NBER WORKING PAPER SERIES

SOME ANSWERS TO THE RETIREMENT-CONSUMPTION PUZZLE

Michael D. Hurd
Susann Rohwedder

Working Paper 12057
http://www.nber.org/papers/w12057

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2006

Research support from the Social Security Administration via a grant to the Michigan Retirement Research Center (UM04-13) and from the National Institute on Aging (P01-AG008291 and P30-AG12815) are gratefully acknowledged. We are also thankful for helpful comments from participants of the Aging Workshop at the NBER Summer Institute 2004, the Sixth Annual Conference of the Retirement Research Consortium, and the Seventh Annual Research Conference at DNB (De Nederlansche Bank). All remaining errors are our responsibility alone. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Some Answers to the Retirement-Consumption Puzzle
Michael D. Hurd and Susann Rohwedder
NBER Working Paper No. 12057
February 2006
JEL No. D91, J26

ABSTRACT

The simple one-good model of life-cycle consumption requires "consumption smoothing." According to previous results based on partial spending and on synthetic panels, British and U.S. households apparently reduce consumption at retirement. The reduction cannot be explained by the simple one-good life-cycle model, so it has been referred to as the retirement-consumption puzzle. An interpretation is that at retirement individuals discover they have fewer economic resources than they had anticipated prior to retirement, and as a consequence reduce consumption. This interpretation challenges the life-cycle model where consumers are assumed to be forward-looking. Using panel data, we find that prior to retirement workers anticipated on average a decline of 13.3% in spending and after retirement they recollected a decline of 12.9%: widespread surprise is not the explanation for the retirement-consumption puzzle. Workers with substantial wealth both anticipated and recollected a decline. Therefore, for many workers the decline is not necessitated by the fall in income that accompanies retirement. Poor health is associated with above-average declines. At retirement time spent in activities that could substitute for market-purchased goods increases. Apparently a number of factors contribute to the decline in spending, which, for most of the population, can be accommodated in conventional models of economic behavior.

Michael D. Hurd
RAND Corporation
1700 Main Street
and NBER
mhurd@rand.org

Susann Rohwedder
RAND Corporation
1700 Main Street
Santa Monica, CA 90407
susannr@rand.org

1. Introduction

The simple one-good model of life-cycle consumption requires "consumption smoothing:" the trajectory of consumption by an individual should be continuous in time. If the trajectory is not continuous, a reallocation of consumption so as to reduce the size of the discontinuity will increase lifetime utility without an increase in the use of resources. However, British households apparently reduce consumption at the ages associated with retirement, and the reduction cannot be explained by the life-cycle model (Banks, Blundell and Tanner, 1998). Households in the Panel Study of Income Dynamics (PSID) sharply reduced several components of consumption at retirement (Bernheim, Skinner and Weinberg, 2001). Because the mechanisms underlying this observed drop in consumption at retirement are not well understood, it has been referred to as the retirement-consumption puzzle.

There are a number of interpretations or explanations for this drop. The most obvious has to do with the cessation of work-related expenses, but it appears that such expenses are not large enough to explain the observed drop in consumption at retirement (Banks, Blundell and Tanner). A second explanation, which is emphasized by Bernheim, Skinner and Weinberg, is that workers do not adequately foresee the decline in income associated with retirement. On reaching retirement they assess their financial resources, and, finding them less than anticipated, reduce consumption. This interpretation is damaging to the life-cycle model, which assumes that economic agents are forward-looking planners. For most workers, retirement is a predicable event, and workers should be assessing continuously their financial situation so that they will not be surprised. They should have saved enough so that they would not have to reduce consumption at retirement.

A third interpretation is that workers under-saved but they were aware they had under-saved: they were not surprised by the inadequacy of their resources. A lack of self control caused the under-saving and the decline in income forced them to reduce consumption. This interpretation is also damaging to the life-cycle model, which assumes that people are both forward-looking and that they follow through on their (optimal) plans. A fourth interpretation is that the timing of retirement is uncertain. Some workers retire earlier than anticipated because of a health event or unemployment, resulting in an unexpected reduction in lifetime resources, and the reduction leads to a concurrent reduction in consumption. Such a reduction in consumption is well within the spirit of the life-cycle model. A final explanation is that retired households have considerably more leisure than working households. The increased leisure can be used to purchase goods more efficiently or to substitute home-produced goods for purchased goods. In this interpretation, spending declines, but actual consumption does not. We note, however, that the increased leisure time could also lead to increases in purchased goods because of complementarities such as spending on travel. If some uses of time are substitutes for market-purchased goods and some are complements, the overall effect is an empirical matter, but we would expect consumption to change at retirement, not that it be smooth.

This paper provides evidence about the retirement-consumption puzzle and several explanations for it. First, based on recollections by retirees both in cross-section and panel data, we find that average spending declined at retirement.[1] The average magnitude of the decline is approximately the same as has been found in previous literature. However, fewer than half of retirees report that spending declined, so that a decline is not the experience of the population in

[1] While we recognize the distinction between consumption and spending, mostly we will use them interchangeably. In practice the empirical results in the literature are based either on food spending or nondurables spending, both of which should be approximately the same as consumption of those items because spending and consumption are almost simultaneous.

general. Second, both in cross-section and in panel workers anticipate that their spending will decline when they retire. This result rules out one explanation for the decline: that there is widespread surprise at the lack of resources at retirement. This explanation would be particularly damaging to the life-cycle model, which, were it verified, would require rethinking of the basic approach to models of intertemporal decision making.

Our third result is that in panel the distributions of anticipated and recollected spending changes are approximately the same. Thus, on average the population that retired between waves had, in this sense, rational forecasts.

