

**Grasshoppers, Ants and Pre-Retirement
Wealth: A Test of Permanent Income
Consumers**

Erik Hurst



Project #: UM04-S1

“Grasshoppers, Ants and Pre-Retirement Wealth: A Test of Permanent Income Consumers”

Erik Hurst
University of Chicago and NBER

September 2004

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98358-5). The opinions and conclusions are solely those of the authors and should not be considered as representing the opinions or policy of the Social Security Administration or any agency of the Federal Government.

Regents of the University of Michigan

David A. Brandon, Ann Arbor; Laurence B. Deitch, Bingham Farms; Olivia P. Maynard, Goodrich; Rebecca McGowan, Ann Arbor; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; S. Martin Taylor, Grosse Pointe Farms; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

Grasshoppers, Ants and Pre-Retirement Wealth: A Test of Permanent Income Consumers

Erik Hurst

Abstract

This paper shows that households who enter retirement with low wealth consistently followed non-permanent income consumption rules during their working years. Using the Panel Study of Income Dynamics (*PSID*), household wealth in 1989 is predicted for a sample of 50-65 year olds using both current and past income, occupation, demographic, employment, and health characteristics. Using the residuals from this first stage regression, the sample of pre-retired households is subsetting into households who save 'lower' than predicted and all other households. By construction, these households had similar opportunities to save; the average household in both these sub-samples are very similar along all observable income and demographic characteristics. It is then shown that households in the low wealth residual sample had much larger declines in consumption upon retirement. Such a result is consistent with the household having inadequately planned for retirement. The panel component of the *PSID* is then used to analyze the consumption behavior of these households early in their lifecycle. It is shown that these low pre-retirement wealth households had consumption growth that responded to predictable changes in income during their early working years. No such behavior was found among the other pre-retired households. Moreover, the low wealth residual households responded both to predictable income increases as well as predictable income declines, a result that is inconsistent with a liquidity constraints explanation. After ruling out other theories of consumption to explain these facts, it is concluded that households who entered retirement with lower than predicted wealth consistently followed near sighted consumption plans during their working lives.

Authors' Acknowledgements

I would like to thank Mark Aguiar, Steve Davis, Jon Gruber, David Laibson, Anna Lusardi, Joe Lupton, Brigitte Madrian, Jim Poterba, Karl Scholz, Matthew Shapiro, Jon Skinner, Mel Stephens, Steve Venti, Paul Willen and Jim Ziliak for helpful comments. Additionally, I would like to thank seminar participants at the University of Chicago's, Graduate School of Business macro lunch, the joint Harvard/MIT public finance workshop, the Harvard macro seminar, the Brown applied economics workshop, the applied economics workshop at the University of Kentucky, the Retirement Session at the 2004 AEA meeting, and the 2003 NBER Aging Summer Institute Workshop. I am indebted to the excellent research of Sonia Oreffice. I would like to acknowledge the financial support received from the Michigan Retirement Research Center. All remaining errors are my own.

It was wintertime, the ants' store of grain had got wet and they were laying it out to dry. A hungry grasshopper asked them to give it something to eat. 'Why did you not store food in the summer like us?' the ants asked. 'I hadn't time', it replied. 'I was too busy making sweet music.' The ants laughed at the grasshopper. 'Very well', they said. 'Since you piped in the summer, now dance in the winter'.

– Aesop's Fable

It is well documented that, conditional on lifetime income, wealth varies dramatically across households entering retirement (Gustman and Juster, 1996; Smith, 1997; Hurst, Luoh and Stafford, 1998; Venti and Wise, 1998 and 2000; Lusardi, 2002). While many authors have attempted to explain this fact (Venti and Wise, 2000; Bernheim, Skinner and Weinberg, 2001; Hurd and Zissimopoulos, 2002), the approach taken in this paper is quite different.¹ In this paper, I directly examine the relationship between households' pre-retirement wealth and their consumption behavior while young. I find that households who entered retirement with much lower than predicted wealth did not follow permanent income consumption rules during their working years; their year-to-year consumption growth responded strongly to predictable income changes. No such behavior was evident in the other group of pre-retired households who had higher wealth conditional on observables. After ruling out other theories of consumption, including the existence of binding liquidity constraints, I conclude that those households who are most likely to under-save for retirement do so, at least in part, because they follow myopic consumption rules during their working years.

Specifically, in the first part of the paper, I segment 50-65 year old households in the 1989 Wave of the Panel Study of Income Dynamics (PSID) by residuals from a regression of observed

¹ There is an additional literature which analyzes the adequacy of pre-retired household wealth by comparing simulated optimal saving behavior from a calibrated lifecycle model to actual household data (Bernheim, 1992 and 1997; Bernheim and Scholz, 1993). The conclusion of these papers is that a large fraction of households have wealth levels that will leave them unprepared to sustain consumption during retirement. More recently, Engen, Gale, and Uccello (1999) and Scholz, Seshadri, and Khitatrakun (2003) maintain that it is possible to account for much of the observed variation in pre-retirement wealth across households using a life-cycle model with heterogeneous earnings shocks and pension coverage. However, both papers conclude that approximately 20% of the pre-retired population under-save for retirement.

household wealth on a vector of current and historical income, employment, demographic and health controls. Doing so, allows me to isolate households with similar opportunities to save. I classify households within the bottom twenty percent of residuals from this first stage regression as having ‘lower than normal’ wealth. In the latter portion of the paper, I justify the use of the twenty percent cutoff. By construction, households with lower than normal pre-retired wealth are very similar to other pre-retired households along income, health, employment, pension and demographic characteristics. However, these households with low pre-retirement wealth residuals experienced a much larger consumption decline upon their subsequent retirement. Their consumption decline during retirement was twice as large, on average, compared to their counterparts with similar pre-retirement income trajectories, but higher pre-retirement wealth. The dramatic decline in consumption at retirement for households with lower-than-normal wealth is consistent with the hypothesis that these households were ill-prepared for retirement.

The innovation of the paper comes next. Using the panel dimension of the *PSID*, I am able to observe the income and consumption behavior of these pre-retired households over a majority of their working years. I then test whether households who appear to violate the Permanent Income Hypothesis (PIH) by accumulating too little wealth to sustain consumption in retirement also violate the PIH while young. Performing standard excess sensitivity tests, I am able to reject that households with lower than normal pre-retired wealth behave as standard permanent income consumers while young. I find that the consumption growth of households with lower than normal pre-retirement savings responds strongly to predictable income changes. The consumption of other pre-retired households does not respond in any way to predictable income changes. Those households entering retirement with little wealth relative to their income, health and demographic trajectories appear to be following rule of thumb consumption plans during their working years.

The fact that consumption responds to predictable income changes for households with little pre-retirement wealth - relative to observables - is not the result of liquidity constraints. Liquidity

constraints may prevent a household from borrowing to smooth their consumption when income is predicted to increase. As a result, their consumption growth may track predictable income growth. However, liquidity constraints do not prevent households from smoothing predictable income declines (Altonji and Siow (1987), Zeldes (1989) and Shea (1995)). If the household realizes that their income is going to decline in the near future, nothing prevents them from saving to smooth their marginal utility of consumption over time. I find that the consumption of households with relatively little retirement savings responds similarly (in terms of magnitude) to both predictable income increases and predictable income declines. These results indicate that liquidity constraints are not the cause of the failure of the permanent income hypothesis during working years for households with little pre-retirement wealth. Furthermore, the substitutability between consumption and leisure is not driving the results. Households with lower than normal wealth entering retirement had consumption profiles while young that responded to predictable income changes even after directly controlling for changes in work hours.

This paper provides a set of facts that describe at least two different types of households. Most households in the population behave according to the Permanent Income Hypothesis. However, there is a segment of households who enter retirement with very low wealth even after controlling for differences in income, demographic, employment and health histories. These same households experience a large consumption decline at the onset of retirement, relative to other pre-retired households. Additionally, these households have consumption profiles that respond to predictable income shocks throughout their working years. Many alternative theories can explain a subset of the above behaviors, but very few theories can jointly explain them all. Specifically, it is shown that the behaviors of these low wealth, pre-retired households are inconsistent with consumption theories such as precautionary savings (Deaton, 1992; Carroll, 1997) or habit formation (Deaton, 1992). These behaviors, however, are consistent with either rule-of-thumb consumption (Campbell and Mankiw, 1989) or hyperbolic consumers (Laibson, 1997). In either case, the households display a lack of planning behavior; under the former

theory, the households are myopic and do not attempt to plan for the future, while under the latter theory, the households attempt to plan, but are incapable of committing themselves to carry out those plans.

As supporting evidence that differences in planning propensities are driving the differences in behavior between the two groups, I find that these households with low wealth entering retirement, conditioned on lifecycle factors, were aware of their near-sighted behavior nearly two decades prior to their retirement. In 1972, questions were asked of all PSID respondents about 1) their propensity to plan for the future, 2) carry out their plans for the future and 3) their propensity to spend their income rather than save it. The answers to such questions are definitely noisy measures of household behavior. However, households who entered retirement with lower than normal savings were much less likely to report that they plan for the future, were much less likely to report that they carry out their plans and were much more likely to report that they spend their income rather than save it.

At the beginning of the paper, a classic fable by Aesop is recounted. In the fable, a sharp distinction is drawn between ants, who saved during their summer (i.e., working years) to sustain consumption during their winter (i.e., retirement), and grasshoppers, who saved little for their future period of low earnings. The results shown in this paper confirm Aesop's proposition; within an economy, consumers are of different types. Some households are forward looking and behave according to the Permanent Income Hypothesis. Others, however, lack either the desire or the ability to plan for the future. When constructing economic models, it is often misguided to assume all households follow similar consumption rules. With respect to retirement saving, while it is true that the majority of households appear to follow permanent income consumption rules, approximately 20% of the population behaves as 'economic grasshoppers'. The consumption of such households closely tracks their income during their working years leaving them with little financial wealth as they enter retirement. Given their lack of planning, such households must sharply decrease consumption – and by extension utility – at the time of retirement.

This paper adds to a growing literature which assesses the effects of planning behavior on household wealth accumulation (Lusardi, 1999, 2002, 2003a; Ameriks, Caplin and Leahy, 2003). Lusardi (2002) uses data from the Health and Retirement Survey (HRS) to show that 1/3 of households nearing retirement report that they have “hardly thought” about retirement. These households have much lower wealth levels than households who reported thinking a lot about retirement. Similarly, Ameriks et. al. (2003) use special surveys of TIAA-CREF investors to show that households who self report spending time financially planning for retirement have higher wealth than households who self report spending little time financially planning. Also, both Bernheim and Garrett (2003) and Lusardi (2003b) find positive causal effects between attending firm sponsored retirement planning seminars and retirement wealth. Collectively, there exists evidence that planning can foster higher savings.

