

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: August 9, 2005

We thank Kerwin Charles, Susann Rohwedder, and seminar participants at Carnegie Mellon/Pittsburgh, Dartmouth, the Federal Reserve Board of Governors, Maryland, Michigan State, Northern Illinois, Washington, Wesleyan, the 2003 NBER Fall Labor Studies meeting, 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.

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

Abstract

Previous research finds a systematic decrease in consumption at retirement, a finding that is inconsistent with the Life-Cycle/Permanent Income Hypothesis if retirement is an expected event. In this paper, we use workers' subjective beliefs about their retirement dates as an instrument for retirement. After demonstrating that subjective retirement expectations are strong predictors of subsequent retirement decisions, we still find a consumption decline at retirement for workers who retire when expected. However, our estimates of this consumption fall are about a third less than those found when we instead rely on the instrumental variables strategy used in prior studies. Finally, we examine several hypotheses that have been put forward to explain the retirement consumption decline. We find little empirical support for these explanations in our data.

JEL Classification. D84, D91, J26

“Old age is the most unexpected of all things that happen to a man.”

- Leon Trotsky (1879-1940)¹

1. Introduction

With a growing number of workers approaching retirement, the preparedness of these households to finance consumption during their retirement 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 standard Life-Cycle/Permanent Income Hypothesis (LCPIH) to examine whether households are saving adequately for retirement. Specifically, if households are rational and foresighted, then their consumption should not change upon expectedly retiring. Contrary to this hypothesis, however, empirical investigations have concluded that household consumption falls at the time of retirement, even when retirement is expected (e.g., Hamermesh 1984; Mariger 1987; Banks, Blundell, and Tanner 1998; Bernheim, Skinner, and Weinberg 2001).²

This fall in consumption at retirement is referred to as the “retirement-consumption puzzle” and has led researchers to call into question the standard rational expectations 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

¹ We are indebted to Eli Katz for suggesting this quote.

² Examining consumption changes at retirement is just one approach to assess the adequacy of retirement savings. Others include comparing the distribution of wealth in a simulated economy with the empirical distribution and calculating the annual payment from converting current household wealth into an annuity and comparing this estimate to current household consumption. For a critical review of this literature, see Engen, Gale, and Uccello (1999).

retirement" (p. 784). Bernheim, Skinner, and Weinberg state that their "... findings are difficult to interpret in the context of the life-cycle model" (p.855).

A number of alternative 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) hypothesize that household bargaining can explain the change in consumption at retirement and find empirical support for their predictions. Angeletos et al. (2001) demonstrate through simulation methods that hyperbolic (rather than geometric) discounting households will have a planned fall in consumption at retirement. Hurd and Rohwedder (2003) and Aguiar and Hurst (Forthcoming) present evidence that incorporating household production decisions into the standard model may explain the puzzle. Moreover, Ameriks, Caplin, and Leahy (2002) and Hurd and Rohwedder find that a substantial percentage of households *expect* their expenditures will decrease upon retirement.

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. Many prior studies examine observed changes in consumption at retirement. To the extent that retirement is caused by an unexpected event such as a job loss or a disability, the observed consumption fall does not refute the LCPIH. Recognizing this fact, some studies exploit the rapid changes in retirement that occur when workers become eligible for government retirement benefits to instrument for retirement. However, the choice of non-linearities in age as an instrument for retirement has two potential problems. First, because older households are generally observed reducing their consumption as they age, it is very important that the age-consumption relationship is properly parameterized. 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 illustrate in the two data sets that we use in our analysis, these relationships are not the same, then age is not a valid instrument for expected retirement.

We use subjective retirement expectations as an alternative instrument to distinguish between expected and unexpected retirements. Expectations questions 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 (2004) 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 expected 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 fall in consumption is 30 to 40 percent smaller in magnitude when using subjective retirement expectations as an instrument relative to using age as an instrument. Therefore, our arguably better methodology provides strong evidence of a fall in consumption at retirement but produces a smaller consumption decrease than the previous literature.

After establishing that consumption falls among workers for whom retirement is an expected event, we then explore the alternative hypotheses that have been put forward. We find little evidence to support either the household bargaining or the household production explanations in our data. We also find that, despite the fact that a sizable number of households expect their expenses to fall at retirement, these expectations are only weakly correlated with the actual consumptions drops that are observed. Thus, our results suggest that there exists a drop in consumption at expected retirements and this drop remains a puzzle given the explanations that have been put forth.

2. Modeling and Estimating the Retirement-Consumption Puzzle

We examine whether consumption falls at retirement in the context of the Life-Cycle/Permanent Income 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, V_t , where

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

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 thought to affect utility (through $\Psi(\cdot)$) such as age and family size, 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 comes from the rational expectations version of the LCPIH pioneered by Hall (1978). For illustrative purposes, assume that the variables modifying consumption (w_t) remain constant and that the interest rate equals the discount rate. Doing so allows us to 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 I_t represents the household's information at time t about its retirement status at period $t+1$, the model imposes the restriction that I_t should be uncorrelated with ε_{t+1} , or that $E_t[\varepsilon_{t+1}|I_t] = 0$.

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.³

To apply this Euler Equation to the retirement-consumption puzzle, consider the estimating 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 does not violate the LCPIH because households may retire for numerous unforeseen reasons such as a job loss or the onset of a disability. In terms of (4), retire_{t+1} may be correlated with the expectations error since some households may retire unexpectedly. 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 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

³ 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.

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

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) and 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 self-reported expectations are clearly known to the individual in each period and are therefore potential instruments for future retirements. The rational expectations assumption that information known at time t is uncorrelated with the expectations error between periods t and $t+1$ 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 econometric requirement that the instruments be 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 $exret_{t+1}$ that equals 1 if the worker expects to retire between time t and $t+1$ and 0 otherwise. The 2SLS estimator of λ when using $exret_{t+1}$ as an instrument is

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

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 because 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, whether their job loss probability is 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 data sets focus on individuals near the usual retirement ages and therefore contain large samples of workers undergoing these transitions. Second, both data sets ask a number of direct questions on retirement expectations. We provide an overview of the key aspects of the data sets here and discuss some of the specific issues in the appendix.

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 period, a total of six waves of information had been collected. We limit our analysis to the first five waves of the RHS because 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 1992.⁴ 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). Age-eligible household members and their spouses (regardless of birth year) were interviewed, resulting in approximately 12,700 initial respondents. As with the RHS, the HRS is 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 expenditures. The RHS collected this information in all waves while the HRS has collected this information in all waves except for wave 4. Food expenditure 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 expenditures is that it is a limited measure of household expenditures. However, a benefit of food expenditures is the fact that food 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 because 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 consumption puzzle studies (Bernheim, Skinner, and Weinberg 2001; Lundberg, Startz, and Stillman 2003), so its use also provides a point of comparison.

In our analysis, the definition of retirement is based on the current labor force status question. This question inquires whether individuals are currently engaged in one of a number of

⁴ See Juster and Suzman (1995) for an overview of the HRS.

activities, including working, unemployed, retired, disabled, and homemaker. We define retirement as those individuals who report being retired.⁵

The most important information available in these surveys for our estimation strategy 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. Survey participants can respond by giving an age (or year) of expected retirement, stating that they will never retire, or stating that they do not know. 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 Because 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 male-headed households. Our main reason for imposing this restriction is that the RHS only collected detailed information on women when a spouse was not 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 wave of our sample periods. We also only consider the first move into retirement and ignore any subsequent movements in and out of retirement. Finally, because retirement expectations are not

⁵ While RHS respondents are only allowed to choose one of the potential labor force activities, HRS respondents may choose as many as they would like. For comparability, we define HRS respondents to be retired if they report being retired and do not report being either unemployed or working.

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.