Accurately forecasting that spending would decline does not establish that workers were consuming optimally. For example, they could have foreseen that income would decline at retirement, yet not have the self-control to save. Upon reaching retirement and lacking adequate savings retirees would be forced to reduce consumption because of the drop in income. We find, however, that at least for a substantial part of the population, a decline in income does not necessitate a decline in spending: Many workers in the upper part of the wealth distribution anticipated similar declines in spending as those in the lower part of the wealth distribution; and many retirees in the upper part of the wealth distribution recollected a decline similar to those at the lower end of the wealth distribution. For these people the decline was not forced by the reduction in income that accompanies retirement.

A fifth result is that health plays an important role in explaining the observed drops in spending. Poor health was an important reason for retirement among 28 percent of retirees. Those retirees reported substantially higher spending declines than those whose health was not an important factor in retirement. An implication is that an unanticipated decline in lifetime resources caused by early retirement could explain a spending decline for part of the population.

Finally, we find that the increase in leisure time accompanying retirement led to an increase in time spent on seven activities that could substitute for market-purchased goods or lead to more efficient shopping. In panel, time spent on activities such as food preparation and shopping increased by about 5.5 hours per week following retirement. This evidence supports the view that the increased ability to engage in home production or more thrifty shopping is an important reason for the observed spending declines.

2. Prior Literature

The literature on this topic began with Banks, Blundell and Tanner (1998) who found that at ages associated with retirement spending declined more rapidly than could be explained by a simple life-cycle model. Their study is based on nondurable consumption in synthetic panel in the U.K. Family Expenditure Survey. They interpret the drop to be the result of "unanticipated shocks occurring around the time of retirement (p. 784)" such as an over-estimate of pension income. According to Bernheim, Skinner and Weinberg spending on food declined sharply at retirement in the Panel Study of Income Dynamics (PSID): the two-year change in log food consumption following retirement averaged –0.14 with the greatest decline among households in the lowest income or wealth quartile. Bernheim, Skinner and Weinberg interpret the decline in food spending (as well as patterns of wealth holdings) to be evidence against models of behavior in which agents are rational and forward-looking. "If households follow heuristic rules of thumb to determine saving prior to retirement, and if they take stock of their financial situation and make adjustments at retirement (so that the adequacy of saving is "news"), then one would expect to observe the patterns documented in this paper (p. 855)." If these interpretations of the

retirement-consumption puzzle are correct, they cast doubt on models of rational forward-looking economic behavior such as the life-cycle model.

Lower spending by those of retirement age in synthetic panel is also found by Blau (2004) in data from the CEX, and by Miniaci, Monfardini and Weber (2003) in the Italian Survey on Family Budgets. A somewhat different result is found by Fisher *et al*. (2005) based on synthetic panels in the CEX. Median spending on food declined by 5.9%, but total spending on consumption, declined by just 2.5% (Table 6 of Fisher *et al*.).

Investigations of food spending in true panel have sometimes corroborated the results of Bernheim, Skinner and Weinberg, but not universally. Smith (2004) divided retirees as observed in the British Household Panel Survey into four groups and found over two groups comprising 57% of the sample that food spending did not decline at retirement. Another group, about 24% of the sample, had a decline in spending but their retirement was involuntary, often associated with unemployment or poor health. This group likely suffered a wealth shock due to early retirement and so would be expected to reduce spending within the framework of the life-cycle model. There remained a group comprising about 19% of retirees. They left the labor force at the normal retirement age, yet experienced a decline of 17% in spending on food.[2] Thus, among the group that retired voluntarily, about 75% experienced no decline in food spending. These findings suggest that any unexplained decline in food spending at retirement is fairly small and not the norm in the U.K. population.

Haider and Stephens (2004) found in the PSID and in the Retirement History Survey that people reduce spending on food when they retire by about 5-10% depending on the specification. In the Health and Retirement Study they found no reduction, and there is no apparent explanation

[2] Group 1 in Table 5 of Smith. However, when controlling for covariates as estimated by regression in Table 6, the decline for this group was just 4%.

for this difference. Haider and Stephens address the issue of the effect of unexpected retirement on food spending by asking whether the decline could be explained by the difference between expected and actual retirement. Controlling for the difference between them, they find that the decline in food spending is reduced by about one third, still leaving an unexplained reduction.

Aguiar and Hurst (2005a) used the Continuing Survey of Food Intake of Individuals, collected by the U.S. Department of Agriculture, to study the fine details of food consumption as well as on food spending. They found that although spending on food declines at retirement, actual consumption as measured by caloric or vitamin intake, or by the quality of food did not decline. Their interpretation is that the extra leisure associated with retirement is used to produce the same food consumption levels but using smaller inputs of market purchased goods. To validate this interpretation, they used scanner data and found that actual prices paid are indeed lower among those aged 65 or over than among those 40-65 (Aguiar and Hurst, 2005b).[3] Apparently, retirees use time to shop more effectively, permitting a reduction in actual spending for the same purchases.

An entirely different approach to the retirement-consumption puzzle is via simulation of a life-cycle model. If the model allows nonseparability between leisure and consumption, then it will, of course, cause a discontinuous change in consumption when hours of work change discontinuously. But even when leisure and consumption are separable, uncertainty can lead to a decline in consumption at retirement as in the models of Blau (2004) and of Benitez-Silva, Buchinsky and Rust (2005). The explanation is that a negative shock, say, to health will lead to unexpected early retirement for some and, therefore, to an unexpected decline in lifetime resources. The discontinuous decline in resources causes a discontinuous decline in

[3] Scanner data are records of actual purchases by households. The purchase codes are linked to price information at the store of purchase. The households are surveyed about demographics and income among other data items.

consumption. However, Blau finds that in a model calibrated to HRS data, the decline in consumption resulting from unexpected events is not large: the drop in spending when retirement is unexpected is only about one percentage point greater than when it is expected.

