Estimating the True Cost of Retirement

David Blanchett, CFA, CFP®

Presented at the Living to 100 Symposium

Orlando, Fla.

January 8–10, 2014

Copyright 2014 by the Society of Actuaries.

All rights reserved by the Society of Actuaries. Permission is granted to make brief excerpts for a published review. Permission is also granted to make limited numbers of copies of items in this monograph for personal, internal, classroom or other instructional use, on condition that the foregoing copyright notice is used so as to give reasonable notice of the Society's copyright. This consent for free limited copying without prior consent of the Society does not extend to making copies for general distribution, for advertising or promotional purposes, for inclusion in new collective works or for resale.

Estimating the True Cost of Retirement

David Blanchett, CFA, CFP®
Head of Retirement Research
Morningstar Investment Management
22 W. Washington St., Chicago, IL
david.blanchett@morningstar.com
(859) 492-5637

Abstract

A common approach to estimating the total amount of savings required to fund retirement is to first apply a generic "replacement rate" to pre-retirement income, such as 80 percent, to get the desired retirement income need. That need is assumed to increase annually at the rate of inflation for the duration of retirement, which is generally assumed to be some fixed period, such as 30 years. Using government data along with a fairly simple market and mortality model, we explore these assumptions to more accurately estimate the true cost of retirement.

We find that the actual replacement rate is likely to vary considerably by retiree household, from under 54 percent to over 87 percent. We note that retiree expenditures do not, on average, increase each year by inflation or by some otherwise static percentage; the actual "spending curve" of a retiree household varies by total consumption and funding level. Specifically, households with lower levels of consumption and higher funding ratios tend to increase spending through the retirement period and households with higher levels of consumption but relatively lower funding ratios tend to decrease spending through the retirement period. When consumption and funding levels are combined and correctly modeled, the true cost of retirement is highly personalized based on each household's unique facts and circumstances, and are likely to be lower than amounts determined using more traditional models.

The author thanks Alexa Auerbach and Hal Ratner for helpful edits and comments.

Estimating the True Cost of Retirement

Estimating how much savings is needed for retirement is a complex calculation. In many cases, advisers or investors estimate the retirement income need by first applying a generic "replacement rate," such as 80 percent, to current or pre-retirement earnings, and assuming the retirement need increases annually by inflation over some fixed retirement period—generally 30 years. A discount rate or a more complex Monte Carlo simulation can then be applied to these cash flows to estimate the total amount of savings required at retirement to achieve success.

These three assumptions—the replacement rate, a constant real consumption level and fixed retirement period—are shortcuts that when combined can overestimate the true cost of retirement for many investors. Through analysis and using government survey data, we explore these assumptions to more accurately estimate the cost of retirement.

We find that:

- While a replacement rate between 70 percent and 80 percent may be a reasonable starting place for most households, the actual overall replacement rate can vary considerably, from under 54 percent to over 87 percent depending on the level of pre-retirement household income and, more importantly, expenses that discontinue after retirement.
- Real retiree expenditures aren't constant and they don't rise (or fall) in nominal terms simply as a function of broad-based inflation or expected health care inflation. The retirement consumption path, or "spending curve," will be a function of the household-specific consumption basket as well as total consumption and funding levels.
- Households with lower levels of consumption and higher funding ratios tend to have real increases in spending through retirement, while households with higher levels of consumption and lower funding ratios tend to see significant decreases. The implication is that households not consuming retirement funds optimally will tend to adjust them during the retirement period, i.e., spending is not constant in real terms.
- While many retirement income models use a fixed time period (e.g., 30 years) to estimate the duration of retirement, modeling the cost over the expected lifetime of the household and incorporating the actual spending curve can result in a required account balance at retirement that can be significantly less than the amount required using traditional models.
- The goal of many retirees is to create income for life, not income over a specific time period. When combined and correctly modeled, the true cost of retirement is highly personalized based on each household's unique facts and circumstances.

In section 1, we review the life-cycle hypothesis and its importance to retirement. In section 2, we review the literature on retirement spending. In section 3, we introduce a replacement rate model to demonstrate how the target household income varies based on different pre- and post-retirement considerations, and in section 4, we use Bureau of Labor Statistics' Consumer Expenditure Survey data to understand the spending habits of retirees and we explore some of the different definitions of inflation. In section 5, we use the dataset to estimate actual changes in consumption for retirees over time. In section 6, we combine the previous findings to better estimate the true cost of retirement, and in section 7 we conclude.

Section 1. Life-Cycle Hypothesis

Before exploring spending habits of retirees, it is important to explore why people save for retirement in the first place. While some forms of saving are required, such as the 6.2 percent employee portion of Social Security tax on earnings, other forms savings, such as in a 401(k) plan, are not. Savings allow a household to transfer consumption over time, i.e., by not consuming those monies today, the household can consume them at some point in the future. There are a number of economic and behavioral theories that have been brought forward to explain this. One of the most prominent is the life-cycle hypothesis (LCH), which was introduced initially by Modigliani and Brumberg (1954).

LCH implies that individuals maximize utility by planning savings and consumption such that lifetime consumption is as smooth as possible. People don't like risk, which is defined as the variability of consumption. The optimal savings and consumption schedule will vary by household and be determined by things like the utility parameters (elasticity of substitution through time, risk aversion), discount rate, mortality risk, expected future compensation and the like.

Consumption smoothing is a relatively simple concept if wages remain constant in real terms over the household's lifetime. For example, if the household earns $50,000 per year in after-tax wages each year while working (adjusted by inflation), the LCH would suggest the target after-tax income should be $50,000 per year during retirement. If we look at actual wages through time, though, we see compensation is not constant over someone's lifetime and tends to increase as someone ages. We see this income growth in figure 1, which includes the average lifetime growth in real wage in panel A and the average annual change in real wage in panel B, for varying levels of education.

