

Preliminary

Are All Americans Saving ‘Optimally’ for Retirement?

William G. Gale
John Karl Scholz
Ananth Seshadri

The Brookings Institution (Gale) and
University of Wisconsin–Madison (Scholz and Seshadri)
1180 Observatory Drive
Madison, Wisconsin 53706-1393

jkscholz@wisc.edu

December 31, 2009

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 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.”² 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 13.2 percent in 1986, but it was negative in 2002, 2005 and 2006, the only negative years since 1932.

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

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.

² 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.

financial preparedness. We also present descriptive evidence on wealth holdings across U.S. birth cohorts and 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 the second part of the paper we extend the dynamic programming approach used in Scholz, Seshadri, Khitatrakun (2006) to assess the adequacy of retirement wealth preparation in 2004, using a sample of Americans born before 1954 (but who are fully retired in 2004 with at least 30 years of market earnings). We examine whether this group of households have accumulated the wealth necessary to maintain pre-retirement living standards in retirement.

The results of this paper are preliminary. With that caveat, our results suggest that roughly three-quarters of the households in our sample had accumulated sufficient resources in 2004 to maintain pre-retirement living standards in retirement. No household characteristics are strongly, systematically correlated with failing to accumulate sufficient wealth. Moreover, there are no strong patterns across the age groups included in the analysis. If these results hold as we incorporate a broader set of households and conduct additional specification checks, they suggest a somewhat less optimistic view about the adequacy of American's retirement preparation than what we found for 1992, the focal year in Scholz, Seshadri, and Khitatrakun (2006). There we found, for households born between 1931 and 1941, that 15.6 percent of the sample had accumulated wealth below their optimal targets in 1992 (and the magnitude of the conditional deficits were somewhat smaller than what we find in 2004).

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, capably discussed in Gale and Sabelhaus (1999), with the NIPA's effort to measure personal saving. One important item is that accrued (and realized) capital gains are excluded from the saving measure. Thus, increases in consumption caused by an appreciating stock market or appreciating housing wealth would 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." They conclude that both the NIPA and Flow of Funds measures substantially overstate the decline in personal saving in the period they study, and they conclude one should not make inferences about the saving behavior of individual households based on the aggregate NIPA or Flow of Funds personal saving rate.

Replacement Rates

The replacement rate – the amount of income in retirement needed to maintain pre-retirement living standards – is a workhorse concept in the financial planning literature. Typical advice suggests that replacement rates should be 70 to 85 percent of pre-retirement income.³

³ 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.

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 than they do during their working years, so saving is a smaller claim on available income. Third, work related expenses generally fall in retirement.

Low income households are thought to need higher replacement rates than high income households because, prior to retirement, they have lower tax burdens, so the difference in tax treatment between Social Security and wages and salaries is less consequential, and they likely devoted less in pre-retirement income to saving than their more affluent counterparts (Dyan, Skinner, and Zeldes, 2004), hence the reduction in taxes and saving experienced in retirement will be less substantial for these households.

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.” Munnell, Webb and Golub-Sass (2007) is perhaps the most ambitious study of wealth accumulation using the replacement rate concept. They conclude 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.

Munnell, Webb and Golub-Sass (2007) make strong assumptions to go from the Survey of Consumer Finances (SCF) – a high quality cross-sectional dataset on household wealth – to an assessment about the adequacy of wealth accumulation at retirement.⁴ They take the components

⁴ 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

of net worth observed for each household in the SCF and extrapolate these to age 65 and then annuitize net worth. This provides the numerator of the statistic they compare to the replacement rate target. The denominator of the statistic is lifetime income, which they need to impute from the SCF cross section. To do this they assume all men and women in the SCF have the median earnings profile (anchored by their observed earnings in 2004), calculated from restricted social security earnings records from the Health and Retirement Study. There is, of course, a great deal of variation in earnings realizations across households. With estimates of the numerator and denominator, they designate a household as being “at risk” if their replacement rate is 10 percent below a target that varies between 65 percent for high-income singles to 85 percent for low-income one-earner couples. The target averages 73 percent across all households.

