Are Retirees Falling Short? Reconciling the Conflicting Evidence

Alicia H. Munnell, Matthew S. Rutledge, and Anthony Webb

September 2014

PRC WP2014-05
Pension Research Council
The Wharton School, University of Pennsylvania
3620 Locust Walk, 3000 SH-DH
Philadelphia, PA 19104-6302
Tel.: 215.898.7620 Fax: 215.573.3418
Email: prc@wharton.upenn.edu
http://www.pensionresearchcouncil.org

The research reported herein was pursuant to a grant from the US 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, the RRC, or Boston College. The authors would like to thank Natalia Orlova and Rebecca Cannon Fraenkel for major contributions to the research and Dina Bleckman for timely research assistance. All opinions, errors, findings, interpretations, and conclusions of this paper represent the views of the authors and not those of the Wharton School or the Pension Research Council.

Are Retirees Falling Short? Reconciling the Conflicting Evidence

Abstract

A fundamental question in the retirement area is whether people will have adequate retirement income to maintain their pre-retirement standard of living. Existing studies offer conflicting assessments; some indicate a serious problem while others present an optimistic view. This chapter attempts to explain why the assessments differ. We find that the optimistic views of retirement preparedness depend crucially on behavioral assumptions that may not reflect real world activity or on consumption levels that are unsustainable in the long run. Thus, our best assessment is that retirees are falling short and will fall increasingly short over time.

Keywords: retirement, retirement income, retirement security, saving, consumption, replacement rate, Social Security, 401(k), reverse mortgage

Alicia H. Munnell
Center for Retirement Research at Boston College
Hovey House
140 Commonwealth Avenue
Chestnut Hill, MA 02467
Email: munnell@bc.edu

Matthew S. Rutledge
Center for Retirement Research at Boston College
Email: rutledma@bc.edu

Anthony Webb
Center for Retirement Research at Boston College
Email: webbaa@bc.edu

Are Retirees Falling Short? Reconciling the Conflicting Evidence

A fundamental question in the retirement arena is whether people will have adequate retirement income to maintain their pre-retirement standards of living. Existing studies offer conflicting assessments. On the one hand, the remarkable stability in the ratio of wealth to income by age from ten Surveys of Consumer Finances indicates that current cohorts have substantially lower retirement assets relative to income in the United States than in the past. In the same vein, the National Retirement Risk Index shows that half of today's working households will not be able to maintain their pre-retirement living standards (Munnell, Webb, and Golub-Sass 2012). Another recent study using the Health and Retirement Study (HRS) and similar assumptions also concludes that about half of pre-retirees are not on track to maintain their pre-retirement consumption (Munnell, Orlova, and Webb 2013). By contrast, Scholz, Seshadri, and Khitatrakun (2006), using the HRS and a life-cycle model of optimal wealth accumulation and decumulation, conclude that the majority of pre-retirees do have an optimal level of wealth. Separately, Hurd and Rohwedder (2013), using HRS consumption data, find that households who retired between 2001 and 2007 experienced only small declines in consumption, suggesting adequate resources.

This chapter attempts to explain why the different approaches yield such different answers. First, we present assumption-free evidence on wealth-to-income ratios from the last ten Surveys of Consumer Finances. These ratios by age have remained virtually unchanged even though people are living longer, Social Security is less generous, 401(k) plans have replaced defined benefit (DB) plans, health care costs have increased dramatically, and interest rates have plummeted. The second section describes two studies that compare projected replacement rates with a consumption-smoothing target. We show that roughly half of US households are unlikely to maintain their pre-retirement standards of living. In a third section we argue that the optimal saving

conclusion that emerges from the Scholz, Seshadri, and Khitatrakun model rests on two key assumptions: (1) households are content with declining levels of consumption in retirement; and (2) households reduce their consumption when children leave home. These assumptions make it much easier for older households to achieve target levels of wealth. To show the impact of these assumptions, we incorporate them in the National Retirement Risk Index. Our results show that these changes reduce the percent of households with inadequate savings produced by the optimal savings model. The fourth section turns to households' ability to maintain consumption after retirement. The questions here are whether the households studied by Hurd and Rohwedder possess sufficient resources to maintain their spending for the remainder of their lives, and whether they maintained their spending as they aged.

Our conclusion is that an optimistic view of US retirement preparedness depends crucially on assumptions about behavior that may not reflect real world activity or consumption patterns that are unsustainable in the long run. Thus, our best assessment is that many retirees will fall increasingly short.

Wealth-to Income Ratios Show Declining Preparedness over Time

While the *adequacy* of current US saving may be open to question, the *trend* in retirement saving relative to income is not. The Federal Reserve's triennial Survey of Consumer Finances (SCF) shows that the ratio of net wealth to income, a good indicator of the extent to which people can replace their earnings in retirement, has remained virtually unchanged at each age from 1983 through 2010.[1] In these ratios, wealth includes all financial assets, 401(k) accumulations, and real estate less any outstanding debt, and income includes earnings and returns on financial assets; importantly, wealth excludes the present expected value of income that the household will

eventually receive from DB pension plans and Social Security.[2] As shown in Figure 2.1, the ratios at each age for each survey lie virtually on top of one another. The outlier is 2010, where the ratios are substantially below those in the other surveys at every age.

Figure 2.1 here

The stability of the ratio reveals a significant decline in retirement preparedness, particularly since five developments should have led to higher ratios of wealth to income. First, life expectancy has increased. Between 1983 and 2010, life expectancy at age 65 rose by 3.5 years for men and 1.8 years for women. As a result, for any given level of income, one would have expected workers to accumulate more wealth in order to support themselves over their longer retirement period. Second, Social Security replacement rates have declined as the Full Retirement Age has been rising from 65 to 67, and the actuarial reduction on benefits claimed early grows larger. Moreover, the growing prevalence of two-earner couples means that fewer households receive the spousal benefit. Third, US retirement plans have shifted over time from DB to defined contribution (DC). Whereas accruals of future benefits under DB plans are not included in SCF wealth, assets in 401(k) plans are included. The shift from unreported to reported retirement assets would have been expected to increase the wealth-to-income ratio. Fourth, health care costs have risen substantially and show signs of further increase. Out-of-pocket expenditures for premiums and copayments under Medicare Part B, the program that covers physician services, have risen from 6.8 percent of the average Social Security benefit in 1983 to 17.0 percent in 2010, and they are projected to climb further in the future.[3] The rising cost of health care should have led to higher wealth-to-income ratios today than in the past. Finally, real interest rates have fallen significantly since 1983, so any given level of wealth now produces less retirement income. If people were

interested in generating a given stream of income, the significant decline in interest rates would have been expected to boost wealth accumulations.

Despite these developments, the stability of wealth-to-income ratios over the 10 SCF surveys between 1983 and 2010 indicates that people are less well prepared than in the past. If they had been over-prepared in the past, they might be fine in the future. But if people were not over-prepared in 1983, then retirees will fall short now.

Failure to Meet Replacement Rate Targets

Some analysts conclude that roughly half of today's working households will fall short in retirement. One approach uses the SCF to construct a National Retirement Risk Index (NRRI); another uses the HRS to calculate a baseline level of retirement income for people in their 50s against which one can measure the impact of alternative strategies such as changing asset allocations, working longer, taking out reverse mortgages, or controlling pre-retirement spending.

The National Retirement Risk Index. The NRRI compares projected replacement rates—retirement income as a percent of pre-retirement income—for today's working-age households to target rates that would allow them to maintain their living standards in retirement. It then calculates the percent at risk of falling short.[4] This calculation assumes that workers retire at 65 (a conservative assumption given that the average retirement age today is lower) and annuitize all their wealth (assumptions that maximize the income from any given amount of wealth).[5] The NRRI was initially constructed using the Federal Reserve's 2004 SCF, and it has been updated to incorporate the 2007 and 2010 surveys. The numbers presented below are based on an update for 2013.

Projecting household retirement replacement rates involves two calculations: estimating how much income households had before age 65, and estimating their post-retirement income. Retirement income is defined broadly to include income from financial assets in 401(k) plans and saved directly (net of non-mortgage debt), housing (net of mortgage debt), DB plans, and Social Security.[6] The approach uses the stable wealth-to-income patterns in the SCF to project where today's younger households are likely to be in terms of future financial and housing wealth. Instead of estimating total wealth directly, however, each component is projected separately. As noted, households are assumed to purchase an inflation-adjusted annuity with their financial wealth and with the proceeds of a reverse mortgage on their homes; retirement income includes imputed rent from the house.[7]

The items that comprise pre-retirement income include labor market earnings, the real return (assumed to be 4.6 percent) on 401(k) plans and other financial assets, and imputed rent from housing, minus mortgage and non-mortgage interest paid. Average pre-retirement lifetime income serves as the denominator for each household's replacement rate. With projections of pre- and post-retirement income, it is possible to calculate the projected replacement rate for each household when the head reaches 65.

