

Is There a Retirement-Consumption Puzzle?
Evidence Using Subjective Retirement Expectations

Steven J. Haider
Department of Economics
Michigan State University
101 Marshall Hall
East Lansing, MI 48824
e-mail: haider@msu.edu

Melvin Stephens Jr.
H. John Heinz III School of Public Policy and Management
Carnegie Mellon University and National Bureau of Economic Research
4800 Forbes Ave.
Pittsburgh, PA 15213
e-mail: mstep@cmu.edu

May 2003
This Version: September 28, 2003

We would like to thank Kerwin Charles, Susann Rohwedder, and seminar participants at Maryland, Michigan State, and the 2003 Retirement Research Consortium Conference for useful comments. Haider acknowledges the financial support from the National Institute on Aging, R03 AG20542-01. Stephens acknowledges the financial support from a Sandell Grant Award from the Center for Retirement Research at Boston University. The research reported herein was supported (in part) by the Center for Retirement Research at Boston College pursuant to a grant from the U.S. Social Security Administration funded as a part of the Retirement Research Consortium. The opinions and conclusions are solely those of the authors and should not be construed as representing the opinions or policy of the Social Security Administration or any agency of the Federal Government, or the Center for Retirement Research at Boston College.

Abstract

Previous research shows a systematic fall in consumption at retirement, a finding that is inconsistent with the life-cycle/permanent income hypothesis. In this paper, we use workers' beliefs about their expected retirement dates as an instrument for retirement. After demonstrating that subjective retirement expectations are strong predictors of subsequent retirement, we still find a systematic fall in consumption for workers who retire when expected. However, our results suggest that this fall in consumption is half as large as that found when we instead rely on the instrumental variables strategy used in prior studies.

JEL Classification. D91, J26

1. Introduction

With a growing number of workers approaching retirement, the preparedness of these households to finance consumption during their retirement years is becoming a topic of increasing concern. Assessing if households are adequately saving for retirement is a difficult task because many factors that affect the optimal level of wealth accumulation, such as tastes, risk preferences, and patience, are hard to quantify. In light of these difficulties, economists have relied upon the rational expectations version of the Life-Cycle/Permanent Income Hypothesis (LCPIH) to judge whether households are adequately saving for retirement. Specifically, if households are rational and foresighted, then their consumption should not change upon leaving the labor force. Contrary to this hypothesis, however, empirical investigations have concluded that household consumption falls at the time of retirement, even for those retirements that are expected (e.g., Hamermesh 1984; Mariger 1987; Banks, Blundell, and Tanner 1998; Bernheim, Skinner, and Weinberg 2001).¹

This fall in consumption at retirement is referred to as the “retirement-consumption puzzle.” The mounting evidence that households do not “smooth” consumption at retirement has led researchers to call into question the standard life-cycle model. In their assessment of this puzzle, Banks, Blundell, and Tanner conclude that their “... evidence strongly suggest that there are unanticipated shocks occurring around the time of retirement” (p. 784). Bernheim, Skinner, and Weinberg state that their “... findings are difficult to interpret in the context of the life- cycle

¹ Examining consumption changes at retirement is just one approach to assess the adequacy of retirement savings. Other approaches include a) comparing the distribution of wealth in a simulated economy with the empirical distribution of wealth and b) calculating the annual payment from converting current household wealth into an annuity and comparing this estimate current household consumption. For a critical review of these literatures, including the method used in this paper, see Engen, Gale, and Uccello (1999).

model” (p.855). Thus, the consensus from these researchers is that households do not behave with the foresight that is key to the life-cycle model.

A number of other explanations for the observed decline in consumption at retirement have been offered in order to rehabilitate the life-cycle model. Lundberg, Startz, and Stillman (2003) suggest that household bargaining concerns between spouses with differing expected lifetimes lead households to adjust their consumption levels upon retirement. Consistent with this hypothesis, they find evidence that consumption decreases at retirement for married couples but do not find such changes for single individuals. Angeletos et al (2001) demonstrate through the use of simulation methods that hyperbolic (rather than geometric) discounting households will have a planned fall in consumption at retirement. Two recent studies, Ameriks, Caplin, and Leahy (2002) with a survey of TIAA-CREF participants and Hurd and Rohwedder (2003) with a mailout supplement to the Health and Retirement Study, find that households *expect* their expenditures to fall upon retirement. In addition, Hurd and Rohwedder and Aguiar and Hurst (2003) suggest that incorporating household production decisions into the standard model may explain the fall in consumption if households shift from using consumption goods that are market intensive to using consumption goods that are time intensive.

In this paper we do not attempt to modify the basic LCPIH to account for the retirement-consumption puzzle, but rather, we return to the question of whether such a puzzle exists. The prediction of the LCPIH is that consumption should not fall *if households retire when expected*. Prior papers have examined observed changes in consumption at retirement. But to the extent that retirement is caused by an unexpected event such as a job loss or disability, the observed fall of consumption at retirement does not refute the LCPIH. Recognizing this fact, some (but not all) previous studies exploited the rapid changes in retirement that occur at the ages when

workers become eligible for government retirement benefits to instrument for retirement. However, the choice of age as an instrument for retirement has two potential problems. First, since older households are generally observed reducing their consumption as they age, it is very important that the age-consumption relationship is properly parameterized. To the extent that the general age-consumption profile is not correctly specified, the use of age as an instrument may violate the exclusion restriction. Second, and more importantly, the implicit assumption when using age as an instrument is that the relationship between age and actual retirement is the same as the relationship between age and expected retirement. If, as we will illustrate in the two datasets that we use in our analysis, these relationships are not the same, then age is not a valid instrument for expected retirement. Thus, while age is an intuitively appealing instrument, it may not be appropriate in this context.

Our goal in this paper is to assess whether the retirement-consumption puzzle does indeed exist. We do so by using an instrument that allows us better to distinguish between expected and unexpected retirements, subjective retirement expectations. In the two datasets that we examine, the Retirement History Survey and the Health and Retirement Study, workers are asked the date at which they expect to retire. Under the rational expectations hypothesis, the responses to this question are valid instruments for future retirement. Equally as important, we demonstrate that these subjective expectations are strong predictors of future retirement, even after allowing for a very flexible age-retirement relationship.

Subjective expectations questions such as the ones we use here have been the focus of a growing area of economic research. This literature finds that subjective expectations are powerful predictors of subsequent outcomes. Hurd and McGarry (1997) and Smith, Taylor, and Sloan (2001) find a strong relationship between subjective survival probabilities and subsequent

mortality. Dominitz (1998) finds that subjective income expectations are good predictors of realized income while Stephens (Forthcoming) finds that subjective job loss probabilities are strongly related to subsequent job displacements. Finally, in an application that is similar in spirit to the test we perform here, Jappelli and Pistaferri (2000) test whether households' subjective income growth is a significant predictor of consumption growth. Consistent with the LCPIH, they find no evidence that expected income growth is correlated with consumption growth even though there is a strong correlation between expected and realized income growth.

Using data from the Retirement History Survey and the Health and Retirement Study, we test whether consumption falls at retirement using subjective retirement expectations as an instrument for retirement. Our first stage results demonstrate that retirement expectations are strong predictors of actual retirement, even after controlling flexibly for the underlying age-retirement relationship. Both reduced form and two-stage least squares estimates of the relationship between consumption changes and expected retirement show that consumption falls at retirement even for workers who retire when expected. While these results reject the LCPIH, the retirement consumption decrease is nearly 50 percent smaller in magnitude when using subjective retirement expectations as an instrument relative to using age as an instrument. Therefore, although our results provide strong evidence of a fall in consumption at retirement, our arguably better methodology produces a smaller consumption decrease than the previous literature.

2. Modeling and Estimating the Retirement-Consumption Puzzle

We examine whether consumption falls at retirement in the context of the life-cycle/permanent income hypothesis model, similar to the prior literature (Banks, Blundell, and

Tanner 1998; Bernheim, Skinner, and Weinberg 2001). In each year, households maximize their utility over the remainder of the life-cycle according to the equation

$$(1) \quad \text{Max } V_t = E_t \left[\sum_{k=t}^T \left(\frac{1}{1+\delta} \right)^{k-t} U(C_k) \Psi(w_k) \right],$$

subject to the asset evolution constraints

$$(2) \quad A_{k+1} = (1+r)(A_k + Y_k - C_k), \forall k = t, \dots, T,$$

where C_t , Y_t and A_t are consumption, income, and assets in year t , respectively; $U(\cdot)$ is the period specific utility function; w_t are variables such as age and family size thought to affect utility (through $\Psi(\cdot)$); r is a constant interest rate; and δ is the subjective discount rate. The resulting Euler Equation that determines the household's optimal allocation of consumption between periods t and $t+1$ is

$$(3) \quad U'(C_t) \Psi(w_t) = \left(\frac{1+r}{1+\delta} \right) E_t [U'(C_{t+1}) \Psi(w_{t+1})].$$

The key intuition behind using this framework to test whether consumption falls at retirement comes from the rational expectations version of the LCPIH pioneered by Hall (1978). Assuming that the variables modifying consumption (w_t) remain constant and that the interest rate equals the discount rate, we can alternatively write (3) as

$$(4) \quad U'(C_{t+1}) = U'(C_t) + \varepsilon_{t+1},$$

where ε_{t+1} is the household's expectation error. Here we have the familiar result that households will "smooth" the marginal utility of consumption between periods. Under the rational expectations assumption, ε_{t+1} should be uncorrelated with any information possessed by the household at time t . In particular, assuming R_t represents the household's information at time t

about its retirement status at period $t+1$, the model imposes the restriction that R_t should be uncorrelated with ε_{t+1} , or equivalently that $E_t[\varepsilon_{t+1}|R_t]=0$.