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. They are also likely to have 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

cohorts. Forecasting problems likely increase with the length of the forecast and the optimal wealth accumulation at younger ages tends to be low, both because of children and upward sloping age-earnings profiles.

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 in that it ignores the role that children play in optimal life-cycle wealth decisions nor do they account for timing of income and wealth shocks, that may affect optimal wealth accumulation if households face credit constraints. Scholz and Seshadri (2009) provide a detailed analysis of the effects children have on optimal wealth accumulation.

To summarize, the replacement rate is a conceptually flawed benchmark for retirement planning. Moreover, the assumptions needed to calculate household replacement rates with some datasets are severe. A natural alternative is to use insights from the life-cycle model, augmented to account for fundamental factors affecting most households, such as demographic changes and uncertainty about future earnings, medical expenses, and longevity. The drawbacks to this approach, done correctly, are computational complexity and the data demands, which among other things, require information on annual earnings over individuals' lifetimes. But the lifecycle model is the appropriate conceptual benchmark, as lifecycle consumption decisions maximize lifetime well-being, subject to lifetime resource constraints. We emphasize the fact that we do not need to assume that people follow the lifecycle model for our approach to be informative. Rather, the lifecycle model provides a rigorous *normative* benchmark – if household wealth exceeds the lifecycle target, they are on-target for being able to maintain their pre-retirement living standards in retirement. Put less intuitively, they will be able to equate the discounted marginal utility of consumption across periods.

Several Lifecycle Papers Suggest Most Americans are Preparing Well for Retirement

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 a model with incorporating uncertain earnings, uncertain 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.

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.


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. Each cohort appears to be mimicking the behavior of the previous cohort, though they appear to be accumulating somewhat more. If people are systematically making important mistakes by accumulating too little for retirement, the mistakes seem to be occurring across all cohorts.

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



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 1941.⁷ 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 compare our

⁷ There are, in fact, over 12,000 households in the 2004 HRS. We restrict the sample here 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, but potential selection issues need further scrutiny. In the model-based calculations described below, we further restrict the sample to those where the primary earner is retired and has at least 30 years of earnings. Details are given below.

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.

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 shows the age ranges that are used when drawing samples for the new HRS cohorts. AHEAD households will clearly be affected by survivorship bias: 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 be in the paid labor market for more than another decade. Moreover, many will

have children who have recently left the household. As emphasized by Scholz and Seshadri (2009), 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 when they retire. 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 2004.⁸ Most households in the oldest three cohorts have retired: for these cohorts, income is higher, the younger the household. 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

⁸ We have imputed earnings for households whose reported earnings are capped by the social security earnings limit. Our approach is described in section 2.

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 are many fewer responses to the retirement comparison question than to the retirement satisfaction question.

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.

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.

2. 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 .⁹ 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

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

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,

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

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

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

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

The first two equations define taxable income for working and for retired households.¹⁰ 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

⁹ 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.

¹⁰In 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.

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 .¹¹ 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.

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

$$U(c) = \begin{cases} c^{1-\gamma} \\ \frac{c^{1-\gamma}}{1-\gamma} \end{cases}, \text{ when } \gamma \neq 1. \text{ In our baseline parameterization, we set the discount factor as}$$

$\beta = 0.96$ 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 and K_j indicates the number of children in the household. This scale implies that a two parent family with 3 children

¹¹Medical expenses are drawn from the Markov processes $\Omega_{jm}(m_{j+1} | m_j)$ for married and $\Omega_{js}(m_{j+1} | m_j)$ 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.

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).¹² 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 directly and through their effects on expected pension and social security benefits. The household model of log earnings (and earnings expectations) is

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

e_j is the observed earnings of the household i at age j in 2004-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

¹²The \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.

equation, and ε_j is a zero-mean i.i.d., normally distributed error term. The estimated parameters are α^l , β_1 , β_2 , ρ , and σ_ε . They are available on request.