Since the SCF does not include wealth from DB pensions and Social Security, the incomes from these two sources must be estimated separately. For DB pension income, our projections are based on income amounts reported in the SCF. For Social Security, benefits are calculated directly based on earnings histories constructed for all members of the household.

To determine the share of the retired population at risk, projected replacement rates are compared with pre-tax target replacement rates. The target permits the household to enjoy the same consumption in each period both before and after retirement. The calculation involves identifying

an age-varying saving rate that results in an accumulation of wealth at age 65 sufficient to generate a level of post-retirement consumption that equals pre-retirement consumption. The calculations incorporate current federal and state income taxes, the Earned Income Tax Credit, Social Security taxes and benefits, house purchase, mortgage borrowing, and loan repayments. The target replacement rate is the ratio of post-retirement income to pre-retirement income associated with the optimal saving rate. The calculation of the level of retirement wealth sufficient to smooth consumption is sensitive to assumptions regarding investment returns, medical and nursing home costs, the scope for economizing in retirement, and the age of retirement. Nevertheless, Skinner (2007) shows that plausible alternatives to our approach yield higher ratios of required retirement wealth to pre-retirement income compared to ours.

The final step in creating the Index is to simply compare each household's projected replacement rate with its target. Those whose projected replacement rates fall more than ten percent below the target are deemed to be at risk of having insufficient income to maintain their pre-retirement living standards. Thus, the Index is the fraction of *all* households that fall more than ten percent short of that target. The percentages of households at risk for three age groups in 2007, 2010, and 2013 appear in Table 2.1. The percent at risk in 2013 declined only slightly from 2010, primarily because the strong performance of the stock market offset the tightening of the reverse mortgage provisions and the continued increase in Social Security's Full Retirement Age. Overall, the conclusions are clear: half of today's working-age households are at risk of not being able to maintain their standard of living in retirement.

Table 2.1 here

Results from the HRS. A second analysis was designed to compare the effectiveness of asset allocation to other levers—controlling spending, working longer, and tapping home equity—in

improving retirement security (Munnell, Orlova, and Webb 2013). That study used data on working households aged 51–64 from the HRS, a nationally representative panel survey of older households conducted every two years since 1992.[8] Determining the retirement readiness of each household involved three steps.

The first identified a target replacement rate that allowed each household to maintain its pre-retirement standard of living. Target replacement rates that vary with marital status and income were drawn from Georgia State University's RETIRE Project (Palmer 2008). These rates were adjusted to reflect a projection that a significant proportion of the sample will have either repaid their mortgages by retirement, or be able to repay all or part of the balances by drawing on financial assets.[9]

The next step was to calculate the projected retirement replacement rate that each household would achieve if it continued on its present course, maintaining its current savings rate and asset allocation, and not taking a reverse mortgage. Total income at retirement in this baseline scenario consisted of Social Security, employer pensions, and income from financial assets. Social Security benefits were based on administrative data and income from current and deferred pensions on self-reports.[10] At retirement, the household was assumed to purchase a nominal joint- or single-life annuity with its financial assets, including 401(k) and IRA balances.

In a third step, the projected replacement rate was compared to the target rate for each household at each age from 60 to 70.[11] If the projected rate was below the target, the household was deemed to fall short in its retirement preparedness. The aggregate result for all households at a given age of retirement was the baseline measure for assessing the impact of the levers.

Figure 2.2 shows the percent of households falling short at each age under both the baseline assumption, and also assuming that the household took advantage of a reverse mortgage. The

baseline results show substantially more households falling short at 65 than the NRRI, but once they take out a reverse mortgage (as do households in the NRRI) the share falling short at 65 drops to 48 percent—close to the NRRI in 2010 (53 percent).[12]

We conclude from these studies that a methodology which identifies a target replacement rate and compares that rate with projected retirement incomes finds almost half of American households are unable to maintain their pre-retirement incomes in retirement. Results from the SCF and the HRS are almost identical.

The Optimal Savings Alternative

In contrast to the target replacement rate approach, Scholz, Seshadri, and Khitatrakun (2006) concluded that most Americans were saving optimally, using HRS data. By this, they meant that households were saving more than enough to smooth lifetime marginal utility of consumption. They defined saving to include 401(k) plan balances and housing equity as well as direct saving. These conclusions were based on a model of optimal wealth accumulation and decumulation over the life cycle that incorporated mortality, labor market, and health cost risk, and income from DB pensions and Social Security. The model was used to calculate the wealth that HRS households should have accumulated by their 50s, given assumed preference parameters along with the health and employment shocks experienced over their lifetimes. These optimal amounts were then compared with the amounts that HRS households actually accumulated. Those authors concluded that only 16 percent of households had less wealth than optimal.[13]

The question is why this optimal savings approach yields such comforting results. Two factors are at play: the cohort analyzed and model assumptions. The two assumptions that stand

out are how households are assumed to consume their accumulated wealth in retirement, and how children affected replacement rate targets.[14]

Cohort analyzed. The optimal savings analysis found that only 16 percent of households saved less than optimally based on the original HRS cohort (age 51–61 in 1992). Interestingly, our NRRI measure for 1992 for those age 51–61 reported only 19 percent at risk. That is, focusing on the same age group in the same year, the two different methodologies yielded the same general conclusion: the vast majority of those households had saved enough to maintain pre-retirement living standards.

Between 1992 and 2004, however, the overall percent at risk increased substantially in the NRRI (see Table 2.2). The three main reasons are a decline in Social Security replacement rates because of a drop in one-earner couples and increases in the Full Retirement Age; a decline in real interest rates; and a shift from DB to DC pensions. The difference in households at risk between the NRRI and the optimization model in 2004 is substantial: the NRRI showed 35 percent of those 50–58 to be at risk, while Scholz-Seshadri (2008) showed only eight percent at risk.[15] To analyze this difference requires examining the implications of the underlying assumptions.

Table 2.2 here

Assumptions underlying optimal saving. As noted, two key assumptions are how much households spend in retirement and how households spend once their children leave home. To evaluate how these affect results, we recalculate the percent at risk of being unable to maintain their living standards during retirement in the NRRI.

Retirement drawdown. The NRRI assumes that people will purchase an annuity when they retire so that they can spend a steady inflation-adjusted amount.[16] By contrast, the optimization model assumes that households undertake an optimal drawdown of unannuitized wealth, carefully trading

off the risk of outliving their wealth against the cost of unnecessarily restricting their consumption. Given the preference parameters in Scholz, Seshadri, and Khitatrakun (2006), households optimally choose higher consumption in their 60s and significantly lower consumption by age 85. This declining consumption reflects the declining probability that the household will be alive.

Using numerical optimization techniques, and adopting Scholz, Seshadri, and Khitatrakun's (2006) preference, mortality, and rate of return assumptions, we compute that a typical married couple undertaking a drawdown of unannuitized wealth will optimally consume 7.9 percent of wealth in the first year of retirement.[17] By contrast, the inflation-indexed annuity purchased by NRRI households provides an income amount of only 5.15 percent of wealth.[18] Thus to meet any target replacement rate, a typical household following the optimal drawdown strategy would only need to accumulate 66 percent (5.2/7.9) of the wealth needed by NRRI households. Of course consumption will decline during the course of retirement, but this is optimizing behavior given the assumed intertemporal elasticity of substitution.[19]

Integrating an optimal drawdown strategy into the NRRI requires two changes. First, as discussed above, wealth annuitization must be replaced by a declining drawdown rate. Second, because consumption falls in retirement, households can consume more during their working years. To equalize pre-and post-retirement consumption, the targets need to be raised (see Figure 2.3), partially offsetting the effect of the declining drawdown rate on the percent saving sub-optimally.

Figure 2.3 here

The speed of the decline in spending also depends on marital status (married couples have the greatest life expectancy and single men the shortest) and the percentage of wealth that is pre-annuitized through employer pensions and Social Security (the higher the percentage, the more

rapid will be the decline). It also depends on the intertemporal elasticity of substitution,[20] which is assumed to equal 1/3, consistent with Scholz and Seshadri. Our procedure involves calculating, for each of 48 household types, a target replacement rate that will permit the household to enjoy the same level of consumption in each period prior to retirement and in the first year of retirement, and an optimally declining level of consumption in retirement. From these 48 household types, targets and drawdown rates are calculated for 12 types of households—low, medium, and high-earner single men, single women, one-earner couples and two-earner couples—by taking weighted averages of renters and homeowners with and without DB pension coverage.

