly correlated with failing to meet the optimal targets. Our new results suggest there may be some identifiable, generally lower SES groups in the economy who may be failing to accumulate the resources needed to maintain living standards in retirement.

While some characteristics may be correlated with undersaving, Figure 2 makes it clear that the problem is, in total, negligible among HRS households. Here we plot the ordered pairs of observed net worth on the horizontal axis and simulated optimal net

worth on the vertical axis. We limit the sample to households with observed and optimal net worth that is less than $1,000,000. If the model perfectly predicted actual wealth accumulation, all observations would cluster on the 45 degree line. In fact, we see overwhelmingly that the observations are below the 45 degree line. This is what was shown in Table 3, where we report that only 3.6 percent of the sample has accumulated less than their optimal target. The curved line gives a cubic spline of the median values of observed and optimal net worth.[13]

Figure 2 provides striking visual evidence that most HRS households have saved at or above their optimal targets. Moreover, it is consistent with the idea that a well-specified lifecycle 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 that the model explains 62 percent of the cross-household variation in wealth (i.e., the R^2 is 62 percent). This R^2, while higher than any alternative reduced form model we examined in our earlier work, is still considerably lower than the R^2 of 86 percent that we generated with just the HRS cohort in 1992. There are three primary differences in our current work, two of which will likely reduce the ability of the model to match observed behavior. First, we have expanded the age groups under investigation. We expect there to be a greater range of income and rate of return shocks across the more heterogeneous groups of households.[14] Second, the data come from 2004 rather than 1992. The 1990s were an extremely strong decade of housing and asset market returns. If rates of return exceed our maintained assumption of 4 percent real, we expect there to be greater dispersion in actual wealth accumulation relative to forecasts, which we expect to lower the ability of our model to match observed behavior. Third, as we have emphasized, our results are preliminary.

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

[14] Splitting the sample by cohort and running the R^2 regression, the R^2 statistics are .40 for the AHEAD cohort, .80 for the CODA cohort, .59 for the HRS cohort, .95 for the war babies cohort, and .84 for the early baby boomers. This suggests that cross-household differences in bequest behavior and medical shocks may lessen the ability of the model to mimic behavior for very old households.