Our summary of these papers is that, as assessed in synthetic panels, there is a decline on average in spending at retirement. But the magnitude of the estimated decline varies between very small as in Fisher *et al*. (2005) to much larger as in Blau (2004) and others. Thus the decline needs to be validated in actual panel data.

The interpretation of the decline in synthetic panel depends on the distribution of the declines in the population: at the extremes, is the average due to a few households having very large declines or is it due to all households have approximately the average decline? The explanation for the latter could be a widespread lack of forward-looking behavior, work-related expenses, substitution of leisure for spending or similar mechanisms, whereas the former would suggest costly low probability risks. However synthetic panel cannot shed light on this issue because it does not track individuals and it cannot associate any decline with personal and financial characteristics such as health and wealth. Finally, these papers do not address the issue of whether any spending changes were anticipated, which could be an important part of the retirement-consumption puzzle.

Food spending as observed in panel declines at retirement but not over all populations at all times. An unanswered question is what causes the difference in measurement. This question is particularly relevant for the comparison between the HRS and PSID in Haider and Stephens (2004). In one comparison, the sample period is approximately the same: 1991-1999 for the PSID and 1992-1996 for the HRS. The questions formats are similar in both surveys and both

surveys are said to be population representative. Yet the log change in food spending at retirement is –0.089 (with standard error of 0.043) in the PSID and 0.005 (0.024) in the HRS.[4]

2. Theoretical Background

In its simplest form the life-cycle model (LCM) with one consumption good specifies that individuals choose a consumption path to maximize expected lifetime utility, and that the instantaneous utility function is unchanging over time. The shape of the optimal consumption path is partially or wholly determined by utility function parameters, the interest rate and mortality risk. The level of the path is determined by the lifetime budget constraint; the difference between the level of consumption and income determines the saving rate and the equation of motion of wealth. Auxiliary assumptions, which are not controversial, are that the marginal utility is continuous in consumption and that marginal utility declines in consumption. A condition for lifetime utility maximization is that marginal utility be continuous in time: were it not continuous a reallocation of consumption across the discontinuity from the low marginal utility state to the high marginal utility state would increase total utility without a greater use of resources. Such a reallocation should continue until there no longer is a discontinuity in marginal utility. Because consumption is monotonic and continuous in marginal utility, an implication is that consumption must be continuous in time. That is, consumption must be smooth over time in a model where utility only depends on consumption.

In a more general model, which recognizes uncertainty, individuals or households experience unanticipated windfall gains or losses to wealth, earnings or annuities, and then re-optimize to a new consumption path, causing a discontinuity in the consumption path. However, wealth, earnings or annuity changes which are foreseeable should cause no change in the

[4] Table 3 in Haider and Stephens (2004).

consumption path because the lifetime budget constraint has not changed. In particular consumption should not change at retirement if retirement occurs as planned.[5] But if retirement occurs sooner than expected, lifetime resources will be less than expected so that consumption will have to be adjusted downward. The obvious example is a stochastic health event that causes early retirement. Negative health shocks leading to early retirement are undoubtedly empirically important, so that we should expect to observe in the population some unanticipated decline in consumption at retirement from these shocks alone.

A second generalization of the LCM specifies that utility depends on more than one good, in particular leisure as well as consumption. Suppose that the within-period utility function is $u(c,l)$. The implications for consumption at retirement depend on whether the utility function is separable; that is, whether the marginal utility of consumption u_c depends on l.

If the utility function is separable, u_c should be continuous in time and consumption will also be continuous. If the utility function is not separable, but retirement is gradual so that l increases slowly, consumption will also change in a continuous manner. But for most workers l increases abruptly by about 2,000 hours per year. A condition of utility maximization is that the marginal utility of consumption, u_c, be the same immediately before and immediately after retirement: the argument is the same as we gave earlier in the context of a single good model of the LCM. Now, however, because of nonseparability and because of the sudden change in l, the LCM *requires* a discontinuous change in consumption.

Some types of leisure are substitutes for the consumption of market purchased goods such as home repairs, some are complements with consumption such as travel, and some are neutral such as watching television. Everyday observation and introspection suggest that we

[5] If some of measured consumption is, in fact, work-related expenses, consumption as measured by spending would drop at retirement, but utility-producing spending would not. This is a measurement issue.

have all types, and it is an empirical question as to which dominates. But the main point is that we would not expect consumption to be smoothed over retirement.

Because of differences in tastes and differences in economic resources we expect heterogeneity across households in whether substitution or complementarity dominates. For example, someone with high wealth may continue to purchase home repairs as before retirement, but spend more on travel with a net effect of an increase in spending. Someone with a high wage rate may have purchased home repairs before retirement but will do them himself after retirement for a net reduction in spending.

This can be illustrated with a three good model. Suppose that utility is given by $u(x, y, l)$, where *x* and *y* are composite goods and *l* is leisure. The optimal path of x will equate u_x before and after retirement and the optimal path of y will equate u_y before and after retirement. If $u_{xl} = 0$ and $u_{yl} = 0$ both x and y will be continuous across retirement. But if x is a substitute for leisure ($u_{xl} < 0$) and y is a complement to leisure ($u_{yl} > 0$), then a discontinuous but anticipated reduction in l will require a decrease in x and an increase in y. Whether total spending increases or decreases would depend on utility function parameters, prices and the levels of spending on each. Even with identical preferences workers facing differing wage rates would change total spending differently at retirement.

We have stated nonseparability in terms of the utility function, but the conclusions are the same in the context of home production. For example, suppose that instantaneous utility is given by $u(f(c,l))$ where f is a production function that uses inputs of purchased goods c and of time l to produce actual consumption which produces utility. Then nonseparability of f will cause a discontinuous change in c when l changes discontinuously at retirement.