We find that imposing an assumption of an optimal drawdown boosts resulting replacement rate targets, but this increase is more than offset by a reduction in the amount of wealth required to finance each dollar of post-retirement consumption. As a result, the percent of those in their 50s at risk in 2004 declined from 35 percent under the original NRRI assumptions to 24 percent after the NRRI is adjusted to reflect optimal drawdown strategies (see Figure 2.4).

Figure 2.4 here

Children. A second important assumption is what happens to household consumption once the children leave home. Under the optimal savings approach, where the marginal utility of consumption varies with the size of the household, households optimally choose lower consumption once the children are gone. As a result, the households have lower replacement rate targets and need to save less for retirement than in cases where consumption remains steady.

Incorporating the impact of children on consumption into the NRRI model requires recalculating the NRRI targets for each of the 48 types of households and collapsing those groups into 12 household types. The new calculation retains the assumption that households reduce their consumption as they age. Adjustments to the targets are based on the equivalence scales from

Scholz and Seshadri (2008), namely $(A_j+0.7K_j)^{0.7}$ where A_j and K_j denote the number of adults and children in the household. The assumption is that children are born when the parents are age 27 and remain in the household for 18 years. Following Scholz and Seshadri (2007), the number of children varies with marital status and income tercile. Figure 2.5 shows the optimal consumption path for a typical household before and after adjustments for children. The new target replacement rate permits the household to enjoy the same level of consumption prior to the arrival of children, after they have left home, and in the first year of retirement, and a higher level of consumption while the children are at home.

Figure 2.5 here

The percentage at risk is then re-estimated using these new targets that reflect both the optimal drawdown and the assumption that consumption varies with household size. The combined impact of the two adjustments reduces the percent at risk in 2004 from the original 35 percent to 11.5 percent, very close to the eight percent reported by Scholz and Seshadri (2008).

The upshot of this discussion is that the standard economic assumptions embedded in the optimization model lead to low wealth accumulation targets and high initial incomes. The notion that households accept declining consumption in retirement implies that households will accumulate much less wealth to maintain their living standards. The assumption that parents reduce their consumption once children leave home means that households have lower saving targets and save more between the emptying of the nest and retirement, yielding few at risk. These two assumptions are the key factors that differentiate the conclusions from the two models.

Which assumptions are more plausible? Do people want steady real consumption in retirement (especially in view of rapidly rising health care costs and expensive long-term care) or will they accept declining payments? Do parents cut back on consumption when the children leave

home, or do they spend the slack in their budgets? No one really knows the answer. While spending does decline as people age, the extent to which this pattern tracks declining income is unclear; people cannot spend what they do not have. On the side of steady consumption, financial planning tools invariably assume that households require level amounts. How households react when children leave home is also not well understood. One study found that household consumption did not decline, and per-capita consumption increased after children left (Coe and Webb 2010). Yet the sample size was small, so the effects of children remain an unresolved issue. Nevertheless, assumptions alter one's assessment of whether households are saving optimally, and one's view of the results depends on the plausibility of the underlying assumptions.

The Sustainability of Initial Retirement Consumption

Another argument in favor of the adequacy of current saving is that some researchers find only modest declines in total spending after retirement (Hurd and Rohwedder 2013). These data come from the HRS and a supplemental survey to the HRS, the Consumption and Activities Mail Survey (CAMS), conducted in 2001, 2003, 2005, and 2007 (and subsequently in 2009 and 2011). For most of their analysis, the authors use panel data from 439 households on actual spending before and after retirement, and on anticipated spending changes prior to retirement and recollections of spending changes after retirement. The authors find only small consumption declines (1–6 percent) post-retirement in total spending, non-durable spending, and food spending.[21] The authors surmised that these declines were consistent with the cessation of work-related expenses, more efficient shopping, or the loss of earnings due to early retirement as a result of a health shock.[22] Nevertheless, there was much heterogeneity. For instance, households in the lowest quartile showed large declines in consumption upon retirement—particularly when health

was a reason for retirement and when they had a short planning horizon. Nevertheless, the overall message was that observed steady consumption profiles suggest that people are well prepared for retirement.

A possible explanation for the absence of any significant drop in consumption at retirement despite apparently inadequate resources is that retirees may initially consume too much, but then experience a reduction in consumption over time. We test this hypothesis in three ways. The first asks whether the Hurd-Rohwedder households possessed sufficient resources in their first year of retirement to maintain the same consumption levels for the remainder of their lives.[23] The second examines whether these households continued to maintain their immediate post-retirement consumption after six to ten years of retirement. The third compares the trajectories of households having sufficient resources to maintain lifetime consumption with those of households having insufficient resources.

To calculate whether the 439 households possessed sufficient financial resources to maintain their first year retirement spending later in life, we take account of financial resources (both inside and outside of retirement plans), Social Security, and employer pension income, net of federal income tax. We also assume that households take a reverse mortgage on their houses. All proceeds of the financial assets and reverse mortgage are then used to purchase an inflation-indexed annuity. Although few households actually do purchase inflation-indexed annuities, the income they could obtain approximates the amounts that households might consume out of their financial assets were they to follow conventional rules-of-thumb patterns (such as the 4-percent rule, Bengen 1994). We note that this income is substantially less than the amount they could enjoy in the first year, if they had decumulated those assets optimally.[24]

Results of this annuitization exercise appear in Figure 2.6. A household with average income and consumption in each of the bottom seven deciles of the income distribution will not have enough money to maintain its first year's consumption.

Figure 2.6 here

A possible concern with these results could be the potential impact of resources and consumption misreporting. Even if all households did save optimally, those who understated their income or overstated their consumption would be classified as under-resourced, while those who overstated their income or understated their consumption would appear to have excess resources. To address this concern, we minimize the assumed mismatch in reporting of income and consumption. This exercise involves sorting households by income and then by consumption, and calculating households' ability to fund the corresponding point of the distribution of spending, at each point of the income distribution. Results show that households in all but two deciles have a consumption-income ratio in excess of one, indicating their inability to maintain initial levels of consumption (see Figure 2.7).

Figure 2.7 here

A second approach is to see what happens to total spending of CAMS respondents in the years after their first observation in retirement. The treatment group is respondents who reported they were not retired at CAMS wave *t* (2001, 2003, or 2005), retired at CAMS wave *t+1,* and remain retired thereafter, a maximum of an additional four surveys. The control group is respondents who reported that they were not retired at wave *t* and *t+1* and remained not retired thereafter. For this exercise, we tighten the definition of retirement from that used in the Hurd-Rohwedder analysis.[25] To be not retired in any wave, the respondent had to be both 'not retired' (the definition used by Hurd-Rohwedder) and 'working;' to be retired at any wave, the respondent

had to be both 'retired' (the definition used by Hurd-Rohwedder) and not 'working.'[26] Our results show a sharp decline in total spending among those who retired (see Table 2.3). By the end of the period, median spending was 23 percent lower than when respondents were working. By contrast, median spending among those who did not retire remained relatively constant.

Table 2.3 here

A possible explanation for the decline in consumption is that it could have been anticipated and hence reflects a high intertemporal elasticity of substitution. Nevertheless, the observed decline is much too large to be explained by optimal decumulation patterns. The average retirement age was 61, and households with the preference parameters assumed by Scholz, Seshadri, and Khitatrakun (2006) would optimally plan for consumption to decline by only eight percent over the subsequent decade. Alternatively, for it to be optimal for consumption to decline by 23 percent, households would have to be almost risk neutral. Furthermore, the two groups are approximately the same average age at time *t* (59 for the retirees and 61 for those who continued to work). Absent implausibly large differences in mortality risk, if it were optimal for the consumption of retired households to decline by 23 percent, it should also be optimal for the consumption of non-retired households to decline by about the same amount. In sum, it seems more likely that retiree households had to adjust consumption to fit within their lifetime budget constraints.

A third exercise compares the trajectories of households having sufficient resources to maintain their consumption for the remainder of their lives with those of households having insufficient resources. This exercise draws on the results of the first exercise and follows those with sufficient and insufficient resources over subsequent waves of the CAMS. We find that the spending for those with insufficient resources declined by 31 percent compared to 16 percent for

those with sufficient resources (see Table 2.4). While the relative response is what one would have expected, the sample sizes are too small to be definitive

Table 2.4 here

Our best guess, overall, is that people tend to maintain their pre-retirement spending when they first retire, but then they cut back sharply thereafter. This reflects the fact that, except for those with very high incomes, HRS respondents in the Hurd-Rohwedder sample lacked adequate resources to maintain their initial levels of consumption throughout their retirement year.

Conclusion

Of all the studies, the evidence we find the most convincing about retirement preparedness is the simple calculation of wealth to income by age from the ten Surveys of Consumer Finances. We show that these ratios remained unchanged over time, despite longer lives, declining Social Security replacement rates, the shift from DB to DC plans, rapidly rising health care costs, and low interest rates. In other words, that people retiring in the future will be less prepared than those in the past, and we are not surprised target replacement rate comparisons show that about half of households will be unable to maintain their standards of living in retirement.