Figure 1. Lifetime real earnings

Panel A. Average lifetime real earnings curve **Panel B. Average annual change in real wage**

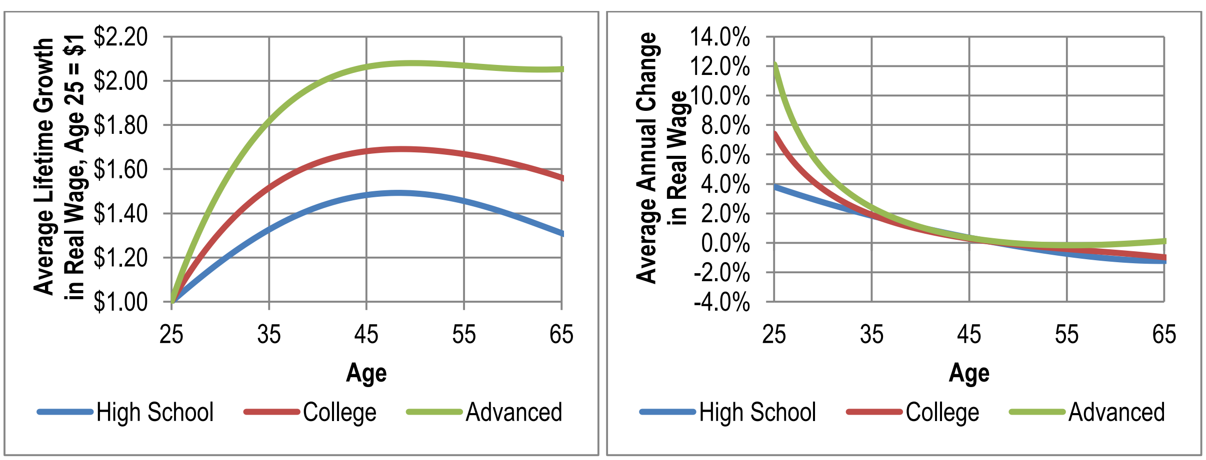

Figure 1 uses data from the Department of Labor's March Current Population Survey (CPS) from 1980 to 2011[1] (32 years). We removed outliers; for example, only workers making at least 75 percent of the federal minimum wage for the respective year were included. We separated workers into three groups: high school (either did or did not graduate), college (either attended some college or graduated from college) and advanced (post undergraduate education). For each of the three, we determined the median[2] compensation for each age, and fit a fourth order polynomial to the data to determine the earnings curve. This created a "smooth" earnings curve for each respective period. We then averaged the growth of each curve to create panels A and B in figure 1.

Figure 1 has important implications from a saving and spending perspective. For example, a college-educated individual will likely be making roughly 50 percent higher wages at retirement than he or she did at age 25. Therefore, a retirement income replacement analysis based on wages at age 30 will likely understate the actual total retirement need. Within the LCH model, there are also important implications about saving for retirement. If an individual is interested in truly smoothing consumption, then it may make sense to delay saving for retirement until age 35, which is when wages are higher.

---

[1] More specifically, the CPS Uniform Data Extract datasets prepared by the Center for Economic and Policy Research.

[2] The median is used versus the average because the average is highly skewed, especially at older ages.

Section 2. Literature Review

There is a growing body of literature exploring the spending habits and tendencies of retiree households. The majority of studies note consumption tends to decline at retirement, an effect commonly referred to as the "retirement consumption puzzle." This is in contrast with what we would expect based on LCH, whereby consumption would remain constant at retirement. The actual amount of the total change in consumption, though, varies materially across past research.

Banks, Blundell and Tanner (1998) was the first study to find a sharp decline in consumption at retirement using U.K. data, while Bernheim, Skinner and Weinberg (2001), using panel data[3] from the Panel Study of Income Dynamics (PSID), also found a drop in consumption at retirement. Also using panel data, Hurd and Rohwedder (2008) found that spending before and after retirement declines at a relatively small rate, from 1 to 6 percent depending on the measure. Research by Aguila, Attansio and Meghir (2011) noted that individuals tend to smooth consumption during the first year of retirement.

Ameriks, Caplin and Leahy (2007) analyzed responses to survey questions answered by TIAA-CREF participants about anticipated changes in spending at retirement among those still working and about recollected spending changes among those who were already retired. They found that the mean anticipated change was −11.3 percent versus the recollected change of −4.6 percent, and that 54.6 percent of their sample anticipated a reduction in spending versus 36.2 percent that recollected a reduction. This suggests the actual reduction in spending for retirees may be less than many forecast.

These findings are similar to others, such as Miniaci, Monfardini and Weber (2003) and Battistin et al. (2007), who use the Italian Survey on Family Budgets, as well as Aguiar and Hurst (2008) and Laitner and Silverman (2005), who use the Consumer Expenditure Survey (CE). In particular, Fisher et al. (2008) find that consumption-expenditures decrease by about 2.5 percent when individuals retire, expenditures continue to decline at about a rate of 1 percent per year after that. In contrast, Christensen (2004) found no evidence of a drop in consumption at retirement in any of the commodity groups using Spanish panel data.

The change in expenditures varies by type. For example, there has been some research that specifically explored food expenditures of retirees. Aguiar and Hurst (2005) find that while food expenditures decline 17 percent at retirement, the quantity and quality of food consumed did not change. In contrast, Haider and Stephens (2007) found in PSID and the Retirement History Survey that people reduce spending on food when they retire by about 5 to 10 percent. Aguila, Attansio and Meghir (2011), using panel data from 1980 through 2000, estimate a 6 percent drop in

[3] For those readers not familiar with panel data, it is a type of survey where the same individual or household is tracked (or measured) over time. Panel data is also referred to as longitudinal data.

food expenditures after retirement although they find no evidence of nondurable spending reduction in other areas. They attribute this decline in food expenditures to the additional time retiree households have to produce food at home and shop for bargains.

Section 3. What is an Appropriate Replacement Rate?

When targeting a retirement income goal, a common rule of thumb is to estimate the "replacement rate." The replacement rate is the percentage of household earnings needed to maintain a similar standard of living during retirement. The replacement rate is typically less than 100 percent of terminal salary because a number of expenses paid by a household decline or disappear when retired. For example, a retired household no longer has to pay Social Security and Medicare taxes or save for retirement. The household may also have a higher standard deduction and receive income (e.g., Social Security) that is taxed more favorably than wages.

One of the most well-known studies on replacement rates is the Aon Consulting "Replacement Ratio Study," most recently updated in 2008. In the study, the authors note that replacement rates vary by income, for example a household with pre-retirement income of $20,000 has a replacement rate of 94 percent versus a replacement rate of 78 percent for a household with pre-retirement income of $90,000. Replacement rates are typically higher for lower income households because they tend to pay lower (or no) taxes.

Similar to the Aon study, we wish to demonstrate how replacement rates vary across different income and expense scenarios. Therefore, we conduct an analysis in which the replacement rate is defined as the total household income in retirement (traditional IRA, Roth IRA, Social Security retirement benefit and taxable account) divided by the pre-retirement household income. We assume that 80 percent of the household account is in pre-tax (i.e., traditional 401(k) and traditional IRA) savings and that the taxable account is large enough to fund the necessary difference.

We assume a married household with no dependents that can claim two exemptions ($3,900 each). The standard deduction is $12,200 before retirement (under the age of 65) and $14,600 afterward. We use 2013 tables and assume the household itemizes deductions if they are larger than the available standard deduction. We assume a state tax rate of 4 percent. We do, however, ignore other potential tax considerations that may affect a retiree, such as health care expenses which may be deductible (if they exceed 7.5 percent of adjusted gross income).