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.¹³ 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

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

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

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 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 consumption (and hence asset accumulation) from the oldest possible age (D) to the beginning of working life (S)

¹³The 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.

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

The Health and Retirement Study (HRS) is sponsored by the National Institute of Aging and conducted by the University of Michigan with supplemental support from the Social Security Administration. 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 1931–41 birth cohort and their spouses, if they were married. Follow up interviews have continued every two years through 2008. As the HRS has matured, new cohorts have been added. We use the 2004 version of the data ($n = 20,129$), which samples households born before 1954.

The Social Security Administration provides detailed earnings records for those HRS respondents who grant explicit permission. This restricted-access data provide a direct measure of earnings realizations and lifetime income. These measures allow for the estimation of household's expectations of future earnings. They also allow us to simulate accurately social security benefits for the respondent and spouse or for the couple, if the benefit would be higher. Two issues arise in using earnings information. First, Social Security earnings records are not

available for 27 percent of the respondents included in the analysis. Second, the Social Security earnings records are top-coded for respondents who earned more than the social security taxable wage caps—for 13 percent of earnings observations between 1951 and 1979. From 1980 through 2003 censoring is much less of an issue because we have access to W-2 earnings records which are very rarely censored.

Dropped Observations

Table 3 shows descriptive statistics for respondents by whether or not they granted permission to release their Social Security earnings history. There are significant differences in 2004 earnings for those not granting permission, these respondents are also less likely to be retired and more likely to be male, black or Hispanic and college graduates. While these differences are statistically significant, most differences are not large. The group that did not grant permission is on average two years younger, 2 percent higher fraction male or black, 3 percent higher fraction Hispanic and slightly better educated (a difference of .18 years on average). The most important differences are earnings in 2004 and the fraction retired. There is no significant difference in net worth, which is the focus of this study. Haider and Solon (2000) briefly examine selection issues that arise with the restricted earnings data from the original HRS cohort and conclude “As far as can be told from observable data, the HRS Social Security earnings sample seems to be reasonably representative.”

Undoing the Top Coding

The second issue is that earnings data from 1951 through 1977 are potentially censored. 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 in each year using