To test this implication of the model, researchers typically assume that the period specific utility function exhibits constant relative risk aversion. Deriving the corresponding marginal utility of consumption and inserting it into (3) yields the familiar first-order approximation for the Euler Equation

$$(5) \quad \Delta \ln C_{t+1} = \frac{1}{\rho}(r - \delta) + \frac{1}{\rho} \Delta \ln \Psi(w_{t+1}) + v_{t+1},$$

where ρ is the coefficient of relative risk aversion.² To apply this Euler Equation to the retirement-consumption puzzle, consider estimating the equation

$$(6) \quad \Delta \ln C_{t+1} = \alpha + \beta \text{retire}_{t+1} + \gamma X_{t+1} + v_{t+1},$$

where X_{t+1} is a vector of time-varying demographic characteristics that are meant to capture $\Delta \ln \Psi(w_{t+1})$ and retire_{t+1} is an indicator for whether or not the household retired between years t and $t+1$. The parameter β measures the observed fall in consumption at retirement. However, a finding that β is non-zero is not a violation of the LCPIH because households may retire for numerous unforeseen reasons such as a job loss or the onset of a disability. The LCPIH only predicts that consumption changes should be uncorrelated with planned retirement behavior.

The endogeneity of retire_{t+1} has been addressed in some previous studies by using an instrumental variables strategy, accomplished through the use of two-stage least squares (2SLS). The first stage involves regressing retirement on a set of instruments that are correlated with retirement but that are not correlated with the error term in (6). In the second stage, predicted

² Work by Carroll (1997) and Ludvigson and Paxson (2001) has shown that equation (5) is a poor approximation because the true Euler Equation is very non-linear. More importantly, this specification may also lead to erroneous rejections of the life-cycle model when it is in fact true. While examining the importance of these biases in explaining the retirement consumption puzzle is of interest, it is not the goal of the current paper.

retirement, $predret_{t+1}$, is constructed for each individual based on the first stage regression, which is then used to replace the endogenous retirement variable in (6). Thus, the estimating equation becomes

$$(7) \quad \Delta \ln C_{t+1} = \alpha + \lambda predret_{t+1} + \gamma X_{t+1} + v_{t+1},$$

and the LCPIH is tested by the null hypothesis $\lambda = 0$. As is typical, the difficulty rests with identifying a valid instrument. However, the rational expectations assumption is very useful in this regard. Under this assumption, all information available at time t (variables dated time t and earlier) as well as any strictly exogenous variables are candidate instruments.

The instrumental variable we use in this paper, subjective retirement expectations, meets the necessary criteria as a valid instrument. These expectations are clearly known to the individual in each period and are therefore potential instruments for future retirements. The rational expectations assumption also satisfies the econometric requirement that the instrument be uncorrelated with the error term in equation (7). Furthermore, as we demonstrate below, these expectations also satisfy the additional econometric requirement that the instruments are highly predictive of subsequent retirement behavior.

The validity of subjective retirement expectations as an instrument for retirement can best be illustrated by examining the 2SLS estimator in this context. Suppose that retirement expectations are captured by a binary variable $exret_t$ that equals 1 if the worker expects to retire between time t and $t+1$ and equals 0 otherwise. The 2SLS estimator of λ when using $exret_t$ as an instrument is

$$\lambda = \frac{E[\Delta \ln C_{t+1} | exret_t = 1] - E[\Delta \ln C_{t+1} | exret_t = 0]}{E[retire_{t+1} | exret_t = 1] - E[retire_{t+1} | exret_t = 0]}.$$

The sample analog of this estimator is known as the Wald estimator (Angrist and Krueger 1999). The numerator of the Wald estimator will equal zero if the average of consumption changes for workers that expect to retire equals the average of consumption changes for workers that do not expect to retire. The LCPIH predicts that the expected consumption change for *both* of these groups should equal zero since households act on their available information at time t to keep their consumption smooth between time t and $t+1$. Notice that this result still holds if differences in retirement expectations across individuals reflect differences in expectations of other outcomes. For example, suppose households reporting that they expect to retire also have a high probability of losing their job while households that do not expect to retire have a low job loss probability. The LCPIH predicts that, knowing their job loss probability to be either high or low, households will adjust their consumption in period t such that the Euler Equation holds, i.e., so that their expected change in consumption equals zero. Thus, using subjective retirement expectations as an instrument provides a test of the LCPIH.

3. The Data

We rely primarily upon two data sets in this paper, the Retirement History Survey (RHS) and the Health and Retirement Study (HRS). The RHS and the HRS are useful in two regards. First, both datasets focus on individuals near the usual retirement ages and therefore contain large samples of workers undergoing these transitions. Second, both of these datasets ask a number of expectations questions including direct questions on retirement expectations.

The Retirement History Study (RHS) began in 1969 and re-interviewed households on a biennial basis until 1979. The initial sample of approximately 11,000 individuals included men and unmarried women born between 1905 and 1911 (ages 58-63 in the initial wave). The survey collected a wide array of information including labor force activities, health experiences, and

demographic details. At the end of the survey, a total of six waves of information had been collected. We only use the first five waves of the RHS since the retirement expectations variables are not asked beyond the fourth wave.

The Health and Retirement Study (HRS) is an on-going longitudinal survey that began in 1992.³ The initial sample consisted of approximately 7,700 households that contained at least one person born between 1931 and 1941 (ages 51-61 in the initial wave); the age-eligible household member and his or her spouse (regardless of birth year) were interviewed, resulting in approximately 12,700 initial respondents. Like the RHS, the HRS is also fielded biennially. The survey collects detailed information on a variety of topics, including demographics, employment, health status, and financial status. We use the publicly available versions of the first five waves of the HRS (1992-2000).

Both surveys collect information on household food consumption. The RHS collected this information in all waves while the HRS has collected this information in all waves except for wave 4. Food consumption information has been used in a number of previous studies testing household consumption behavior (Hall and Mishkin 1982; Zeldes 1989; Shea 1995). The main drawback to using food consumption is that it is a limited measure of household expenditures. However, a benefit of food consumption is the fact that it is a non-durable good, which means that changes in food expenditures should be closely linked to changes in household utility. It is difficult to measure the utility changes associated with changes in durable good expenditures since households can receive service flows from past purchases of these items. Furthermore, food consumption is either the main or a component of the main consumption measure in two previous retirement savings puzzle studies (Bernheim, Skinner, and Weinberg 2001; Lundberg, Startz, and Stillman 2003), so its use also provides a point of comparison.

³ See Juster and Suzman (1995) for an overview of the HRS.

Information on retirement status is collected at two separate points in all waves during each survey. First, both studies ask a general labor force status question, which allows individuals to choose from several possible activities: working, unemployed, retired, disabled, and homemaker.⁴ Second, these studies also ask workers to report their current retirement status as being either fully retired, partly retired, or not retired at all. Since most prior research has defined retirement using the labor force status question, we will also use that definition of retirement throughout the majority of our analysis.

For our estimation strategy, the most important information available in these surveys is the questions regarding each individual's expected age of retirement. In the initial wave of both surveys, individuals who have not yet retired are asked when they expect to retire.⁵ Survey participants can respond by giving an age (or year) of expected retirement, stating that they will never retire, or respond that they do not know the date at which they expect to retire. Although the HRS only asks this question during the first wave, the RHS continues to ask workers their expected retirement age through the fourth wave.⁶ Since households continually receive new information that may cause them to alter their expectations, the availability of updated expectations allows us to use even more precise information in our analysis with the RHS.

For our analysis, we limit our attention to men. Our main reason for imposing this restriction is that the RHS only collected detailed information on women when a spouse was not

⁴ When reporting current labor force status, the RHS only allowed individuals to choose one activity while the HRS allows individuals to choose multiple activities. For comparability between the two surveys, we consider all individuals in the HRS who report that they are retired as retired in our analysis unless they also report that they are either working or unemployed.

⁵ In the first wave of the RHS, the question is, "At what age do you expect to stop working at a regular job?" The RHS question in the later waves is, "Do you expect ever to stop working for a living? At what age?" The HRS question is, "When do you think you will retire (completely)? [Respondent can then provide year or age.]"