In our view, studies showing that people are saving optimally hinge crucially on assumptions that imply people are willing to accept declining consumption as they age and sharply reduce their consumption when the children leave home. Applying these patterns to the National Retirement Risk Index reduces the percent at risk to levels very similar to the optimal savings studies. Yet the key question is whether these assumptions seem plausible. While bits and pieces of evidence exist on both sides of these assumptions, answers to this question are not yet conclusive. Moreover, the fact that consumption does not decline early in retirement ignores the

fact that many people do not have the resources to continue consuming at that pace over their entire retired lives.

Technical Appendix

This appendix describes the methodology used to calculate three replacement rate targets: (1) the target used in the National Retirement Risk Index, which assumes that the marginal utility of consumption does not vary with household size and that the household purchases an annuity on retirement; (2) a revised target that retains the assumption that the marginal utility of consumption does not vary with household size, but that households undertake an optimal decumulation of unannuitized wealth in retirement; and (3) a further revised target that assumes that the marginal utility of consumption varies with household size, and that households undertake an optimal decumulation of unannuitized wealth in retirement.

The NRRI target. The NRRI target assumes that the household's goal is to accumulate sufficient wealth to generate a level of post-retirement consumption that equals consumption immediately before retirement. The household achieves this goal by choosing an age-varying savings rate. The target replacement rate is the ratio of post-retirement income to pre-retirement income associated with the optimal savings strategy. Pre-retirement income equals labor market earnings, imputed rent, and investment returns, minus mortgage and loan interest paid, all averaged over ages 20 to 65. Post-retirement income equals income from Social Security, employer pensions, and an inflation-indexed annuity, plus imputed rent. The household is assumed to purchase an inflation-indexed annuity with its financial assets plus the proceeds of a reverse mortgage. The calculations include federal, state (Massachusetts), and Social Security taxes based on 2006 law, including the Earned Income Tax Credit and the favorable tax treatment accorded to income from Social Security.

Targets are calculated for 48 types of households—those in the bottom, middle, and top tercile of the income distribution who are single men, single women, or one- or two-earner couples,

with or without defined benefit pensions, and who are homeowners or renters. Weighted averages are calculated to yield targets for three income terciles for single men, single women, and one- and two-earner couples.

The calculations assume that both housing and financial assets yield a historical real return. The secondary earner joins the household at age 25 with zero assets. Married couples, single men, and single women face annuity rates of 4.69, 5.22, and 4.90 percent, respectively, corresponding to the income payable on inflation-indexed annuities for members of the 1956 birth cohort at 2004 interest rates and expense loads. At age 30, homeowners purchase a house valued at twice their age-50 earnings with the aid of a 30-year mortgage at a real interest rate of 2.23 percent. At retirement, homeowners can borrow 53 percent of the value of the house on a reverse mortgage.

Low, middle, and high earners experience wage increases derived from Clingman and Nichols (2004). The calculations assume that young households are able to borrow if their desired consumption exceeds their net income. They also assume that, during their working lives, households optimally choose a level of consumption that increases at the rate of one percent a year.

Target adjustment for optimal accumulation. If the marginal utility of consumption does not vary with age or health status and if the rate of interest equals the rate of time preference, households that choose not to purchase an annuity will optimally choose declining consumption in retirement.[27] The rate of decline will be governed by mortality risk and the intertemporal elasticity of consumption. If the initial level of consumption, expressed as a percent of wealth at retirement, is higher than the income obtainable on an inflation-indexed annuity, households choosing an optimal decumulation of unannuitized wealth will need to accumulate less wealth than households purchasing an annuity. They will therefore be able to enjoy higher pre-retirement income and will face higher initial replacement rate targets. Assuming zero mortality before

retirement, the optimal target is one associated with an age-varying savings rate that permits the household to enjoy the same consumption in the first period of retirement as it enjoyed during its working life, with an optimal decline in consumption thereafter.

The optimal savings rate and decumulation path is the one that maximizes:

$$\sum_{t=20}^{T}\sum_{m=1}^{m=3} B^{t-20}\rho_{m,t}U(C_{t,m}) \qquad (1)$$

where *m* is marital status (married, surviving male, or surviving female), $\rho_{m.t}$ is the probability of being in marital status *m* at time *t*, and $C_{t,m}$ is consumption at time *t* in marital state *m*. Mortality is assumed to be zero prior to age 65, and households are assumed to have population mortality for the 1956 birth cohort after age 65. Following Scholz, Seshadri, and Khitatrakun (2006), the rates of interest and time preference are both assumed to equal four.

The household's utility function is of the following form:

$$U_m(C_t^m, C_t^f) = \frac{(C_t^m + \lambda C_t^f)^{1-\gamma}}{1-\gamma}, U_f(C_t^f, C_t^m) = \frac{(C_t^f + \lambda C_t^m)^{1-\gamma}}{1-\gamma} \qquad (2)$$

where λ measures the jointness of consumption, C_t^m, C_t^f denote the consumption of the husband and wife at time *t*, and γ is the coefficient of risk aversion, assumed to be three, following Scholz, Seshadri, and Khitatrakun (2006). When λ equals one, all consumption is joint. When λ equals zero, none of the household's consumption is joint. We assume that λ equals 0.5.

Target adjustment for children. Scholz, Seshadri, and Khitatrakun (2006) assume that households attempt to smooth the marginal utility of per-adult equivalent consumption, calculated as follows:

$$\frac{(A_j + K_j)^{0.7}}{(A_j)^{0.7}} \qquad (3)$$

where A_j and K_j are the number of adults and children in the house. In their model, the number of children and parental ages at which the children are born varies with income decile. Children are incorporated into the NRRI by assuming that households choose an age-varying savings rate and a replacement rate target that permits the household to enjoy constant per-adult equivalent consumption during its working life and an optimal decline in consumption in retirement. Children are assumed to arrive at age 27 and stay for 18 years. Bottom-tercile households are assumed to have the average number of children of households in the second income decile, middle-tercile households the average for the fifth decile, and top-tercile households the average for the ninth decile. The savings plan and target replacement rate are those that maximize equation (1), with the utility function modified as follows:

$$U_m(C_t^m, C_t^f) = \frac{(\delta C_t^m + \delta\lambda C_t^f)^{1-\gamma}}{1-\gamma} \frac{(\delta C_t^f + \delta\lambda C_t^{fm})^{1-\gamma}}{1-\gamma} \qquad (4)$$

where δ is $\frac{(A_j)^{0.7}}{(A_j + K_j)^{0.7}}$

If the optimal strategy is to save zero financial assets, the target is set at the level that the household will achieve if it saves nothing, but arrives at retirement debt-free.

Appendix Table 2.1 reports weighted average targets for low-, middle-, and high-tercile single men, single women, one-earner couples and two-earner couples calculated under the assumption that households: (1) annuitize; (2) undertake an optimal drawdown of unannuitized wealth; and (3) choose higher consumption when the children are at home.

Appendix Table 2.1 here

References

Banks, J., R. Blundell, and S. Tanner (1998). 'Is There a Retirement-Savings Puzzle?' *American Economic Review*, 88(4): 769–788.

Bengen, W. P. (1994). 'Determining Withdrawal Rates Using Historical Data,' *Journal of Financial Planning*, 7(4): 171–180.

Bernheim, B. D., J. Skinner, and S. Weinberg (2001). 'What Accounts for the Variation in Retirement Wealth among US Households?' *American Economic Review*, 91(4): 832–857.

Clingman, M. and O. Nichols (2004). 'Scaled Factors for Hypothetical Earnings Examples under the 2004 Trustees Report Assumptions,' *Actuarial Note*, 2004.3. Washington, DC: US Social Security Administration.

Hurd, M. D. and S. Rohwedder (2011). 'Economic Preparation for Retirement,' Working Paper 17203. Cambridge, MA: National Bureau of Economic Research.

Hurd, M. D. and S. Rohwedder (2013). 'Heterogeneity in Spending Change at Retirement,' *Journal of the Economics of Ageing*, 1(2): 60–71.

Hurd, M. D. and S. Rohwedder (2008). 'The Retirement Consumption Puzzle: Actual Spending in Panel Data,' Working Paper 13929. Cambridge, MA: National Bureau of Economic Research.

Hurst, E. (2008). 'The Retirement of a Consumption Puzzle,' Working Paper 13789. Cambridge, MA: National Bureau of Economic Research.

Juster, F. T. and R. M. Suzman. (1995.) 'An Overview of the Health and Retirement Study,' *Journal of Human Resources*, 30(supplement), S7–S56.