II. Segmenting Pre-Retired Households By Wealth Residuals

The purpose of this section is to identify pre-retired households with lower than ‘normal’ wealth. Household wealth, at the time of retirement, is a function of economic factors (income, demographics, health shocks, interest rates) and individual decision factors (saving propensities, portfolio allocation). In order to explore household planning behavior, households who had similar opportunities to save over their lifetime are compared. Using the 1989 Panel Study of Income Dynamics (*PSID*), where there are multiple decades of past income, demographic and health data for each household, pre-retired households are segmented by whether they have higher or lower wealth than other households who experienced similar economic, demographic, and health histories.

A. Segmenting Pre-Retired Households

To start, a cross section of pre-retired households who were in the *PSID* during the 1989 survey were examined. Pre-retired households are defined to be households with a non-retired head between the age of 50 and 65. The analysis year of 1989 was chosen purposefully. While

the *PSID* has collected income, employment and demographic information in all survey years since its inception in 1968, information on wealth and savings were only asked at five-year intervals starting in 1984 through 1999. Since 1999, the *PSID* has tracked household wealth every other year. Consequently, cross-sectional studies of wealth using *PSID* data are limited to the years of 1984, 1989, 1994, 1999, and 2001 where the 2001 data is the most recent *PSID* data made available. Given the nature of the tests that I perform below, I need to follow the pre-retired households backwards in time (to observe their consumption behavior during their working years) as well as to follow the households into the future (to observe their consumption behavior around the period when they subsequently retire). For these reasons, 1989 was chosen as the year in which pre-retired households were segmented by their conditional wealth levels.² Additionally, as discussed in the Data Appendix, certain pension questions used in the wealth prediction equation were asked only in 1989. Given the sample design of the *PSID*, nearly all 1989 pre-retired households had one family member participate in the *PSID* every year since the survey's inception in 1968. As a result, there are almost twenty years of income, employment, demographics and health data for each pre-retired household in the 1989 *PSID*.

The analysis sample was restricted to include only those households who had positive 1989 wealth. Given that pre-retired households are well into their lifecycle, this requirement was not overly restrictive. Less than 4% of non-retired 50-65 year olds in the 1989 *PSID* had zero or negative wealth. The positive wealth restriction was imposed so that the log of wealth can be used as the dependent variable in subsequent regressions.

The *PSID* wealth supplements contain information on the household's investment in real estate (including main home), vehicles, farms, businesses, stocks, bonds, mutual funds, saving and checking accounts, money market funds, certificates of deposit, government savings bonds, Treasury bills, IRAs, bond funds, cash value of life insurance policies, valuable collections for

² The main results of the paper carry through if the 1984 wealth data is used to segment pre-retired households by their wealth residuals.

investment purposes, and rights in a trust or estate, mortgage debt, credit card debt, and other outstanding collateralized and non-collateralized debt. The measure of wealth used in this paper is the sum of all of the above asset measures less all of the above debt measures. For a full discussion of the *PSID* wealth data, see Hurst, Luoh and Stafford (1998).

The *PSID* wealth supplements have one major drawback when used to assess retirement savings. Through the late 1990s, the *PSID* did not explicitly ask households about their wealth in either private or public pensions. In 1989, however, this problem can be partially alleviated. The 1989 *PSID* respondents were asked questions about their expectations of the percentage of their pre-retirement yearly labor earnings that would be replaced by all household pension plans (including social security) during retirement. Using this information, one can account for expected differences in pension coverage across households.

To identify households who saved little given their economic opportunities, the following regression was estimated:

$$W_{i,1989} = \varphi_0 + \varphi_1 X_{i,1989} + \varphi_2 Z_{i,historical} + \eta_{i,1989}, \quad (1)$$

where $W_{i,1989}$ is the log of household i 's net wealth in 1989, $X_{i,1989}$ is a vector of household i 's 1989 income, employment, demographic, and health controls, and $Z_{i,historical}$ is a vector of household i 's historical income, employment, demographic and health controls. The error term, $\eta_{i,1989}$, represents the portion of current household log wealth that is unexplained by the X and Z controls.

The Data Appendix details the specific controls used in estimating regression (1). Briefly, $X_{i,1989}$ includes the current age of the household head, age squared, dummies for the household head's race, marital status, educational attainment, occupation, industry and family composition, a quadratic in household current total labor income, dummies for the household head and "wife's" current health and employment status, and the household's self reported expectation of

their pension replacement rate.³ $Z_{i,\text{historical}}$ includes a quadratic in average household labor income between 1980 and 1987, a quadratic in average household labor income between 1974 and 1979, the change in labor income between 1980 and 1988, the coefficient of variation of income over 1975 and 1989, and health and unemployment shocks experienced by the head and the wife between 1980 and 1988. In some specifications, health and unemployment shocks experienced by the household in the 1970s were included. These variables provided no additional explanatory power to the regression and, as a result, were omitted from the base specification.

The residuals from (1), $\eta_{i,1989}$, provide a measure of whether the household has saved more or less than households with similar economic, demographic, employment, and health trajectories. The adjusted R-squared from (1) was 0.53 indicating that the controls included captured a majority of the variation in wealth across households. Figure 1 presents the distribution of 1989 wealth residuals for the sample of 1989 pre-retired households. Any classification of households into two groups based on these wealth residuals is in some sense arbitrary. To begin, I segment households with the lowest 20% of residuals as having low “normalized” wealth. These households correspond to the proverbial economic grasshoppers discussed above. My comparison group will be all other pre-retired households in the sample (the proverbial economic ants).

The 20th percentile cutoff is chosen given that: 1) Hall and Mishkin (1982) find that about twenty percent of the population appears to be rule of thumb and 2) both Engen et al. (1999) and Scholz et al. (2003) find that about 20% of households under save for retirement. However, in the sections that follow, I explore the robustness of my results when the cutoff is redefined as the 10th, 30th, 40th or 50th percentile of the residual distribution. For my sample, based on the first stage regression, the corresponding cutoffs of the log pre-retirement wealth residuals for the 10th, 20th, 30th, 40th and the 50th percentiles of the wealth residual distributions are, respectively, -1.32,

³ The *PSID* surveys its respondents in the spring of the year. During the 1989 survey year, households are asked about their wealth (spring 1989) and about the previous year of income (1988 income). As a result, when predicting 1989 wealth, 1988 income is the appropriate income measure.

-0.73, -0.36, -0.12, and 0.12. As we will see in the following sections, the twenty percent cutoff is well justified.

These residuals, however, should not be strictly interpreted as a measure of a household's planning propensity because many of the controls in equation (1) are a function of the extent to which individuals are forward looking. For example, education affects income (the opportunity to save). However, education is also the result of household forward looking behavior. The fact that wealth differs across households despite the inclusion of these controls suggests that I have isolated groups of households who have different saving propensities above and beyond the extent that these saving propensities are correlated with the controls included in X and Z . The alternative would be to include only controls which are totally exogenous to the propensity to plan. However, there are very few controls that are truly exogenous. Even age, given its relationship to health, is potentially a proxy for household planning behavior. Using very few controls in equation (1) would leave the results discussed below open to the criticism that the sample selection procedure isolated households who had different opportunities to save. By including controls that may also proxy for household planning propensity into equation (1), however, my estimates below may underestimate the amount of 'grasshopper' behavior in the sample. Given this, the estimates below should be treated as lower bounds. In Section IV, the same analysis is performed on samples split using actual wealth levels as opposed to the wealth residuals. In that section, the results of the two approaches are compared. Foreshadowing the results, the *qualitative* conclusions of the paper are not sensitive to the control variables included in equation (1). The reason for the similarity in results is discussed in Section IV.

Table 1 presents descriptive statistics for the two samples of pre-retired households where the sample is split based on the first stage wealth residuals estimated from (1). Aside from

wealth, the two samples look very similar along income, demographic and health histories.

Given the sample selection procedure, this result should not be surprising.⁴

B. Differences in Subsequent Retirement Behavior By Wealth Residual Groups

If differences in planning ability across households were driving the difference in normalized pre-retirement wealth across households, we would expect to see subsequent differences in retirement behavior. Households who accumulated too little wealth for retirement should react to the realization that they were ill prepared by: 1) reducing their consumption in retirement and/or 2) delaying the time of their retirement (or working a part time job after retiring).

There are potentially two interrelated drawbacks to comparing the subsequent retirement behavior of the households with low wealth residuals and all other pre-retired households. First, given that the 1999 *PSID* is the most recent data available for public use, it is not possible to observe all households actually retiring. A household who was 50 in 1989 will likely not retire until the mid 2000s. Second, and potentially more important, those who retire early may be a selected sample. One may imagine that those households with very low wealth would delay their retirement relative to other households. There is no information, however, to suggest that the households in the two pre-retirement wealth residual groups retire at different ages. Panel A of Table 2 shows that the average age of retirement, conditional on retiring, is between 62 and 63 years old for both those with low and high first stage 1989 pre-retirement wealth residuals. Self-reported retirement status is used to define household retirement behavior. Furthermore, similar percentages of both groups were observed actually retiring before 1999 (43% of ants and 38% of grasshoppers). Given the incentives in public and private pension systems and the fact that the average household in each group is similar along occupation, education and income dimensions, it is not surprising that the retirement age is similar between the two groups. However, as we

⁴ Unless otherwise specified, all dollar values reported in the paper are in 1989 dollars.

show below, the propensity to take a part-time job in retirement does differ between the two groups.

The only measure of consumption, aside from housing expenditures, that the *PSID* directly asks its respondents about is their food consumption.⁵ Specifically, in all years between 1970 and 1987 and all survey years between 1990 and 1999, households were asked to report the amount that that they spent on food at home and food away from home during the previous month. Consistent with the hypothesis that our group of low residual households are ill prepared for retirement, it is found that such households have much larger declines in consumption upon retirement. Table 2 shows *the level of food consumption* averaged over the three years preceding retirement, *the level of food consumption* averaged over the three years after retiring, and the mean and median percentage decline in consumption associated with retirement for those households that retired. The percentage decline compares the three year average consumption prior to retirement to the three year average after retirement for each household who retired and then averaged the percentage declines over all households.