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. These figures correspond to households that contribute at least one first-differenced observation to the analysis.⁷ 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 not to be married, are 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 respond with the modal (i.e., most likely) date of expected retirement rather than the mean date. 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 provide modal responses to the retirement expectations question. In addition, they find that, when predicting retirement using regressions that include a large number

⁷ Since consumption data is not available in wave 4 of the HRS, each household can contribute at most two first-differenced observations (between waves 1 and 3) to the consumption regressions. The HRS sample described in Table 1 corresponds to the set of first-differenced observations that can be created from all five waves of the HRS from individuals that contribute at least one observation to the consumption regressions. This larger HRS sample is used in the analysis of the retirement expectations variables. The sample used in the HRS consumption regressions is comprised of 4,045 first-differenced observations from 2,617 individuals.

of observable characteristics, expected age of retirement is a very strong predictor of the actual age of retirement.⁸

In Figure 1, we use data from the RHS to demonstrate the relationship between expected wave of retirement and actual wave of retirement. 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.⁹ Of the workers who report an expected wave of retirement, the majority of them expect to retire by the fourth wave.¹⁰ 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 the 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.¹¹ Three main patterns can be observed in these figures. First, the modal wave of realized retirement corresponds to the subjective expected retirement 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.

⁸ Other aspects of retirement expectations variables have been analyzed. Deviations between expected and actual retirement ages are correlated with wealth and health changes as well as marital transitions (Anderson, Burkhauser, and Quinn 1986; Disney and Tanner 1999; Loughran, Panis, Hurd, and Reti 2001; Benitez-Silva and Dwyer 2003). Benitez-Silva and Dwyer (2003) find that retirement expectations in the HRS are consistent with the rational expectations hypothesis.

⁹ Since our subsequent analysis examines changes in consumption and retirement status between survey waves, we examine the distribution of retirement expectations by survey wave. The distribution of expected retirement ages are presented in Appendix Figures 1a and 1b for the RHS and HRS, respectively.

¹⁰ 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).

¹¹ Workers who leave the RHS before retiring are excluded from the last three panels of Figure 1.

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 retirement date after the available sample period. Panels B and C of Figure 2 show that HRS respondents also appear to provide 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 data sets.

The correlation between retirement expectations and realizations could simply be 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 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. Aside from estimating the additional explanatory power contained in the subjective retirement expectations responses, equation (8) also represents the first stage for equation (7).

¹² These covariates are included since, as we discuss below, they are the same regressors used in the consumption change equations.

Panel A of Table 2 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 33 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.

Because 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 2 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 43 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 little impact on the coefficient estimates.

Finally, Panel C presents the results for the HRS using wave 1 expectations. In comparison to the results using wave 1 expectations in the RHS, the correlation between expected and observed retirement is not as large in the HRS. Nonetheless, the estimates are still highly significant and do not change across different specifications.¹⁴

¹³ The reported standard errors for the regression coefficients in Table 2 as well as in all subsequent tables are robust to arbitrary forms of correlation within individuals over time.

¹⁴ If younger individuals have noisier retirement expectations, then the fact that the initial age of respondents in the HRS is 51 to 61 while the initial RHS ages is 58 to 63 may be responsible for the smaller estimates in the HRS. However, when the HRS is limited to respondents who are at least age 58 in wave 1, the results do not change.

One possible concern with the construction of our expected retirement variable $exret_{t+1}$ is that combining an array of beliefs about the expected date of retirement into a dichotomous variable, especially “never” and “don’t know”, may improperly characterize the relationship between retirement expectations and realizations. To examine this issue, we modified equation (8) by replacing $exret_{t+1}$ with a set of “years until expected retirement” indicators which correspond to the number of years from the survey date until the respondent expects to retire. In addition, we are able to place “never” and “don’t know” responses into their own categories.

Figure 3 presents the results from estimating this adjusted version of equation (8) using the most recent retirement expectations response in the RHS. The solid line presents the “raw” correlations where we do *not* include age or any other covariates aside from the years until expected retirement indicators. The dashed line shows the results where we include the covariates discussed above as well as a full set of age dummies in the adjusted equation (8). The respondents we code as having $exret_{t+1}$ equal to one are those expecting to retire within two years of the survey date. As the results in the Figure show, these individuals have similar high probabilities of retiring by the next survey. The large decrease in the probability of retiring between those who say they will retire in two years and those who will retire in three years is consistent with a distinct change between the two groups categorized by $exret_{t+1}$. Finally, the low retirement probabilities, especially when adjusted for age, among respondents expecting to retire no sooner than three years after the survey date as well as respondents answering “never” or “don’t know” suggest that combining this group of workers should not be problematic for our analysis. Nonetheless, we will also discuss results in which we use this more flexible specification for retirement expectations.

Figure 3 also illustrates the degree to which actual retirement deviates from retirement expectations. For example, the solid line in the Figure indicates that less than 70 percent of respondents who expect to retire within one year of the survey date are in fact retired by the next survey date (which occurs two years after the survey). Also, roughly fifteen percent of respondents who state that they will never retire do indeed retire by the next survey wave. In fact, when we calculate the share of actual retirements that are due to workers who expected to retire between the waves using the most recent expectations in the RHS, only about half (51 percent) of retirements are expected.¹⁵ Thus, while retirement expectations are strong predictors of retirement, a large share of retirements are nonetheless unexpected.

While our analysis is based on surveys with two years between waves, many other papers that examine consumption changes at retirement use data from surveys with only a single year between waves. Therefore, understanding the accuracy of retirement expectations at a one year interval is relevant for comparing our analysis to much of the literature. Appendix Figure 2 presents a comparable set of results to those found in Figure 3 except that retirement is now defined as having retired in the year following the survey.¹⁶ Less than 60 percent of individuals who expect to retire within one year are in fact retired within that time frame. Furthermore, we calculate that less than 40 percent of respondents who retire within one year of the survey had in fact expected to retire by that date. These results suggest that there is no reason to believe that using data from surveys with a one year interval between waves would produce appreciably different results from those found here.

¹⁵ Based on wave 1 expectations in the RHS, 34 percent of retirements are expected. In the HRS, based on wave 1 expectations and using all five survey waves as in Table 2, 32 percent of retirements are expected.

¹⁶ Individuals who exit the labor force are asked when they left their last job. The RHS reports this date as the number of months since April of the previous survey year. Unfortunately, the RHS does not report the interview date. However, interviews were generally conducted between March 1 and July 1 of a survey year. To ensure that we allow all of the respondents at least a one year window, we define a retirement as occurring within one year if the respondent reports having left their last job no more than 15 months since April of the last survey year. Date of retirement is missing for seven percent of retirements between survey waves.

5. The Change in Consumption at Observed Retirement

Before examining 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 (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 (6) is not evidence of a failure of the LCPIH because these regressions pool households retiring both expectedly and unexpectedly.

Panel A of Table 3 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 because 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 4, 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 (2004) 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 data set, the Panel Study of Income Dynamics (PSID). 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 retirement. We construct comparison samples from the PSID for both the RHS and the HRS.¹⁷ Figure 4 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 3. 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 data sets. The results of regressing the change in log family income on the same regressors used in Panel A of Table 3 are shown in Panel B of the Table. The results show that family income falls fairly similarly across all the data sets and time periods. Although the income drop at retirement in the HRS is smaller than the estimate for the contemporaneous PSID sample, the decrease is very similar to that found in the RHS.

¹⁷ The appendix provides additional information on the PSID samples including a discussion about why the calendar years for the PSID samples used here do not perfectly align with those of the RHS and HRS.

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 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 data sets is also present in the RHS but not in the HRS. As we discussed in the introduction, however, these results alone do not refute the LCPIH because 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 4 presents the estimates of equation (9). When RHS wave 1 expectations are used to construct $exret_{t+1}$, we reject the null hypothesis for two of the three specifications (Panel A of the Table). When only a linear term in age is included in the model (column 1), the estimate of ϕ is negative and significant. When the potential non-linearities between age and consumption changes are accounted for by using a quadratic specification for age, the estimate of ϕ remains significant. However, when a complete set of age dummies is included in the analysis, the estimated value of ϕ becomes smaller and marginally significant.