Figure 2: Scatterplot of Optimal and Actual Net Worth, 2004

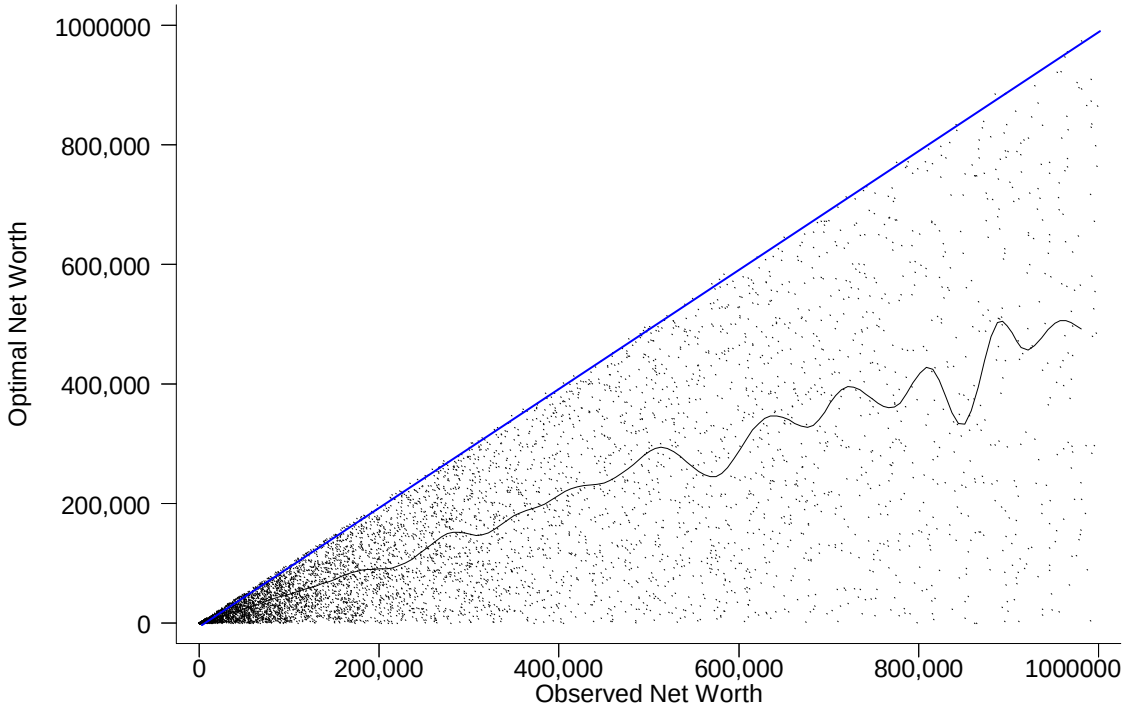

As is clear from Figure 2, many households in the HRS are accumulating more wealth than their optimal targets. There are several reasons why this may be the case. Households may have received a rate of return on net worth that exceeds the 4 percent (real) return we assume in the model. Households may anticipate life expectancy that exceeds the life-table estimates that we use. Households may have purposeful bequest intentions. Or households may anticipate future reductions in the generosity of social security or they may anticipate that out-of-pocket medical expenses, perhaps for end-of-life nursing home expenditures, will be larger than we assume. In this preliminary draft, we have not pursued additional analyses that might illuminate the importance of these factors. Instead, we take a first step by showing some characteristics that are correlated with the difference between observed net worth and optimal net worth (which we call the median wealth surplus).

[16] Income deciles are defined within the cohort (so the decile cutoffs are lower, for example, for early baby boomers, who have not completed their work years, than for households in the HRS cohort).

The results are given in Table 4. High income households are more likely to exceed their optimal net worth target.[16] Households in the ninth (tenth) decile of their cohort income distribution have a median wealth surplus that is $78,000 ($183,000) higher than households in the first decile. Retired households and married households have larger median wealth surpluses, and the magnitude of this surplus increases with educational attainment. In contrast, African American households are likely to have lower median wealth surpluses. None of the cohort indicators are statistically significant. We think it is plausible that higher income, more highly educated households may wish to leave bequests or may have higher desired precautionary saving, perhaps in anticipation of substantial late-in-life medical expenses. Hence, it is perhaps not surprising to see high SES households exceed their optimal targets by a substantial amount.

4. Conclusions

There is a considerable amount of discussion in the popular media and in policy and academic writing that Americans are doing a poor job of preparing for retirement. This perception is reinforced by recent low (and sometimes negative) personal saving rates in the National Income and Product Accounts, and perhaps recent disruptions in the housing, credit, and stock markets, and slow economic growth. But efforts to assess the adequacy of wealth accumulation require an objective standard to reach conclusions. The workhorse standard has been the replacement rate. But replacement rates are a conceptually flawed measure. We argue instead that the lifecycle model provides a natural, normative tool for assessing the adequacy of wealth accumulation.

Table 4: Correlates of Median Wealth Surplus			
	Coefficient	**Standard Error**	**T-statistic**
Constant	-35,193	123,240	-0.29
Income decile 2	-3,552	14,884	-0.24
Income decile 3	-4,871	12,557	-0.39
Income decile 4	-6,537	10,749	-0.61
Income decile 5	177	17,321	0.01
Income decile 6	5,678	14,739	0.39
Income decile 7	27,128	23,924	1.13
Income decile 8	52,349	31,255	1.67
Income decile 9	77,757	19,271	4.03
Income decile 10	182,722	20,943	8.72
CODA cohort	13,943	17,901	0.78
HRS cohort	23,124	23,377	0.99
WB cohort	8,559	41,420	0.21
EBB cohort	-9,849	53,504	-0.18
Retired	16,946	4,945	3.43
Age of head	507	1,434	0.35
African American	-22,191	3,300	-6.72
Other non-white	-19,977	15,159	-1.32
Married	30,977	9,637	3.21
GED	5,340	13,760	0.39
HS grad	22,975	5,238	4.39
Some college	27,982	11,720	2.39
College or More	121,748	16,078	7.57
One child	1,102	15,052	0.07
Two children	8,593	14,445	0.59
Three children	2,283	10,366	0.22
Four children	-3,955	19,911	-0.20
Five children	-6,237	10,275	-0.61
Six or more children	-8,095	15,067	-0.54