The mean consumption levels prior to retirement for both groups of households were quite similar (second row of Table 2). However, after retirement, those with low pre-retired wealth residuals consumed \$2,900 of food per year while those with the higher wealth residuals consumed over \$3,700 of food per year. The average decline in consumption for the low wealth residuals households was 11%, while the other group, on average, only decreased their consumption during retirement by 3%. The median decline in consumption at retirement showed a similar pattern: low wealth pre-retired households experienced nearly a 20% consumption decline compared to a 11% decline for the other households. The fact that the average household

⁵ Many authors examining consumption in the *PSID* use the ‘Skinner’ measure of consumption (Skinner, 1987). The Skinner consumption measure optimally weights food consumption with measures of housing expenditures to come up with a total measure of consumption. However, when examining changes in consumption, all the time series variation in the Skinner consumption measure comes from either the variation in food consumption or the variation in housing expenditures. Given that housing expenditures may be directly related to a household’s level of wealth (because of liquidity constraints in the housing market), it is inappropriate to use the Skinner consumption measure when estimating consumption Euler equations when the samples are split based on wealth. For this reason, in this paper, I only focus on food consumption and do not include housing expenditures in my *PSID* consumption measure.

experiences a consumption decline during retirement is consistent with almost all existing empirical work (see, Banks, Blundell and Tanner, 2000; Bernheim, Skinner and Weinberg, 2001, Aguiar and Hurst, 2003; Hurd and Rohwedder, 2003).⁶ Furthermore, Bernheim et al. (2002) also finds that low wealth households experience much larger consumption expenditure declines at the onset of retirement.

The results in Table 2 are based on very few low wealth residual households who were observed to subsequently retire. As seen in Table 2, the average decline in consumption between the two samples is not statistically different from each other at standard levels. The question is whether the lack of significance is due to low power or is it because there is no actual difference in behavior between the two groups. To explore the robustness of these results, the exact same analysis was performed using 1984 wealth to split the sample.⁷ Doing so allowed many more households who actually retired between 1984 and 1999 to be observed. The results are shown in Appendix Table A1 and are nearly identical to the results shown in Table 2. Low residual households experienced a consumption decline that was nearly twice as large as the other pre-retired households. However, these average and median declines are statistically different from each other at standard levels of significance. The results from Table 2 and Appendix Table A1 document large differences in consumption declines between the two groups at the time of retirement.

While the average household in each group retire at similar ages, there is evidence that the low wealth residual households take less leisure in retirement. Households in the bottom 20% of residuals and the remaining 80% of residuals both work similar average annual hours during the three years prior to retirement (1,631 hours vs 1,806 hours, p -value of difference = 0.11).

However, during the three years after retirement, low wealth households work 413 hours annually

⁶ Both Hurd and Rohwedder (2003) and Aguiar and Hurst (2003) provide evidence that a large portion of the decline in expenditure – particularly food expenditure - associated with retirement for the median household is the result of a switch towards home production. Aguiar and Hurst show that while food expenditure falls for the average households by 10-20%, food consumption, measured via food diaries, does not change at the onset of retirement.

⁷ For a sample of non-retired, 50-65 year olds in 1984, 1984 log wealth was regressed on the full set of variables used to segment the 1989 pre-retired households, except all control variables were lagged 5 years.

(on average) compared to 240 annually for high wealth households (p -value of difference = 0.05).

The average decline in work hours is 11 percentage points less for low wealth residual households (75% decline vs. 86% decline, p -value of difference = 0.03). These work hours result from the retired household working a part-time job. So, even though the low wealth residual and high wealth residual households retire at a similar age, the low wealth residual households are much more likely to work after their self reported age of retirement. These results can be reconciled given institutional incentives for households to retire between the ages of 63 and 65.

III. Testing For Differences in PIH Behavior Across Wealth Residual Groups

The large decline in consumption at retirement for low wealth households seems at odds with the standard permanent income hypothesis (PIH) model (Modigliani and Brumberg, 1954; Friedman, 1957).⁸ Given that the date of retirement is largely forecastable, forward looking households should accumulate enough wealth so as to sustain consumption during retirement. In this section, I explicitly test whether the low wealth residual households behaved as permanent income consumers in the decades prior to retirement. In other words, do low wealth pre-retired households seem to violate the PIH throughout their lives or only at the time of retirement?

According to the PIH with perfect capital markets and patient consumers, expected income growth between period t and $t+1$ should not have statistical power in predicting consumption growth between period t and $t+1$ (Hall 1978). Any predictable future changes in the household's income stream should already be incorporated into the household's current consumption plan. Linking household retirement wealth with their consumption behavior while young is the innovation of this paper.

⁸ If consumption and leisure are substitutes, the PIH model could predict a large consumption decline at retirement. This proposition is tested in Section V.

A. *Empirically Testing for Permanent Income Consumption Behavior*

In this sub-section, I outline the standard PIH model as presented in Zeldes (1989). I then use this model to test for differing consumption behavior between groups of households defined by their pre-retirement wealth residuals.

Assume that households solve the following maximization problem:

$$\begin{aligned} \max_{C_{ikt}} \quad & u(C_{ikt}, \Theta_{ikt}) + E_t \sum_{s=t+1}^T \left(\frac{1}{1 + \delta_k} \right)^{(s-t)} u(C_{iks}, \Theta_{iks}) \quad (2) \\ \text{s.t.} \quad & X_{ik,t+1} = (1 + r_{ik,t+1})(X_{ikt} - C_{ikt}) + Y_{ik,t+1} \\ & Y_{ikt} = P_{ikt} V_{ikt} \\ & P_{ikt} = g_{ik} P_{ik,t-1} N_{ikt} \\ & u(C_{ikt}, \Theta_{ikt}) = \frac{C_{ikt}^{1-\rho_k}}{1-\rho_k} \exp(\Theta_{ikt}), \quad \rho_k > 1; \end{aligned}$$

where i indexes households, k indexes household type (i.e., whether or not the household belongs to the low pre-retirement wealth residual group) and t indexes time; C_{ikt} , X_{ikt} , and Y_{ikt} are, respectively, household i 's consumption, cash on hand for consumption, and household income in period t ; $r_{ik,t+1}$ is the household specific after tax interest rate between years t and $t+1$ and δ_k is the discount rate that pertains to a household of type k . The household's utility function is of the Constant Relative Risk Aversion form with a time invariant coefficient of relative risk aversion, ρ_k , and a time invariant time discount rate, δ_k . Both ρ and δ are allowed to differ by household type but are assumed constant within each type. Utility of the household is also dependent upon the household's tastes in period t , Θ_{ikt} . Household income can be decomposed into two parts; a permanent component (P_{ikt}) and a transitory component (V_{ikt}). Permanent income in the current period is equal to permanent income in the previous period multiplied by a nonstochastic growth factor (g_{ik}), specific to the household, and a stochastic shock (N_{ikt}). The stochastic components to income $\{N_{ikt}, V_{ikt}\}$ are assumed to be independently and identically distributed jointly lognormally

with zero means and variances of the underlying distributions equal to zero and $\{\sigma_{ik,N}^2, \sigma_{ik,V}^2\}$, respectively.

The Euler Equation to the above optimization problem can be estimated with the familiar specification (see, among others, Shapiro (1984), Zeldes (1989) and Lawrance (1991)):⁹

$$\dot{C}_{ik,t+1} = \frac{-\ln(1 + \delta_k)}{\rho_k} + \frac{\omega_{ik}^2}{2\rho_k} + \frac{\ln(1 + r_{ik,t+1})}{\rho_k} + \frac{(\Theta_{ik,t+1} - \Theta_{ikt})}{\rho_k} + \varepsilon_{ik,t+1}, \quad (3)$$

where $\dot{q}_{ik,t+1} = \Delta \ln q_{ik,t+1}$, for any variable q , $\varepsilon_{ik,t+1}$ is the mean zero expectations forecast error, and ω_{ik}^2 is the variance of the forecast error. The law of iterated expectations implies that $\varepsilon_{ik,t+1}$ is uncorrelated with any variable known at time t (Hall, 1978).

Similar to Zeldes (1989), household tastes are defined according to the following:

$$\Theta_{ikt} = b_{0k} age_{ikt} + b_{1k} age_{ikt}^2 + b_{2k} \ln(famsize_{ikt}) + \tau_{ik} + \mu_t + \xi_{ik,t+1} \quad (4)$$

where age_{ikt} is the age of the household head in year t and $famsize_{it}$ represents the number of members in the household in year t . The effects of age and family size on the taste shifter are allowed to differ by household type, k . The unobservable (to the econometrician) component of the taste shifter includes a family fixed component which is constant over time for any family within a type (τ_i), an aggregate component that is constant across types and families but varies across time (μ_t), and a remaining component that is orthogonal to the other two (ξ_{ikt}).^{10, 11}

Substituting (4) into (3), one gets:

$$\dot{C}_{ik,t+1} = \lambda_{0k} + \lambda_{1k} \ln(1 + r_{ik,t+1}) + \lambda_{2k} \dot{famsize}_{ik,t+1} + \lambda_{3k} age_{ikt} + \mu_{t+1} - \mu_t + \varepsilon_{ik,t+1}^* \quad (5)$$

⁹ The solution to this model is discussed in Zeldes (1989). Also, see Lawrance (1991) for a discussion of estimating this consumption Euler equation when consumption is measured with error.

¹⁰ Innovations to ξ_{ikt} are assumed to be persistent such that $E_t[\xi_{ik,t+1} - \xi_{ikt}]$ equals zero.

¹¹ Allowing a component which varies by type over time (ψ_{kt}) did not alter the results presented below in any way.

where $\varepsilon_{ik,t+1}^* = \varepsilon_{ik,t+1} + (\xi_{ik,t+1} - \xi_{ikt})/\rho_k$ and has mean zero. The constant, λ_{0k} , can be expressed as $(\delta_k - \omega_{ik}^2/2 + b_{0k} + b_{1k})/\rho_k$. The coefficient λ_{1k} in (5) is equal to $(1/\rho_k)$.¹²

As outlined in the previous section, (5) will be jointly estimated for the two different sub-populations of households - those with low first stage pre-retired wealth residuals and all other households. Formally, the following equation allows for the parameters of (5) to differ accordingly between the two groups of households:

$$\begin{aligned} \dot{C}_{ik,t+1} = & \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1 + r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1 + r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} \\ & + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi D_{Year} + \varepsilon_{ik,t+1}^* \end{aligned} \quad (6)$$

where $D_{<20}$ is a dummy variable equal to 1 if the household has a first stage wealth residual (defined in the previous section) in the lowest twenty percent of the wealth residual distribution. Including the low residual dummies and these dummies interacted with the interest rate, age and family size allows for preference parameters (δ and ρ), as well as the impact of taste shifters (b_1 and b_2), to differ across the two groups. D_{Year} is a vector of year dummies which are included to account for aggregate shocks which affect both types of households.