In this discussion we have simplified the problem by assuming that retirement is given exogenously. Whether retirement is chosen does not affect the discontinuity in consumption when leisure and consumption are not separable provided the increase in leisure is discontinuous. As an empirical matter a substantial majority of retirement is from full-time to completely out of the labor force (Rust, 1990) and there are good reasons for such a sharp transition.[6]

3. Data

Our data come from the Health and Retirement Study (HRS) and from a supplemental survey to the HRS, the Consumption and Activities Mail Survey (CAMS).[7] The HRS is a biennial panel. Its first wave was conducted in 1992. The target population was the cohorts born in 1931-1941 (Juster and Suzman, 1995). Additional cohorts were added in 1993 and 1998 so that beginning with the 1998 wave it represented the population from the cohorts of 1947 or earlier. The HRS interviewed about 20,000 subjects in 13,100 households in the year 2000 wave. A random sample of 5,000 households (38.2 percent of all households interviewed in HRS 2000) was asked to participate in CAMS. The questionnaires for CAMS wave 1 were sent out in October, 2001.[8] In married or partnered households it was sent to one of the spouses, chosen at random. There were 3,866 responses in the CAMS wave 1, which corresponds to a total response rate of 77.3 percent.[9]

---

[6] For example, a defined benefit pension plan can have such strong incentives to retire that workers within a wide range of tastes for retirement will all retire. Most firms will not allow a gradual reduction in work hours, so that a worker who would like to retire gradually will be forced to change employers and possibly occupations (Hurd, 1996).

[7] The HRS is sponsored by the National Institute of Aging (grant number NIA U01AG09740) and conducted by the University of Michigan.

[8] See Hurd and Rohwedder (2005) for a more extensive description of CAMS.

[9] Although the response rate was high there was some differential non-response by demographic characteristics. HRS has supplied weights to account for non-response, and most of our analyses will use them. Because our main variables are about household spending we will use the household weights. We have conducted parallel analyses using unweighted data and the results are very little different.

In October, 2003, CAMS wave 2 was sent to the same households. The structure of the questionnaire was almost the same so as to facilitate panel analysis. The response rate in CAMS wave 2 was 78.3 percent.[10]

CAMS has three main topics: Part A is about activities or uses of time; Part B collects data on spending, including anticipations and realizations about changes in spending at retirement; and Part C asks information about marital status and labor force participation.[11]

The main variables of interest in our analysis are anticipated and recollected spending changes at retirement which come from the following question sequence in the CAMS questionnaire:[12]

[10] Response rates are lower bounds in that they are not adjusted for mortality or undeliverable questionnaires.
[11] In wave 1 of CAMS section C included in addition questions about prescription drug usage.
[12] The CAMS questionnaires for wave 1 and 2 are accessible online at http://hrsonline.isr.umich.edu/meta/2001/cams/qnaire/cams01abc.pdf for CAMS wave 1 and http://hrsonline.isr.umich.edu/meta/2003/cams/qnaire/cams2003.pdf for CAMS wave 2.

Excerpt from the CAMS Questionnaire:

Question B38 in CAMS wave 1; question B44 in CAMS wave 2.

We would like to understand more about spending in retirement.

Are you retired?

_____ Yes ➔ **Complete BOX A** No ➔ **Complete BOX B**

BOX A – Retired:	**BOX B – Not Retired**:
a. How did your TOTAL spending change with retirement? _____ Stayed the same ➔ **Go to c** _____ Increased _____ Decreased	**d**. How do you expect your TOTAL spending to change with retirement? _____ Stay the same ➔ **Go to f** _____ Increase _____ Decrease
b. By how much? _____%	**e**. By how much? _____%
c. For the items below, check (✓) whether the spending increased, decreased or stayed the same in retirement:	**f**. For the items below, check (✓) whether you expect spending to increase, decrease or stay the same in retirement:

We link these responses to the rich information obtained on the same respondents in repeated HRS core interviews. For example, we obtain information on demographics and socio-economic status from HRS 2000. We also make use of the panel nature of the HRS to obtain from prior waves information such as the year a CAMS respondent retired, self-rated health in the wave immediately before and after a respondent's retirement, and reasons for retirement.[13]

In the empirical analysis of this paper we use responses where the qualitative and the quantitative information on the anticipated or recollected spending decline is not missing. [14] We use data on retirees between the ages of 50 and 74 who retired no earlier than age 50 and no later than age 70. For those not yet retired we use the responses of those aged between 50 and 69.[15] We make these restrictions to find changes in spending when retirement is "normal" and to reduce recall bias due to retirement having occurred many years in the past.

4. Declines in Spending in the HRS: Anticipations and Realizations

In this section we provide empirical evidence about three issues. First, we find average declines in spending at retirement of about the magnitude that has been found in previous literature. However, we also find that only about half the population experienced a decline. Then, we address two key questions that have not yet been resolved in the literature in a satisfactory manner: While still working did workers anticipate that spending would decline at retirement?

[13] Most HRS core variables are obtained from The RAND HRS Data file which is an easy to use longitudinal data set based on the HRS data. It was developed at RAND with funding from the National Institute on Aging and the Social Security Administration.

[14] Response rates were quite high in both waves of CAMS: The lead-in question about retirement status was answered by 93.8 (95.8) percent of CAMS respondents to wave 1 (wave 2); the qualitative question about change in spending at retirement was answered by 93.5 (93.6) percent, retired and not retired respondents combined. Questions B38b and B38e (combined) were answered by 77.4 percent in wave 1 and 78.2 percent in wave 2.

[15] In addition we exclude 8 observations where the respondent filled out both sides of the box, that is answered both the questions for the retired and for the not retired, and the answers differed from each other.