⁶ Although the HRS contains similar expectation questions in later waves, the skip patterns are such that many fewer individuals respond to the questions. For example, Loughran, Panis, Hurd, and Reti (2001) show that the response rate in Wave 2 is about three-quarters less than the response rate in Wave 1. The HRS also elicits information about the subjective probability individuals will retire by ages 62 and 65 in all waves. However, Loughran, Panis, Hurd, and Reti (2001) demonstrate that these questions were not consistently asked of the appropriate populations.

present. We also impose the restriction that each respondent had to be working at the initial wave so that they a) can potentially enter into retirement during the survey period and b) will be eligible to answer the question on expected retirement. We use an unbalanced sample in that we include observations from individuals who leave the surveys prior to the final (most recent) wave. For our analysis, we only consider the first move into retirement and ignore any subsequent movements in and out of retirement. Finally, since retirement expectations are not elicited from workers who have left the labor force, our analysis is restricted to observations up to and including the wave of retirement.

We provide basic descriptive statistics for both samples in Table 1. All dollar amounts are adjusted to 2001 dollars using the annual Personal Consumption Expenditure (PCE) deflator. The differences in the observable characteristics, such as the sample becoming more educated and more likely to not be married, are broadly consistent with well-known secular trends of the last three decades.

4. Retirement Expectations and Realizations

Prior research has found that subjective retirement expectations are strong predictors of subsequent retirement behavior. Bernheim (1989) examines the relationship between retirement expectations and realizations using the RHS. He finds that respondents appear to give the modal (i.e., most likely) date of retirement rather than the mean date to the retirement expectations question. Across all of the expected retirement dates he examines, roughly two-thirds of men retire within one year (before or after) of their expected date. Loughran, Panis, Hurd, and Reti (2001) find that retirement expectations are strong predictors of retirement in the HRS. Using two waves of the British Retirement Survey, Disney and Tanner (1999) find evidence similar to Bernheim's in that respondents appear to give modal responses to the retirement expectations

question. In addition, they find that when predicting retirement using regressions that include a large number of observable characteristics, expected age of retirement is a very strong predictor of the actual age of retirement.

In Figure 1, we use data from the RHS to demonstrate the relationship between expected wave of retirement and actual wave of retirement.⁷ The expectations question is from the first wave of the survey. Panel A of Figure 1 shows the distribution of expected retirement waves in the RHS. Of the workers who report a wave of retirement, the majority of them expect to retire by the fourth wave.⁸ Given that respondents in the RHS are ages 58-63 at the initial interview, nearly all of them will be eligible for the Social Security normal retirement age by wave four. More interesting, however, is fact that approximately one-third of the workers in the RHS report that they will never retire. Furthermore, about one-eighth of workers do not know when they will retire.

The remaining three panels of Figure 1 show the relationship between workers' wave one expected retirement ages and their actual retirement dates.⁹ Three main patterns can be observed in these figures. First, the modal wave of realized retirement corresponds to the expected wave for all expected retirement waves. Second, the accuracy of the expectations is stronger for workers expecting to retire at waves closer to the initial survey date. Third, the timing of observed retirements for those who say they will never retire and those who do not know when they will retire are very similar.

As shown in Figure 2, similar patterns are found in the HRS. Since workers in the HRS are younger than those in the RHS, a higher fraction of HRS respondents report an expected

⁷ Bernheim (1989) performs a similar analysis except he compares expected year of retirement with actual year of retirement.

⁸ Because retirement expectations are elicited by age of retirement, we determine the expected wave of retirement by assuming that workers expect to retire on the day upon which they reach that age (i.e., their birthday).

⁹ Workers who leave the RHS before retiring are excluded from the last three panels of Figure 1.

retirement date after the available sample period. Panels B and C of Figure 2 show that HRS respondents also appear to be giving modal responses to the expected retirement question. As can be seen in Panel D of the Figure, workers who state that they will never retire have subsequent dates of retirement that are comparable to workers who report that they do not know when they will retire. Thus, the strong relationship between retirement expectations and subsequent retirement dates persists in both datasets.

To further illustrate the strength of the retirement expectations variables, Table 2 shows the fraction of workers retiring in each wave conditional on whether or not a worker reports that he planned to retire by that wave.¹⁰ We define workers as expecting to retire if they responded that they expect to retire in the year of survey, the year after the survey, or two years after the survey. Workers that give any other response, including “never” or “don’t know”, are classified as not expecting to retire. The first two columns of Table 2 use wave 1 expectations in the RHS to compare the fraction of workers who actual retire between those workers who expected to retire and those workers who did not. Between waves 1 and 2, 58 percent of workers who expected to retire did so while 11 percent of workers who did not expect to retire did in fact retire. Moving to later waves, the fraction of workers who expected to retire between the waves and subsequently retired remains roughly constant.¹¹ However, the fraction of workers who did not expect to retire but do increases at the later waves. Columns 3 and 4 show the results using the workers’ most recent wave expectations in the RHS to measure whether retirement is expected. As anticipated, the most recent expectations responses are better predictors of

¹⁰ The results in Table 3 restrict the sample to observations that also have non-missing consumption data between the two waves in order to use the same sample as in the subsequent regression analysis.

¹¹ At first glance, this result may appear to be at odds with the results in Figure 1 that suggest expectations get less accurate for later expected waves. However, remember that the results in Table 2 are based upon those workers who have yet to retire by the time they reach these later waves while Figure 1 is based on all workers no matter which wave they retire. Thus, the results between Table 2 and Figure 1 are consistent.

retirement status than the expectations responses given at wave one. The final columns present the results for the HRS based upon wave 1 expectations. The results are qualitatively similar to those based upon wave 1 expectations in the RHS. The noticeable difference is that the fraction actually retiring is lower in the HRS than in the RHS, presumably due to the HRS representing a younger sample.

While retirement expectations are strongly correlated with retirement, it could be that they are simply reflecting the strong age-retirement relationship. To determine the additional information contained in the expected retirement variables, we estimate the equation

$$(8) \quad \text{retire}_{t+1} = \pi_0 + \pi_1 \text{exret}_{t+1} + \pi_2 X_{t+1} + u_{t+1},$$

where retire_{t+1} is an indicator for whether the worker retired between waves t and $t+1$, exret_{t+1} is an indicator for whether the worker expected to retire between waves t and $t+1$, X_{t+1} contains age at time t , the change in household size between t and $t+1$, and wave dummies, and u_{t+1} is an error term.¹² In other words, this equation represents the first stage for equation (7).

Panel A of Table 3 presents the results of estimating (8) where exret_{t+1} is based upon the response to the retirement expectations question in wave 1 of the RHS.¹³ We present results using three specifications to control for age. When using only a linear age term (column 1), workers who expect to retire between waves are 34 percent more likely to retire than workers who do not. Using a quadratic in age has no impact on this estimate (column 2). Finally, including a complete set of age dummies to fully capture the age-retirement relationship has a negligible impact on the estimate. Across all three specifications, the point estimate is strongly significant with t-statistics exceeding 20 across the columns.

¹² These covariates are included since, as we discuss below, they are included when estimating the consumption changes equations.

¹³ The reported standard errors for the regression coefficients in Table 3 as well as in all subsequent tables are robust to arbitrary forms of correlations within individuals over time.

Since workers update their expectations as they receive new information, the correlation between retirement expectations and actual retirement decisions should be strengthened when the worker's most recent retirement expectations are used in place of the worker's initial (wave 1) expectations. The results in Panel B of Table 3 confirm this intuition in the RHS. When using the most recent (wave t) expectations to construct $exret_{t+1}$, workers who expect to retire are 42 percent more likely to retire than workers who do not expect to do so. As with the results in Panel A, alternative methods to control for the relationship between age and retirement have a minimal impact on the magnitude of the correlation between expected and observed retirement.

Finally, Panel C presents the results for the HRS using wave 1 expectations. Consistent with Table 2, the correlation between expected and observed retirement is not as large in magnitude in the HRS as in the RHS. Nonetheless, the estimates are still highly significant and do not change across different specifications.

5. The Change in Consumption at Observed Retirement

To answer the question of whether expectations can explain the retirement-consumption puzzle, we first document the extent to which consumption falls at retirement in the RHS and the HRS. We estimate equation (6) where our primary interest is on the estimate of β , the coefficient on *observed* retirement. The vector of observable characteristics, X_{t+1} , contains age at time t , the change in household size between t and $t+1$, and wave dummies. While this list of covariates is rather sparse, it is consistent with the controls used in numerous studies examining changes in consumption. As we discussed above, finding a drop in consumption at retirement when estimating equation (6) is not evidence of a failure of the LCPIH since these regressions pool households retiring both expectedly and unexpectedly. In fact, when we calculate the share

of retirements due to workers who expected to retire between the waves using the most recent expectations in the RHS, less than half (48.7 percent) of retirements are expected.