Munnell, A. H., A. Webb, and L. Delorme. (2006.) 'A New National Retirement Risk Index,' *Issue in Brief*, 48. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, A. H., N. Orlova, and A. Webb. (2013.) 'How Important Is Asset Allocation to Financial Security in Retirement?' in O. S. Mitchell and K. Smetters, eds., *The Market for Retirement Financial Advice.* Oxford, UK: Oxford University Press, pp. 89–106.

Munnell, A. H., A. Webb, and F. Golub-Sass. (2012.) 'The National Retirement Risk Index: An Update,' *Issue in Brief,* 12-20. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Palmer, B. A. (2008.) '2008 GSU/Aon RETIRE Project Report,' Research Report Series 08-1. Atlanta, GA: J Mack Robinson College of Business, Georgia State University.

Scholz, J. K. and A. Seshadri (2008). 'Are All Americans Saving "Optimally" for Retirement?' Presented at the 10th Annual Joint Conference of the Retirement Research Consortium in Washington, DC, August 7–8.

Scholz, J. K. and A. Seshadri (2007). 'Children and Household Wealth,' Working Paper 2007-158. Ann Arbor, MI: University of Michigan Retirement Research Center.

Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006). 'Are Americans Saving "Optimally" for Retirement?' *Journal of Political Economy*, 114(4): 607–643.

Skinner, J. (2007). 'Are You Sure You're Saving Enough for Retirement?' *Journal of Economic Perspectives*, 21(3): 59–80.

US Board of Governors of the Federal Reserve System. *Survey of Consumer Finances,* 1983-2010. Washington, DC.

Venti, S. F. and D. A. Wise (2002). 'Aging and Housing Equity,' in O. S. Mitchell, Z. Bodie, P. B. Hammond, and S. Zeldes, eds., *Innovations in Retirement Financin*g. Philadelphia: University of Pennsylvania Press, pp. 254–281.

Endnotes

[1] For a selection of research papers using the SCF, see http://www.federalreserve.gov/econresdata/scf/scfindex.htm.

[2] The exact definition of income in the SCF includes wages, investment income, interest and dividend income, capital gains or losses, unemployment payments, alimony, welfare, pension income, and some other less common income; it is essentially all pre-tax income that comes into a household in a given year.

[3] This information comes from unpublished data from the Centers for Medicare and Medicaid Services, Office of the Actuary.

[4] For a fuller discussion of the National Retirement Risk Index, see Munnell, Webb, and Delorme (2006).

[5] Under a reverse mortgage, a homeowner borrows against the equity in his house and receives money from a lender. Unlike a home equity loan, no loan payments or interest are due until the individual dies, moves out, or sells the house. When one of these events occurs, the borrower or the estate is responsible for repaying the loan in full.

[6] The NRRI does not include income from work, since the goal of the Index is to assess the retirement security of households once they make a complete exit from the labor force.

[7] For couples, the annuity provides the surviving spouse two-thirds of the base amount. While inflation-indexed annuities are not popular with US consumers, they provide a convenient tool for converting a lump sum of wealth into a stream of income. And while inflation-indexed annuities provide a smaller initial benefit than nominal annuities, over time they protect a household's purchasing power against the erosive effects of inflation.

[8] For an overview of the HRS, see Juster and Suzman (1995).

[9] The target replacement rates used in the HRS study differ from those in the NRRI. These differences reflect not only the treatment of mortgage debt, referred to above, but also differences in the denominator. The denominator in the NRRI replacement rate target includes imputed rent and investment returns minus interest paid and is averaged over the household's working life. The denominator in the HRS study is labor market earnings averaged over the previous 10 years.

[10] Social Security benefits are calculated from actual or imputed earnings records when self-reported data are missing.

[11] For more details on the methodology, see Munnell, Orlova, and Webb (2013).

[12] This exercise assumes a nominal annuity, whereas the NRRI assumes the purchase of an inflation-indexed annuity. Nominal annuities provide a higher initial income than inflation-indexed annuities, increasing the likelihood that households will achieve their targets, albeit at the risk of falling below their targets later in retirement.

[13] An updated version of the analysis for 2004 showed even lower levels of households at risk (Scholz and Seshadri 2007).

[14] Two issues not considered here are the treatment of the house and the need for precautionary saving. The authors assume that housing wealth is available to fund consumption, so the household is in the same situation as a renter. To avoid implausibly rapid declines in non-housing consumption, households would have to move to a smaller and smaller house. Clearly people do not do this in the real world (Venti and Wise 2002), so the assumed decline in total consumption produces severe pressure on non-housing consumption. If housing represented one third of total consumption, non-housing consumption would decline by 45 percent from age 61 to 85. The other issue is the assumption that households invest in a risk-free asset yielding a real 4-percent return. It has not been possible to earn a risk-free 4 percent return since around the year 2000, so to earn

a 4-percent return households would need to assume investment risk. Under CRRA utility, households optimally respond to investment risk by undertaking precautionary saving and, averaged over draws of investment returns, optimally accumulate more than they would have accumulated in the absence of investment risk. The assumption of a 4-percent risk-free investment return leads to an overstatement of the percentage of households who have saved at least optimally.

[15] The eight percent is a weighted average of the 5.2 percent for households born from 1942 to 1947 and the 10.2 percent for those born from 1948 to 1953.

[16]Although households can also purchase nominal annuities that provide a declining real income, at plausible preference parameters and prevailing annuity prices, they will prefer the level consumption obtainable from a real annuity. 'Mortality credits,' the reallocation of resources from those who die to those who survive, enable annuities to pay higher returns than equivalent unannuitized investments, particularly at older ages. This additional return makes annuity purchasers more willing than non-purchasers to forego consumption at younger ages in return for additional consumption at older ages.

[17] This calculation ignores medical costs and assumes that one half of the household's wealth is pre-annuitized Social Security wealth. Scholz, Seshadri, and Khitatrakun (2006) assume a real return of 4 percent, a coefficient of risk aversion of three, and a rate of time preference of 3 percent. The calculation is based on population average mortality for the 1936 birth cohort when calculating the optimal drawdown rate for the Scholz, Seshadri, and Khitatrakun (2006) sample, who were mostly born from 1931–41. Subsequent birth cohorts would consume somewhat smaller percentages, reflecting their greater life expectancy.

[18] The NRRI calculations assume market expense loads, mortality improvements based on Social Security Administration projections, and that interest rates revert to levels prevailing in 2004. The

5.15 percent annuity rate is for households born before 1952. Subsequent birth cohorts are projected to face lower annuity rates.

[19] The assumed intertemporal elasticity of substitution, under the assumption of constant relative risk aversion, equals the inverse of the coefficient of risk aversion. If the assumed intertemporal elasticities of consumption were lower, post-retirement wealth paths would be flatter, requiring the household to accumulate more wealth to generate a given amount of first year retirement income. If, at the extreme, the households had an intertemporal elasticity of substitution of zero, they would optimally spend only the interest of 4 percent and would require 29 percent more wealth at retirement than the NRRI households.

[20] The intertemporal elasticity of substitution measures how responsive consumption growth is to the real interest rate.

[21] The finding of only a small drop contradicts earlier studies that identified a retirement-consumption puzzle (Hurst 2008; Bernheim, Skinner, and Weinberg 2001; Banks, Blundell, and Tanner 1998). It was considered a puzzle because the mechanism underlying the drop was not well understood.

[22] The authors also contend that the decline is consistent with the optimal decline in consumption implied by the traditional life cycle model, which they assert should equal the mortality rate, assuming the rate of time preference equals the interest rate. In fact, the prediction of the life-cycle model for households exhibiting constant relative risk aversion is that the optimal decline equals the mortality rate divided by the coefficient of risk aversion. This calculation implies that for a typical 61 year-old married couple with a coefficient of risk aversion of three, the decline in consumption over a two-year period would be one percent.

[23] Consumption includes both property taxes and purchases of automobiles and other durables, including automobile finance charges.

[24] On the other hand, it is more than the amount that they would consume were they to undertake an 'optimal' decumulation as in Hurd and Rohwedder (2011), who find that most newly retired households possess sufficient amounts.

[25] Hurd and Rohwedder (2013) impose the additional restriction that the respondent remained retired in the following wave if observed in that following wave.

[26] This adjustment was necessary because many of those retired were working and some of those working characterized themselves as retired. Respondents are also categorized as working if they report they are in part-time/less than part-time employment.

[27] This also sets aside the utility of leisure.