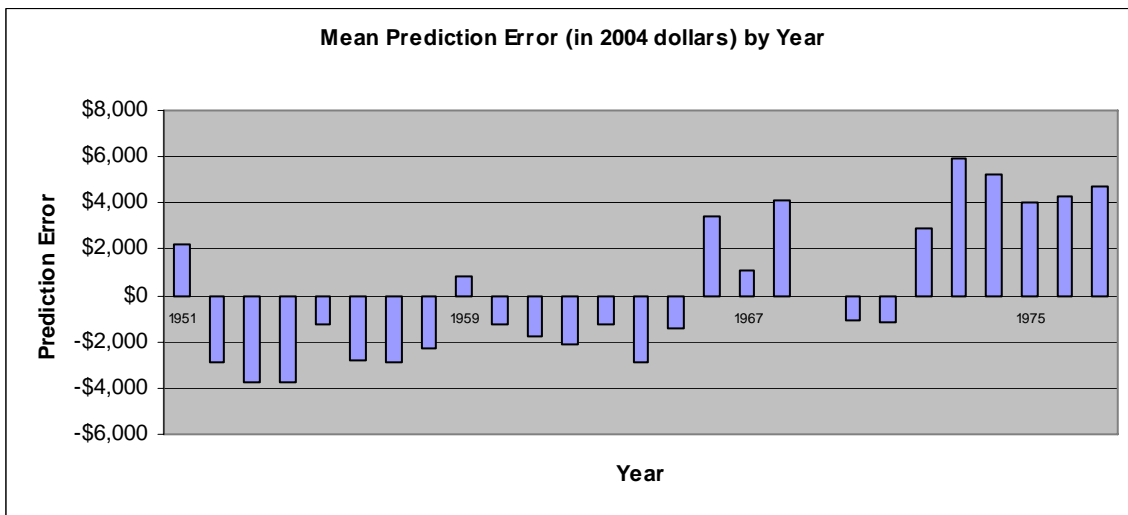
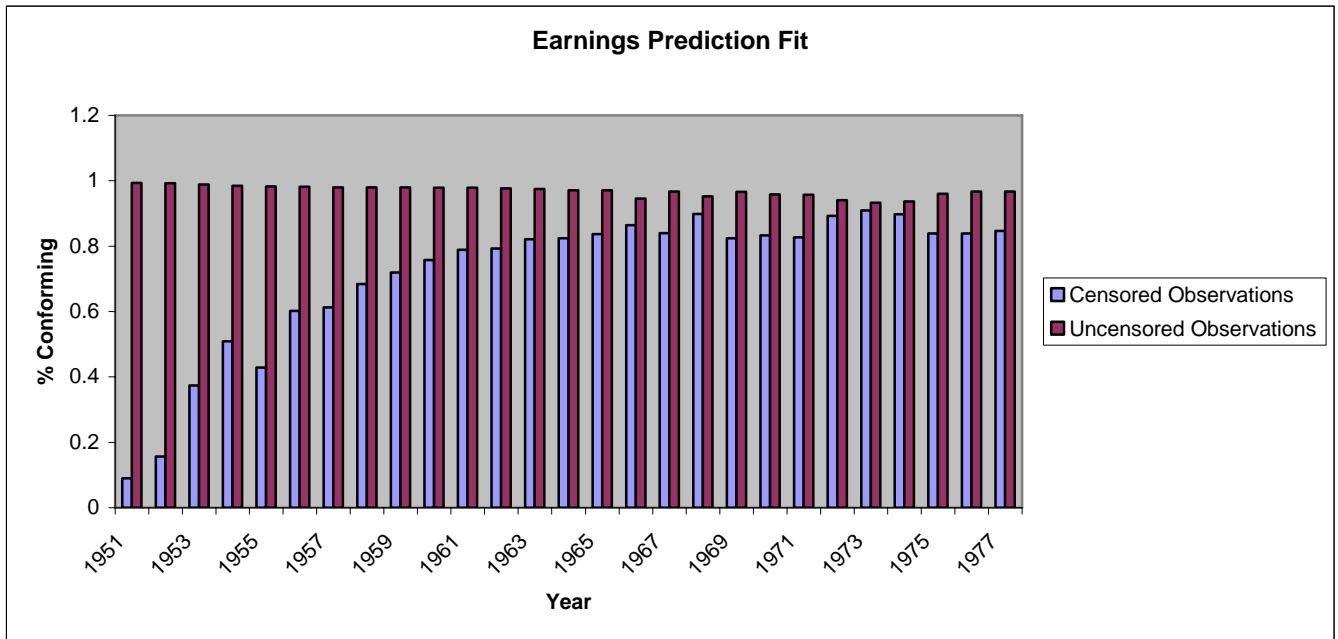
Tobit regressions of earnings on indicator variables for marital status, census regions, race and ethnicity, birth year, gender and education group. To add a dynamic element to the earnings imputations, we include variables for the household's position in the aggregate earnings distribution in each of the preceding 4 years. We replace capped earnings in cases where the predicted earnings from the regression exceed capped earnings.