Panel A of Table 4 presents the results of estimating (6) separately for the RHS and the HRS samples. The estimate in the RHS (column 1) is consistent with previous results: consumption is significantly reduced at retirement. Surprisingly, however, we do not find evidence of a consumption decline at retirement in the HRS (column 2).

This differential result between the RHS and the HRS could be due to many reasons. One potential reason is that food consumption data are very noisy. However, the result in the HRS cannot simply be due to measurement error since we find evidence of the decline at retirement in the RHS. In addition, when we examine the food consumption by household size for both data sets in Figure 3, the food consumption levels are remarkably similar across both surveys. Further casting doubt on this explanation, when using the same measure of food consumption, Stephens (Forthcoming) finds that job loss significantly reduces consumption by 15 percent in the HRS over the same time period.

A second potential explanation for the differences is that the 1990s could have been substantively different than the 1970s, perhaps because of unexpected increases in wealth due to higher stock market returns in the latter period. If a large number of retirements were induced by these increases in wealth, then time specific factors may explain the differences between the surveys. To examine whether there exist important time period effects, we make use of a third dataset, the Panel Study of Income Dynamics (PSID).¹⁴ The PSID is useful because it spans the time period of both the RHS and the HRS and it collects information on food consumption and

¹⁴ The PSID is an on-going panel dataset that began in 1968. The original sample consisted of approximately 5,000 families, and these families and their off-spring have been followed ever since. The PSID interviewed households annually until 1997 at which time it began interviewing on a biennial basis.

retirement. From the PSID, we construct comparison samples for both the RHS and the HRS.¹⁵ Figure 3 demonstrates that the food consumption levels for both comparison samples are remarkably similar to the food consumption levels in the RHS and HRS.

The results of estimating (6) on the PSID comparison samples are shown in columns 3 and 4 of Table 4. In both PSID samples, consumption falls by roughly 9 percent at observed retirement and the estimated decreases are both significant. Thus, the results in both of the PSID samples are very similar to the RHS estimate, leaving the HRS estimate as the anomalous result.

As another method of examining the possibility of differential time period effects, we examine changes in total family income at retirement across all of the datasets. The results of regressing the change in log family income on the same regressors used in Panel A of Table 4 are shown in Panel B of the Table. The results show the family income falls by at least 20 percent across all of the samples including the HRS. Furthermore, it is interesting to note that the estimated income decline at retirement is nearly identical in the datasets that correspond to the same time periods; the estimated decline is 0.25 log points in the RHS and the PSID-70s samples while the consumption drop is over 0.30 log points in the HRS and the PSID-90s. Thus, differences in income changes at retirement cannot explain the differences across studies.

A third potential explanation for the divergence between the RHS and the HRS is that, as noted by Chamberlain (1984), estimates of rational expectations models are inconsistent in short panels because rational expectation errors have an expected value of zero as the number of time periods increases, not as the number of cross-sectional observations increases. Because food consumption is not available in the fourth wave of the HRS, our examination of consumption changes between the first three waves of the survey may exacerbate this problem. To examine

¹⁵ We use PSID sample members who were 53-62 year olds in 1970 to compare to the RHS and a 53-62 year olds in 1991 to compare to the HRS. We then follow each comparison sample biennially for five waves. Statistics for the PSID comparison samples are shown in Appendix Table 1.

this possibility, we restrict the second PSID sample to span 1991-1995 to approximate the HRS panel in terms of time period and length of sample. The results for this sample (column 5) are not statistically different than the HRS result in column 2, although the point estimate is negative.

Overall, we take these estimates as illustrating that the negative correlation between consumption and retirement found in other datasets is also present in the RHS but not in the HRS with the latter result likely reflecting sampling variability. As we discussed in the introduction, however, these results alone do not refute the LCPIH since observed retirements may be correlated with unexpected events that cause households to change their consumption. Therefore, in order to test the LCPIH, we turn our focus onto the consumption response to expected retirements.

6. The Change in Consumption at Expected Retirement

The primary parameter of interest in our analysis is λ , which is the coefficient on predicted retirement in equation (7). Before presenting our structural estimates of this parameter from using our instrumental variables estimation strategy, we first present the results of estimating the reduced form model

$$(9) \quad \Delta \ln C_{t+1} = \alpha + \phi \text{exret}_{t+1} + \gamma X_{t+1} + v_{t+1}$$

where we include exret_{t+1} directly into the Euler Equation in lieu of predret_{t+1} . Notice that this reduced form estimate is actually a direct test of the rational expectations LCPIH. Since expectations at time t should be uncorrelated with expectations error between future periods, a test of the null hypothesis $\phi = 0$ provides a simple test of the model. In fact, ϕ is the numerator of the Wald estimator.

Table 5 presents the estimates of equation (9). When RHS wave 1 expectations are used to construct $exret_{t+1}$, we cannot consistently reject the null hypothesis (Panel A of the Table). When only a linear term in age is included in the model (column 1), the estimate of ϕ is marginally significant. When the potential non-linearities between age and consumption changes are accounted for by using a more flexible specification for age, the estimates of ϕ are insignificant. In addition, when wave 1 expectations are used in the HRS as shown in Panel C, the estimates are insignificant across all three specifications.

The model is rejected, however, when the most recent retirement expectations are used as the basis for $exret_{t+1}$ as shown in Panel B of Table 5. Individuals who expect to retire between waves t and $t+1$ have a significant fall in their consumption relative to those workers who do not. The results remain stable across alternative specifications for age. Thus, these estimates reject the LCPIH. We interpret these results as strong evidence that consumption falls at retirement even among those workers who expect to retire.

While the reduced form estimates in Table 5 are evidence of a decrease in consumption for workers who retire when expected, the structural (2SLS) results of equation (7) presented in Table 6 estimate the magnitude of the consumption decline. Consistent with the results in the previous table, we find a negative but insignificant impact of retirement on consumption when using wave 1 expectations in the RHS (Panel A). The point estimates in the HRS are slightly positive, although the confidence intervals around these estimates are relatively large (Panel C). The results using the most recent expectations responses in the RHS are negative and significant. Our estimates of λ in equation (7) are stable across all three specifications, even when we allow for the most general relationship between age and consumption (column 3). Our estimates imply that consumption falls by 8 to 9 percent when workers retire as expected. Overall, our results in

Tables 5 and 6 find strong evidence of a decline in consumption at retirement, even when workers retire as expected.

Alternative Instrument Specification

Our specification for the relationship between retirement expectations and observed retirements is a simple binary indicator for whether or not the worker expects to retire by the next survey wave. However, this specification ignores much of the information contained in the subjective retirement response since workers can report the age (or year) in which they expect to retire, if they expect to never retire, or if they do not know when they will retire. A more flexible specification for this relationship between may more precisely estimate the change in consumption when workers retire as expected since it exploits more variation across workers. To more fully exploit the workers' subjective retirement information, we use a set of dummies variables based on the worker's most recent retirement expectations to capture the number of years from the survey date until the date that the worker expects to retire. We also include separate dummy variables for whether worker expects to never retire or the worker does not know when they will retire.

The estimated relationship between actual and expected retirement using this alternative specification is presented in the first two columns of Table 7. Relative to workers who expect to retire eight or more years from the survey date, workers who expect to retire in either the survey year or one year following the survey are more than 50 percent more likely to retire. The estimate for workers who expect to retire two years following the survey is more than 40 percent more likely to retire. These three groups comprise the set of workers where $exret_{t+1}=1$ when expectations are specified as a binary variable in the previous specifications. The impact drops

substantially for workers who expect to retire three years following the survey and continues to decline for workers whose expected retirement date is further in the future. Using a complete set of age dummies in lieu of a linear age term has a negligible effect on the estimates.

The last two columns of Table 7 show the second stage estimates when the alternative instrument specification is implemented. The point estimates are nearly identical to the analogous results in Table 6. Thus, the more flexible instrument specification yields nearly the same results as the simple binary instrument used in the main analysis.

7. Using Age as an Instrument for Retirement

Previous studies have used age as an instrument for retirement.¹⁶ Noting that there are sharp changes in the likelihood of retirement at ages when workers become eligible for government retirement benefits and that these ages are known in advance, researchers have exploited this non-linear relationship between age and retirement status as a source of variation in retirement that is uncorrelated with the error term in the Euler Equation. Since the availability of these benefits are known in advance to households and the benefit value should be easily forecasted, the rationale behind using age as an instrument is that the increase in retirement induced by the age specific benefit eligibility should not represent new information to households and therefore be uncorrelated with the error term in equation (7).

While the use of age as an instrument is intuitively appealing, we note two potential concerns with this approach. First, the rapid change in retirement status by age may be correlated with changes in the marginal utility of consumption at these ages. If these changes are not captured by the variables in X_{t+1} but are correlated with the non-linearity in age, then the

¹⁶ The method of Bernheim, Skinner, and Weinberg amounts to using a complete set of dummies variables for each age covered in their sample (54 to 70). Banks, Blundell, and Tanner use lagged regressors such as past retirement status as instruments. However, to the extent that these lagged values are age dependent, they are implicitly using age as an instrument. Thus, the arguments discussed here also apply to their study.

exclusion restriction will be violated and render age an inappropriate instrument. For example, age is usually entered linearly as a regressor when the first-difference of consumption is the dependent variable (or, equivalently, age is specified as a quadratic when using consumption in levels). The exclusion restriction will be violated if this parameterization is inadequate to capture rapid changes in the age-consumption profile around the retirement age.