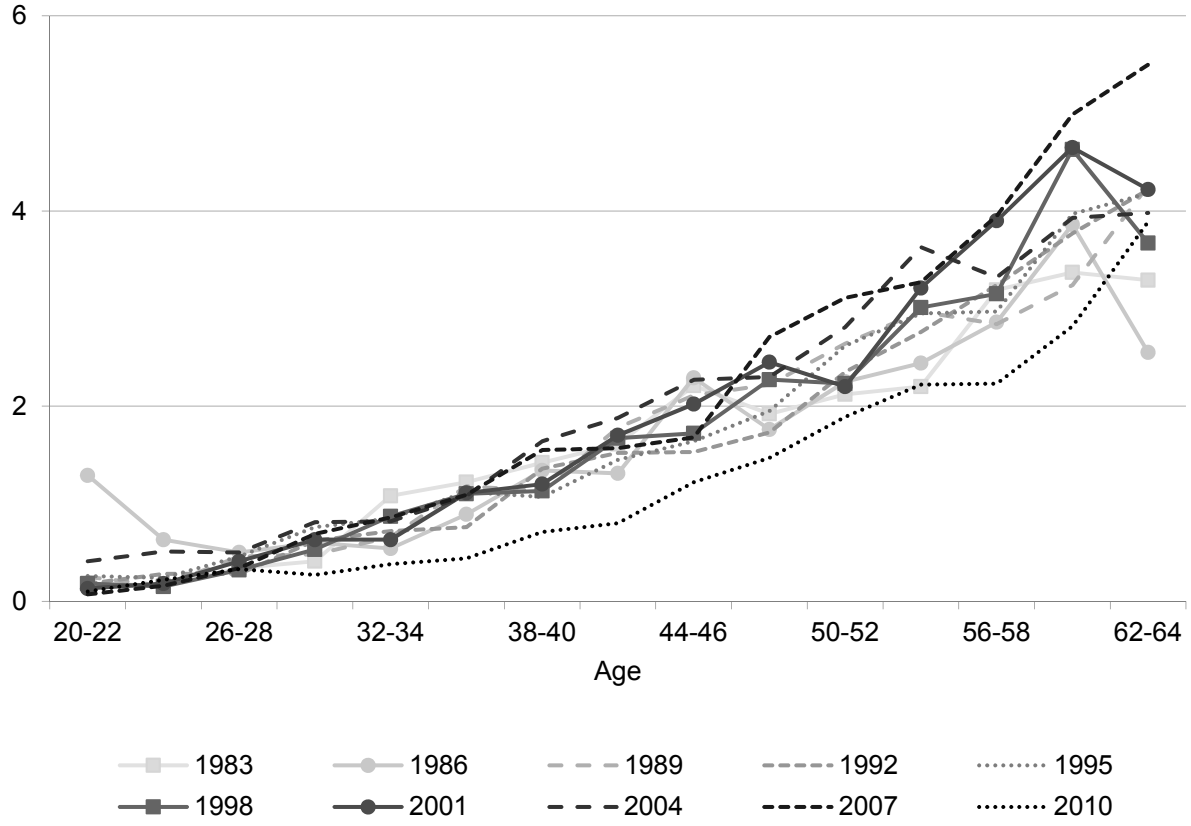

Figure 2.1. Ratio of wealth to income by age from the Survey of Consumer Finances, 1983-2010. *Source:* Authors' calculations based on U.S. Board of Governors of the Federal Reserve System, Survey of Consumer Finances, 1983-2010.

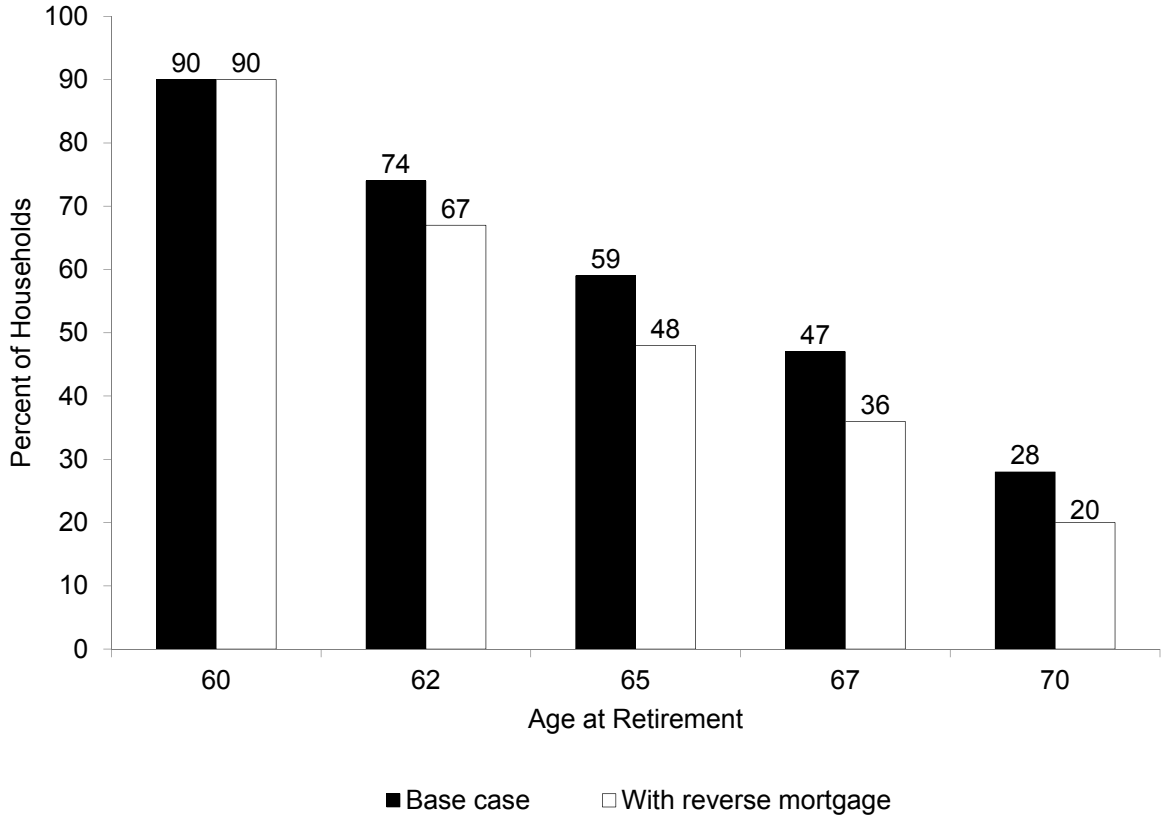

Figure 2.2. Percent of all households falling short of target by age at retirement, base case, and with a reverse mortgage. *Source:* Munnell, Orlova, and Webb (2013).

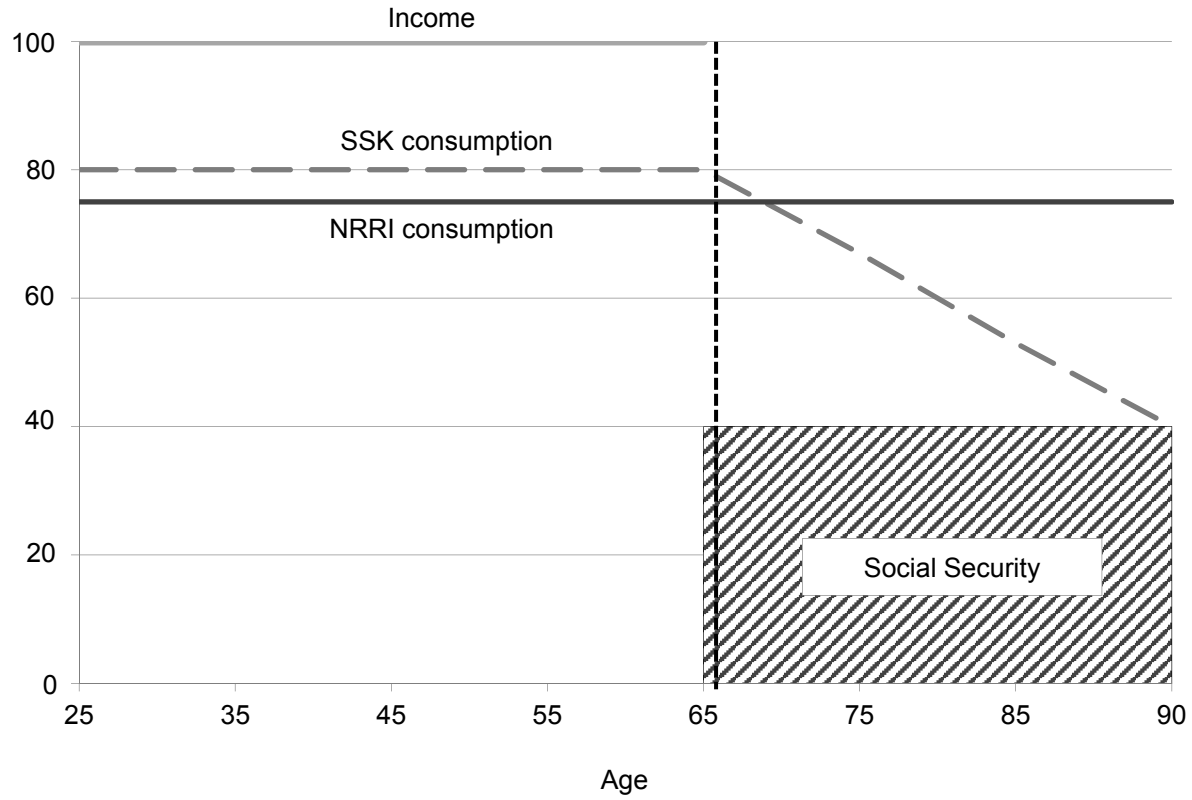

Figure 2.3. Illustrative consumption by age, SSK, and NRRI as percent of income. *Notes*: SSK refers to Scholz, Seshadri, and Khitatrakun (2006) computations, and NRRI refers to National Retirement Risk Index. *Source:* Authors' computations.