When the most recent retirement expectations from the RHS are used as the basis for $exret_{t+1}$ as shown in Panel B of Table 4, the model is rejected across all three specifications for age. 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, regardless of the specification for age. Thus, these estimates reject the LCPIH. We interpret these results as strong evidence that households do not fully adjust their consumption to all available information.

In Panel C of Table 4, we examine the consumption response to wave 1 retirement expectations in the HRS. Although we did not find a retirement consumption decline for the HRS in Table 3, there remains the possibility that positive wealth shocks may have led the observed change to be insignificant during this time period. However, it still may be the case that retirement consumption changes among those who expect to retire may differ from the

observed changes. Across all three specifications in Panel C, we find no difference in consumption changes among those who expect to retire relative to other households.

While the reduced form estimates in Table 4 for the RHS are evidence of a decrease in consumption for workers who retire when expected, the structural (2SLS) results of equation (7) presented in Table 5 estimate the magnitude of the consumption decline. Consistent with the results in the previous table, we find a negative and significant impact of retirement on consumption when using wave 1 expectations in the RHS (Panel A). These estimates imply that consumption falls by 7 to 10 percent upon retirements that are expected. The corresponding point estimates in the HRS based upon wave 1 expectations vary between being positive and negative, although the confidence intervals around these estimates are relatively large (Panel C). Finally, 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 10 to 11 percent when workers retire as expected. Overall, we interpret these results as suggesting that a retirement-consumption puzzle exists for individuals who expect to retire.

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, as we discussed in section 4, this specification ignores the fact that individuals 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. To exploit this diversity of retirement expectations, we also estimated a version of equation (7) where the first stage employs the same

specification used to generate the results shown in Figure 3. Our second-stage estimates from this alternative specification are nearly identical to the analogous results reported in Table 5.¹⁸

Using Age as an Instrument for Retirement

To account for the endogeneity of the retirement decision, previous studies examining the retirement-consumption puzzle have used age as an instrument.¹⁹ 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. Because the availability of these benefits are known in advance to households and the benefit value should be easily forecasted, 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 (7).

While the use of age as an instrument is intuitively appealing, we note two 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 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,

¹⁸ In the alternative specification that includes age linearly, the second-stage coefficient for the retirement variable is -0.102 (standard error of 0.023), as compared to -0.109 (standard error of 0.024) reported in Table 5. In the alternative specification that includes age dummies, the second-stage coefficient for the retirement variable is -0.091 (standard error of 0.025), as compared to -0.099 (standard error of 0.026) reported in Table 5. The F-statistic for the instruments in the first stage is well over 100 for both alternative specifications.

¹⁹ Bernheim, Skinner, and Weinberg use separate first-stage equations for each individual age covered in their sample (54 to 70) to predict retirement status. This approach is identical to a single first-stage regression pooling across all ages and including dummy variables for each age along with interactions between these age dummies and the other covariates. Banks, Blundell, and Tanner use lagged regressors such as past retirement status as instruments. To the extent that these lagged values are age dependent, they are implicitly using age as an instrument via a non-linear transformation. Thus, the arguments discussed here also apply to their study. Aguiar and Hurst use age dummies to instrument for retirement.

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 vary systematically 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 5 suggests that a systematic bias does indeed exist. The solid lines the panels show the fraction of workers at each age, conditional on not having retired yet, retire by the next survey wave in the RHS and HRS. 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, conditional on having not yet retired.²⁰ Prior to normal retirement, the solid and dashed lines are nearly identical. 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 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.

²⁰ The long-dashed lines in the Figure are based on the expected date of retirement in the first survey wave. The short-dashed line in Panel A of the Figure are based upon the worker's most recent expected retirement date.

To compare our estimates using retirement expectations as an instrument to prior research, we estimate the model using age as an instrument in the same RHS samples used above.²¹ 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. We present estimates of models that 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.

The results of using age as an instrument for retirement are presented in Table 6. The results show that consumption is significantly reduced by 16 percent when workers retire as "expected" when the linear age term is included. When a quadratic in age is included, the point estimate implies that consumption drops by 14 percent at expected retirement, but the standard errors increase dramatically so that the estimate is now marginally significant.²² The F-statistics at the bottom of Table 6 indicate that the age dummies are jointly significant predictors of retirement. Note, however, that these F-statistics are much smaller than the analogous values for the expectations dummy variables in Table 6.

Compared to the results where expectations are used as an instrument (Tables 5 and 6), the estimated consumption decrease at retirement when using age as an instrument is substantially larger in magnitude. Moving from using age as an instrument to using the expectations variables reduces the point estimates by 30 to 40 percent. The standard errors around the estimates in Table 6, however, are just large enough to include the previous point estimates within one standard deviation. These larger estimates are likely due to the instrument

²¹ Since the HRS results have been consistently insignificant, we focus on the RHS for the remainder of our analysis.

²² We also estimated these models including the post-retirement observations as has been done in previous studies. The results for the linear and quadratic specifications resulted in point estimates of -0.163 (with a standard error of 0.047) and -0.202 (0.089), respectively. Thus, we do not believe that our slightly different sample construction (relative to the rest of the literature) has any impact on the analysis we present. If anything, these additional estimates imply that the bias from using age as an instrument is larger than suggested by the results in Table 6.

treating the retirements at ages beyond normal retirement age as expected retirements. In addition, the standard errors on the structural estimates are substantially smaller when retirement expectations are 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.

7. Why Does Consumption Fall at Expected Retirement?

Our results suggest that better distinguishing between expected and unexpected retirements can explain about a third of the drop in consumption when compared to previous methods. Nevertheless, we still find that the puzzle persists: consumption falls at retirement, contrary to the basic Life-Cycle/Permanent Income Hypothesis. We now turn to assessing several of the hypotheses that have been put forward to explain the puzzle.

Lundberg, Startz, and Stillman (2003) suggest that household bargaining between spouses with differing expected lifetimes can explain the drop in consumption at retirement. Consistent with their hypothesis, they present evidence that consumption drops at retirement for married couples but remains constant for single individuals. Not only can we replicate their main analysis with the RHS, but we can also account for the endogeneity of retirement.²³

In Table 7, we follow the approach of Lundberg et al. and estimate the models separately for married and single households.²⁴ Panel A of the Table is comparable to the analysis reported

²³ Lundberg et al. do not instrument for retirement in their analysis.

²⁴ A household is considered single if the male head is single during both waves covered by the first-differenced observation and is considered married if the head is married for both waves. As an alternative specification, we divided households based on wave 1 marital status. The results when using this alternative method are unchanged, which is not surprising given that only 318 observations were dropped based on the first definition.

in Table 3 of Lundberg et al. (2003). Our OLS results in this Panel suggest that the drop in consumption is *larger* among single households. When we instrument for retirement using most recent expectations (Panel B), the point estimate for single households is still larger than for married households although the former result is imprecisely estimated. Thus, we do not find empirical support for the Lundberg et al. household bargaining prediction in our data.

Hurd and Rohwedder (2003) and Aguiar and Hurst (Forthcoming) suggest that incorporating household production decisions into the basic LCPIH may explain the consumption decrease at retirement. Specifically, upon retiring, households might shift from using consumption goods that are market-good intensive to those that are time intensive because the opportunity cost of time is lower. For example, individuals may choose to substitute food prepared at home for food in restaurants more frequently during retirement.