Were anticipations and realizations consistent on average? We have unique data that allow us to address these issues directly.

4.1 Spending Change at Retirement in CAMS

Table 1 shows recollected spending change at retirement as reported by retirees in CAMS wave 1.[16] The observations are classified according to the age at which the respondent retired as found in prior waves of HRS. The average recollected change was a decline of about 14%. There is some variation with age showing a larger decline among those who were 50-54 at the time of their retirement. Such early retirement is often associated with a health event that causes workers to retire unexpectedly early, so the larger decline in spending is not surprising.

Because we have data on individual changes, we can examine the distribution of change, which cannot be done in synthetic panels. Despite the similarity of the average reduction with synthetic panel results, just 49% recollected a decline (not shown in table): the median reduction was zero. Therefore the drop in spending is not a population-wide phenomenon:

4.2 Anticipated Changes in Spending at Retirement

An interpretation of the decline in spending is that economic resources at retirement are less than anticipated prior to retirement either because saving was inadequate (Bernheim, Skinner and Weinberg) or because pension income was over-estimated (Banks, Blundell and Tanner). If that is the case, workers would not have predicted that spending would decline at retirement for the obvious reason that they were as yet unaware of the soon-to-be-discovered

[16] That part of the literature which uses synthetic panel to measure spending change refers to nondurable spending as in Banks, Blundell and Tanner, whereas the CAMS questions refers to total spending. According to the analysis of Blau (2004) based on the CEX the results are practically identical whether one uses spending on nondurables or on total spending.

shortfall. We investigate this interpretation using data on expectations about spending change at retirement by CAMS respondents who were not retired. Table 2 shows the average and median anticipated change in spending at retirement among those working in wave 1 of CAMS. The average anticipated decline was about 20%.[17] The greatest anticipated declines are among the younger workers, which is consistent with their understanding that they face a risk of early unexpected retirement.

If there is a correlation between retirement and expectations about spending change at retirement, the expectations of workers who are in their 60s will not reflect population expectations because some have already retired. Because there are few retirements before age 55, the average anticipated decline among workers in their early 50s is approximately a population average. We conclude that a decline is widely anticipated.

Because the medians are similar to the means, the average decline is not simply the result of a relatively small number of outliers; in fact about 66 percent of workers anticipated a decline (not shown).

Ameriks, Caplin and Leahy (2002) found in a sample of TIAA-CREF participants that among those not retired 55% expected lower spending in retirement and among those already retired 36% experienced lower spending. These percentages are somewhat lower than what we find. The most important difference between the TIAA-CREF sample and our sample is that the TIAA-CREF sample is much wealthier and does not represent the entire population.

4.3 Anticipated and Recollected Spending Change in Panel

[17] The average anticipated decline is greater than the average recollected decline reported in Table 1. There is no reason they should be the same as the anticipations and recollections come from different populations.

Using data from CAMS wave 1 and CAMS wave 2 we observe 201 retirements between waves among those in the age range 50 to 69. Table 3 shows that among those where we have complete data in both waves, the anticipated reduction in spending was 13.3% and the recollected actual reduction was 12.9%. Even though the number of observations is just 146, the recollected change is close to the change in the entire wave 1 CAMS sample, 13.8% (Table 1), and the standard error is small enough to rule out exceptionally large declines or declines of zero. The 95% confidence interval is 9.5% to 17.1%. The age pattern shows that both the anticipated and actual spending decline is greatest at younger ages when a health shock would cause a larger deviation in retirement from normal retirement age: the resulting shortfall in lifetime resources should lead to a greater spending decline. The median actual spending change was zero. In total, 44.5 percent recollected a spending change (not shown). This percentage is close to percentage of retirees in CAMS wave 1 who recalled a decline in spending (49%) as stated in Section 4.1. The panel results verify that the actual decline in spending is not population-wide.

The distributions of anticipated and actual spending change are similar as shown in Figure 1. The figure shows the cumulative distributions of percent spending change at retirement. The distribution of anticipations is marginally to the left of the distribution of recollections and both show a greater tendency toward declines rather than increases.

5. Explanations for the observed drop in spending

While our results based on CAMS wave 1 and on the panel show that an actual spending drop at retirement is limited to somewhat less than half the population, it is of considerable interest to understand why spending dropped among that group. In this section we provide evidence about three possible mechanisms: the affected population had little savings so that

spending was forced to decline when income declined; a health shock precipitated an unexpected early retirement; greater leisure in retirement was partially used in home production or more efficient shopping.

Spending Declines and Wealth

Models of intertemporal optimization require that agents act on the information they have about future events. Retirement with its accompanying decline in income is forecastable for most people. A widespread failure to save adequately could be signaled by a decline in consumption that is forced by the drop in income. However, as shown in Table 4 even those in the highest wealth quartile anticipated and experienced a spending decline.[18] Furthermore there is no systematic pattern in the actual decline as a function of wealth quartiles. The only noticeable differences in the entries are the substantially lower anticipated decline among those in the lowest wealth quartile and the smaller recollected decline in the 2nd quartile.