In earlier work, we showed the original HRS cohort, those born between 1931 and 1941, had overwhelmingly met or exceeded their optimal wealth target in 1992. But an extremely important unanswered question from our earlier work remained: do our findings apply to households born in other cohorts? Munnell, Webb and Golub-Sass (2007) argue that the answer is "no." But for reasons discussed earlier, we think their work is not the last word on this topic.

In this paper we present preliminary evidence from the HRS making use of all existing waves of the data. This includes two cohorts older than the original HRS cohort (the AHEAD and CODA) and two cohorts younger than the original HRS cohort (the War Babies and Early Baby Boomers). The approach used in our analysis is data intensive. It requires social security earnings histories and data on fertility, because one cannot develop a suitable measure of pre-retirement living standards without knowing the annual flow of resources that households received during their working lives and the composition of households when income arrives. Through the procedures the HRS has established for researchers to gain access to HRS respondents' social security earnings records, we are able to acquire the necessary data for a broader set of HRS cohorts and apply our earlier methodology.

Our preliminary evidence is striking. Only 3.6 percent of HRS households have net worth that is below their optimal targets. Conditional on having accumulated too little, the magnitude of the deficits is small. There is some evidence that younger households (those in the War Babies and Early Boomer cohorts) are less likely to have met their targets. But even in the early boomer cohort, only 10.2 percent of households are below their targets. The median deficit, conditional on not meeting the target, is $16,306. Thus, we think there is very little evidence that Americans, at least for those born before 1954, are preparing poorly for retirement.

These results increase our confidence that Americans are, by in large, preparing sensibly for retirement, given the existing generosity of social security, Medicare, and pension arrangements. We have additional work to do to explore the robustness of our results, particularly the ability of households to comfortably weather unusually large out of pocket medical expenses (recall, households in the model are hit with out of pocket medical expenses after retirement – the magnitude of these shocks are estimated using eight HRS waves of data). But we see little in the descriptive data or our model-based analyses that leads us to think that households are making large, systematic errors in their financial preparation for preparation.

References

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

Court, David, Diana Farrell, and John E. Forsyth. 2007. "Serving Aging Baby Boomers," *The McKinsey Quarterly*. Number 4, pages 103-113.

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

Gale, William G. and John Sabelhaus. 1999. "Perspectives on the Household Saving Rate." *Brookings Papers on Economic Activity* 1: 181–214.

Haider, Steven and Gary Solon. 2000. "Nonrandom Selection in the HRS Social Security Earnings Sample," mimeo, http://www-personal.umich.edu/~gsolon/workingpapers/nonresp.pdf.

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

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

Love, David A., Michael Palumbo, and Paul A. Smith. 2008. "The Trajectory of Wealth in Retirement." Board of Governors of the Federal Reserve System Research Paper Series.

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

Munnell, Alicia H., Anthony Webb, and Luke Delorme. 2006. "A New National Retirement Risk Index," An Issue in Brief, Center for Retirement Research at Boston College, June, No. 8.

Munnell, Alicia H., Anthony Webb, and Francesca Golub-Sass. 2007. "Is There Really A Retirement Saving Crisis? An NRRI Analysis." An Issue in Brief, Center for Retirement Research at Boston College, August, No. 7-11.

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.* August, 607-643.

Scholz, John Karl and Ananth Seshadri. 2007. "Children and Household Wealth." Mimeo, <http://www.ssc.wisc.edu/~scholz/Research/MRRC_Scholz_Seshadri.pdf>