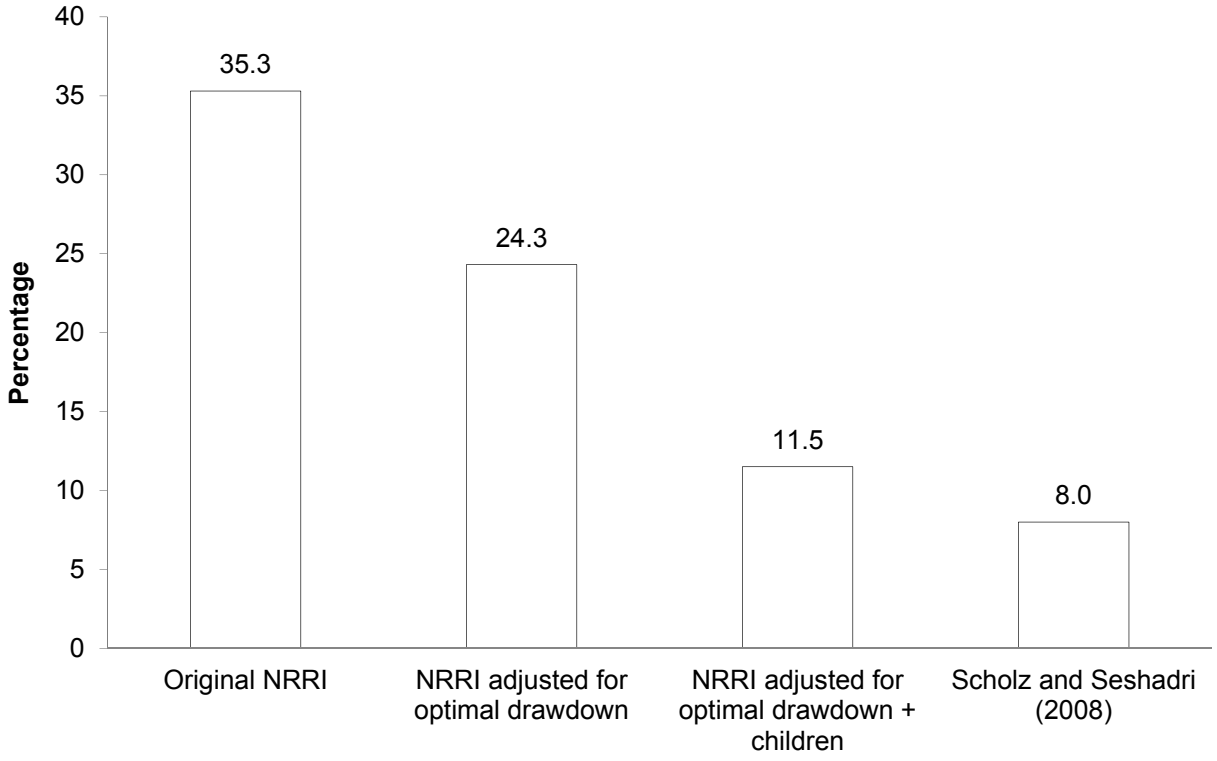

Figure 2.4. Percentage of households age 51–61 at risk, 2004. *Note:* NRRI refers to National Retirement Risk Index. *Source:* Authors' computations.

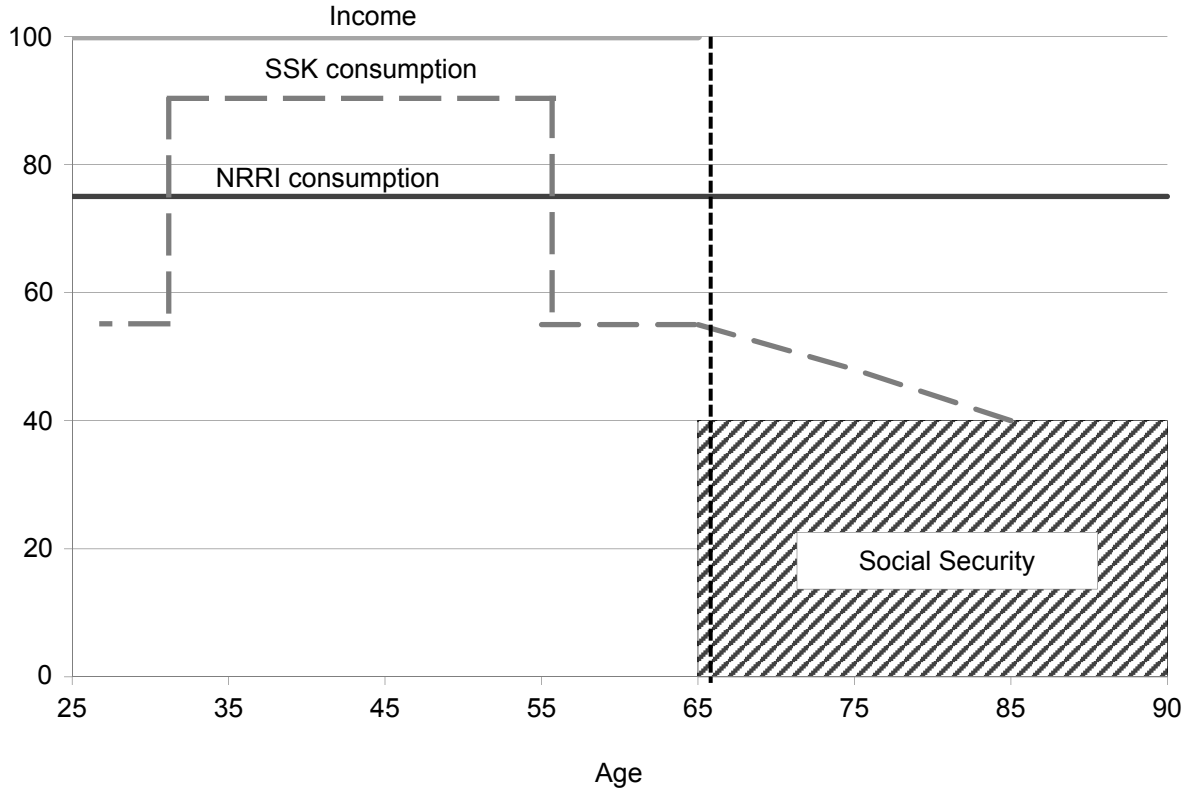

Figure 2.5. Illustrative consumption by age, SSK, and NRRI as percent of income, with children. *Notes:* SSK refers to Scholz, Seshadri, and Khitatrakun (2006) computations, and NRRI refers to National Retirement Risk Index. *Source:* Authors' computations.