We assume the household ceases to pay Medicare and Social Security taxes upon retirement and its goal is to have the same total after-tax income when retired. The additional incremental expenses factored into the analysis are pre-tax and post-tax expenses, each of which are treated as a percentage of terminal salary. The pre-tax expenses are most likely to be things such as a traditional 401(k) or traditional IRA deferral but could also be things like company-sponsored insurance premiums. The post-tax expenses are most likely to be things like a Roth 401(k) or Roth IRA deferral but could also be costs associated with working, such as purchasing clothes and commuting to work, that will no longer be realized upon retirement. Additional post-tax expenses, such as college tuition for children, mortgage payments, etc., may be additional expenses paid while working but not for the entire retirement period.

We assume the household consists of a primary worker and spouse and the spouse makes half as much as the primary worker. Spousal income is an important consideration since total household Social Security benefits will be based on either the primary worker's earnings (half) or the spousal benefit, whichever is greater. We assume both members retire at age 65.

In table 1, we present four household profiles and examine the replacement rate that results as we vary pre-tax and post-tax retirement expenditures. Again we assume retirement is funded by a traditional IRA, a Roth and Social Security. Although a "rule of thumb" replacement rate of 70 to 80 percent is clearly reasonable, it isn't ideal and, moreover, it is clear the replacement rate is sensitive to the proportion of pre-tax expenses to post-tax expenses—in fact, the range expands to 54 to 87 percent.

Table 1. Initial target replacement rates as a percentage of pre-retirement income

$25,000 Primary / $12,500 Spouse							
		Pre-Tax Expenses as a % of Income					
		0%	3%	6%	9%	12%	15%
Post-Tax %	0%	87%	84%	82%	79%	76%	74%
	3%	84%	81%	79%	76%	73%	71%
	6%	81%	78%	76%	73%	70%	68%
	9%	78%	75%	73%	70%	67%	65%
	12%	75%	72%	70%	67%	64%	62%

$100,000 Primary / $50,000 Spouse							
		Pre-Tax Expenses as a % of Income					
		0%	3%	6%	9%	12%	15%
Post-Tax %	0%	84%	81%	78%	75%	72%	70%
	3%	80%	77%	74%	71%	69%	66%
	6%	76%	73%	70%	68%	65%	62%
	9%	72%	70%	67%	64%	61%	59%
	12%	69%	66%	63%	61%	58%	55%

$50,000 Primary / $25,000 Spouse							
		Pre-Tax Expenses as a % of Income					
		0%	3%	6%	9%	12%	15%
Post-Tax %	0%	87%	84%	81%	78%	72%	69%
	3%	84%	81%	78%	72%	69%	66%
	6%	80%	77%	71%	68%	65%	62%
	9%	77%	71%	68%	65%	62%	58%
	12%	70%	67%	64%	61%	58%	55%

$150,000 Primary / $75,000 Spouse							
		Pre-Tax Expenses as a % of Income					
		0%	3%	6%	9%	12%	15%
Post-Tax %	0%	84%	81%	79%	76%	73%	70%
	3%	80%	77%	75%	72%	69%	66%
	6%	76%	73%	71%	68%	65%	62%
	9%	72%	69%	67%	64%	61%	58%
	12%	68%	65%	63%	60%	57%	54%

Section 4. Do Retirement Income Needs Rise With Inflation?

In the previous section, we explored how replacement rates can vary depending on pre-retirement income and expenses; in this section, we explore the second assumption in estimating retirement cost—whether retirement income needs rise with inflation. First, we use data from the Consumer Expenditure Survey to explore how actual expenditures differ for households of varying ages. Then, we use the RAND Health and Retirement Study (HRS) dataset to understand how consumption changes over time.

Consumption Profiles

We use the Consumer Expenditure Survey for this section from the Bureau of Labor Statistics website,[4] in particular the 2011 datasets. For each household, the age is defined either as the age of the reference person for a single household, or the average of the reference person and the spouse if it is a two-person household. For expenditures, we focus on the primary categories used to estimate total expenditures (code TOTEXPPQ). We focus specifically on clothing (APPARPQ), charitable contributions (CASHCOPQ), food (FOODPQ), entertainment (ENTERTPQ), health care (HEALTHPQ), housing (HOUSPQ), insurance and pensions (PERINSPQ), and transportation (TRANSPQ), and combine the remaining expenditure groups: alcoholic beverages (ALCBEVPQ), personal care (PERSCAPQ), reading (READPQ), education (EDUCAPQ) and tobacco (TOBACCPQ).

Figure 2 contains the average percentage of total expenditures devoted to these different categories for different household ages. We see two prominent changes in relative expenditures for older retirees: the relative amount spent on insurance and pensions decreases significantly at older ages, while the relative amount spent on health care increases significantly at older ages.

---

[4] http://www.bls.gov/cex/pumdhome.htm

Figure 2. Changing expenditures over time

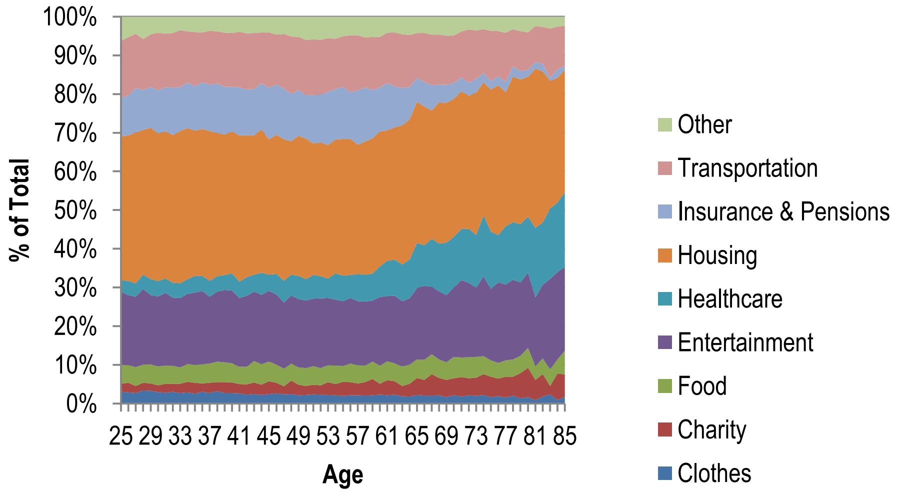

Source: Consumer Expenditure Survey, Morningstar

These different consumption baskets are reflected in the types of indexes created by the Bureau of Labor Statistics to track inflation. The most commonly cited definition of inflation is the change in the Consumer Price Index (CPI) for Urban Consumers, or CPI-U. There are alternative definitions of the CPI that exist as well. For example, the CPI-W, the Consumer Price Index for Urban Wage Earners and Clerical Workers, is the inflation rate for Social Security retirement benefits. An alternative inflation proxy for older workers is the experimental Consumer Price Index for Americans 62 Years of Age and Older, often referred to as the Consumer Price Index for the Elderly (CPI-E).