To test whether household consumption responds to predictable changes in income, the following regression can be estimated:

$$\begin{aligned} \dot{C}_{ik,t+1} = & \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1 + r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1 + r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} \\ & + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi D_{Year} + \beta_1 \overset{\text{Predict}}{Y}_{ik,t+1} + \beta_2 D_{<20} \overset{\text{Predict}}{Y}_{ik,t+1} + \varepsilon_{ik,t+1}^* \end{aligned} \quad (7)$$

where $\overset{\text{Predict}}{Y}_{ik,t+1}$ is the predictable component of income growth rate between t and $t+1$ estimated simultaneously with (7).^{13,14} If households are not sufficiently ‘impatient’, the Permanent Income

¹² Given that changes in family size are planned in advance, $\Delta \ln \dot{famsize}_{it}$ is assumed to be uncorrelated with $\varepsilon_{i,t+1}^*$. The following results were also run omitting changes in family size as a control but with the growth in per capita consumption as the dependent variable. The results were unchanged.

¹³ This procedure to test for the excess sensitivity of consumption is standard in the literature. See Browning and Lusardi (1996) and the cites within.

Hypothesis predicts that consumption growth between periods t and $t+1$ should be unaffected by forecastable changes in income between periods t and $t+1$.¹⁵ Any predictable change in income should already be included in the household's consumption plan. If either β_1 or β_2 is positive and significant, predictable income growth has statistical power in predicting consumption growth and the standard Permanent Income Hypothesis with no liquidity constraints and patient consumers can be rejected.¹⁶

In order for the two stage least squares estimation of (7) to yield unbiased estimates of β_1 and β_2 , both the predictable income growth components and the dummy indicating the bottom 20% of the pre-retirement wealth distribution have to be independent of the regression error term. In the following empirical work, I will instrument for a household's predictable component of income growth using four lags of income growth, excluding the first lag.¹⁷ By definition, these lagged variables are orthogonal to the error term, $\varepsilon_{i,t+1}^*$.

The standard assumption made about the error structure when estimating consumption Euler equations is that $E_t[\varepsilon_{ik,t+1}^* | \Omega_t] = 0 \forall k$ and $\forall t$, where Ω_t is all information known at time t .

However, given the above sample selection procedure, $E_t[\varepsilon_{ik,t+1}^* | D_{<20}]$ need not equal zero within each wealth residual group. For example, it is possible that households with low wealth residuals repeatedly received poor income draws throughout their lifetime. In other words, these households may have been perpetually unlucky. Once the subsequent period came, they realized bad income growth, adjusted their consumption downward, and continued to expect high income

¹⁴ It has been suggested that the wealth residual itself, as opposed to the wealth dummy, $D_{<20}$, should be interacted with the predictable component of income in (7). Given the theory, this would be inappropriate. The relationship between the wealth residual and the response of consumption to predictable income changes is non-linear. If the permanent income hypothesis is correct, the coefficient on the predictable income change should be zero for most households and positive for others. One could, in principle, include higher powers of the residual interacted with the predictable component of income growth to capture the non-linear relationship. For simplicity, and ease of exposition, the wealth dummy approach is used in this paper. In Section IV and in Table 5, the non-linear relationship between the wealth residuals and the violation of the PIH is explored more fully.

¹⁵ 'Impatient' households are classified as households who wish to borrow, all else equal, in the current period. Formally, households are deemed 'impatient' if the following condition is satisfied: $\gamma_k (r_{i,t+1} - \delta_k) + (\rho_k/2) \sigma_{i,N}^2 < g_i - \sigma_{i,N}^2/2$. This impatience condition is necessary to generate buffer stock saving behavior (Carroll, 1997). Below, I rule out buffer stock saving behavior as an explanation for the results presented in this and previous sections.

¹⁶ We will specifically test for the existence of liquidity constraints in the following section.

¹⁷ The first stage regression showing the validity of the instruments is discussed in the following section.

growth in the following period. Such households who were persistently unlucky (or persistently over optimistic) could end up in retirement with little wealth. Furthermore, their expectations about future income growth would not be ascertained from lagged income growth. In each period, they expected higher future income growth and realized lower income growth. Their expectations, on average, and their realizations, on average, would not be equal for perpetually unlucky households. Such a situation could cause households with negative realizations of ε^* in years prior to 1989 to have low wealth in 1989.

However, under rather general assumptions, estimating (7) will still yield unbiased estimates of β_1 and β_2 . The reason for this is that (7) includes a dummy variable indicating differences in type across households, $D_{<20}$. If the low residual group had persistently low-income realizations (relative to their expectations), the lagged income growth controls will persistently under-predict their expected income growth. These households will be consistently revising their consumption downward with each ‘unlucky’ income draw. As a result, their consumption growth will be lower, on average, than the other group of pre-retired households. If this underestimate is constant over time within the low wealth residual grouping, the differences in expectations will be captured by the inclusion of $D_{<20}$ into the estimating equation. In other words, the estimation of (7) will yield valid estimates of β_1 and β_2 if $\text{Cov}[\varepsilon^*_{i,t+1}, D_{<20} \overset{\bullet}{Y}_{i,t+1} \mid D_{<20}] = 0$. The two groups could have different, non-zero ex-post mean realized shocks. However, in order for (7) to be valid, the shocks must be i.i.d *within* each group. This assumption is not much different than the standard assumption that $\varepsilon^*_{i,t+k}$ is i.i.d. for the full sample. Therefore, the inclusion of $D_{<20}$ allows the two groups to have different mean expected shocks, on average. As long as the shocks are i.i.d. within each group, the instruments for the predictable component of income growth will be orthogonal to the error term within each group.

Two further comments can be made about the identification assumptions. First, we can directly test to see if the income processes differ between the two groups by regressing current

income growth on lags of income growth and those lags interacted with $D_{<20}$. If the two groups have different income processes, the interaction terms would enter significantly. There is no evidence that the two groups have different income processes. The coefficients on the interaction terms in such a regression were essentially zero and no interacted term was statistically different from zero. Furthermore, the joint test of significance on the interacted terms could not reject that the two income processes were similar. Second, this result should not be surprising. By definition, the low wealth residual households and the high wealth residual households are very similar along all observables including actual income shocks (unemployment spells) and expected income shocks (education/occupation/industry interactions) (see Table 1). Given the way the sample was split in regression (1), there should be no reason to believe that the expected income processes differ between the two groups of pre-retired households.¹⁸

B. Estimation Results

Equation (7) is estimated on data from 1975–1987 for the sample of *PSID* households who were ‘pre-retired’ in 1989. Formally, consumption growth is defined as the change in log annual food expenditures between year t and year $t+1$.¹⁹ See the Data Appendix for a full discussion of the creation of household real consumption growth and household real after-tax interest rates.

When estimating (7), I instrument for the predictable component of household labor income growth using 4 lags of household labor income growth, excluding the first lag. If household labor income growth follows an autoregressive or moving average process, past labor income growth will have predictive power in determining expected future labor income growth. As noted above,

¹⁸ For all estimations reported in Tables 3-5, the income processes were estimated separately for each group. All the regressions were re-estimated forcing both groups to have the same income process. As expected, given the above discussion, there was no difference in results between the two procedures.

¹⁹ Food consumption in the *PSID*, used by many authors to estimate consumption Euler equations, is a good measure to test household consumption behavior. First, food consumption has little aspects of durability. Second, because households can substitute away from eating in restaurants or from buying more expensive brands, food consumption will be sensitive to changes in income. To the extent there is habit formation in food consumption or if food consumption responds little to income changes, the estimation of (7) will be biased against finding significant coefficients on β_1 and β_2 .

I allow the income processes to differ accordingly between the two different groups of pre-retired households. A first stage regression of current household labor income growth on four lags of household labor income growth shows that the lags have strong predictive power both for households who have wealth residuals in the top 80% (F -statistic = 10.7, p -value < 0.01) and for households who have 1989 wealth residuals in the bottom 20% (F -statistic = 4.2, p -value < 0.01).

Table 3 shows the results from estimating equation (7). If both the low residual and high residual groups followed standard PIH consumption rules, then both β_1 and β_2 would equal zero. If the low residual wealth groups followed a similar consumption plan as the high wealth residual group, β_2 would equal zero, regardless of the value of β_1 . Table 3 reports that β_1 , the coefficient on the predictable change in income for the whole sample, is negative and not statistically different from zero. β_2 , however, is large, positive, and statistically different from zero. Households who had little pre-retirement wealth relative to their lifecycle characteristics responded positively to predictable income changes. The model predicts that the marginal propensity to consume out of predictable income changes is 56 percentage points higher for households with lower than normal pre-retirement wealth (t -statistic = 2.0). The net response to predictable income changes for the low residual group ($\beta_1 + \beta_2$) is positive (an estimated marginal propensity to consume of 0.40) and statistically different from zero (p -value 0.06).

In summary, the results of this section show that households who display behavior that is inconsistent with the permanent income hypothesis as they transition into retirement also display behavior that is inconsistent with the permanent income hypothesis during their working lives.

IV. Robustness Specifications

The results in section III suggest that pre-retired households residing in the lowest twenty percent of the normalized wealth distribution have different consumption behavior during their working years than other households with similar economic histories. Is the difference in consumption behavior due to liquidity constraints? How sensitive are the results to the choice of

a 20% wealth residual cutoff? Are the findings in section III robust to changes in the sample selection criteria? In this section, all of these questions are explored.

A. Empirically Testing for the Existence of Binding Liquidity Constraints

If market imperfections prevent households from borrowing when expected income growth is positive, predictable income growth will have statistical power in predicting consumption growth. The lower the household wealth, the more likely the household will be liquidity constrained. However, as noted by Altonji and Siow (1987), Zeldes (1989) and Shea (1995), it is possible to empirically test whether the existence of liquidity constraints is driving the rejection of the Permanent Income Hypothesis in micro data. Liquidity constraints prevent a household from borrowing but do not place any restrictions on a household's ability to save. As a result, the consumption growth of liquidity constrained households should only respond to predictable *increases* in income, but not predictable income *declines*. If households truly expect their income to decline, they could save a percentage of their income today so as to fund consumption in the future, leaving their discounted marginal utility of consumption unchanged.