Table 3 (Dollar Amounts in 2004 Dollars)			
Variable	Mean Value		Difference in Mean p-value
	Release	Did Not Release	
2004 earnings	\$16,780	\$24,766	0.000
Net worth	\$469,667	\$491,913	0.420
Mean age (years)	65.77	63.17	0.000
Mean education (years)	12.72	12.90	0.002
Fraction male	0.45	0.47	0.021
Fraction black	0.09	0.11	0.000
Fraction Hispanic	0.06	0.09	0.000
Fraction couple	0.66	0.67	0.782
No high school diploma	0.24	0.20	0.000
High school diploma	0.48	0.49	0.957
College graduate	0.13	0.15	0.001
Post college education	0.10	0.11	0.125
Fraction self-employed	0.10	0.12	0.005
Fraction partially or fully retired	0.52	0.36	0.000
Observations	14,598	5,531	
<p>SOURCE: Authors' calculations from the 2004 HRS. The table is weighted by the 1992 HRS person level analysis weights. Two sided difference in mean test: $H_0 : \mu_R - \mu_{NR} = 0$.</p>			

The predictions exceed the capped amounts for 79 percent of the capped observations. Of the uncensored observations, 97 percent of the predictions are below the cap. Predictions in early years suffer from the lack of earnings history as the predictions condition on location in the earnings distribution over the preceding 4 years.

Another way to test the accuracy of the imputation method is to impose a cap, \hat{c} that is lower than the actual cap, C . It is then possible to follow the imputation procedure and check its

accuracy against the true earnings data in the interval $[\hat{c}, C]$. The second chart shows the mean prediction error in 2004 dollars for each year from 1951 to 1977 when an artificial cap is imposed 10 percent lower than the actual cap. The error appears biased downward in early years and biased upward in later years while predictions are typically within \$4,000 of true annual earnings.



Steps to Develop the Analysis Sample

Our analysis starts with the RAND HRS Data

(<http://hrsonline.isr.umich.edu/modules/meta/rand/index.html>, the site requires registration with the HRS), which aggregates HRS data for respondents and spouses across waves into a single analysis file with consistent variable definitions across waves. We add core HRS measures to the RAND data, including information about child ages, defined contribution pension benefits from past and current jobs, and defined benefit pension coverage from past and current jobs. The restricted access Social Security earnings data are added and we impute earnings for censored observations as described above. We transform the data from respondent level into household level and add household-level pension data. We calculate Social Security benefits for each member of the household based on the earnings history for respondents, respondents' living spouses and deceased spouses.

The initial 2004 HRS sample has 8,513 households that provide Social Security earnings histories. As explained earlier, these earnings histories are a necessary input to our optimal wealth calculations. Besides requiring Social Security earnings data, we make one additional strong sample restriction. In this draft we focus on households where the primary earner is retired (and the self-reported year of retirement is 2004 or earlier) and at least one adult in the household has at least 30 years of Social Security earnings. The resulting sample has 2,996 households.

There is no compelling reason for this sample restriction. Our suspicion is that those who do not meet the 30-years of earnings condition were more likely than others to have some period of disability or other sources of income not covered by the social security system. Because we do not have information on transfer payments prior to 1992 (the first wave of the HRS) or on

earnings outside the Social Security system, we will not have good measures of lifetime earnings. Our sample restrictions also minimize the importance of earnings imputations. But subsequent versions of the paper will include all households with restricted social security earnings data.

Optimal Wealth Accumulation Across HRS Cohorts

Table 4 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. Recall that each household is retired (in this paper draft). 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 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.

The first column of Table 4 shows the median level of optimal wealth for households in each HRS cohort. We do not report data for Early Baby Boomer households, since it is very unusual for households in this cohort to be retired and few have 30 years of earnings. The optimal amounts peak for the HRS cohort, and then it falls as households age, as households presumably decumulate wealth in retirement. Actual accumulation for retired households in our

analysis subsample, (as opposed to all households, which is given in Table 1), is shown at the bottom of the table.

Table 4: Optimal and Actual 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	\$128,718	\$343,219	23.3	\$23,713
CODA	174,420	412,539	25.0	32,280
HRS	228,405	485,439	26.9	37,595
War Babies	112,789	305,340	28.1	24,302
Full Sample	188,835	434,293	25.9	32,260

	Median Actual Wealth	Mean Actual Wealth
AHEAD	\$200,000	\$584,197
CODA	271,600	468,435
HRS	355,000	637,077
War Babies	212,000	388,883
Full Sample	300,600	584,197

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 War Babies (they are young), 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. Overall, 25.9 percent of households have net worth that is below their optimal targets. Conditional on not meeting the target, the magnitude of the deficit is \$32,260. Recall, 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.