In table 2, we contrast the differences in the weights among the eight major expenditure groups for these three price indexes. Not surprisingly, we see the weights for things like medical care are higher in the CPI-E (versus the CPI-U), while things like education, apparel and transportation are lower. From December 1982 to December 2012, the average annual change in the CPI-E has been 3.07 percent versus 2.92 percent for CPI-U, therefore, the costs of goods for retirees (as defined by the CPI-E) have increased by approximately 5 percent more, per year, relative to general inflation (CPI-U). If this relationship persists and general inflation (CPI-U) is expected to be 3.0 percent per year, retiree inflation would be 3.15 percent per year. This difference would become increasingly important over longer retirement periods, which is likely a concern for retirees given increasing mortality levels.

Table 2. Different consumer price indexes

	Expenditure Weights			Δ from CPI-U	
Expenditure group	**CPI-U**	**CPI-W**	**CPI-E**	**CPI-W**	**CPI-E**
Apparel	3.5%	3.6%	2.4%	0.1%	-1.1%
Education and communication	6.7%	6.7%	3.8%	0.0%	-2.9%
Food and beverages	15.0%	15.7%	12.8%	0.7%	-2.2%
Housing	40.2%	39.2%	44.5%	-1.0%	4.3%
Medical care	6.9%	5.6%	11.3%	-1.3%	4.4%
Other goods and services	5.3%	5.1%	5.4%	-0.2%	0.1%
Recreation	5.9%	5.5%	5.3%	-0.4%	-0.6%
Transportation	16.5%	18.7%	14.5%	2.2%	-2.0%

The increase in medical care is the largest difference between the CPI-U and CPI-E. Even with social programs like Medicare, medical costs are a significant concern to retirees, especially since expenses like long-term care costs are not covered under the program. Medical inflation, defined as the Consumer Price Index for All Urban Consumers: Medical Care, obtained from the Federal Reserve Bank of St. Louis (FRED), has averaged +5.42 percent per year from 1948 to 2012, versus +3.63 percent for the CPI-U. Therefore, the increase in medical costs has been approximately 50 percent higher than general inflation.

The relationship between general inflation (CPI-U) and medical inflation is included in figure 3. We see a relatively strong relationship historically, with a coefficient of determination (R^2) of 59.07 percent. As of June 19, 2013, the Federal Reserve Bank of Cleveland (Cleveland Fed) was forecasting a 10-year expected inflation rate of 1.55 percent. If we use the results of the ordinary least-squares regression in figure 3, the forecasted medical inflation rate would be approximately 4.0 percent per year.

Figure 3: General inflation (CPI-U) versus medical inflation

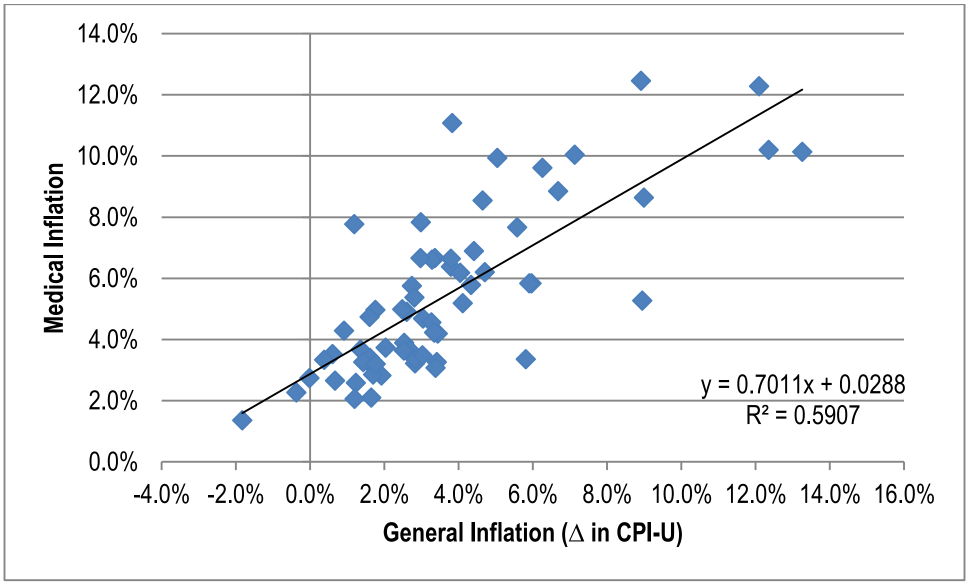

Medical costs are likely to affect retirees differently. Many retirees will have the majority of their medical expenses covered by Medicare, while some may incur significant out-of-pocket expenses for items not covered by Medicare, such as long-term care expenses. To better understand the potential impact of varying levels of medical expenses on household expenditures, we conduct an additional analysis where we segment the households into three groups based on the total level of expenditures (the low income group is defined as households with total expenditures in the 95th to 65th percentile, the mid income group is defined as households with total expenditures in the 65th to 35th percentile, and the high income group is defined as households with total expenditures in the 35th to 5th percentile).

We find no meaningful difference in the medical costs as a percentage of total expenditures among the three income groups either at the median or 95th percentile (highest one in 20) total expenditure levels. The median percentage of medical expenses increases from approximately 5 percent of total expenditures at age 60 to 15 percent by age 80. The 95th percentile, which is the group that has the highest costs in one of 20 households, increases from approximately 25 percent at age 60 to approximately 35 percent by age 80. These findings are important since they suggest medical expenses affect households similarly from a total cost perspective.

Section 5. Consumption Changes Over Time

In the previous section, we explored the changing consumption profiles for households at different ages. In this section, we seek to examine the actual changes in total consumption (or expenditures) for a retiree household over time. While the Consumer Expenditure Survey includes data on total consumption, it is cross-sectional (or longitudinal) and there is no reliable dataset that links changes in household consumption over time. Therefore, to estimate the changes in consumption for retirees, we use the RAND HRS dataset, a panel household survey (combining both cross-sectional and longitudinal data) specifically focused on the study of retirement and health among individuals over the age of 50 in the United States. The RAND HRS is a user-friendly version of a subset of the HRS. It contains cleaned and processed variables with consistent and intuitive naming conventions, model-based imputations and imputation flags, and spousal counterparts of most individual-level variables.