Second, the fraction of workers unexpectedly retiring may systematically vary by age. The first stage of the 2SLS methodology assigns the actual retirement experiences of the current retirees to be their present expectations. In other words, the implicit assumption when using age as an instrument is that the observed fraction of workers retiring at each age is equal to the current expectations of those still working. If this assumption does not hold, then the 2SLS estimate when using age as an instrument will be contaminated by this systematic bias.

The evidence presented in Figure 4 suggests that a systematic bias does indeed exist. The solid lines in Panels A and B of the Figure show the fraction of workers at each age, conditional on not having yet retired, who retire by the next survey wave in the RHS and HRS, respectively. These lines illustrate the standard result: retirement rates increase with age, peak at the Social Security normal retirement age, and remain relatively high at subsequent ages. The dashed lines illustrate the fraction of workers at each age who expect to retire by the next survey wave, again conditional on having not yet retired.¹⁷ Prior to normal retirement, the solid and dashed lines are nearly identical. At each age after normal retirement age, however, the fraction of workers retiring exceeds the fraction expecting to retire. When age is used as an instrument, the 2SLS estimator falsely treats these later ages as having a relatively high fraction of expected

¹⁷ Both surveys ask the expected date of retirement in the first survey wave. These data are illustrated with the long-dashed lines in the Figure. The Retirement History Survey continues to ask the question through the first four survey waves. The short-dashed line in Panel A of the Figure illustrates the results based upon the worker's most recent expected retirement date.

retirements. Thus, the 2SLS estimate will be biased with the direction of the bias depending upon the correlation of consumption changes with these unexpected retirements.

To compare our estimates using expectations as an instrument to prior research, we estimate the model using age as an instrument in our samples. Since the HRS results have been consistently insignificantly, we focus on the RHS for the remainder of our analysis. The age instruments are a full set of dummy variables for the worker's age at wave t in the t to $t+1$ change used in the analysis. For the RHS these age dummies range from 58 to 69. Our models also include either a linear or a quadratic specification for age so that only the non-linearities in age identify the consumption change at "expected" retirement. We also present two sets of analyses based upon whether we do or do not include post-retirement observations. In our analysis using retirement expectations, we do not use post-retirement observations because the instrument is not defined for these periods. However, previous studies have included post-retirement observations in the analysis when age is used as an instrument.

The results of using age as an instrument for retirement are presented in Table 8. The first two columns of the Table use the same sample as in the earlier tables. The results show that consumption is significantly reduced by 13 percent when workers retire as "expected" when the linear age term is included. When a quadratic in age is included, the point estimate becomes more negative magnitude but the standard errors increase dramatically so that the estimate is now marginally significant. In the final two columns of Table 8, we include post-retirement observations in the analysis. The reason for adding these variables back into the analysis is that prior studies had included these observations although we are unable to do so in our main analysis since retirement expectations are not asked of workers who have left the labor force. When these observations are included, the estimates show that consumption falls by nearly 15

percent at retirement and these results are not affected by the choice of the age controls in consumption equation.

Compared to the results where expectations are used as an instrument (Tables 6 and 7), the estimated consumption decrease at retirement when using age as an instrument are substantially larger in magnitude. Moving from using age as an instrument to using the expectations variables reduces the point estimates by nearly 50 percent. The standard errors around the estimates in Table 8, however, are just large enough to include the previous point estimates within one standard deviation. These larger estimates are likely due to the instrument treating the large number of retirements at ages beyond normal retirement age as expected retirements. In addition, note that the standard errors on the structural estimates are substantially smaller when expected retirement is used as an instrument as opposed to age. This difference is likely due to the additional power of the expectations variable in predicting subsequent retirement. Overall, we believe the comparison between the two sets of structural estimates is consistent with our arguments that not only are the subjective retirement expectations a better instrument but that the use of age as an instrument for retirement will lead to biased estimates.

8. Conclusions

A number of previous studies have found a significant drop in household consumption at retirement. If individuals are behaving in accordance with the rational expectations version of the Life-Cycle/Permanent Income Hypothesis, then this consumption decline is a puzzle if it is observed for workers who retire when expected. As we mentioned in the introduction, a number of alternative hypotheses have been offered to explain this consumption decline at retirement. Our goal in this paper is not to refute any of these possible explanations. Rather, we have investigated whether the consumption drop at retirement can be explained by the use of valid

instruments for expected retirement. Using age as an instrument for expected retirement, prior researchers find that consumption declines even when retirements are expected, a finding that has been labeled the retirement-consumption puzzle. As we have documented here, however, there is strong evidence that the age profile of actual retirements and expected retirements are not the same which means that age is not a valid instrument for expected retirement.

In this paper, we use subjective retirement expectations as instruments for expected retirement. Our results show that workers' expectations are strong predictors of their subsequent retirement patterns. We find that retirement expectations are significantly correlated with consumption changes. This result rejects the LCPIH since any information known to the household should be uncorrelated with the expectations error in the Euler Equation. When using retirement expectations as an instrument for retirement, we find that consumption falls by 8 to 9 percent for workers who retire as expected. Thus, even when a valid instrument is used for retirement, there is still evidence of a significant fall in consumption at retirement. However, this estimated retirement consumption decline is nearly half the magnitude that one would find when using the methods found in previous studies.

How should we interpret this fall in consumption with retirement? As noted at the outset, a number of alternative explanations have been offered. Although our analysis does not validate nor rule out any specific hypotheses, we can provide some suggestive evidence. For example, Figure 4 demonstrates that workers past the Social Security retirement age systematically underestimate their likelihood of retirement. This finding is consistent with retirement systematically being correlated by the arrival of unanticipated shocks, as suggested by the Banks, Blundell, and Tanner quote in the introduction.¹⁸ On the other hand, other studies have presented evidence that individuals *expect* their expenditures to decline upon retirement (Hurd and

¹⁸ Comparisons between expectations and realizations by age in Disney and Tanner also show these underestimates.

Rohwedder 2003; Ameriks, Caplin, and Leahy 2002). Interestingly, the RHS also asks working respondents whether their expenses will be more, less, or about the same after retirement, and our tabulations show that three times as many respondents answer that expenses will be less as compared to those who answer more.¹⁹ Such *ex ante* reports of a consumption fall suggest that the basic LCPICH model might not be sufficiently rich to describe consumption changes around retirement. Because our results suggest that the fall in consumption indeed exists, more work is needed to uncover the exact mechanisms behind this finding.

¹⁹ Of the individuals who respond to the question, 13% answer that expenses will be more, 39% less, and 48% about the same.

References

- Aguiar, Mark, and Erik Hurst (2003). "Consumption vs. Expenditure," University of Chicago Manuscript.
- Anderson, Kathryn H., Richard V. Burkhauser, and Joseph F. Quinn (1986). "Do Retirement Dreams Come True? The Effect of Unanticipated Events on Retirement Plans," *Industrial and Labor Relations Review*, 39:4, 518-26.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman and Stephen Weinberg (2001). "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15:4, 47-68.
- Angrist, Joshua D. and Alan B. Krueger (1999). "Empirical Strategies in Labor Economics," in *Handbook of Labor Economics, Volume 3*, eds. Orley C. Ashenfelter and David Card, Amsterdam: Elsevier Science Publishers BV.
- Ameriks, John, Andrew Caplin, and John Leahy (2002). "Retirement Consumption: Insights From a Survey," National Bureau of Economic Research Working Paper No. 8735.
- Banks, James, Richard Blundell, and Sarah Tanner (1998). "Is There a Retirement-Savings Puzzle?" *American Economic Review*, 88:4, 769-88.
- Bernheim, B. Douglas (1989). "The Timing of Retirement: A Comparison of Expectations and Realizations," in *The Economics of Aging* (ed. David A. Wise). Chicago and London: University of Chicago Press.
- 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.
- Carroll, Christopher D. (1997). "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)" National Bureau of Economic Research Working Paper No. 6298.
- Chamberlain, Gary (1984). "Panel Data," in *Handbook of Econometrics, Volume II*, eds. Zvi Griliches and Michael Intriligator. Amsterdam: Elsevier Science Publishers BV.
- Disney, Richard and Sarah Tanner (1999). "What Can We Learn From Retirement Expectations Data?" The Institute For Fiscal Studies Working Paper Series No. W99/17.
- Dominitz, Jeff (1998). "Earnings Expectations, Revisions, and Realizations," *The Review of Economics and Statistics*, 80:3, 374-88.
- Engen, Eric M., William G. Gale, and Cori E. Uccello (1999). "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, 1999:3, 65-187.
- Hall, Robert E. (1978). "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86:6, 971-87.
- Hall, Robert E. and Frederic S. Mishkin (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 50:2, 461-81.
- Hamermesh, Daniel S. (1984). "Consumption during Retirement: The Missing Link in the Life Cycle," *Review of Economics and Statistics*, 66:1, 1-7.