We explore the potential importance of such substitution by presenting results separately for food at home and food away from home. Table 8 shows the results using the change in the log of consumption as the dependent variable (as in our analysis above) as well as using the change in the level of consumption. We include the latter specification because of the substantial number of zeroes for the disaggregated expenditure categories, especially food away from home. The results suggest that there are statistically significant drops in both food at home and food away. Whereas the drops are similar when the dependent variable is specified in levels (just over \$6 in the IV results), the results are much bigger for food away when it is specified in logs.²⁵ Given that we find a decrease for both types of food expenditures, our results do not support a prediction of the home production hypothesis that the retirement consumption drop is generated

²⁵ Lundberg et al. find that the drop in consumption at retirement is the same for both food at home and food away from home (in logarithms), and Bernheim et al. find that the decrease for food away from home is actually smaller in magnitude than for food at home. Thus, although studies have found varying results regarding whether food away or food at home drops more at retirement, all studies find that food at home drops as well.

by a strict substitution from food away to food at home. However, our data do not allow us to rule out the possibility that households are substituting between market- and time-intensive types of grocery items that would account for the decline in food at home expenditures.

Finally, Ameriks, Caplin, and Leahy (2002) and Hurd and Rohwedder (2003) find that households *expect* their expenditures to fall at retirement. Although these findings are not directly informative as to the mechanism causing expenditures to decline at retirement, they provide an important insight that any explanation of the decline must capture. All households in the RHS who state that they expect to retire are asked whether when they retire, “[w]ill your *expenses* be more, less, or the same as now?” (emphasis added). Although the wording of this question differs slightly, we believe it provides a useful point of comparison to the previous studies. Among households who expect to retire in the first wave of the RHS, we find qualitatively similar results to those reported in previous studies: 39 percent expect their expenses will fall, 13 percent say they expect their expenses to increase, and 48 percent expect their expenses will remain the same.²⁶

Unlike these prior studies, we can examine whether these consumption change beliefs are correlated with the actual consumption changes occurring at retirement. We re-estimate our basic models but interact the retirement variable with indicators for whether individuals expect expenses to be more, less, or the same upon retirement.²⁷ We present these results in Table 9. Examining the OLS results in column 1, we find that the observed change in consumption is

²⁶ The RHS further asks these individuals why they expect their expenses to change. For those who believe that their expenses will be less, the reasons most often given include a decline in work expenses (45 percent) and voluntary adjustments (30 percent). To the extent that these work expenses and voluntary adjustments are reflected in food costs (perhaps entertaining or eating out), then the decline in food expenditures at retirement would not be due to unanticipated shocks.

²⁷ Because individuals who report they do not expect to retire are not asked how they expect their expenses to change at retirement, we include a separate interaction for these individuals. We use the most recent report of expected expense change to define these interactions. We limit the sample to observations reporting both the wave 1 expected change and the most recent wave expected change.

essentially identical across all categories of expected retirement consumption changes. The IV results using most recent retirement expectations (column 2) suggest that the drop in consumption is somewhat more concentrated among those who expect their expenses to be less at retirement, but we still find a retirement consumption decline for all three expected consumption change groups.²⁸ Thus, differing beliefs about changes in expenses upon retirement cannot explain the retirement consumption decline since it is present for all expected consumption change responses and the actual consumption changes only weakly conform to the expected change patterns.

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. 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. Agents' expectations are an important component of numerous economic models,

²⁸ In generating the IV results, we employ a procedure wherein we predict retirement using most recent retirement expectations, interact these predictions with the expected retirement consumption change, and then include these interactions in the consumption equation. The reported standard errors are computed as the standard deviations of the bootstrapped distributions of the coefficients using 200 replications and re-drawing households with replacement.

including the Life-Cycle/Permanent Income Hypothesis, and our paper adds to the growing literature that uses subjective expectations responses in survey data to test economic hypotheses. 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 because 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 7 to 11 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 about a third less in magnitude, and more precisely estimated, than we find when we follow previous studies and use age as an instrument for retirement.

We have also examined a number of the alternative hypotheses that have been offered to explain the retirement consumption decrease. Our results do not support the household bargaining hypothesis put forth by Lundberg et al. because we find similar if not larger consumption decreases for single households relative to married households. Examining the changes in food at home and food away separately, we find no evidence that there is a direct substitution between these two types of food expenditures, a result that has also been found by Bernheim et al. and Lundberg et al. Although this finding does not rule out the household production explanation that has been put forward, it reduces the scope for its importance. Finally, two prior studies have pointed to survey responses where individuals report that they expect expenditures to fall at retirement. However, our results suggest that such responses may not be so important given that the retirement consumption decline is not concentrated among those who expect their expenses to be less after retiring.

How should we interpret this fall in consumption at retirement? Although our analysis has not been able to validate any specific hypotheses, we have used expectations data to demonstrate the extent to which a puzzle exists and we have been able to examine several explanations for the puzzle. Our results both suggest that more work is needed to uncover the exact mechanisms behind this finding and demonstrate that subjective retirement expectations will be useful in this endeavor.

Appendix

The Retirement History Survey

The RHS interviewed 8,132 men in the first wave. We begin by dropping those men who do not work in the first wave and who do not show up at least twice, leaving 5,761 men. We then follow these men until the end of the sample period (wave 5), until they retire, or until they leave the sample. We also drop those respondents with missing retirement information. This initial sample provides 14,670 first-differenced observations. Of these first-differenced observations, we then eliminate those that include a wave where an individual is a boarder or where the necessary age or household size information is missing, leaving 14,120 observations. Finally, we drop all remaining observations that do not contain the appropriate food expenditure information or expectation information, leaving 11,510 observations on 5,060 individuals.

The food expenditure information is collected by asking questions regarding several underlying sources: food at the grocery store, food delivered to the home, food at other types of stores (e.g., bakeries, fruit stands, and dairy stores), food at work, and food eating out. All questions refer to expenditures over the last week, except the food out questions that refer to the last month. As discussed in detail in Hurd (1992), some of the variables are apparently recorded in cents rather in dollars, as suggested by the codebook. We have adjusted all food expenditures at home in waves 2 and 4 to account for this inconsistency. Also discussed in Hurd, some of the individual observations also appear to be recorded in cents. For the results reported in the paper, we follow Hurd and attempt to individually identify these observations and adjust them from cents to dollars.²⁹ In addition, there is a follow-up question to the food-at-home questions that asks whether the expenditures represent the “usual” amount. If the respondent indicates that the given amount does not represent the “usual” amount, an alternative total is inquired about.

²⁹ These adjustments affect under 200 person-year expenditure variables.

When such an alternative total is provided, we use this amount; approximately 15 percent of the respondents per wave provide an alternative “usual” amount.³⁰

Total household income is collected by approximately 50 questions about various underlying income sources in each wave. Although most people answer most questions, missing data is a problem because total income cannot be computed directly if a person does not respond to any one of the numerous income questions. We impute missing data for each of the underlying sources using the “predictive mean matching” method (Little 1988), the same method used by the University of Michigan for HRS imputations. This method imputes a missing value for an individual by using an actual outcome value from a “similar” person within the data set, where similarity is based on the prediction from a regression. Our imputations are based on gender, race, education, employment status, age, home ownership, and household size.

The Health and Retirement Study

The HRS interviewed 4,596 men in the first wave who were born between 1931 and 1941 (excluding the “overlap sample” individuals). We begin by dropping those men who did not work in the first wave and who do not show up at least twice, leaving 3,255 men. We then limit our HRS sample to those individuals who have valid consumption data in order to contribute at least one first-difference observation to the sample between waves 1 and 3. This restriction limits the sample to 2,866 men. We then follow these men until the end of the sample period (Wave 5), until they retire, or until they leave the sample. We also drop those with missing retirement information. This initial sample provides 8,933 first-differenced observations. Of these first-differenced observations, we then eliminate those that include a wave the necessary age or household size information is missing, leaving 8,811 observations. Next, we drop all individuals that do have the appropriate retirement expectations information, leaving 8,086

³⁰ This adjustment is empirically unimportant.

observations from 2,167 men. We use this sample for assessing the retirement expectations variables in section 4. Finally, for use in the consumption regressions, we drop all individuals who do not have the appropriate food expenditure information, leaving 4,045 observations. This much smaller sample size is due the food expenditure questions not being asked in wave 4 and the food expenditure questions being placed in an experimental module that participants are allowed to bypass in wave 3.