We use the RAND HRS data for spending and match each household to the RAND Consumption and Activities Mail Survey (CAMS), a supplement to the HRS. CAMS was first mailed in September 2001; therefore, to match the two series, we use the five available waves: 2001, 2003, 2005, 2007 and 2009.

As opposed to using all available households, we apply a number of filters. To be included in our analysis, we require the total household spending be greater than $10,000 for each of the five surveys and a consumption change of no greater than 50 percent (in absolute terms) between any two of the five surveys. We do this to create a cleaner dataset, under the assumption households that complete the survey each year and do not have significant changes in consumption are likely more reliable indicators of actual retirees. These filters reduce our sample to 591 households, which is 10.9 percent of the total number of households available in the CAMS series.

For our analysis, we exclude households if any member of the household classifies himself or herself as “not retired.” We test the real growth in consumption by reducing the change in consumption by inflation (CPI-U) over the two-year period between surveys. Once the average annual real change for each household has been estimated for each age, we average the changes for each age group. Similar to our aggregation methodology for the CE data, the age for a single household is based on the age of that household individual, while the age for a married household is the average age of the two spouses.

Figure 5 includes the annual real (inflation-adjusted) change in consumption for retirees ages 60 to 90. Our results are bound between these two ages to ensure a large enough sample of retirees at each age (we generally seek a minimum of 30 households for each age). We include the results of a second order polynomial regression for the entire age range as well as from ages 65 to 75. We include this smaller age range (age 65 to 75) because in future tests we are forced to only consider that limited range for sample size reasons.

Figure 5. Annual real change in consumption for retirees

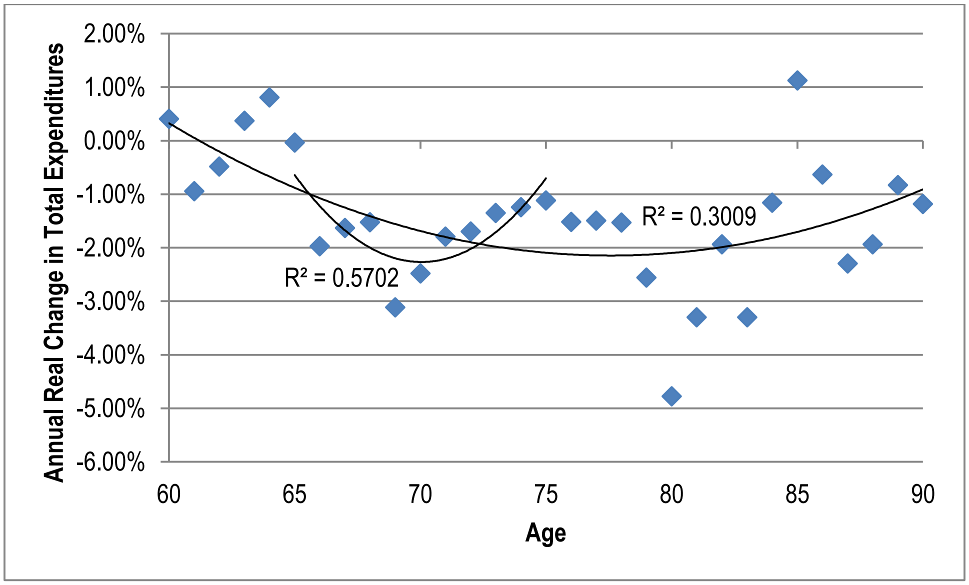

While research on retirement spending commonly assumes consumption increases annually by inflation (implying a real change of 0 percent), we do not witness this relationship within our dataset. We note there appears to be a "retirement spending smile" whereby the expenditures actually decrease in real terms for retirees throughout retirement and then increase toward the end. Overall, however, the real change in annual spending through retirement is clearly negative.

What is less clear from figure 5 is whether the change in expenditures (i.e., consumption) is by choice or by need. It may be that average expenditures decrease because the average retiree did not save enough for retirement and is therefore forced to reduce consumption not out of want but out of need. To better understand this dynamic, we further refine our sample into four groups, based on consumption and total household net worth. The approximate median consumption in our sample is $30,000 per year and the approximate net worth is approximately $400,000.

Our proxy for net worth includes a secondary residence (an aggregated value within the dataset), as well as the estimated total value of pensions and Social Security received by the household. We estimate the value of pensions and Social Security by calculating the mortality-weighted net present value of the future payments, in which we assume a discount rate of 2 percent for Social Security benefits (since these are assumed to increase with inflation) and a 4 percent discount rate for pensions (which are assumed to be nominal). We use the Gompertz Law of Mortality to estimate mortality, as described by Milevsky (2012). Within our Gompertz model, the model lifespan of 88 years and dispersion coefficient of 10 years are fitted based on the unisex mortality from the Society of Actuaries 2000 Annuity Table.

Households with consumption less than $30,000 and a net worth below $400,000 in an initial year (each of the four potential linked survey values are viewed independently) are assumed to be “low spend, low net worth” households. Those with consumption greater than $30,000 and a net worth above $400,000 in an initial year would be “high spend, high net worth” households. The remaining two groups therefore are “low spend, high net worth” and “high spend, low net worth.”

Breaking down the households into these four groups helps us better understand how consumption changes for a household given both its level of consumption and its available resources. Households in which spending and net worth are the same, either low/low and high/high would roughly be considered to consuming optimally, i.e., their consumption is roughly consistent with their resources. In contrast, households where spending and net worth are not the same, either high/low or low/high, would be consuming suboptimally, either too much (high/low) or not enough (low/high). We contrast the changes in spending habits of these two groups in figure 6.

Figure 6. Impact of the amount of consumption and net worth on the average real change in consumption