- Hurd, Michael D. and Kathleen McGarry (1997). "The Predictive Validity of Subjective Probabilities of Survival," National Bureau of Economic Research Working Paper No. 6193.
- Hurd, Michael D. and Susann Rohwedder (2003). "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending and Retirement," National Bureau of Economic Research Working Paper No. 9586.
- Jappelli, Tullio and Luigi Pistaferri (2000). "Using Subjective Income Expectations to Test for Excess Sensitivity of Consumption to Predicted Income Growth," *European Economic Review*, 44:2, 337-58.
- Juster, F. Thomas and Richard Suzman (1995). "An Overview of the Health and Retirement Study," *Journal of Human Resources*, 30:S, S7-S56.
- Loughran, David, Constantijn Panis, Michael Hurd, and Monika Reti (2001). "Retirement Planning," RAND Manuscript.
- Ludvigson, Sydney, and Christina H. Paxson (2001). "Approximation Bias in Linearized Euler Equations," *Review of Economics and Statistics*, 83:2, 242-56.
- Lundberg, Shelly, Richard Startz, and Steven Stillman (2003). "The Retirement-Consumption Puzzle: A Marital Bargaining Approach," *Journal of Public Economics*, 87:5-6 1199-1218.
- Mariger, Randall P. (1987). "A Life-Cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results." *Econometrica*, 55:3, 533-57.
- Shea, John (1995). "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis," *American Economic Review*, 85:1, 186-200.
- Smith, V. Kerry, Donald H. Taylor, Jr., and Frank A. Sloan (2001) "Longevity Expectations and Death: Can People Predict Their Own Demise?" *American Economic Review*, 91:4, 1126-34.
- Stephens Jr., Melvin (Forthcoming). "Job Loss Expectations, Realizations, and Household Consumption Behavior," *The Review of Economics and Statistics*.
- Zeldes, Stephen P. (1989). "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97:2, 305-46.

Table 1: Sample Characteristics

	RHS	HRS
Interview years	1969, 1971, 1973, 1975, 1977	1992, 1994, 1996, 1998, 2000
Age range in first wave	58 to 63	51 to 61
Persons	3,706	3,093
Person-Observations (first differences)	14,210	10,599
Married in first wave	91.5%	85.8%
Education	10.0	12.6
Percent retiring during the sample period	59.6%	34.5%

Notes: All results are unweighted. Because food consumption questions were not asked in the 1998 wave of the HRS, all of the HRS analysis that directly uses consumption data is limited to the 1992 – 1996 waves.

Table 2: The Relationship between Expected and Actual Retirement

	RHS: Wave 1 Retirement Expectations		RHS: Most Recent Wave Expectations		HRS: Wave 1 Retirement Expectations	
	Expect to Retire	Do Not Expect to Retire	Expect to Retire	Do Not Expect to Retire	Expect to Retire	Do Not Expect to Retire
Fraction Retiring Between Waves:						
1 and 2	0.576 (.022)	0.105 (.005)	0.576 (.022)	0.105 (.005)	0.378 (.030)	0.037 (.004)
2 and 3	0.584 (.020)	0.261 (.009)	0.689 (.017)	0.213 (.009)	0.365 (.030)	0.087 (.006)
3 and 4	0.592 (.030)	0.343 (.012)	0.722 (.020)	0.254 (.012)	0.382 (.032)	0.101 (.007)
4 and 5	0.686 (.039)	0.389 (.017)	0.690 (.027)	0.327 (.018)	0.288 (.032)	0.162 (.009)

Notes: Each cell in this table reports the fraction of individuals who retired between the waves among those who have not retired by the first of the two waves. Standard errors are reported in parentheses.

Table 3: Predicting Retirement Using Retirement Expectations

	(1)	(2)	(3)
Age Controls	Age	Age, Age ²	Age Dummies
A. RHS: Wave 1 Expectations			
Expect to Retire	0.347 (0.014)	0.332 (0.014)	0.311 (0.015)
R-squared	0.141	0.146	0.152
N	9,472	9,472	9,472
B. RHS: Most Recent Wave Expectations			
Expect to Retire	0.448 (0.012)	0.439 (0.012)	0.425 (0.012)
R-squared	0.220	0.224	0.229
N	9,472	9,472	9,472
C. HRS: Wave 1 Expectations (Using All Five Waves)			
Expect to Retire	0.229 (0.016)	0.229 (0.016)	0.219 (0.017)
R-squared	0.104	0.106	0.116
N	9,360	9,360	9,360

Notes: The dependent variable is an indicator for whether the worker retired between the waves. All regressions include the change in household size and indicators for the survey wave. Standard errors are robust to arbitrary forms of correlations within individuals over time.

Table 4: OLS Estimates of Changes in Consumption and Income at Retirement

	RHS 1969-1977 (1)	HRS 1992-1996 (2)	PSID 1970-1980 (3)	PSID 1991-1999 (4)	PSID 1991-1995 (5)
A. OLS: dependent variable is change in log consumption					
Retired	-0.098 (0.012)	0.019 (0.025)	-0.101 (0.031)	-0.089 (0.043)	-0.051 (0.061)
Age	-0.005 (0.002)	0.000 (0.002)	-0.002 (0.003)	0.002 (0.004)	-0.001 (0.006)
Dhhsz	0.093 (0.009)	0.058 (0.013)	0.108 (0.025)	0.154 (0.045)	0.198 (0.057)
R-square	0.053	0.011	0.047	0.047	0.097
N	9,472	3,766	716	549	333
B. OLS: dependent variable is change in log household income					
Retired	-0.279 (0.015)	-0.312 (0.045)	-0.253 (0.039)	-0.373 (0.102)	-0.337 (0.153)
Age	0.002 (0.003)	-0.004 (0.003)	-0.001 (0.004)	0.012 (0.017)	0.017 (0.028)
Dhhsz	0.054 (0.008)	0.023 (0.019)	0.017 (0.025)	0.093 (0.093)	0.083 (0.138)
R-square	0.040	0.021	0.074	0.019	0.010
N	9,386	3,740	716	549	333

Notes: The dependent variable is the change in the log of consumption (Panel A) and the change in the log of household income (Panel B) between waves. All regressions include dummy variables for the survey wave. Standard errors are robust to arbitrary forms of correlations within individuals over time.

Table 5: Reduced-Form Estimates of the Impact of Expected Retirement on Consumption

Age Controls	(1) Age	(2) Age, Age ²	(3) Age Dummies
A. RHS: Wave 1 Expectations			
Expect to Retire	-0.021 (0.014)	-0.019 (0.014)	-0.008 (0.015)
R-squared	0.045	0.046	0.047
N	9,472	9,472	9,472
B. RHS: Most Recent Wave Expectations			
Expect to Retire	-0.038 (0.012)	-0.037 (0.012)	-0.031 (0.012)
R-squared	0.046	0.046	0.047
N	9,472	9,472	9,472
C. HRS: Wave 1 Expectations			
Expect to Retire	0.008 (0.026)	0.011 (0.026)	0.003 (0.026)
R-squared	0.011	0.011	0.015
N	3,766	3,766	3,766

Notes: The dependent variable is the change in log consumption between the waves. All regressions include the change in household size and indicators for the survey wave. Standard errors are robust to arbitrary forms of correlations within individuals over time.

**Table 6: IV Estimates of the Impact of Expected Retirement on Consumption
Using Expected Retirement as an Instrument**

Age Controls	(1) Age	(2) Age, Age ²	(3) Age Dummies
A. Instrument: RHS Wave 1 Expectations			
Retired	-0.061 (0.040)	-0.058 (0.043)	-0.026 (0.048)
N	9,472	9,472	9,472
B. Instrument: RHS Most Recent Wave Expectations			
Retired	-0.084 (0.027)	-0.083 (0.028)	-0.074 (0.029)
N	9,472	9,472	9,472
C. Instrument: HRS Wave 1 Expectations			
Retired	0.029 (0.089)	0.039 (0.092)	0.009 (0.093)
N	3,766	3,766	3,766

Notes: The dependent variable is the change in log consumption between the waves. All regressions include the change in household size and indicators for the survey wave. Standard errors are robust to arbitrary forms of correlations within individuals over time.