The food expenditure information is collected via questions regarding several underlying sources of food. For those receiving food stamps last month, the survey inquires about the quantity of food stamps received and the additional amount of money spent in the last month on food at the store, food delivered to the door, and food eating out. For those not receiving food stamps, the survey asks about the same three underlying sources, but instead asks about expenditures in an average week.

We use the HRS-provided imputations for the income variables.

Comparability of the RHS and HRS

Although the labor force status questions are different in the RHS and HRS, both surveys ask a second direct question about retirement. This second question asks workers to report their current retirement status as being fully retired, partly retired, or not retired at all. Defining retirement based on this alternative question does not affect our results.

The wording of the expected retirement questions varies slightly between the RHS and the HRS. 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)?”

Differences in the wording of the expectations questions across the two surveys are reflected in differences in the information available in the data files. The RHS data reports the expected age of retirement since survey participants are prompted to give an expected age response. The HRS question allows households to give either an expected age *or* year of retirement response. The HRS data, however, only reports the expected year of retirement that is converted from expected age if an age response is given. Thus, we cannot compute with precision the age at which a person expects to retire in the HRS. This imprecision can be observed in Panel B of Appendix Figure 1 as an elevated retirement hazard at ages 63 and 66. However, this imprecision has little impact on the results in the paper because the analysis is based on expected wave of retirement, not expected age of retirement.

The Panel Study of Income Dynamics

The PSID is an on-going panel data set that began in 1968. The PSID interviewed households annually until 1997, at which time it began interviewing on a biennial basis. The original sample consisted of approximately 5,000 families, and the families and their off-spring have been followed ever since. Of these original families, approximately 60 percent sampled from a nationally representative sampling frame while the remaining households were taken from a set of low-income households. We limit the observations in our analysis to households originating from the nationally representative sample.

We create two PSID samples: the PSID-70s sample that spans 1970 to 1980 to coincide with the RHS and the PSID-90s sample that spans 1991 to 1999 to coincide with the HRS. The PSID samples do not directly match the years in the RHS and HRS due to data availability issues. We start the PSID-70s sample in 1970 since the PSID does not collect any food consumption data in 1973 (hence eliminating the use of odd number years for the PSID-70s

sample) and only collects food at home consumption data in 1968. We begin the PSID-90s sample in 1991 since, as mentioned above, the PSID only surveys households in odd-numbered years beginning in 1997. For each of the PSID samples, we draw male household heads that are 53-62 years old in the first year of that sample. We then follow each comparison sample biennially for five waves. Statistics for the PSID comparison samples are shown in Appendix Table 1.

References

- Aguiar, Mark, and Erik Hurst (Forthcoming). "Consumption vs. Expenditure," *Journal of Political Economy*.
- 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.
- Benitez-Silva, Hugo and Debra S. Dwyer (2003). "What to Expect When You Are Expecting Rationality: Testing Rational Expectations Using Micro Data," SUNY-Stony Brook Manuscript.
- 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 Grilliches 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. (1992). "Wealth Depletion and Life-Cycle Consumption by the Elderly." *Topics in the Economics of Aging*. Chicago and London: The University of Chicago Press.
- 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.
- Little, Roderick (1988). "Missing-Data Adjustments in Large Surveys," *Journal of Business and Economic Statistics*, 6:3, 287-301.
- 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 (2004). "Job Loss Expectations, Realizations, and Household Consumption Behavior," *The Review of Economics and Statistics*, 86:1, 253-69.
- 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	5,060	2,617
Person-Observations (first differences)*	11,510	8,086
Married in first wave	91.8%	87.2%
Years of Education	10.0	12.6
Percent retiring during the sample period	56.9%	40.0%

Notes: All results are unweighted. *For all analysis of food expenditure information in the HRS, we can only use the 1992, 1994, and the 1996 waves of the HRS because food expenditures were not asked in the 1998 wave, leaving 4,045 observations.

Table 2: Predicting Retirement Using Retirement Expectations

Age Controls	(1) Age	(2) Age, Age ²	(3) Age Dummies
A. RHS: OLS, independent variable is wave 1 expectations			
Expect to Retire	0.330 (0.013)	0.317 (0.013)	0.297 (0.014)
R-squared	0.149	0.154	0.160
N	11,510	11,510	11,510
B. RHS: OLS, independent variable is most recent wave expectations			
Expect to Retire	0.425 (0.011)	0.415 (0.011)	0.403 (0.011)
R-squared	0.224	0.229	0.234
N	11,510	11,510	11,510
C. HRS: OLS, independent variable is wave 1 expectations			
Expect to Retire	0.254 (0.018)	0.254 (0.018)	0.243 (0.018)
R-squared	0.126	0.129	0.140
N	8,086	8,086	8,086

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 3: Changes in Consumption and Income at Retirement
OLS Regressions

	(1)	(2)	(3)	(4)	(5)
	RHS	HRS	PSID	PSID	PSID
	1969-1977	1992-1996	1970-1980	1991-1999	1991-1995
A. OLS, dependent variable is change in log consumption					
Retired	-0.090	0.005	-0.101	-0.089	-0.051
	(0.011)	(0.024)	(0.031)	(0.043)	(0.061)
Age	-0.005	0.000	-0.002	0.002	-0.001
	(0.002)	(0.002)	(0.003)	(0.004)	(0.006)
Dhhsz	0.100	0.061	0.108	0.154	0.198
	(0.009)	(0.013)	(0.025)	(0.045)	(0.057)
R-square	0.057	0.010	0.047	0.047	0.097
N	11,510	4,045	716	549	333
B. OLS, dependent variable is change in log household income					
Retired	-0.258	-0.230	-0.253	-0.373	-0.337
	(0.015)	(0.041)	(0.039)	(0.102)	(0.153)
Age	0.007	-0.005	-0.001	0.012	0.017
	(0.003)	(0.003)	(0.004)	(0.017)	(0.028)
Dhhsz	0.049	0.020	0.017	0.093	0.083
	(0.007)	(0.019)	(0.025)	(0.093)	(0.138)
R-square	0.031	0.013	0.074	0.019	0.010
N	11,419	4,012	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 4: The Impact of Expected Retirement on Consumption
Reduced-Form Estimates**

Age Controls	(1) Age	(2) Age, Age ²	(3) Age Dummies
A. RHS: OLS, independent variable is wave 1 expectations			
Expect to Retire	-0.033 (0.012)	-0.029 (0.012)	-0.021 (0.013)
R-squared	0.052	0.052	0.053
N	11,510	11,510	11,510
B. RHS: OLS, independent variable is most recent wave expectations			
Expect to Retire	-0.046 (0.010)	-0.044 (0.010)	-0.040 (0.011)
R-squared	0.053	0.053	0.054
N	11,510	11,510	11,510
C. HRS: OLS, independent variable is wave 1 expectations			
Expect to Retire	0.005 (0.026)	0.008 (0.026)	-0.002 (0.026)
R-squared	0.010	0.010	0.014
N	4,045	4,045	4,045

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 5: 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. RHS: IV using wave 1 expectations			
Retired	-0.099 (0.036)	-0.093 (0.038)	-0.071 (0.042)
N	11,510	11,510	11,510
B. RHS: IV using most recent wave expectations			
Retired	-0.109 (0.024)	-0.106 (0.025)	-0.099 (0.026)
N	11,510	11,510	11,510
C. HRS: IV using wave 1 expectations			
Retired	0.017 (0.078)	0.023 (0.082)	-0.006 (0.082)
N	4,045	4,045	4,045

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: The Impact of Expected Retirement on Consumption
Using Age as an Instrument**