Because wealth is subject to considerable observation error which would result in classification error, we display similar results but stratified by education (Table 5). Those with more economic resources, as reflected in their higher education, reported declines in spending, and their anticipated declines were larger than their realized declines. The table also shows that those with the least education had similar declines to those with more education. However the least educated had the greatest difference between anticipations and recollections. This discrepancy is similar to the discrepancy in the lowest wealth category in Table 4. An implication is that a small fraction of the least well-off and/or least well-educated misperceive what their spending change will be. Despite the misperception, actual spending decline is not

[18] Quartile boundaries are for singles: p25 = $15,200; p50 = $74,500; p75 = $258,500; and for couples: p25 = $ 72,000; p50 = $186,000; p75 = $407,000.

greater than for others, suggesting that the misperception was due to an over-estimate of spending needs in retirement.[19]

Spending Declines and Health

Workers who have worse health are more likely to have a health shock and are more likely to retire early (McClellan, 1998). If exposed to an unexpected reduction in lifetime resources, they will reduce spending at retirement. Table 6 classifies the CAMS respondents who retired between waves 1 and 2 according to their health status prior to retirement. While working, respondents anticipated spending changes that were independent of health status, but actual spending change varied substantially with prior health status. For example, those in fair or poor health just before retirement recollected a spending decline of 18.4% compared with the spending decline of just 10.3% among those in excellent or very good health prior to retirement. These results suggest the larger declines may have been due to a health shock among those in worse health and that the early retirement led to a decline in spending.

We can study health transitions more directly in the full CAMS sample based on a measure of health transitions at retirement using panel information derived from multiple prior HRS waves. Table 7 shows health transitions from the wave preceding retirement to the wave after retirement. The population that enters this table is CAMS respondents who were retired in CAMS wave 1. The method of calculating this table is to examine prior HRS waves as far back as wave 2 (1994) to find the wave in which the respondent had transited from working to fully or partially retired. Health prior to retirement is reported in the preceding HRS wave, and health after retirement is reported in the subsequent wave. The table shows that 679 (74%) reported

[19] Part of the discrepancy between expectation and realizations for the least well off/or least well educated could also be due to this group being more affected by health shocks. However, in that case one would expect this group to experience larger declines in spending than the remainder of the population, which is not the case.

their health to be excellent, very good or good following retirement and that their average spending decline was 12.0%. Spending declined by 18.7% among the 233 whose health was fair or poor following retirement. Most maintained their health status at retirement: for example, 151 (69%) of those in fair or poor health prior to retirement remained in fair or poor health. This group comprises about 17% of the table, and they had the largest decline in spending: 20.9%. Because those in worse health are more likely to have a health shock than those in better health, some in that group probably retired unexpectedly early. Others may have anticipated an early retirement because of their health but the exact timing was uncertain.

Among those whose health improved at retirement (67 observations) spending declined by 12.9%, approximately the same as the spending decline of those whose health was good or better both before and after retirement. Approximately 12% of those with good or better health prior to retirement had fair or poor health following retirement. Their spending decline was 14.6%. We conclude that worse health and a transition to worse health are associated with larger spending declines, indicating that health risk is likely responsible for some of the decline in spending at retirement.

Another piece of evidence linking health to spending change comes from a question that HRS asks the newly retired. Respondents are asked their reasons for retirement, including the importance of health. Table 8 shows for the entire CAMS sample that about 29% of retirees mentioned health as an important reason for retirement.[20] They were much more likely to be in fair or poor health prior to retirement and they recollected that spending declined by 20.5%.

[20] We coded responses of "very important" or "moderately important" as health was an "important reason for retirement." Most respondents affirmed the first category. Similarly, we coded responses of "somewhat important" and "not important at all" as health was "not an important reason for retirement." Among these, most respondents (about 90%) affirmed that health was not important at all. The sample in Table 8 is reduced from the sample in Table 1 because of missing data on reasons for retirement.

Among those where health was not an important reason for retirement just 13% said that pre-retirement health was fair or poor and the average decline in spending was 11.0%.

Overall there was a small decline in health at retirement as measured by the two categories: 24% had fair or poor health prior to retirement compared with 26% after retirement. Among those where health was an important reason for retirement, the decline in health was about seven percentage points.

While these results strengthen the argument that health is an important factor in spending declines, they do not isolate the causes: Spending could decline because of a health shock leading directly to retirement.[21] Poor health could be associated with greater uncertainty about the timing of retirement. Or there could be other factors correlated with poor health. But because a large fraction of retirees report either that health was not an important factor in retirement or that their health was excellent, very good or good after retirement, we conclude that a health shock was important for some but not for all.

Change in Available Leisure Time

If leisure and consumption are not separable then a discontinuous change in leisure should lead to a discontinuous change in consumption in order to smooth the marginal utility of consumption. In section A of CAMS, the respondent was asked about his or her use of time. Many categories of time-use would neither be complements nor substitutes with market purchased goods. For example, "walking," or "watching TV" would seem to interact very little with market purchased inputs. We chose seven activities as shown in Table 9 that might be

[21] In results not shown, spending declines by about two percentage points more among those who say that health was an important reason for retirement and who report a transition to fair or poor health, but this difference is not economically important.

substitutes for market purchased goods or services.[22] If home production is partly responsible for any drop in spending, the amount of time spent on activities such as cooking, cleaning, home maintenance and the like should increase at retirement.

For those who retired between waves 1and 2 of CAMS, the table shows the change in hours per week spent on activities that we have classified as possible substitutes for market purchased goods. The changes are calculated as wave 2 hours spent in the activities minus wave 1 hours. For example, the individuals who retired between the waves spent 1.42 hours more per week in yard work and gardening at wave 2 than they did in wave 1. Overall the time spent in market substitutes increased by about 5.5 hours per week.

6. Conclusions

In models that are based on forward-looking optimal behavior by individuals, it is reasonable for the model to admit that some people will make optimization mistakes. In addition, uncertainty can cause apparent mistakes by individuals even though *ex ante* they made optimal choices. However, these mistakes should be symmetric in the sense that mistakes on one side of the optimum should be approximately balanced by mistakes on the other side of the optimum. Thus, in the absence of macro shocks, systematic deviations from the optimum at the population level would cast doubt in the validity of the model.