**Table 7: Estimates of the Impact of Retirement on Consumption with the RHS
Alternative Instrument Specification**

	First Stage		Second Stage	
	Dep Var = $retire_{t+1}$		Dep Var = $\Delta \ln C_{t+1}$	
	(1)	(2)	(3)	(4)
Retired			-0.081	-0.072
			(0.026)	(0.028)
Years Until Expected Date:				
Year of Survey	0.533	0.534		
	(0.040)	(0.040)		
One Year	0.579	0.548		
	(0.026)	(0.027)		
Two Years	0.475	0.434		
	(0.026)	(0.027)		
Three Years	0.136	0.147		
	(0.025)	(0.026)		
Four Years	0.069	0.070		
	(0.025)	(0.026)		
Five Years	0.066	0.066		
	(0.025)	(0.027)		
Six Years	0.022	0.028		
	(0.025)	(0.027)		
Seven Years	0.056	0.039		
	(0.025)	(0.027)		
Eight or More Years (Excluded)	---	---		
Never Retire	0.051	0.042		
	(0.021)	(0.021)		
Don't Know When	0.138	0.127		
	(0.024)	(0.024)		
F-Test (Expectation Variables)	163.2	137.9		
Age Controls	Age	Age	Age	Age
		Dummies		Dummies
N	9,472	9,472	9,472	9,472

Notes: All regressions include the change in household size and indicators for the survey wave. Standard errors are robust to arbitrary forms of correlations within individuals over time.

**Table 8: IV Estimates of the Impact of Expected Retirement on Consumption
Using Age as an Instrument**

	(1)	(2)	(3)	(4)
	Age	Age, Age ²	Age	Age, Age ²
Additional Age Controls				
Include Post-Retirement Observations?	No	No	Yes	Yes
Instrument: RHS Age Dummies (Ages 59-69)				
Retired	-0.132 (0.062)	-0.151 (0.087)	-0.155 (0.055)	-0.154 (0.103)
First Stage F-test (For Excluded Age Variables)	30.8	31.8	39.2	39.9
N	9,472	9,472	14,210	14,210

Notes: The dependent variable is the change in log consumption between the waves. All regressions include the change in household size and indicators for the survey wave. Standard errors are robust to arbitrary forms of correlations within individuals over time.

Appendix Table 1: Characteristics of the PSID Comparison Samples

	PSID-70s	PSID-90s
Interview years	1970, 1972, 1974, 1976, 1978	1991, 1993, 1995, 1997, 1999
Age range in first wave	53 to 62	53 to 62
Persons	252	235
Person-Observations (First Differences)	1,111	878
Married in first wave	94.4%	91.9%
Education	10.3	12.2
Percent retiring during era	71.0%	54.0%