Panel A. Matched spending and worth **Panel B. Mismatched spending and worth**

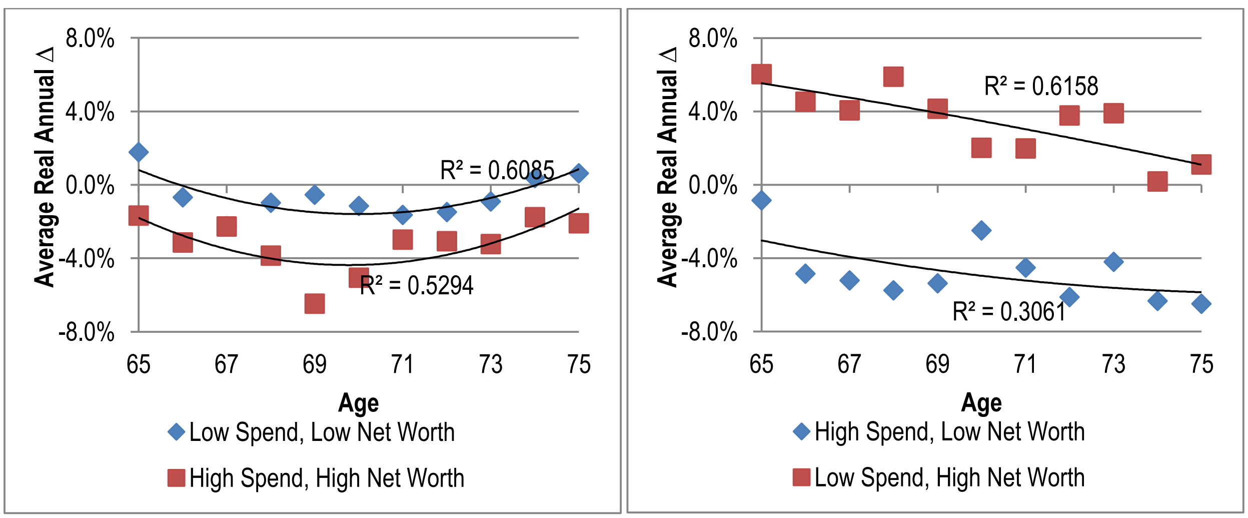

We find the “matched” groups with similar levels of spending and net worth have relatively similar average real changes in expenditures from ages 65 to 75. We note that the lower spending households also tend to see lower decreases in spending over time. This may be due to the fact a higher percentage of household spending is on nondiscretionary items for the lower income household when compared to the higher income household. It’s also important to note that households with lower levels of consumption (low spend, low net worth) tend to have real increases in spending, as denoted by the blue diamonds above the zero mark in panel A, that are greater than households with higher levels of consumption (red squares all in negative territory).

There is a much greater difference in the change in real spending for the mismatched household. We see those households that are overfunded and not spending optimally (the “low spend, high net worth” group) actually tend to increase consumption as they move from age 65 to age 75 but at a decreasing rate. In fact, the real increase by age 75 for these households approaches 0 percent. In contrast, those households that are underfunded and spending too much tend to see considerable declines in consumption. While there are a number of different potential explanations for this spending decline, it may be brought on by the realization that the household spending is not expected to be sustainable over the lifetime of that household.

Section 6. Estimating How Much a Household Should Save for Retirement

Up to this point, we have explored important considerations when estimating the "cost" of retirement. In this section, we want to extend the model to better understand the implications of how much someone has to save for retirement. To do so, we will assume the retiree household has first determined the appropriate total after-tax, post-retirement expenditures required from a portfolio consistent with section 3. To start, we build a "retirement spending curve" that incorporates our expectations about consumption based on our previous analysis.

Retirement Spending Curve

We are not the first to estimate the impact of a consumption path during retirement that increases by some value other than inflation. For example, research by Bernicke (2005), using data from the 2002 Consumer Expenditure Survey, noted that older households tend to spend less than younger households. This decreased level of consumption increases the initial available withdrawal rate when compared to the traditional inflation-adjusted Monte Carlo simulation. Zolt (2013) introduces a dynamic withdrawal adjustment based on whether the portfolio is ahead of or behind target at any point during retirement based on withdrawal findings from Blanchett and Frank (2009). In both cases, the authors note that the required retiree savings decreases when lower inflation rates are used for predicting the lifetime retiree household income need.

From our analysis, we create equation 1, which tells us the change in real annual spending (ΔAS) as a function of age (Age) and the after-tax total expenditure target ($ExpTar$) of a retiree. To take into account that higher-income households spend a higher percent of income on medical costs than lower-income households and are therefore more affected by the higher medical inflation rate, we create a curve in equation 1 that differs from the curve in figure 5. In this curve, we increase the average annual spending by approximately 0.5 percent per year for households that spend over $85,000 per year. Our selection of $85,000 was subjective and higher than the breakpoints in the previous analysis. We use a 0.5 percent increase to approximate the potential future impact of increases in health care costs as a percentage of total costs, especially since the compounded impact of this change may be material for younger retirees or those who are still working. Both of these changes were relatively subjective.

$$\Delta AS = .00008(Age^2) - (.0125 * Age) - .0066\ln(ExpTar) + 54.6\% \qquad [1]$$

In figure 7, we use equation 1 to create various "spending curves" for retirees with different levels of initial total retirement spending goals: $25,000, $50,000 and $100,000. In panel A of figure 7, we demonstrate how the annual real change in spending (based on equation 1) increases at a greater rate (or decreases at a slower rate) for the lower total target expenditure level (e.g., $25,000 versus $100,000). This is consistent with panel A in figure 6. In panel B of figure 7, we show the annual target income (in real terms) over the lifetime for 65-year-old retiree. A retirement spending curve that assumed the annual income need increased annually by inflation, which is the most common assumption when estimating retirement needs, would result in a 0 percent change for each age in panel A and a $1 constant need in panel B. However, using the spending curves based on actual retiree expenditures, we see that the total need decreases in panel B throughout retirement.

Figure 7. Retirement income targets

Panel A. Annual real change in consumption **Panel B. Lifetime real income target, age 65 retiree**

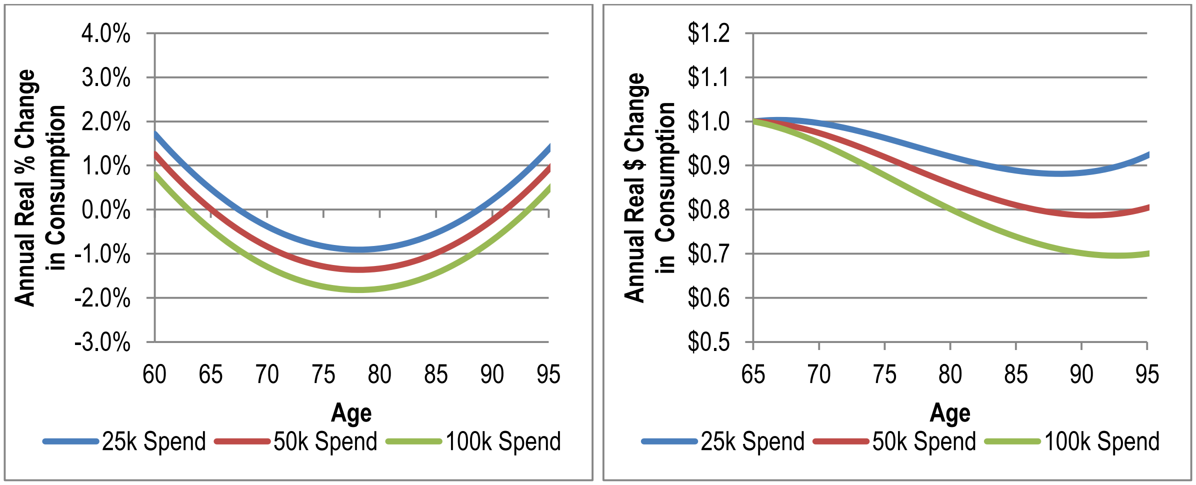

To determine the impact of different retirement spending curves on the cost of retirement, we conduct different simulations. Our first simulation looks at the probability of an initial withdrawal rate lasting 30 years given a constant real spending need as well as the 25k, 50k, and 100k spending curves noted in figure 7. The term "initial withdrawal rate" is used to note the initial amount withdrawn from the portfolio, where the amount is increased by some amount going forward. The constant real spending curve assumes the need increases annually by inflation. The three spending curves result in changes to the initial withdrawal amount based on equation 1, which is displayed visually in figure 7.