To test whether liquidity constraints are causing the rejection of the Permanent Income Hypothesis, the following equation can be estimated:

$$\begin{aligned} \dot{C}_{ik,t+1} = & \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1 + r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1 + r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} \\ & + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 \dot{age}_{ikt} + \alpha_7 D_{<20} \dot{age}_{ikt} + \phi D_{Year} + \beta_3 D_{Up} \dot{Y}_{ik,t+1}^{Predict} \\ & + \beta_4 D_{Down} \dot{Y}_{ik,t+1}^{Predict} + \beta_5 D_{<20} D_{Up} \dot{Y}_{ik,t+1}^{Predict} + \beta_6 D_{<20} D_{Down} \dot{Y}_{ik,t+1}^{Predict} + \varepsilon_{ik,t+1}^* \end{aligned} \quad (8)$$

where D_{Up} and D_{Down} are dummy variables indicating whether the household's predicted income growth between t and $t+1$ was positive or negative, respectively. Table 4 reports the results of estimating equation (8). If liquidity constraints were driving the failure of the permanent income hypothesis for the low wealth residual households, the consumption of the household would respond to predictable income increases but not predictable income declines. In other words, β_5

would be positive and β_6 would be equal to zero. Table 4 shows that the low residual households respond equally to predictable income increases and to predictable income declines. Such a result suggests that liquidity constraints are not driving the rejection of the permanent income hypothesis.

This should not be surprising. By definition, the two groups of households had earned similar levels of labor income over their lives. If one group was liquidity constrained while young, their consumption over that time period should have been lower, all else equal. This implies that liquidity constrained households will show up in retirement with higher wealth, not lower wealth. The fact that a household has low normalized wealth at retirement suggests against liquidity constraints while young.

B. Alternate Specifications of Low Wealth Residual Households

The results in Tables 3 and 4 compare the consumption behavior of those households with the lowest twenty percent of 1989 pre-retirement wealth residuals to the consumption behavior of all other pre-retired households. There is nothing magical about the twenty percent cutoff. One question of interest is how robust are the findings in Tables 3 and 4 to alternate definitions of ‘lower than normal’ pre-retirement wealth. Table 5 explores this question.

Instead of defining the low pre-retirement wealth residual cutoff as an absolute cutoff, the responsiveness to predictable income changes are allowed to differ by residual deciles. Specifically, households are grouped together based upon having pre-retirement residuals (estimated from the first stage regression discussed in section II) in the following ranges: 0-10%, 10%-20%, 20%-30%, 30%-40%, 40%-50%, and greater than 50%. Such a specification allows for a non-linear consumption response to predictable income changes depending on the wealth residuals. Additionally, with such a specification one can see whether the responsiveness of consumption growth to predictable income shocks increases as the wealth residuals become more negative.

Table 5 reports the results of this specification. Households with wealth residuals in the bottom 10% of the distribution respond the strongest to predictable income changes; the marginal propensity to consume out of predictable income shocks is 0.67 (0.93-0.26) with a p -value = 0.07. Households with residuals in the 10th-20th percentile also respond positively to predictable income changes, although the responsiveness was much smaller and not statistically different from zero; 0.08 (0.34 – 0.26), p -value = 0.34. It is worth noting that those in the 10th-20th percentile of the wealth residual distribution did respond differently than those in the top ½ of the wealth residual distribution (coefficient = 0.34, p -value = 0.05). Households in the other wealth residual deciles (20th-30th, 30th-40th, or 40th-50th) did not respond in any way that was economically or statistically different than households with higher wealth residuals. It is concluded that somewhere between 10-20% of the pre-retired wealth distribution have lower than normal pre-retirement wealth and have consumption profiles that respond to predictable changes in income. In other words, the twenty percent cutoff used in the early parts of the paper was well-justified.

C. Redefining the Sample Splitting Equation

As discussed in Section II, given the sample selection procedure, the results may be biased against finding grasshopper behavior. For example, the X and Z controls in regression (1) are likely correlated with poor planning on the part of the household. Households who do not go to college may plan less for the future than households who go to graduate school. By including education dummies as controls when estimating (1), some differences in planning propensities across households will be purged from the residuals. To examine the impact of including such controls in the sample selection equation, two samples are isolated using the actual wealth distribution as opposed to the residual wealth distribution. This procedure is the same as including no X or Z controls when estimating (1). As noted in section II, including many controls has a distinct advantage. The PIH theory says that households who experience different economic, demographic and health trajectories should have different consumption and saving

patterns entering retirement. Using the X and Z controls creates households who are similar along all the X and Z dimensions but differ in wealth. However, it can still be illustrative to examine the results when no controls are used when splitting the sample.

To start, equation (7) is re-estimated with $D^*_{<20}$ is replaced by $D_{<20}$, where $D^*_{<20}$ is defined to include all pre-retired households in the bottom 20% of the actual 1989 log wealth distribution (as opposed to the normalized log wealth residual distribution). The overlap of households in both $D^*_{<20}$ and $D_{<20}$ is 76%. This implies that the actual wealth levels of most households in the left tail of the wealth distribution cannot be explained by the X and Z controls. Given the high overlap of households in the normalized and actual wealth groupings, it is not surprising that the results do not change when (7) is re-estimated using $D^*_{<20}$ instead of $D_{<20}$. The estimate of β_1 was -0.22 (standard error = 0.16) and the estimate of β_2 was 0.85 (p -value = 0.09). Using the same procedure to get the results presented in Table 5, it is found that between 20 and 30% of households with the lowest wealth residual respond to predictable income changes. In other words, slightly more households behave as the proverbial grasshopper using the method with no first stage controls.

V. Interpretation

The results in sections II – IV show that there is a link between household retirement wealth and their consumption while young. Specifically, a set of facts are presented which describe two different types of households. First, there is a segment of households who enter retirement with very low wealth even after controlling for differences in income, demographic, employment and health histories. Second, these same households experience a large consumption decline at the onset of retirement, relative to other pre-retired households. Third, these households have consumption profiles that respond to predictable income shocks throughout their working years. Differences in observables do not explain differences between these groups. These households, by construction, are nearly identical along income, employment, demographic, and health

dimensions. While it is found that a majority of households behave as predicted by the standard permanent income hypothesis, a subset of households (roughly 10-20% of the population) have consumption and saving behavior that is inconsistent with such a theory. The goal, then, is to explain the behavior of this latter group of households.

Many alternative theories can explain a subset of the above facts, but very few theories can jointly explain them all. For example, the results described above for low wealth residual households are inconsistent with models of consumption habits or precautionary savings. While both of those theories could predict excess sensitivity in consumption to predictable income changes during a household's working years, neither would predict such a sharp decline in consumption upon retirement. To the contrary, standard models of consumption habits (Deaton, 1992) predict that consumption will decline less slowly during retirement for those households with habit preferences, compared to PIH households, all else equal. Furthermore, there is nothing to suggest that models of habit persistence in consumption would lead households to have extremely low levels of wealth upon retirement. Precautionary models of saving may predict low pre-retirement savings along with the excess sensitivity of consumption, but such households would still smooth their consumption across the period of retirement (Gourichas and Parker, 2000). Also, differences in time preferences across households cannot alone generate the above findings. It is true that households with a high, constant time discount rate will enter retirement with little wealth, relative to households with similar income profiles and a lower time discount rate. However, such a household will still smooth the marginal utility of consumption over time. These households will have consumption profiles that would not respond to predictable income changes, including retirement.

A strong substitutability between leisure and consumption could reconcile the main results outlined above. Leisure is high in both retirement and periods of job loss. Households who treat consumption and leisure as substitutes would optimally plan for lower consumption in retirement and, as a result, would save less during their working years. Furthermore, as the household

became retired, and leisure increased, we would expect a sharp decline in consumption if the household was smoothing total utility and not just consumption. Lastly, these households would optimally choose to have a positive correlation between predictable income movements and consumption while young, if the predictable income changes were associated with a change in leisure (such as job loss). When leisure is low, consumption would be high and when leisure is high, consumption would be low.²⁰ To explore whether the substitutability between leisure and consumption is causing the failure of the permanent income tests, changes in total hours worked by both the household head and wife (if present) were included directly into the estimation of regressions (7) and (8). This procedure is similar to that used by Attanasio and Browning (1995) to test for the substitutability of consumption and leisure. The results reported in Tables 3-5 were essentially unchanged with the inclusion of work hours into the estimation equation. Such a finding suggests that the substitutability between consumption and leisure is not causing the failure of the permanent income hypothesis documented above.

There are two possible types of behavior that are consistent with all of the results above. First, households who follow myopic (rule-of-thumb) consumption rules would have consumption that closely tracks income (Campbell and Mankiw, 1987). If households do not plan for the future they will end up in retirement with little wealth, be forced to take a consumption decline upon retirement, and will have consumption profiles that respond to predictable income changes during their working years. Second, a theory of time inconsistent preferences can match the above facts. Such households may want to plan for the future, but are incapable of doing so; these households have relatively high discount rates over short horizons and relatively low discount rates over longer horizons. This discount structure sets up a conflict between today's preferences (which prefer current consumption), and the preferences that will be held in the future (which prefer that they had deferred consumption in the past). A household may realize that retirement is coming, yet in each period the household would choose to postpone

²⁰ See Attanasio and Browning (1995) for a discussion.

saving for retirement until the next period. Eventually, the household could enter retirement with little accumulated wealth, and as a result, consumption would eventually have to decline.

Furthermore, consumers with time inconsistent preferences could have consumption growth that responds to both predictable income increases and predictable income declines (Laibson, 1997) and could have consumption that falls sharply at the time of retirement (Angeletos et al., 2000).

While both the rule of thumb and the time inconsistent preference theories can reconcile the behavior of households with low wealth residuals, it is not possible to disentangle the two theories. Specifically, it is not possible to distinguish whether these identified household are completely myopic with respect to their consumption decisions or whether they would like to plan for the future, but are incapable of doing so. However, certain types of consumers with time inconsistent preferences can be ruled out. Laibson (1997) and Angelitos et. al. (2000) draw distinctions between ‘sophisticated hyperbolic consumers’ and ‘naïve hyperbolic consumers’. The former group is aware of their time inconsistent preferences while the latter is not. Sophisticated hyperbolic consumers will take steps to commit themselves to saving. Such households, knowing their desire to reduce saving and increase consumption, will shift their savings toward illiquid assets (i.e., real estate, business or pension).