These results differ somewhat from 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 (in 1992 dollars).¹⁴ Hence, we find a somewhat larger fraction of households below their optimal targets in this preliminary analysis than we did in our earlier analysis. The magnitude of the median deficit, conditional on having one, is also somewhat higher. There is not a great deal of variation across cohorts in the fraction of the sample that is undersaving.¹⁵

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 (these results are not shown, but are available on request). No covariate is significantly correlated with having wealth less than the optimal target (covariates include dummy variables for income decile, cohort, education, race, children, marital status, and age in years). This result is similar to our earlier results (SSK, 2006) where no covariate besides marital status was significantly correlated with failing to meet the optimal targets. In our earlier results, single people were much more likely than married couples to have wealth below their optimal targets.

¹⁴ The CPI-U is higher by a factor of 1.346 in 2004 relative to 1992.

¹⁵ Readers should not place much weight on the “War Babies” results, since their average age is 59. Those in this cohort that are retired may be more likely than usual to have had some negative shock that forces them out of the labor market. Consequently, we might expect to see a somewhat larger fraction of these households with wealth below their optimal targets.

As is clear from Table 4, 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).

The results are given in Table 5. High income households are more likely to exceed their optimal net worth target.¹⁸ Households in the ninth (tenth) decile of their cohort income distribution have a median wealth surplus that is \$73,745 (\$119,324) higher than households in the first decile. The surpluses increase with educational attainment. African American households are likely to have lower median wealth surpluses, while households in the HRS cohort tend to have higher surpluses. 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.

¹⁸ 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).

Table 5: Correlates of Median Wealth Surplus

	Coefficient	Standard Error	T-statistic
Constant	-59,145	37,029	-1.60
Income decile 2	6,325	2,852	2.22
Income decile 3	6,797	3,832	1.77
Income decile 4	7,446	5,639	1.32
Income decile 5	28,433	5,995	4.74
Income decile 6	38,853	9,393	4.14
Income decile 7	41,606	8,208	5.07
Income decile 8	46,933	10,899	4.31
Income decile 9	73,745	16,224	4.55
Income decile 10	119,324	14,923	8.00
CODA cohort	5,966	6,380	0.94
HRS cohort	15,906	7,423	2.14
WB cohort	11,364	12,417	0.92
EBB cohort	10,237	17,532	0.58
Age of head	732	441	1.66
African American	-7,315	2,494	-2.93
Other non-white	-6,824	4,552	-1.50
Married	2,464	4,326	0.57
GED	-2,701	3,789	-0.71
HS grad	8,976	4,021	2.23
Some college	9,408	5,234	1.80
College or More	71,786	12,266	5.85
One child	2,062	4,928	0.42
Two children	4,735	7,140	0.66
Three children	879	6,479	0.14
Four children	-1,055	5,593	-0.19
Five children	468	6,816	0.07
Six or more children	-4,395	4,747	-0.93

3. 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 preparing for retirement. This perception is reinforced by low (and sometimes negative) personal saving rates in the National Income and Product Accounts for selected years over the past decade, 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.

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 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 a broader age range of households. 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).¹⁹ 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 suggests that 25.9 percent of HRS households have net worth that is below their optimal targets. Conditional on having accumulated too little, the magnitude of

¹⁹ We again emphasize the sample restriction in this draft that the primary earner in each household is retired and at least one household member has at least 30 years of earnings in the restricted social security data. These restrictions will be relaxed in subsequent drafts.

the deficits is \$32,210. While the results suggest that some households will need to ratchet their living standards downward in retirement, most Americans are, by in large, preparing sensibly, 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. 2009. "Children and Household Wealth." Mimeo, <http://www.ssc.wisc.edu/~scholz/Research/Children.pdf>