Additional Age Controls	(1) Age	(2) Age, Age ²
RHS: IV using age dummies (ages 59-69)		
Retired	-0.158 (0.060)	-0.141 (0.083)
First Stage F-test (For Excluded Age Variables)	33.7	35.0
N	11,510	11,510

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: The Impact of Retirement on Consumption by Marital Status
RHS Respondents**

	(1)	(2)
	Married Households	Single Households
A. OLS		
Retired	-0.083 (0.011)	-0.158 (0.049)
R-Squared	0.050	0.088
N	10,331	861
B. IV using most recent wave expectations		
Retired	-0.093 (0.027)	-0.167 (0.124)
N	10,331	861

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

**Table 8: The Impact of Retirement on Consumption by Food Category
RHS Respondents**

	(1)	(2)	(3)	(4)
	Food At Home	Food Away From Home	Food At Home	Food Away From Home
	Logs	Logs	Levels	Levels
A. OLS				
Retired	-0.035 (0.012)	-0.336 (0.028)	-3.38 (1.37)	-7.88 (0.68)
R-Squared	0.059	0.056	0.012	0.025
N	11,397	7,601	11,510	11,510
B. IV using most recent wave expectations				
Retired	-0.054 (0.029)	-0.301 (0.064)	-6.23 (3.36)	-6.56 (1.69)
N	11,397	7,601	11,510	11,510

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

**Table 9: The Impact of Retirement on Consumption by Expected Consumption Change
RHS Respondents**

	(1)	(2)
	OLS	IV using Most Recent Expectations
Interaction of Retired with:		
Expect Expenses to be More	-0.081 (0.024)	-0.073 (0.042)
Expect Expenses to be Less	-0.089 (0.017)	-0.124 (0.035)
Expect Expenses to be Same	-0.082 (0.017)	-0.065 (0.036)
Expect to Never Retire	-0.089 (0.021)	-0.078 (0.075)
N	11,372	11,372

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 for the OLS regression are robust to arbitrary forms of correlations within individuals over time. Standard errors for the IV regression are the standard deviations of the bootstrapped distributions that were calculated from 200 replications in which households were drawn with replacement.

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%
Years of Education	10.3	12.2
Percent retiring during era	71.0%	54.0%