In the simple one-good life-cycle model, a widespread drop in consumption at retirement looks like a widespread or population-level systematic mistake. One explanation is that most of the population systematically over-estimated what their resources would be following retirement: on observing their actual income they were either compelled to reduce consumption or re-

[22] Section A of CAMS has 31 categories of time-use, but they were not chosen for an analysis of how activities interact with spending. An unfortunate result is that CAMS is lacking time-use categories that may be complements to spending.

optimization required that they reduce consumption. In either case most of the population experienced a loss in lifetime utility because of the discontinuity in marginal utility at retirement. They would not have foreseen the decline in consumption because ex ante they did not foresee that their income would be inadequate. A second explanation is that individuals did foresee that they would reach retirement with inadequate resources, which would force a reduction in consumption, but a lack of self control prevented them from reducing consumption prior to retirement. Both explanations are damaging to the life-cycle model. First, at the population level, retirement is not a surprise, and on average economic resources should not be a surprise. Second, models of optimal behavior are neither empirically valid nor useful unless the economic agents have the ability to make the choices that they deem optimal conditional on their constraints.

The results we have presented in this paper are evidence against both of these interpretations. First, at the population level, individuals are not surprised by the drop in consumption at retirement: indeed about two-thirds of workers anticipated a drop. Second, while some workers may have foreseen that consumption would have to drop in a sub-optimal manner, but lacked the self-control to reduce consumption while still working, this explanation cannot be valid at the population level: workers in the third and fourth wealth quartiles, whose consumption would not be forced to drop at retirement, also anticipated a decline in consumption.

If surprise at economic resources and/or lack of self control do not explain the drop in consumption at retirement, what does? Widespread health shocks that cause an unexpectedly early retirement could, in principle, cause the drop. In fact, in a young population where everyone has a positive probability of a health shock which would lead to unexpectedly early

retirement, everyone should with some positive probability anticipate a decline in consumption at retirement. There is no symmetric increase in consumption at retirement because with each passing year that a worker escapes the health shock, consumption will increase a little. Then workers who remain in the labor force unexpectedly long will have gradually increased consumption while still working so that consumption before and after retirement will be approximately the same.

Even though in principle health shocks could cause widespread consumption declines at retirement, their empirical importance is not great enough to explain fully the recollected declines in consumption: indeed about two-thirds of retirees say that health was not important at all in their retirement decision.

Work-related expenses could explain part of the decline. We have not addressed this issue in this paper because we do not have data on what spending can be confidently attributed to work-related expenses.[23] Most researchers believe that work-related expenses cannot be large enough to explain the average declines in the population. Furthermore, there must be considerable heterogeneity in work-related expenses because many retirees recollect that spending did not change at retirement.

The use of the increased leisure time to reduce spending while maintaining constant marginal utility is a plausible explanation for part of the population. Indeed, following retirement, activities that are plausible substitutes for spending increased by 5.5 hours per week in our panel data. But heterogeneity in time use is surely important: because of differing tastes and economic situations, some may spend large budget shares on goods that are complements to leisure and some may spend large budget shares on goods that are substitutes for leisure. Indeed

[23]Work-related expenses are difficult to define because most relevant categories can be both work-related and not work-related, such as transportation, clothing, dining and eating out. This difficulty is also mirrored in the fact that reliable estimates of work-related expenses are hard to find in the literature, most likely for exactly this reason.

some 11 percent of our sample indicated that spending increased following retirement, which indicates that complements dominate for this group.

We conclude that no single explanation is likely to be important for the large fraction of the population that anticipated and recollected a decline in consumption at retirement. Rather each explanation makes a contribution to the decline, and it will be an objective of future research to quantify their respective contributions. At this point, however, evidence points to a mix of explanations that, for most of the population, can be accommodated by conventional life-cycle economic models.

References

Aguiar, Mark and Erik Hurst, 2005a, "Consumption vs. Expenditure," *Journal of Political Economy*, 113(5), pp. 919-948.

Aguiar, Mark and Erik Hurst, 2005b, "Lifecycle Prices and Production," NBER Working Paper No. 11601.

Ameriks, John, Andrew Caplin and John Leahy, 2002, "Retirement Consumption: Insights from a Survey," NBER Working Paper No. 8735.

Banks, James, Richard Blundell and Sarah Tanner, 1998, "Is There a Retirement-Savings Puzzle?" *American Economic Review*, 88 (4), pp. 769-788.

Benitez-Silva, Hugo, Moshe Buchinsky and John Rust (2005), "Is There a Decline in Consumption after Retirement, and if so, Is It Inconsistent with the Predictions of the Life Cycle Model," presented at the 7th Workshop of the RTN-AGE Project, Venice, IT.

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), pp. 832-857.

Blau, David, 2004, "Retirement and Consumption in a Life Cycle Model," typescript, Department of Economics, University of North Carolina at Chapel Hill.

Consumption and Activities Mail Survey, 2001 and 2003, Early Release. Supplemental Mail Survey to the Health and Retirement Study, funded by National Institute on Aging at NIH (U01 AGO9740).

Fisher, Jonathan, David S. Johnson, Joseph Marchand, Timothy M. Smeeding and Barbara Boyle Torrey, 2005, "The Retirement-Consumption Conundrum: Evidence from a Consumption Survey," Retirement Research Center at Boston College, WP 2005-14.

Haider, Stephen J. and Melvin Stephens, 2004, "Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations," NBER Working Paper No. 10257.

Health and Retirement Study, Waves 1 through 7, public use datasets. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG09740). Ann Arbor, MI, (various years).

Hurd, Michael D., 1996, "The Effect of Labor Market Rigidities on the Labor Force Behavior of Older Workers," in David Wise, ed., *Advances in the Economics of Aging,* Chicago: University of Chicago Press, pp. 11-58.