There is no evidence among the low pre-retirement wealth residual households to suggest that they are trying to commit themselves to save. Actually, the results suggest the opposite. Only 56% of the pre-retired, low wealth residual households (those in the bottom 20% of the wealth residual distribution) own any ‘illiquid’ assets, where liquid assets are defined as the sum of housing, other real estate and business equity. The comparable number for the other pre-retired households was over 90%. Furthermore, the median fraction of household wealth in these illiquid assets is much smaller for those households with low pre-retirement wealth residuals (0.23 vs. 0.66). If the households with low pre-retirement wealth residuals have time inconsistent preferences, they are not taking actions to commit themselves to save for the future.

In conclusion, there exists at least two different types of consumers in the population; those that plan for the future and those that do not plan for the future. In future work, it would be useful to understand whether such households simply do not plan for the future (are myopic) or do plan for the future, but are unable to commit themselves to save (have time inconsistent preferences).

VI. The Self-Awareness of Economic Grasshoppers

There is one final piece of evidence that suggests pre-retired households in 1989 who had lower than normal wealth did so because of poor planning. In 1972 and 1975, the *PSID* asked its respondents to self assess many of their socio-economic characteristics. Some of the questions asked of *PSID* respondents included: whether they get angry easily, whether they are concerned about failure, whether they feel life will work out, whether they have control over life, and whether they are satisfied with themselves. There were three questions asked as part of these supplements that directly pertained to a household's consumption-savings tradeoff and the household's willingness to plan for the future. Specifically, the questions were:

1. Are you the kind of person that plans his life ahead all the time or do you live more from day-to-day?
2. When you make plans ahead, do you usually carry out things the way you expected?
3. Would you rather spend your money and enjoy life today or save more for the future?

The first question was asked in both 1972 and 1975, while the second and third questions were only asked in 1972. The questions above are inherently vague about what they were intended to measure. I do not want to claim that the answers to such questions are in any way indicative of the household's actual behavior. I do, however, believe that the answers to such questions may provide some suggestive evidence about how the two groups of pre-retired households studied above assess their own behavior.

As noted earlier, most of the pre-retired *PSID* households in 1989 were in the *PSID* since its inception in 1968 and, as a result, provided answers to the self-assessment questions in both 1972

and 1975. The results are reported in Table 6. Responses were compared between households with low normalized pre-retirement wealth residuals (bottom 20%) and all other households. The sample split is identical to the one described in Table 1. As expected, households with lower than normal pre-retirement wealth were much less likely to classify themselves as planners in 1972 (46% for the low wealth residual households vs. 59% for the higher wealth residuals, p -value of difference = 0.06). The results are even more striking in 1975. Households with lower than normal pre-retirement wealth in 1989 only reported themselves to be a planner 38% of the time (compared to 56% of the time for the other 1989 pre-retired households). In 1972, only 54% of the low wealth residual households said that they were likely to carry out plans as expected. The other pre-retired households in 1989 said that they carried out their plans 67% of the time. All these differences were significant at the ninety-four percent level of confidence.

Perhaps the most interesting question is the one that most directly assesses the household's consumption decisions. Of the households with the lowest 1989 pre-retirement wealth residuals, 60% of them reported in 1972 preferring spending money today (as opposed to saving it for the future). The comparable number for the other 1989 pre-retired households was only 40%. While only suggestive, it appears that those households who behave most like the proverbial 'economic grasshopper' are aware of their grasshopper tendencies. Such households report being less likely to plan for their future, less likely to carry out plans conditional on making them, and more likely to spend their money today rather than save it for the future. One should not forget that these self-assessments were made almost two decades prior to when their pre-retirement wealth was measured.

There are two other questions in the early *PSID* surveys that are worth reporting. The first question asked households whether or not they had any positive savings. The second asked if the household had accumulated savings greater than two months of income. These questions were asked both in 1972 and 1975. The responses to these questions for our 1989 pre-retired households are also reported in Table 6. Not surprisingly, households who under-saved entering

retirement were low savers throughout their early working years. For example, in 1972, only 30% of those pre-retired households with wealth residuals in the bottom twenty percent of the wealth distribution had two months worth of accumulated savings. Over ½ of the other pre-retired households had at least two months of accumulated savings. Households who have lower than normal wealth entering retirement were much less likely to have had any significant amount of saving early in their lifecycle.

VI. Conclusion

In this paper, it is shown that households who enter retirement with low wealth consistently followed non permanent income consumption rules during their working years. Using the Panel Study of Income Dynamics (*PSID*), household wealth in 1989 is predicted for a sample of 50-65 year olds using both current and past income, occupation, demographic, employment, and health characteristics. Using the residuals from this first stage regression, the sample of pre-retired households is subsetting into households who save ‘lower’ than predicted and all other households. By construction, these households had similar opportunities to save; the average household in both these sub-samples are very similar along all observable income and demographic characteristics. It is then shown that households in the low wealth residual sample had much larger declines in consumption upon retirement. Such a result is consistent with the household not adequately planning for retirement.

In the main part of the paper, the panel component of the *PSID* is used to analyze the consumption behavior of these households early in their lifecycle. It is shown that these low pre-retirement wealth households had consumption growth that responded to predictable changes in income during their early working years. No such behavior was found among the other pre-retired households. Moreover, the low residual households responded both to predictable income increases as well as predictable income declines, a result that is inconsistent with a liquidity constraints explanation. It is also shown that the set of facts presented about households with

low pre-retirement wealth residuals are inconsistent with theories of habit formation or precautionary saving. Lastly, these results are not caused by the households treating consumption and leisure as substitutes.

There are two theories that can explain all of the behaviors associated with low wealth residual households. First, households who follow myopic (rule-of-thumb) consumption rules would have consumption that closely tracks income. If households do not plan for the future they will end up in retirement with little wealth, be forced to take a consumption decline upon retirement, and will have consumption profiles that respond to predictable income changes during their working years. Second, the results could be reconciled if the household had time inconsistent preferences. Such households may want to plan for the future, but are incapable of doing so. There is no evidence, however, that these households had taken steps to commit themselves to saving for the future. While both the rule of thumb and the time inconsistent preference theories can reconcile the behavior of households with low wealth residuals, it is not possible to disentangle the two theories. Specifically, given PSID data, it is not possible to distinguish whether these identified household are completely myopic with respect to their consumption decisions or whether they would like to plan for the future, but are incapable of doing so. Regardless, the near sighted behavior of both types of households leaves them ill prepared to sustain consumption through retirement.

At the beginning of the paper, a classic fable by Aesop is recounted. In the fable, a sharp distinction is drawn between ants, who saved during their summer (i.e., working years) to sustain consumption during their winter (i.e., retirement), and grasshoppers, who saved little for their future period of low earnings. The results shown in this paper confirm Aesop's proposition; within an economy, consumers are of different types. Some households are forward looking and behave according to the Permanent Income Hypothesis. Others, however, lack either the desire or the ability to plan for the future. When constructing economic models, it is often misguided to assume all households follow similar consumption rules. With respect to retirement saving, while

it is true that the majority of households appear to follow permanent income consumption rules, approximately 20% of the population behaves as ‘economic grasshoppers’. In future work, it would be useful to understand the reasons why households do not plan for the future. Exploring the origins of such households could also be fruitful. Are households born of a given type or do they evolve as their life progresses?

References

- Aguiar, Mark and Erik Hurst (2003). "Consumption vs. Expenditure", University of Chicago Working Paper.
- Altonji, Joseph and Aloysius Siow (1987). "Testing the Response of Consumption to Income Change with (Noisy) Panel Data." *Quarterly Journal of Economics*, 102(2): 293-328.
- Ameriks, John, Andrew Caplin, and John Leahy (2003). "Wealth Accumulation and the Propensity to Plan," *Quarterly Journal of Economics*, 118(3), 1007-48.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001). "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15(3), Summer, 47-68.
- Attanasio, Orazio and Martin Browning (1995). "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review*, 85(5): 1187-1237.
- Banks, James, Richard Blundell and Sara Tanner (1999). "Is There a Retirement-Savings Puzzle?" *American Economic Review* 88(4): 769-88.
- Bernheim, Douglas (1992). "Is the Baby Boom Generation Preparing Adequately for Retirement? Technical Report." Unpublished paper. New York. Merrill Lynch.
- Bernheim, Douglas (1997). "The Adequacy of Personal Retirement Saving: Issues and Options." In *Facing the Age Wave*, eds, David Wise. Stanford, CA. Hoover Press.
- Bernheim, Douglas and Daniel Garrett (2003), "The Effects of Financial Education in the Workplace: Evidence from a Survey of Households," *Journal of Public Economics*, 87, pp. 1487-1519.
- Bernheim, Douglas and Karl Scholz (1993). "Private Saving and Public Policy." In *Tax Policy and the Economy*, vol 7, eds, James Poterba. Cambridge and London: MIT Press.
- Bernheim, 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.
- Browning, Martin and Annamaria Lusardi (1996). "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature* 34(1): 1797-1902.
- Carroll, Christopher (1997). "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112(1): 1-56.
- Campbell, John and Gregory Mankiw (1989). "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." *NBER Macroeconomics Annual* 4(1): 185-247.
- Deaton, Angus (1992). *Understanding Consumption*. Oxford: Clarendon Press.
- Dynan, Karen (2000). "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 90(3): 391 – 406.

- Engen, Eric M., William G. Gale, and Cori Uccello (1999). "The Adequacy of Household Saving", *Brookings Papers on Economic Activity*, 1999:2, 65-165.
- Friedman, Milton (1957). *A Theory of the Consumption Function*. Princeton: Princeton University Press.
- Gourinchas, Pierre-Olivier and Jonathan Parker (2000). "Consumption Over the Life Cycle," *Econometrica*, 70(1): 47-89.
- Gustman, Alan and F. Thomas Juster (1996). "Income and Wealth of Older American Households," in Eric Hanushek and Nancy Maritato, Assessing Knowledge of Retirement Behavior, Washington D.C.: National Academy Press.
- Hall, Robert (1978). "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86(6): 971-987.
- Hall, Robert and Frederic Mishkin (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica* 50(2): 461-481.
- Hansen, Lars Peter and Kenneth Singleton (1983). "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy* 91(2): 249-265.
- Hurd, Michael and Susann Rohwedder (2003). "The Retirement Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement," NBER Working Paper #9586.
- Hurd, Michael and Julie Zissimopoulos (2002). "Saving for Retirement: Wage Growth and Unexpected Events", *Rand Working Paper*, July.
- Hurst, Erik, Ming-Ching Luoh and Frank Stafford (1998). "Wealth Dynamics of American Families: 1984-1994." *Brookings Papers on Economic Activity* 1998(1): 267-337.
- Laibson, David (1996). "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112(2): 443-477.
- Lawrance, Emily (1991). "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* 99(1): 55-77.
- Lusardi, Annamaria (1999). "Information, Expectations, and Savings", in *Behavior Dimensions of Retirement Economics*, Henry Aaron, eds. New York: Brookings Institution/Russel Sage Foundation, 81-115.
- Lusardi, Annamaria (2002). "Explaining Why So Many Households Do Not Save", Mimeo, Dartmouth College.
- Lusardi, Annamaria (2003a). "Planning and the Effectiveness of Retirement Saving", Mimeo, Dartmouth College.