Figure 1: RHS Wave 1 Retirement Expectations and Realizations

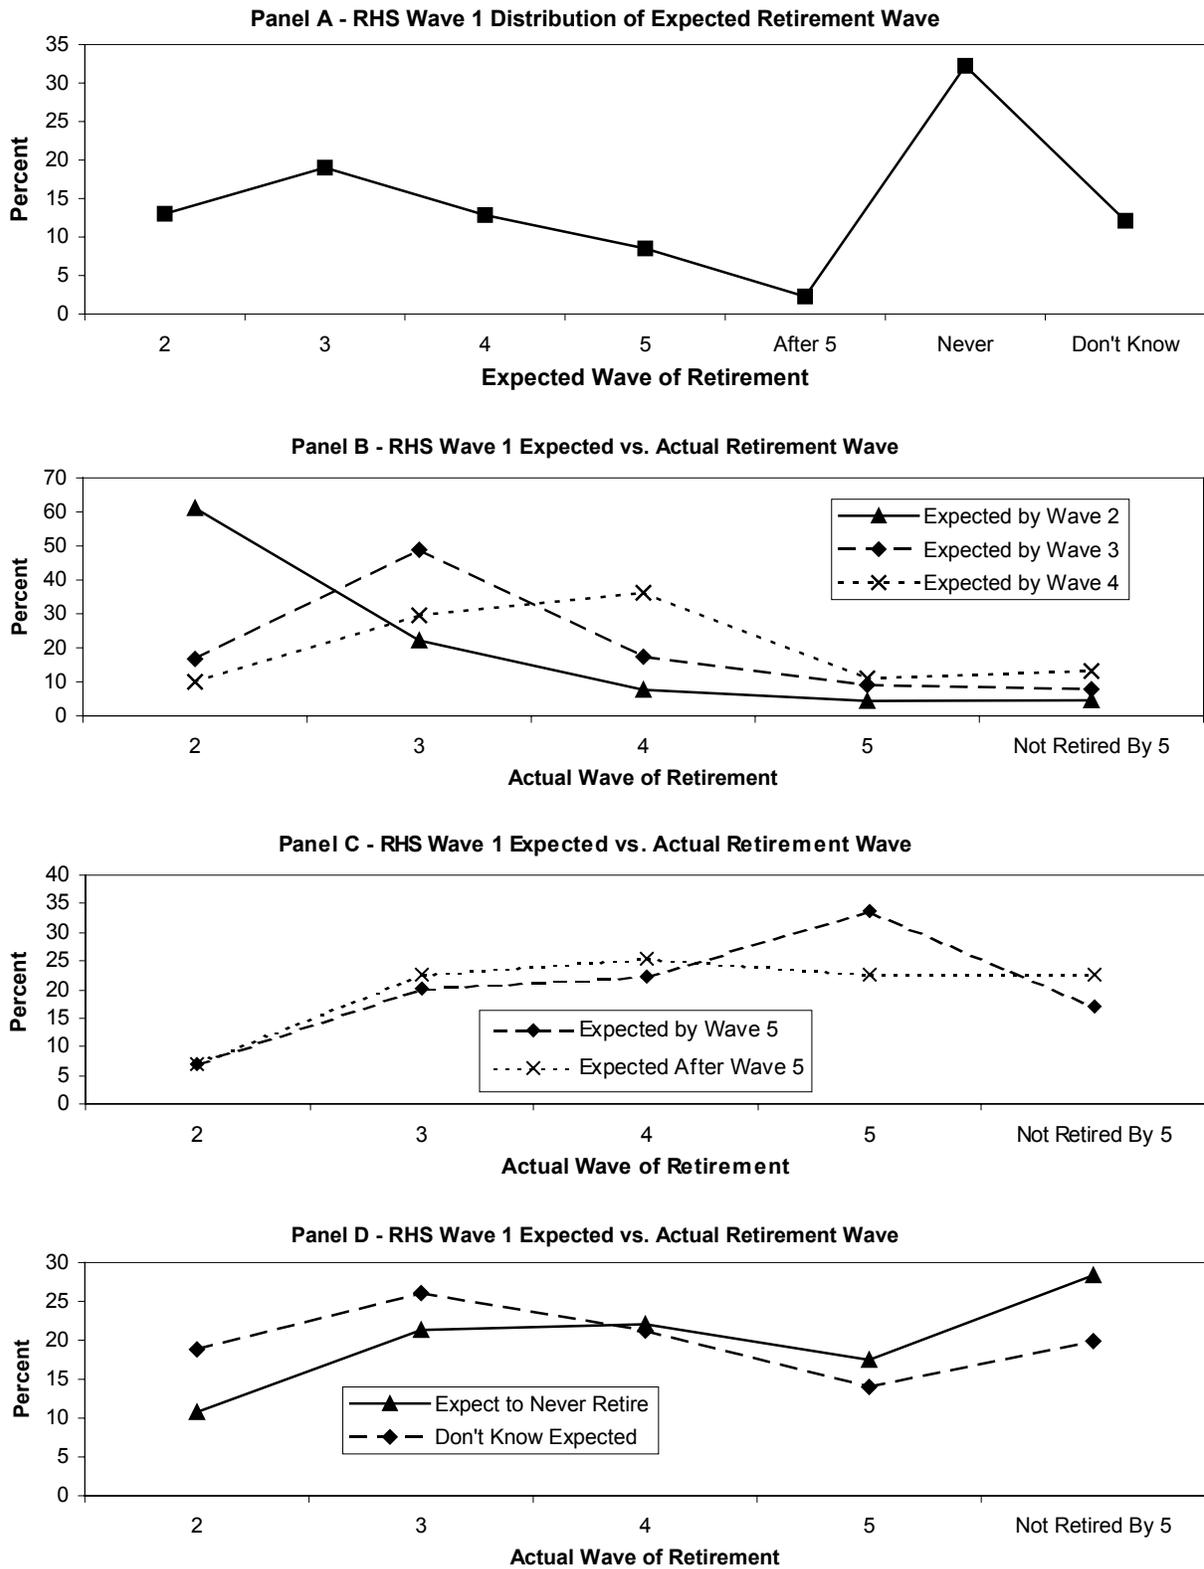


Figure 2: HRS Wave 1 Retirement Expectations and Realizations

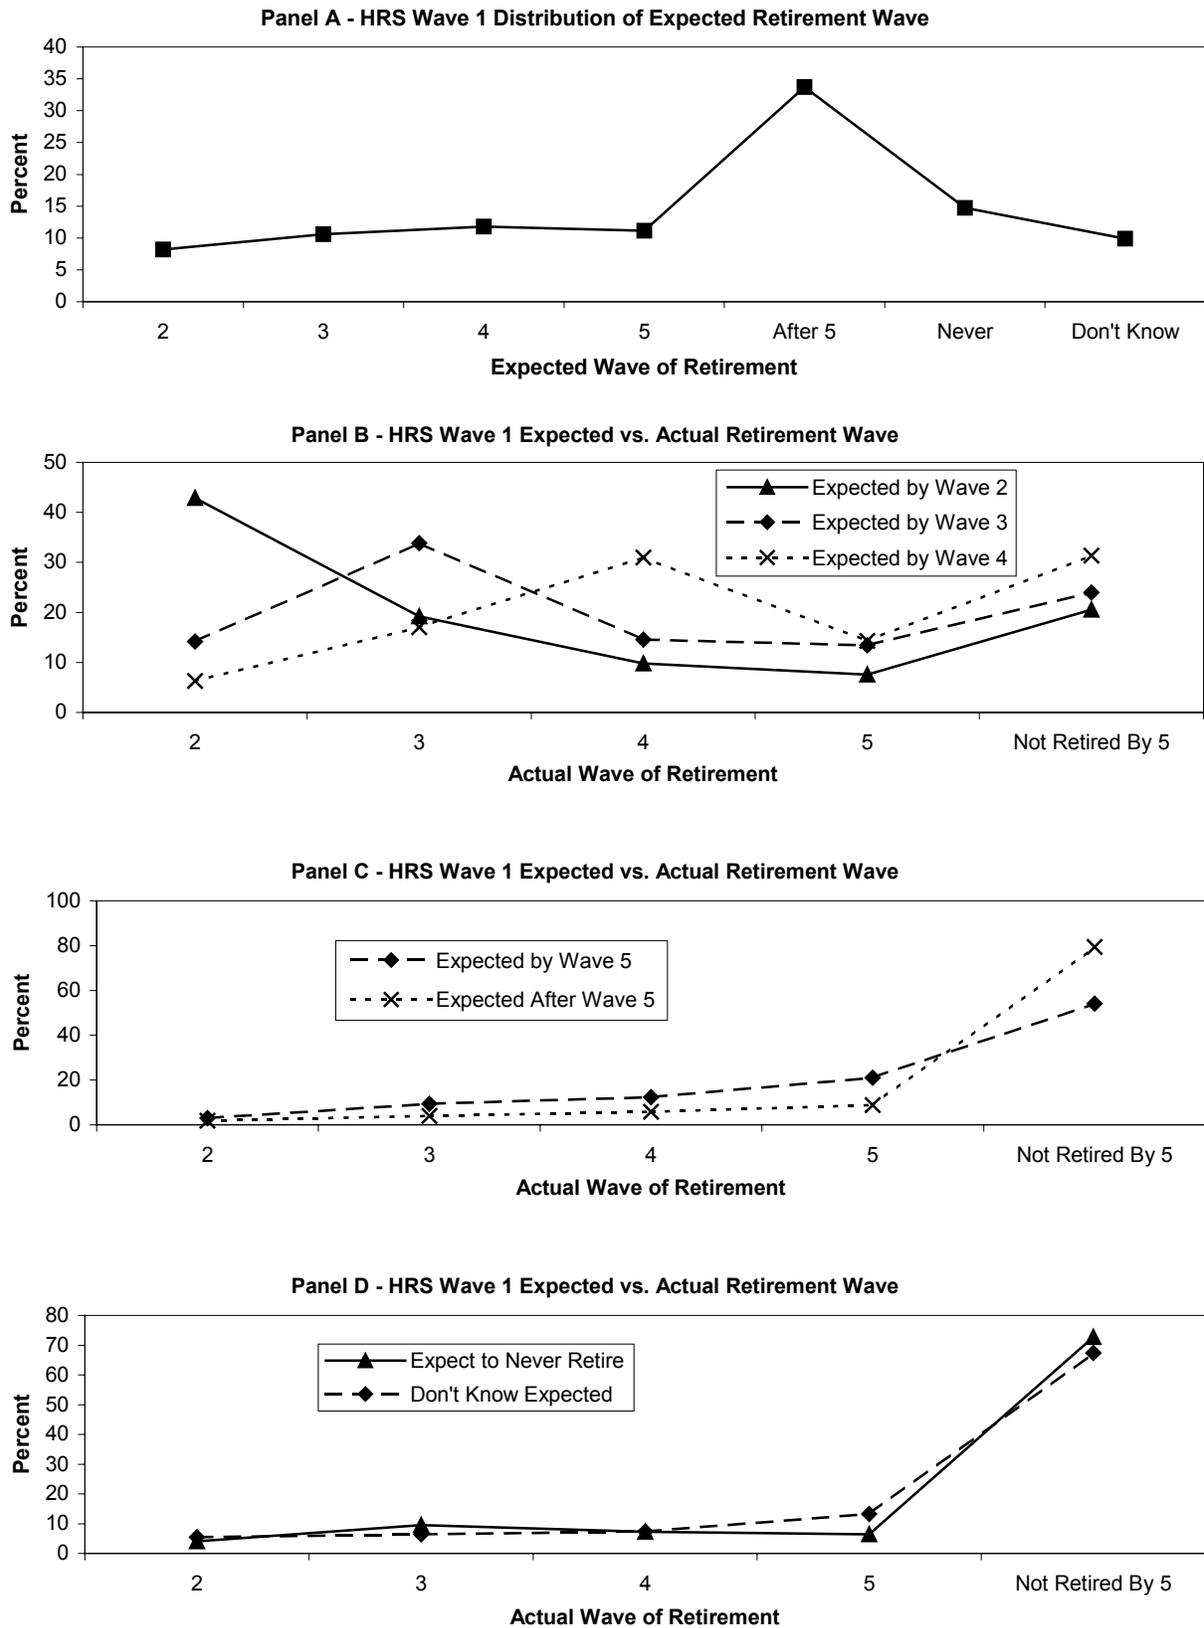
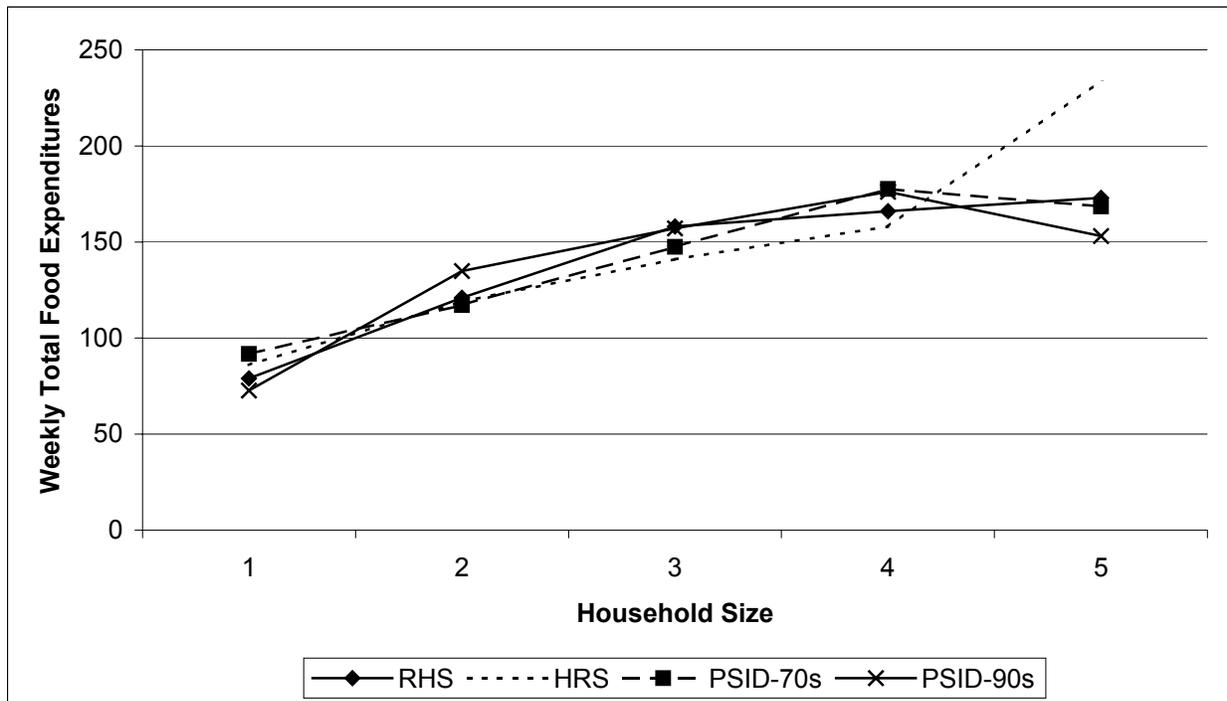


Figure 3: Weekly Food Expenditures by Household Size



Notes: These figures graph the mean weekly food expenditures by household size. All dollars are deflated to 2001, using the annual Personal Consumption Expenditure (PCE) deflator.

**Figure 4: Fraction of Workers Retiring By the Next Wave
Observed vs. Expecting to Retire**