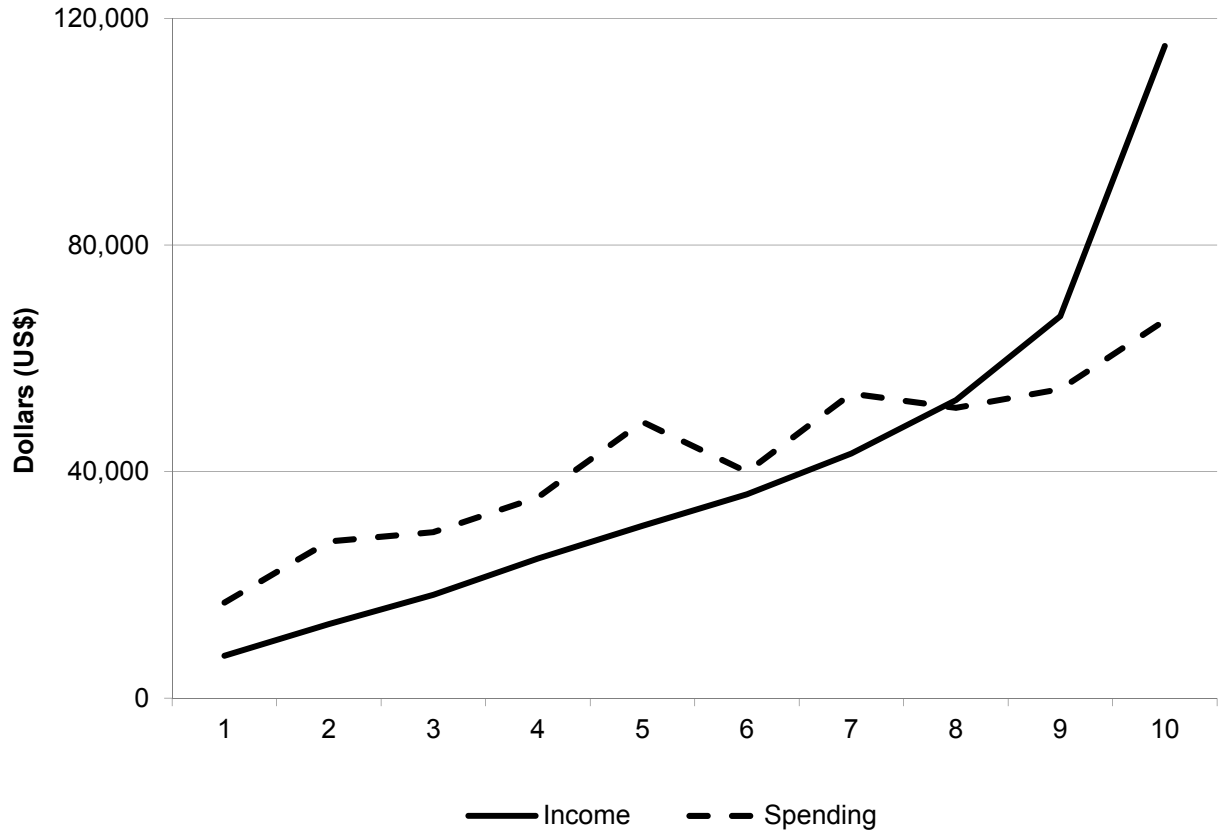

Figure 2.6. Mean income and spending by income decile at time of retirement for Hurd-Rohwedder (2011) sample of CAMS households. *Source:* Authors' calculations.

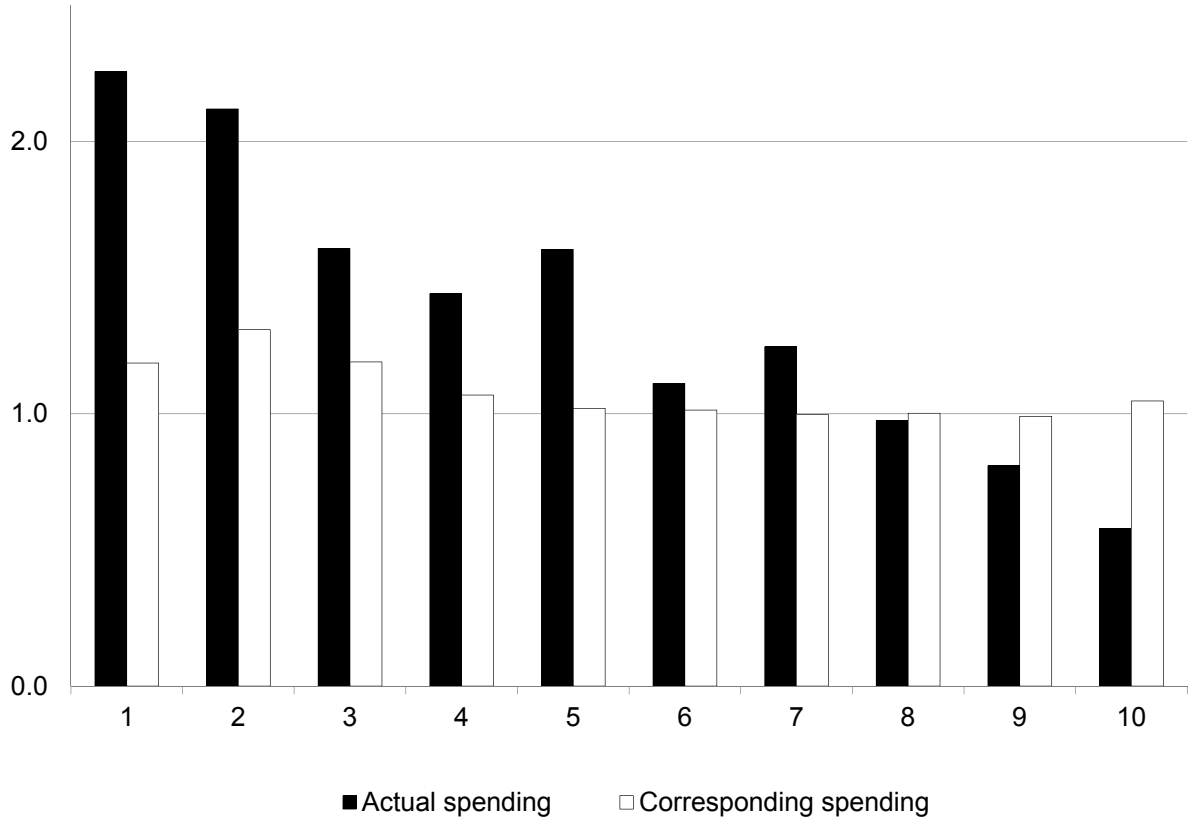

Figure 2.7. Average actual and corresponding consumption-to-income ratio at time of retirement by income decile for Hurd-Rohwedder sample of households. *Source:* Authors' calculations.

Table 2.1. Percent of households 'at risk' at age 65 by age group over time

Age group	2007	2010	2013
All	44%	53%	52%
30-39	53	62	59
40-49	47	55	54
50-59	32	44	44

Source: Authors' calculations.

Table 2.2. Percent 'at risk' using the NRRI versus the 'optimal saving' approach: 1992 and 2004

Age group	1992		2004	
	NRRI	Optimal savings	NRRI	Optimal savings
All groups	36	--	43	--
51–61	19	16	35	5

Note: The NRRI result for 2004 is for households age 50–58.
Source: Authors' calculations.

Table 2.3. Median respondent spending by CAMS retirement status for respondents age 50 to 70

Observation period	Not retired at time *t*, retired at time *t+1* and thereafter		Not retired throughout	
	Consumption	Sample size	Consumption	Sample size
t	23,100	275	25,900	1,424
t+1	25,300	275	25,600	1,424
t+2	21,000	208	24,500	898
t+3	21,000	193	24,300	677
t+4	18,300	123	24,200	291
t+5	17,700	71	25,600	148
Percent change				
From *t* to *t+1*	9.5		−1.2	
From *t* to *t+5*	−23.4		−1.2	

Notes: Respondents are classified as retired if they also reported that they were not working. Conversely, they are classified as not retired if they also reported that they were working. The comparison is between individuals who were not retired at t and retired at *t+1* and those who were not retired at both waves. The consumption of respondents is reported at subsequent waves if their retirement status is the same as that at *t+1*, irrespective of their retirement status at other waves subsequent to *t+1*. The difference in the percentage change from *t* to *t+5* between the two groups is statistically significant at the 5-percent level.
Source: Authors' calculations.

Table 2.4. Median respondent spending by retirement status for respondents age 50 to 70

Observation period	Not retired at time *t*, retired at *t+1* and thereafter			
	Insufficient	Sample size	Sufficient	Sample size
t	25,300	132	21,100	143
t+1	27,300	132	23,400	143
t+2	20,100	101	22,400	107
t+3	20,400	82	21,300	111
t+4	18,400	51	18,200	72
t+5	17,700	31	17,800	40
Percent change				
From *t* to *t+1*	7.9		10.9	
From *t* to *t+5*	-30.9		-15.6	

Note: See note for Table 2.3. The increase in the 'sufficient' sample size at *t+3* is the result of the inclusion of respondents who were retired at waves *t*, *t+1*, and *t+3*, but who were not retired at *t+2*.

Source: Authors' calculations.

Appendix Table 2.1. NRRI target replacement rates by household type and drawdown alternative

	Low income			Middle income			High income		
	Annuitization	Optimal drawdown	Optimal drawdown w/ children	Annuitization	Optimal drawdown	Optimal drawdown w/ children	Annuitization	Optimal drawdown	Optimal drawdown w/ children
One-earner couple	77	90	71	69	78	65	66	74	55
Two-earner couple	70	82	65	65	76	60	65	74	55
Single male	69	79	63	65	73	58	63	69	50
Single female	74	85	65	65	74	58	64	69	50

Note: The target replacement rates are weighted averages.
Source: Authors' calculations.