- Annamaria Lusardi (2003b). "Financial Education and Saving", in *Pension Design and Decision-Making Under Uncertainty*. Eds, Olivia Mitchell and Steven Utkus. Pension Research Council, Philadelphia.
- Modigliani, Franco and Richard Brumberg (1954). "Utility Analysis and the Consumption Function: An Interpretation of the Cross-Section Data," in *Post-Keynesian Economics*. Eds., Kenneth Kurihara. New Brunswick, NJ. Rutgers U. Press: 388-436.
- Scholz, Karl, Anath Seshadri, and Surachai Khitatrakrun (2003). "Are Americans Saving "Optimally" for Retirement?", Mimeo, University of Wisconsin.
- Shapiro, Matthew (1984). "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data." *Economic Letters* 14(1): 93-100.
- Shea, John (1995). "Union Contracts and the Life Cycle-Permanent Income Hypothesis." *American Economic Review* 85(1): 186-200.
- Skinner, Jonathan (1987). "A Superior Measure of Consumption from the Panel Study of Income Dynamics," *Economic Letters*, 23(2): 213-16.
- Smith, James P. (1997). "Wealth Inequality Among Older Americans," *Journal of Gerontology*, 52(B), May: 74-81.
- Venti and Wise (2000). "Choice, Chance and Wealth Dispersion at Retirement," *NBER Working Paper #7521*.
- Venti, Stephen and David Wise (1998). "The Cause of Wealth Dispersion at Retirement: Choice or Chance?" *American Economic Review* 88(2): 185-191.
- Zeldes, Stephen (1989). "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2): 305-346.

Data Appendix

In this appendix, the controls included in the estimation of equation (1) are discussed. After which, the construction of the after tax interest rate and the household's consumption growth rate, included in the estimation of equations (7) and (8) are detailed.

Regression equation (1) included a vector of 1989 household controls, $X_{i,1989}$, and a vector of historical, pre-1989 controls, $Z_{i,historical}$. The specific controls included in the X vector were:

- a quadratic in the household head's age,
- a dummy equal to 1 if the household head is black,
- a dummy equal to 1 if the household head is currently married,
- the number of children in the family unit,
- the number of children between the ages of 13-17 (living within or away from the household),
- the number of children between the ages of 18 and 21 (living within or away from the household),
- a dummy equal to 1 if the head had less than a high school degree,
- a dummy equal to 1 if the head had exactly a high school degree,
- a dummy equal to 1 if the head had some college education (but no college degree) (household heads with four or more years of college was the excluded group),
- dummy variables indicating the census region where the household resides,
- a vector of one digit occupation dummies,
- a vector of one digit industry dummies,
- a dummy variable equal to 1 if the head was unemployed,
- and, a quadratic in household labor income, where household labor income is defined as the sum of labor earnings of both the head and spouse, if a spouse was present.

Additionally, the PSID asked the following question of both heads and spouses to ascertain individual health status: "Would you say your (Head's/Wife's) health in general is excellent, very good, good, fair, or poor?". If the individual responded that their health was fair or poor, the

individual was coded as having “bad health”. 1989 “bad health dummies” were included for both the head and the wife.

Lastly, two questions were asked in 1989 about household pensions. First, households were asked the following question about their pension replacement rates:

“We're interested in how much of earnings will be replaced by pensions. Thinking of your (and your (wife's/"WIFE'S")) total pension benefits when you (both) retire, including Social Security, how will they compare with your (and your (wife's/"WIFE'S")) pre-retirement earnings--I mean, about what percent of your pre-retirement earnings will they be?”

The response to this question was included as part of the X vector of controls. Secondly, equation (1) also included the amount that households directly contributed to their pension during the previous 5 years. In other specifications (not reported), occupation, industry and education controls were interacted. These interactions had no effect on the results and, as a result, were omitted from the main specification discussed in the text.

The Z vector included pre-1989 controls including:

- a dummy variable equal to 1 if the household head had any “bad health” between 1980 and 1988,
- a dummy variable equal to 1 if the spouse had any “bad health” between 1980 and 1988,
- the number of years of “bad health” reported for both the head and the spouse between 1980 and 1988,
- a dummy variable equal to 1 if the household head become divorced anytime between 1980 and 1988,
- a quadratic in average family labor income between 1980 and 1988,
- the change in family labor income between 1980 and 1988,
- the coefficient of variation of income between 1975 and 1989,
- a dummy variable indicating whether the household head was ever unemployed between 1980 and 1988,
- and the total duration of unemployment spells (in weeks) of the household head between 1980 and 1988.

For completeness, income and employment controls from the 1970s were included in many of the regression specifications. These controls added little to the predictive power of regression (1).

Including income and employment controls from the 1970s did not alter the reported results in any way.

When estimating equation (7), a measure of consumption growth and the after tax interest rate must be constructed. Consumption growth is defined as the percentage increase in annual food expenditures between year t and year $t+1$. The *PSID* annually collects information on the cost of food consumed at home, the amount spent away from home in restaurants and the value of food purchased with food stamps. Aside from the potential measurement error in the reporting of consumption data, researchers also must deal with the potential uncertainty surrounding the time period to which the respondents' answers refer. Questions dealing with the amount of non-food stamp consumption the household undertakes in an average week are asked primarily during April through August of each year, with the median interview date occurring in June. To obtain an estimate of annual food consumption flow, the *PSID* editors multiply this response by 52. A potentially important question is the time frame used by consumers in determining average consumption since that will determine the appropriate timing of the interest rate, current consumption, consumption growth and the relevant income growth. All food consumption reported during the early summer of year t is assumed to apply to all of year t .

Formally, a household's after-tax interest rate can be expressed as $r_{i,t+1} = r_{t+1}^{t\text{bill}} (1 - \tau_{i,t+1}) - \pi_{t+1}^e$, where $r_{t+1}^{t\text{bill}}$ equals the average nominal rate on a one year treasury starting in June of year t , τ_{t+1} equals household i 's marginal tax rate in year $t+1$, and π_{t+1}^e is the expected inflation rate between t and $t+1$. The *PSID*, during this sample period, reports the household's marginal tax rate computed on the basis of detailed income data and the relevant tax tables. It is assumed that households have perfect foresight with respect to their future marginal tax rates. Because inflation is unpredictable, the household in period t may not know the real rate that it will face between t and $t+1$. Using the actual inflation rate between t and $t+1$ may bias estimated coefficients because of the potential correlation with the error term. To account for this, most

researchers use the actual inflation rate when calculating the return and instrument using lagged inflation rates. Instead, I calculate the household's expected borrowing rate using the Livingston Inflation Forecast, where the forecast runs from June of year t to June of year $t+1$. The Livingston Survey of Inflation Forecasts is maintained by the Philadelphia Federal Reserve. All values were converted into 1989 dollars. This methodology for computing after-tax interest rates and consumption growth rates is similar to most authors who use the *PSID* data to estimate consumption Euler equations (see Shapiro, 1984; Zeldes, 1989; and Shea, 1995).

**Table 1: Descriptive Statistics for Two Sub Samples of Pre-Retired Households:
Bottom 20% of 1989 Wealth Residuals and Top 80% of 1989 Wealth Residuals**

	I	II	III
	<i>1st Stage</i>	<i>1st Stage</i>	<i>p-value of</i>
	<i>Residual</i>	<i>Residual</i>	<i>difference</i>
	<i>Top</i>	<i>Bottom</i>	<i>Column I and</i>
<i>Selected Income and Demographic Variables</i>	<i>80 Percent</i>	<i>20 Percent</i>	<i>II</i>
<u><i>Wealth Distribution and Portfolio Composition</i></u>			
25 th percentile of Household Wealth	\$28,900	\$2,014	<0.01
Median Household Wealth	\$83,150	\$8,275	<0.01
75 th percentile of Household Wealth	\$205,500	\$29,954	<0.01
% Owning Home	0.77	0.43	<0.01
% Owning Stocks	0.31	0.12	<0.01
% Owning Business	0.19	0.03	<0.01
<u><i>Demographics</i></u>			
Age of Head in 1989	57	57	0.30
Dummy: Marital Status in 1989	0.68	0.58	0.02
Dummy: Divorced Anytime 1980 – 1988	0.11	0.12	0.50
Dummy: Race of Head in 1989 (Black = 1)	0.10	0.13	0.35
Dummy: Education in 1989 12 years or less	0.48	0.55	0.22
Dummy: Have Children Aged 1– 5 in 1989	0.03	0.01	0.09
Dummy: Have Children Aged 6 – 13 in 1989	0.06	0.05	0.65
Dummy: Have Children Aged 14 - 20 in 1989	0.20	0.20	0.91
Number of People in Household	2.4	2.2	0.25
<u><i>Labor Income and Labor Income Variability</i></u>			
Mean Family Labor Income in 1989	\$39,947	\$39,790	0.98
Mean Family Labor Income: 1980-1989	\$37,793	\$35,260	0.49
Change in Family Labor Income: 1980 to 1989	\$4,349	\$6,949	0.62
% Heads Unemployed in 1989	0.04	0.03	0.66
% Heads Unemployed Anytime 1980-1988	0.21	0.23	0.60
Median Coefficient of Variation of Income: 1975-1989	0.51	0.53	0.21
<u><i>Health Shocks</i></u>			
% With Head ‘Bad Health’ in 1989	0.19	0.25	0.10
% With Head ‘Bad Health’ Anytime 1980-1988	0.33	0.39	0.14
% With Wife ‘Bad Health’ in 1989	0.11	0.09	0.52
% With Wife ‘Bad Health’ Anytime 1980-1988	0.22	0.24	0.68
<u><i>Retirement Pension</i></u>			
% of 1989 Income Replaced During Retirement	0.54	0.52	0.66
<u><i>Consumption</i></u>			
Mean Household Food Consumption: 1980-1987	\$4,125	\$3,901	0.25
Mean Household Food Consumption: 1975-1979	\$3,481	\$3,264	0.13
Med. Coefficient of Variation of Consumption: 1975-87	0.37	0.41	0.13

Notes: The sample was split using the residuals from a first stage regression of 1989 household log wealth on a vector of household observables. The sample included all non-retired households aged 50-65 in the 1989 wave of the PSID (819 households). All dollar values are in 1989 dollars.