Figure 1: RHS Wave 1 Retirement Expectations and Realizations

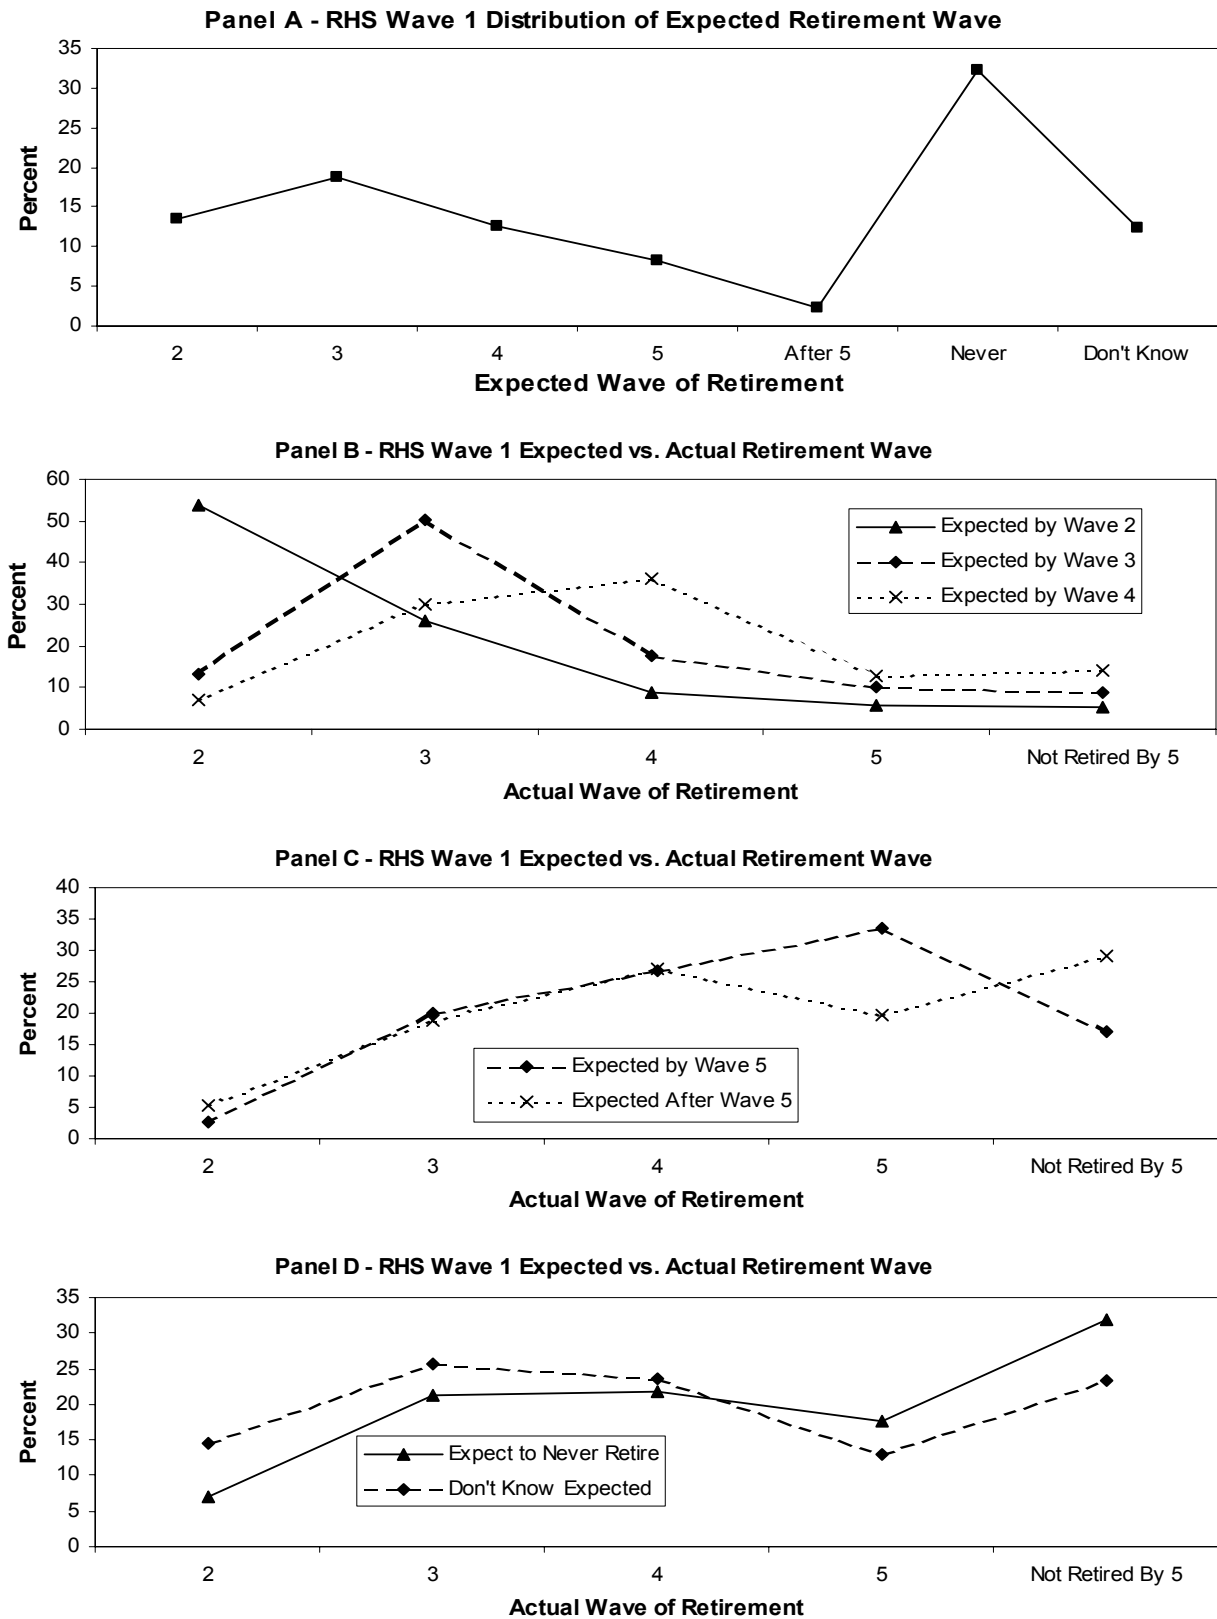
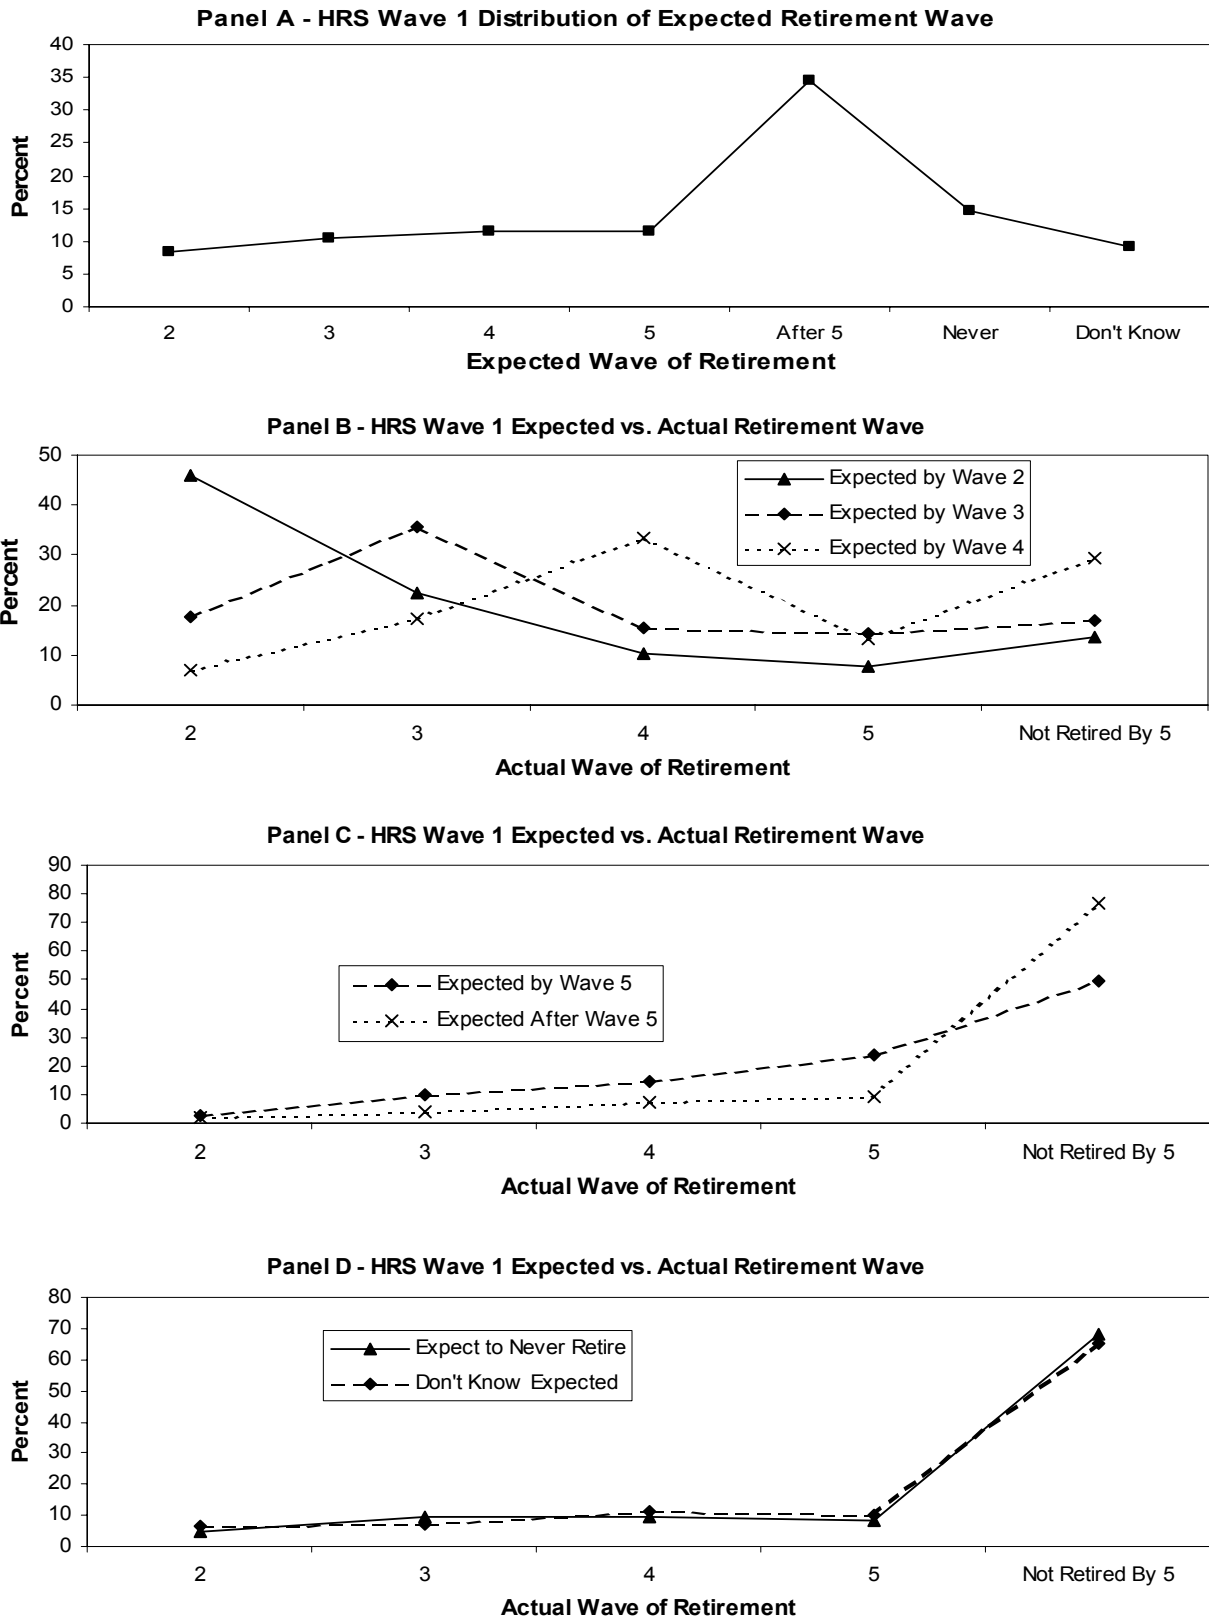