The analysis is based on a portfolio with a 40 percent equity allocation, which is assumed to have a 3.0 percent real return and a standard deviation of 10 percent. The return of the portfolio can roughly be decomposed into a stock return of 9.0 percent, a bond return of 4.0 percent, inflation of 2.5 percent and assumed fees of approximately 0.5 percent. The assumed standard deviation of stocks is 20 percent versus 7 percent for bonds with a correlation of zero between the two asset classes. These numbers are based approximately on Ibbotson's 20-year forecasted capital market assumptions for the year 2013.

Each test scenario is based on a 10,000-run Monte Carlo simulation. For the first simulation, we determine the probability a given withdrawal strategy, based on the different spending curves, survives a 30-year period. We test initial withdrawal rates from 2.0 to 8.0 percent in 0.2 percent increments. The results are included in figure 8.

Figure 8. Retirement income targets

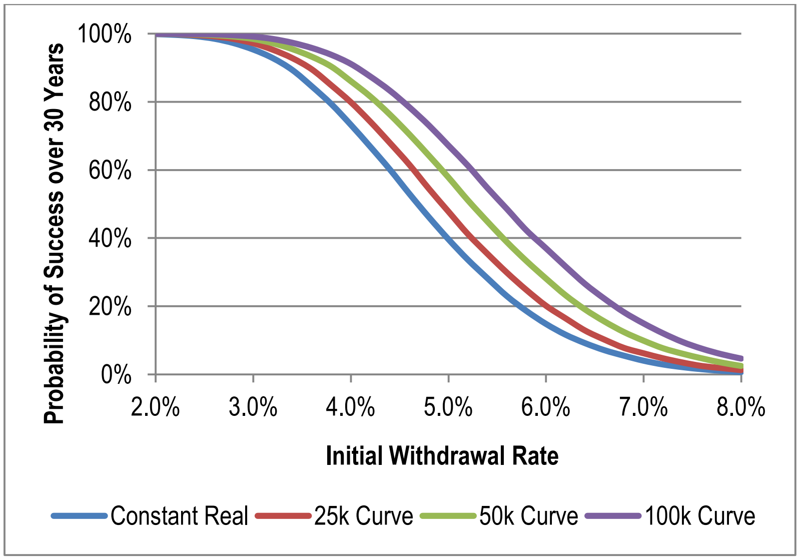

As we expected, the probabilities of success increase across the different initial withdrawal rates when using the spending curves versus assuming a constant real withdrawal amount increase. For example, a 4.0 percent initial withdrawal rate has a 73.3 percent probability of success using a constant real strategy (where the withdrawal increases each year by inflation), while the 25k curve has an 79.9 percent chance of success, the 50k curve has an 86.0 percent and the 100k curve a 91.1 percent.

For the second simulation, we incorporate life expectancy. Here, failure is defined as running out of money while any member of the household is still alive. The differences between modeling for a fixed period (assuming a death date) and modeling for conditional mortality have been noted by Blanchett and Blanchett (2008), among others.

For this simulation, we assume the retirement need doesn't change after age 95. We do this because in our primary RAND HRS dataset we do not have enough data to forecast increases in consumption past age 95. When estimating mortality, we use the Gompertz law of mortality, as described by Milevsky (2012). Our model life span is 86 for males and 90 for females, and we use a dispersion coefficient of 11 for males and 9 for females. These are based on mortality from the Society of Actuaries 2000 Annuity Table.

For the simulation, we test retirement periods of 20, 25, 30, 35 and 40 years. We also include a life expectancy test, where success is determined by the portfolio's ability to maintain the withdrawal during the lifetime of the household, based on either a 65-year-old male, a 65-year-old female or a couple both age 65. The results of the different scenarios are included in table 3.

Table 3. Probabilities of success for various initial withdrawal rates, retirement period and spending curves

Withdrawal Increases Annually by Inflation									
		Period Certain (Years)					During Lifetime (Age 65)		
		20	25	30	35	40	Male	Female	Joint
Initial Withdrawal Rate	3.0%	100.0%	98.9%	95.4%	89.6%	81.9%	98.3%	97.5%	96.3%
	4.0%	97.8%	88.4%	73.3%	58.5%	46.9%	90.7%	87.0%	81.5%
	5.0%	85.6%	61.0%	39.7%	26.3%	17.5%	76.4%	68.6%	57.4%
	6.0%	59.0%	29.9%	14.8%	8.1%	4.8%	60.0%	49.3%	34.5%

$25k Initial Goal Curve									
		Period Certain (Years)					During Lifetime (Age 65)		
		20	25	30	35	40	Male	Female	Joint
Initial Withdrawal Rate	3.0%	100.0%	99.4%	97.2%	92.6%	86.2%	98.9%	98.3%	97.5%
	4.0%	98.5%	91.9%	79.9%	66.2%	53.3%	92.9%	90.0%	85.6%
	5.0%	88.8%	68.1%	47.8%	32.7%	22.7%	80.0%	73.1%	63.2%
	6.0%	65.0%	36.5%	20.2%	11.1%	6.8%	63.6%	53.4%	39.3%

$50k Initial Goal Curve									
		Period Certain (Years)					During Lifetime (Age 65)		
		20	25	30	35	40	Male	Female	Joint
Initial Withdrawal Rate	3.0%	100.0%	99.7%	98.5%	95.6%	91.6%	99.3%	99.0%	98.5%
	4.0%	99.2%	94.8%	86.0%	75.4%	64.2%	95.1%	93.0%	89.9%
	5.0%	92.5%	75.9%	57.8%	42.8%	31.7%	84.2%	78.5%	70.3%
	6.0%	72.7%	45.3%	28.1%	17.2%	11.0%	68.4%	59.2%	46.3%