Table 2
Average Age of Retirement and Consumption Decline Upon Retirement,
By 1989 Wealth Residual Groups

	I	II	III
	<i>1st Stage</i>	<i>1st Stage</i>	<i>p-value of</i>
	<i>Residual</i>	<i>Residual</i>	<i>difference</i>
	<i>Top</i>	<i>Bottom</i>	<i>Column I and</i>
	<i>80 Percent</i>	<i>20 Percent</i>	<i>II</i>
<i>Average Age of Retirement</i>	62.4	62.3	0.99
<i>Average Yearly Food Consumption Prior to Retirement</i>	\$4,045	\$3,700	0.39
<i>Average Yearly Food Consumption After Retirement^a</i>	\$3,744	\$2,913	<0.01
<i>Average Percentage Point Decline in Consumption At Retirement^b</i>	-0.03	-0.11	0.14
<i>Median Percentage Point Decline in Consumption At Retirement^c</i>	-0.11	-0.19	0.40

^a *Yearly Food Consumption Prior to Retirement* is defined as the household's food consumption averaged over the three years prior to retirement.

^b *Yearly Food Consumption After Retirement* is defined as the household's food consumption averaged over the three years after retirement.

^c *Average Percentage Point Decline in Consumption At Retirement* is computed for each household and then averaged over households.

Notes: Sample includes all non-retired households aged 50-65 in the 1989 wave of the PSID, who were subsequently retired between 1990 and 1999. The sample size included 267 high wealth residual households (column I) and 51 low wealth residual households (column II). The wealth residuals are defined as in Table 1. Retirement Age is defined as the first year that the household self-reports being retired. To compute the consumption statistics in this table, the top/bottom 5% of outliers in the change in consumption distribution were truncated. All dollar amounts reported in 1989 dollars.

Table 3: Responsiveness of Early Life Consumption to Predictable Income Changes, Segmenting By 1989 Pre-Retired Wealth Residuals

Variable	Coefficient (Standard Error)
Predictable Income Growth, All Households (β_1)	-0.16 (0.15)
Predictable Income Growth, Low Wealth Residual Households (β_2)	0.56 (0.28)

Notes: This table reports the estimates of the response of household consumption growth to predictable income changes (equation (7) in text). β_1 is the coefficient on predictable income changes for the full sample. β_2 is the coefficient on predictable income changes for households with low first stage wealth residuals. All other estimated coefficients were suppressed. Households with low first stage wealth residuals are defined to be pre-retired households between the age of 50 and 65 who had wealth residuals in the bottom 20% of the wealth residual distribution. See footnote to Table 1 for additional details. Households for which their one year consumption growth was in excess of 50% or less than -33% were excluded from the sample. The household was only excluded from the sample for that given year. These sample restrictions left 4,668 observations, based on 727 households. Standard errors (in parentheses) were adjusted for within household heterogeneity. The equation was estimated using two-stage least squared where the predictable component of household labor income growth was instrumented for using four lags of household labor income growth (excluding the first lag). The income processes was estimated separately for low wealth residual and other wealth residual households.

Table 4: Responsiveness of Early Life Consumption to Predictable Income Changes, Testing For Liquidity Constraints

Variable	Coefficient (Standard Error)
Predictable Income Increase, All Households (β_3)	-0.11 (0.25)
Predictable Income Decline, All Households (β_4)	-0.01 (0.29)
Predictable Income Increase, Low Wealth Residual Households (β_5)	0.35 (0.14)
Predictable Income Increase, Low Wealth Residual Households (β_6)	0.31 (0.17)

Notes: This table reports the estimates of the response of household consumption growth to predictable income increases and to predictable income declines (equation (8) in text). β_3 is the coefficient on predictable income increases for the full sample. β_4 is the coefficient on predictable income declines for the full sample. β_5 and β_6 are the coefficients on predictable income increases and predictable income declines for households with low first stage wealth residuals. All other coefficients were suppressed. Households with low first stage wealth residuals are defined to be pre-retired households between the age of 50 and 65 who had wealth residuals in the bottom 20% of the wealth residual distribution. See footnote to Table 1 for additional details. Households for which their one year consumption growth was in excess of 50% or less than -33% were excluded from the sample. The household was only excluded from the sample for that given year. These sample restrictions left 4,668 observations, based on 727 households. Standard errors (in parentheses) were adjusted for within household heterogeneity. The equation was estimated using two-stage least squared where the predictable component of household labor income growth was instrumented for using four lags of household labor income growth (excluding the first lag). The income processes was estimated separately for low wealth residual and other wealth residual households.

**Table 5:
Responsiveness of Early Life Consumption to Predictable Income Changes, Additive
Definitions of 1989 Low Wealth Residual**

Variable	Estimated Coefficient
Predictable Income Change, All Households	-0.26 (0.16)
Predictable Income Change, Households with 0-10 th percentile wealth residuals	0.93 (0.45)
Predictable Income Change, Households with 10-20 th percentile wealth residuals	0.34 (0.17)
Predictable Income Change, Households with 20-30 th percentile wealth residuals	0.14 (0.19)
Predictable Income Change, Households with 30-40 th percentile wealth residuals	-0.29 (0.25)
Predictable Income Change, Households with 40-50 th percentile wealth residuals	-0.18 (0.29)

Notes: This table reports the estimates of the response of household consumption growth to predictable income changes. The predictable income changes are interacted with five different dummy variables corresponding to the household's 1989 wealth residual. All five dummy variables interacted with predictable income changes were included in the regression simultaneously. The five dummy variables indicated whether the household had 1989 wealth residuals in the: 0-10th percentile, the 10th-20th percentile, the 20th-30th percentile, the 30th – 40th percentile, and the 40th – 50th percentile. All other coefficients from the regression were suppressed. Households for which their one year consumption growth was in excess of 50% or less than -33% were excluded from the sample. The household was only excluded from the sample for that given year. These sample restrictions left 4,668 observations, based on 727 households. Standard errors (in parentheses) were adjusted for within household heterogeneity. The equation was estimated using two-stage least squared where the predictable component of household labor income growth was instrumented for using four lags of household labor income growth (excluding the first lag). The income processes was estimated separately for each wealth residual group.

Table 6
Household Response to Historical Saving Questions and to Subjective Self-Assessment of Planning and Spending Behavior:
By 1989 Wealth Residual Groups

	I	II	III
	<i>1st Stage</i>	<i>1st Stage</i>	<i>p-value of</i>
	<i>Residual</i>	<i>Residual</i>	<i>difference</i>
	<i>Top</i>	<i>Bottom</i>	<i>Column I</i>
1972 and 1975 Saving and Subjective Self Assessment Questions.	<i>80 Percent</i>	<i>20 Percent</i>	<i>and II</i>
<i>1972 Variables</i>			
% of households who describe themselves as being a planner (as opposed to living day-to-day). ^a	0.59	0.46	0.06
% of households who describe themselves as likely to “carrying out plans”. ^b	0.67	0.54	0.06
% of households who describe themselves as being a spender (as opposed to a saver). ^c	0.41	0.60	<0.01
% of households in 1972 with any positive savings.	0.80	0.66	0.02
% of households in 1972 with accumulated savings greater than two months of income.	0.52	0.30	<0.01
<i>1975 Variables</i>			
% of households who describe themselves as being a planner (as opposed to living day-to-day). ^a	0.56	0.38	0.01
% of households in 1975 with any positive savings.	0.80	0.74	0.25
% of households in 1975 with accumulated savings greater than two months of income.	0.48	0.29	<0.01
Sample Size	500	127	

^a PSID question reads: “Are you the kind of person that plans his life ahead all the time or do you live more from day-to-day?”

^b PSID question reads: “When you make plans ahead, do you usually carry out things the way you expected?”

^c PSID question reads: “Would you rather spend your money and enjoy life today or save more for the future?”

Notes: All non-retired households between the age of 50 and 65 in the 1989 PSID who were in the sample during either 1972 or 1975 (627 households). Subjective self assessment questions were asked only in the 1972 and 1975 waves of the PSID. The 1972 set of questions were more extensive. As a result, the ‘carry out plans’ and the ‘saver versus spender’ questions were not asked in 1975. Household response to the question were on a 1-5 scale, with 5 being agree strongly with the statement and 1 being disagree strongly with the statement. Households are considered to respond positively to a question if they answered a 4 or a 5. Wealth residuals are defined as in Table 1.

Appendix Table A1
Average Age of Retirement and Consumption Decline Upon Retirement,
By 1984 Wealth Residual Groups

	I	II	III
	<i>1st Stage Residual Top 80 Percent</i>	<i>1st Stage Residual Bottom 20 Percent</i>	<i>p-value of difference Column I and II</i>
<i>Average Age of Retirement</i>	62.7	62.8	0.99
<i>Average Yearly Food Consumption Prior to Retirement^a</i>	\$4,238	\$4,093	0.61
<i>Average Yearly Food Consumption After Retirement^b</i>	\$3,703	\$3,122	0.02
<i>Average Percentage Point Decline in Consumption At Retirement^c</i>	-0.08	-0.16	0.07
<i>Median Percentage Point Decline in Consumption At Retirement</i>	-0.13	-0.22	0.08

^a *Yearly Food Consumption Prior to Retirement* is defined as the household's food consumption averaged over the three years prior to retirement.

^b *Yearly Food Consumption After Retirement* is defined as the household's food consumption averaged over the three years after retirement.

^c *Average Percentage Point Decline in Consumption At Retirement* is computed for each household and then averaged over households.

Notes: Sample includes all non-retired households aged 50-65 in the 1984 wave of the *PSID*, who were subsequently retired between 1985 and 1999. The sample size included 421 high wealth residual households (column I) and 91 low wealth residual households (column II). The wealth residuals are defined as in Table 1, except using 1984 *PSID* households instead of 1989 *PSID* households. Retirement Age is defined as the first year that the household self-reports being retired. To compute the consumption statistics in this table, the top/bottom 5% of outliers in the change in consumption distribution were truncated. All dollar amounts reported in 1989 dollars.