Figure 2: HRS Wave 1 Retirement Expectations and Realizations



**Figure 3 – RHS Fraction Retiring Before the Next Survey Wave
By Most Recent Expected Retirement Year**

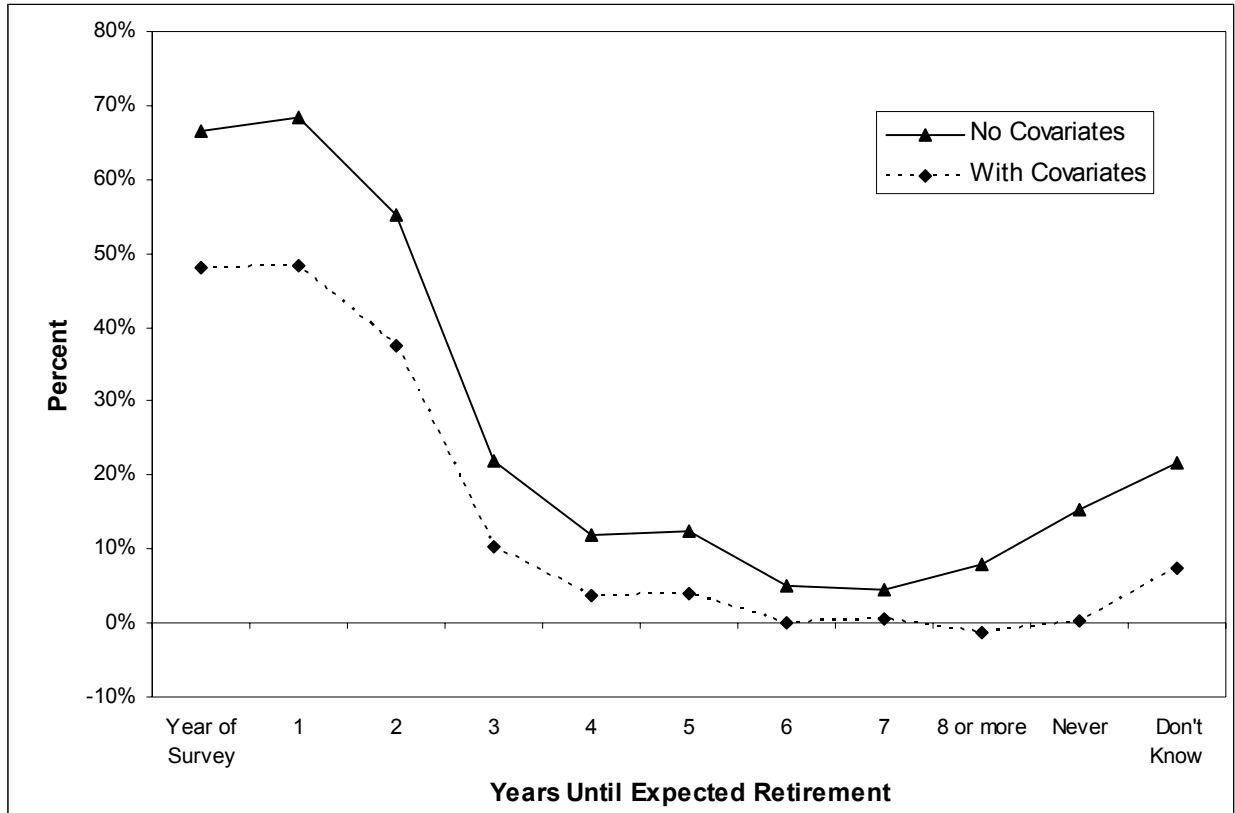

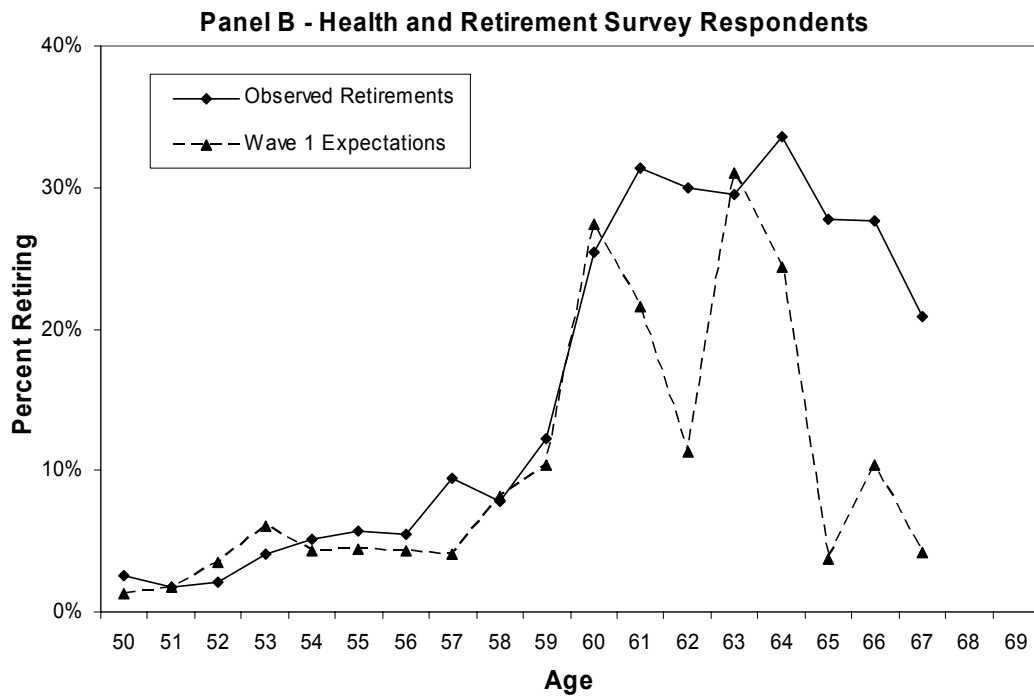



Figure 4: Weekly Food Expenditures by Household Size

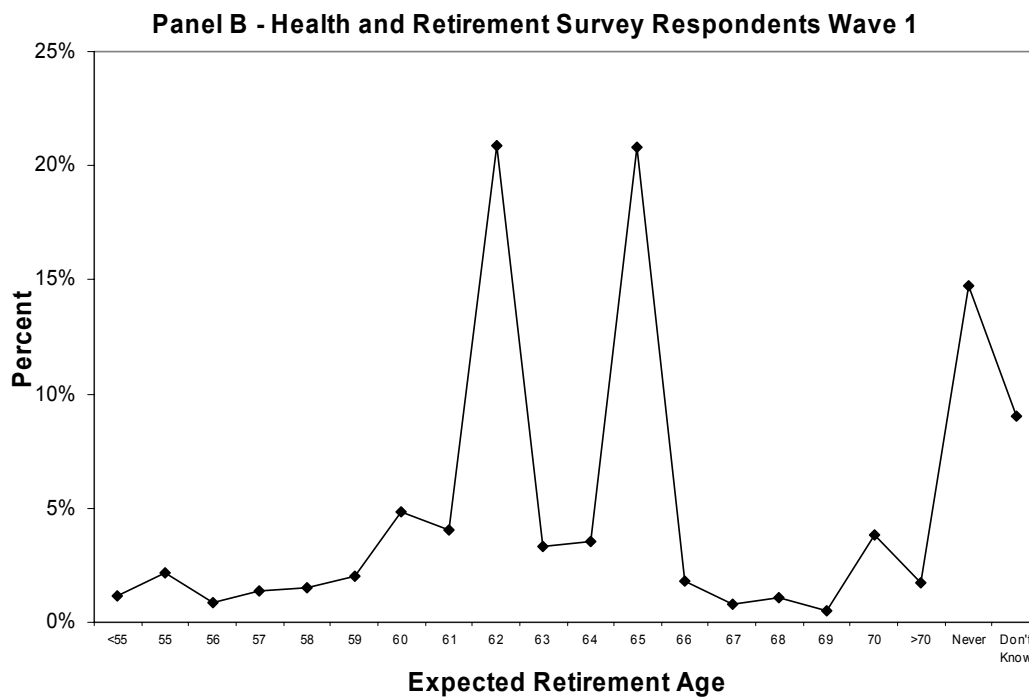



Notes: These figures graph the mean weekly food expenditures by household size. All dollars are deflated to 2001, using the annual Personal Consumption Expenditure (PCE) deflator.

**Figure 5: Fraction of Workers Retiring By the Next Wave
Observed vs. Expecting to Retire**



Appendix Figure 1 – Distribution of Expected Retirement Ages



**Appendix Figure 2 – RHS Fraction Retiring Within One Year of the Previous Survey Wave
By Most Recent Expected Retirement Year**