Hurd, Michael D. and Susann Rohwedder, 2005, "The Consumption and Activities Mail Survey: Description, Data Quality, and First Results on Life-Cycle Spending and Saving," typescript, RAND.

Juster, F. Thomas, and Richard Suzman, 1995, "An Overview of the Health and Retirement Study," *The Journal of Human Resources*, 30. Supplement, pp. S7-S56.

McClellan, Mark, 1998, "Health Events, Health Insurance and Labor Supply: Evidence from the Health and Retirement Survey," in D. Wise, ed., *Frontiers in the Economics of Aging*, 1998, Chicago: University of Chicago Press, pp. 301-346.

Miniaci, Raffaele, Chiara Monfardini and Guglielmo Weber, 2003, "Is There A Retirement Consumption Puzzle In Italy?" IFS Working Paper WP03/14, Institute for Fiscal Studies, London.

RAND HRS Data, Version E. Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa

Monica, CA (April 2004).

Rust, John, 1990, "Behavior of Male Workers at the End of the Life Cycle: An Empirical Analysis of States and Controls," in David Wise, ed., *Issues in the Economics of Aging,* Chicago: University of Chicago Press, pp. 317-379.

Smith, Sarah, 2004, "Can The Retirement Consumption Puzzle Be Resolved? Evidence from UK Panel Data," IFS Working Paper WP04/07, Institute for Fiscal Studies, London.

Table 1. Recollected percent change in spending at retirement (weighted)

Retirement age	N	Mean	Std. Err.	Median
50-54	154	-16.2	2.1	-1.0
55-59	301	-12.7	1.5	0.0
60-64	592	-13.2	0.9	0.0
65-69	237	-14.7	1.5	-5.0
70-74	22	-14.9	4.6	-20.0
All	1,306	-13.8	0.7	0.0

Note: Current age greater than or equal to 50 and less than 80.
Source: Authors' calculations based on CAMS wave 1

Table 2. Anticipated percent change in spending at retirement (weighted)

Age	N	Mean	Std. Err.	Median
50-54	158	-19.4	1.5	-20.0
55-59	289	-23.1	1.2	-25.0
60-64	263	-16.1	1.4	-20.0
65-69	105	-16.7	1.9	-10.0
All	815	-20.1	0.7	-20.0

Source: Authors' calculations based on CAMS wave 1

Table 3. Anticipations and recollections of percent change in spending at retirement. Panel

		Anticipations			Recollections		
Age in wave 1	N	Mean	Std.err. of mean	Median	Mean	Std.err. of mean	Median
50-54	18	-15.8	5.8	-20.0	-15.4	5.1	-11.5
55-59	38	-19.6	3.2	-20.0	-16.7	4.4	-10.0
60-64	65	-9.4	3.2	0.0	-11.9	3.0	0.0
65-69	25	-12.3	3.7	0.0	-7.6	3.7	0.0
All	146	-13.3	1.9	-10.0	-12.9	2.0	0.0

Source: Authors' calculations based on CAMS wave 1 and wave 2

Table 4. Anticipated and recollected change in spending at retirement (percent) by wealth quartile, panel.

	Wealth quartile in HRS 2000				
	Lowest	2	3	Highest	All
Anticipated while working	-5.4	-17.8	-14.5	-15.9	-13.3
Recollected after retirement	-14.6	- 9.0	-16.1	-11.7	-12.9

Source: Authors' calculations. N = 146

Table 5. Anticipated and recollected change in spending at retirement (percent) by education, panel.

	Education				
	Less than high school	High school	Some college	College	All
Anticipated while working	-6.7	-12.7	-17.6	-13.9	-13.3
Recollected after retirement	-10.2	-13.9	-14.3	-11.1	-12.9
Observations	*21*	*56*	*36*	*33*	*146*

Source: Authors' calculations. N=146

Table 6: Anticipated and recollected change in spending at retirement (percent) by health status before retirement, panel

	Health status in HRS 2000			
	Fair or poor	Good	Excellent or very good	All
Anticipated while working	-13.9	-13.4	-13.2	-13.3
Recollected after retirement	-18.4	-15.0	-10.3	-12.9
Observations	*22*	*41*	*83*	*146*

Source: Authors' calculations

Table 7. Recollected change in spending (percent) and number of observations. Health transitions at retirement

Health status prior to retirement	Health status after retirement		
	fair or poor	good or better	All
fair or poor	-20.9	-12.9	-18.4
	151	*67*	*218*
good or better	-14.6	-12.0	-12.3
	82	*612*	*694*
All	-18.7	-12.0	-13.7
	233	*679*	*912*

Note: number of observations in italics

Table 8. Health status before and after retirement, and recollected change in spending at retirement

Importance of health for retirement	N	Fraction in fair or poor health		Percent change in spending	
		Before retirement	After retirement	Mean	Median
Health important	265	0.51	0.58	-20.5	-20.0
Health not important	647	0.13	0.12	-11.0	0.0
All	912	0.24	0.26	-13.8	0.0

Note: Based on HRS 2000 question G138a-1: "I am going to read you a list of reasons why some people retire. Please tell me whether, for you, these were very important reasons for retirement, moderately important, somewhat important, or not important at all." Same question was asked in all waves. We use the latest observation available for the retirees in the sample, using all available waves 1 through 7.
Source: Authors' calculations.

Table 9. Changes in hours per week at retirement, panel.

House cleaning	0.55	Washing, ironing	0.16
Yard work/gardening	1.42	Shopping	0.91
Food preparation	1.47	Finances	0.08
Home improvements	0.83		
Total substitutes 5.42 hours per week			

Note: Additional restriction on sample is that the same person must have responded to the time use section in both waves of CAMS.
Source: Authors' calculations. N=130

Figure 1. Distribution of Panel Responses