$100k Initial Goal Curve									
		Period Certain (Years)					During Lifetime (Age 65)		
		20	25	30	35	40	Male	Female	Joint
Initial Withdrawal Rate	3.0%	100.0%	99.9%	99.2%	97.6%	95.1%	99.6%	99.5%	99.2%
	4.0%	99.5%	96.7%	91.1%	82.7%	74.3%	96.8%	95.3%	93.2%
	5.0%	94.9%	82.3%	67.2%	53.4%	42.6%	87.9%	83.4%	76.9%
	6.0%	79.2%	54.5%	37.1%	25.3%	17.3%	73.4%	65.2%	53.7%

In table 3, we note the relative safety of a given initial withdrawal rate can vary considerably based on the assumed spending curve and the retirement period (either the number of assumed years or a life expectancy model). Using the constant real model, a 4.0 percent initial withdrawal rate has a 73.3 percent probability of success over a 30-year period. This period is generally assumed to represent the retirement horizon for a joint couple. Note, though, the probability of success for a 4.0 percent initial withdrawal rate using the constant real model increases to 81.5 percent over the expected mortality of a joint couple (male and female both age 65). Moreover, the success rate for the joint couple climbs even higher to 89.9 percent if one assumes the $50,000 spending curve rather than the constant real model. Another way of looking at the results is that the 4.0 percent initial withdrawal scenario over 30 years under the constant real model has the same approximate probability of success (70.3 percent versus 73.3 percent) as the 5.0 percent initial withdrawal scenario with the $50,000 spending curve over the expected mortality of the couple.

A 5.0 percent initial withdrawal rate results in a 20 percent reduction in the amount of savings required to fund a retirement goal when compared to the traditional 4.0 percent initial withdrawal rate. This may seem counterintuitive, but if we assume a retiree household requires $40,000 of income per year from a portfolio, using the 5.0 percent rule, the necessary balance at retirement is $800,000 ($40,000/0.05 = $800,000) versus $1 million if a 4.0 percent initial withdrawal rate is used. This 5.0 percent initial withdrawal amount can likely be further increased if the retiree is willing to take on the potential risk of future reductions in spending by implementing a more dynamic withdrawal strategy.

Section 7. Conclusions

In this paper, we use various government survey data and perform an analysis to more accurately estimate the cost of retirement. We note that while a replacement rate between 70 and 80 percent is likely a reasonable starting place for most households, the actual replacement goal can vary considerably based on the expected differences between pre- and post-retirement expenses. We also find that retiree expenditures do not, on average, increase each year by inflation and that the actual “spending curve” of a retiree household also varies by total consumption, whereby households with lower levels of consumption tend to have real increases in spending that are greater than households with higher levels of consumption.

When combined, these findings have important implications for retirees, especially when estimating the amount that must be saved to fund retirement. While many retirement income models use a fixed time period (e.g., 30 years) to estimate the duration of retirement, modeling the cost over the expected lifetime of the household, along with incorporating the actual spending curve, results in a required account balance at retirement that can be 25 percent less than the amount required using traditional models. In summary, a more advanced perspective on retiree spending needs can significantly change the estimate of the true cost of retirement.

References

Aguiar, Mark, and Erik Hurst. 2005. “Consumption Versus Expenditure.” *Journal of Political Economy* 113 (5): 919–48.

———. 2008. “Deconstructing Lifecycle Expenditure.” NBER Working Paper No. 13893, National Bureau of Economic Research, Cambridge, MA.

Aguila, Emma, Orazio Attanasio, and Costas Meghir. 2011. “Changes in Consumption at Retirement: Evidence From Panel Data.” *Review of Economics and Statistics* 93 (3):1094–99.

Ameriks, John, Andrew Caplin, and John Leahy. 2007. “Retirement Consumption: Insights from a Survey.” *Review of Economics and Statistics* 89 (2): 265–74.

Aon Consulting. 2008. “Replacement Ratio Study: A Measurement Tool for Retirement Planning.” Seventh update, Aon Consulting, London. http://www.aon.com/about-aon/intellectual-capital/attachments/human-capital-consulting/RRStudy070308.pdf.

Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber. 2007. "The Retirement Consumption Puzzle: Evidence From a Regression Discontinuity Approach." Research Paper Series 27/07, University Ca'Foscari of Venice, Department of Economics.

Banks, James, Richard Blundell, and Sarah Tanner. 1998. "Is There a Retirement-Savings Puzzle?" *American Economic Review* 88 (4): 769–88.

Bernicke, Ty. 2005. "Reality Retirement Planning: A New Paradigm for an Old Science.” *Journal of Financial Planning* 18 (6): 56–61.

Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review* 91 (4): 832–57.

Blanchett, David M., and Brian C. Blanchett. 2008. "Joint Life Expectancy and the Retirement Distribution Period." *Journal of Financial Planning* 21 (12): 54–60.

Blanchett, David M., and Larry R. Frank. 2009. “A Dynamic and Adaptive Approach to Distribution Planning and Monitoring.” *Journal of Financial Planning* 22 (4): 52–66.

Christensen, Mette Lunde. 2004. "Demand Patterns Around Retirement: Evidence from Spanish Panel Data." Economics Discussion Paper Series EDP-0809, University of Manchester, Economics, School of Social Sciences.

Fisher, Jonathan D., David S. Johnson, Joseph Marchand, Timothy M. Smeeding, and Barbara Boyle Torrey. 2008. “The Retirement Consumption Conundrum: Evidence From a Consumption Survey." *Economics Letters* 99 (3): 482–85.

Haider, Steven J., and Melvin Stephens Jr. 2007. “Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations.” *Review of Economics and Statistics* 89 (2): 247–64.

Hurd, Michael D., and Susann Rohwedder. 2008. “The Retirement Consumption Puzzle: Actual Spending Change in Panel Data.” NBER Working Paper No. 13929. National Bureau of Economic Research, Cambridge, MA.

Laitner, John, and Dan Silverman. 2005. "Estimating Life-Cycle Parameters from Consumption Behavior at Retirement." NBER Working Paper No. 11163. National Bureau of Economic Research, Cambridge, MA. Cambridge, MA: National Bureau of Economic Research. March. http://papers.nber.org/papers/w11163.

Milevsky, Moshe A. 2012. *The 7 Most Important Equations for Your Retirement: The Fascinating People and Ideas Behind Planning Your Retirement Income.* Mississauga, Ontario: Wiley.

Miniaci, Raffaele, Chiara Monfardini, and Guglielmo Weber. 2003. "Is There a Retirement Consumption Puzzle in Italy?" IFS Working Paper No. 03/14, Institute for Fiscal Studies (IFS), London.

Modigliani, Franco, and Richard Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." In *Post-Keynesian Economics*, edited by Kenneth K. Kurihara. New Brunswick: Rutgers University Press.

Zolt, David M. 2013. "Achieving a Higher Safe Withdrawal Rate with the Target Percentage Adjustment." *Journal of Financial Planning* (January): 51–59.