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

Erik Hurst
University of Chicago
erik.hurst@gsb.uchicago.edu

(*Preliminary Version*)

February 2003

Abstract

In this paper, it is shown that households who enter retirement with lower than 'normal' wealth do so because they had consistently followed near-sighted consumption rules during their working years. Using the Panel of Income Dynamics (*PSID*), household wealth in 1989 is predicted for a sample of 50-65 year old non-retired households using both current and past income, occupation, demographic and health characteristics. Using the residuals from this first stage regression, the sample of pre-retired households can be subsetted 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 is identical 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. It appears that retirement came as more of a surprise to these households. In the main part of the paper, I use the panel component of the *PSID* and 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. After ruling out other theories of consumption to explain these facts, it is concluded that households who entered retirement with lower than predicted wealth followed some near-sighted consumption rule of thumb early in their working lives.

* I would like to thank Mark Aguiar, Steve Davis, Anna Lusardi, Joe Lupton, Brigitte Madrian, Mel Stephens, and Paul Willen for helpful comments. Additionally, I would like to thank seminar participants at the University of Chicago's, Graduate School of Business macro lunch. I am indebted to the excellent research of Sonia Oreffice.

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

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 wealth, conditional on lifetime income, varies dramatically across households entering retirement (Gustman and Juster, 1996; Smith, 1997; Hurst, Luoh and Stafford, 1998; Venti and Wise, 1998 and 2000; Lusardi, 2001). While many authors have attempted to explain this variation in wealth across pre-retired households (Venti and Wise, 2000; Bernheim, Skinner and Weinberg, 2001; Hurd and Zissimpopoulos, 2002), the approach taken in this paper is quite different. After isolating households near retirement that have saved little given their lifetime income, employment, demographic and health trajectories, I examine the consumption behavior of these households while they were young. I find that these 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 optimizing consumption rules as an explanation for my set of findings, I conclude that those households who are most likely to under-save for retirement do so, at least in part, because they persistently follow near-sighted consumption rules during their working lives.

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 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 identical to other pre-retired households along income, health, employment, pension and demographic characteristics. However, these households experience a consumption decline during their subsequent retirement that is twice as large compared to their counterparts with similar pre-retirement income trajectories, but higher pre-retirement wealth. For households with lower than normal pre-retirement wealth, retirement appears to be a surprise.

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. Performing standard excess sensitivity tests, I am able to reject that households with lower than normal pre-retirement wealth behave as standard permanent income consumers. According to the Permanent Income Hypothesis (PIH) with perfect capital markets, 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. I find, however, 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 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, but nothing prevents a household from saving to smooth expected income declines (Zeldes, 1989; Shea, 1995). I find that the consumption of households with relatively little retirement savings

responds similarly to both predictable income increases and predictable income declines. These results indicate that liquidity constraints are not driving the failure of the permanent income hypothesis during working years for households with little pre-retirement wealth.

There are many alternate theories of consumption that can explain either the excess sensitivity of consumption to predictable income changes during a household's working years or the little pre-retirement wealth compared to life cycle income, health and demographics, or a large drop in consumption upon retirement, but very few theories can jointly explain the three facts. For example, theories of habit formation (see Deaton, 1992; Dynan, 2000) or precautionary saving (see Deaton, 1992; Carroll, 1997) may be able to explain the excess sensitivity of consumption during the household's working years but are unable to explain the large drop in consumption upon retirement. Additionally, I explicitly show that the substitutability of consumption and leisure are not driving the results. Including changes in work hours into the estimation of the consumption Euler equation for households during their working years leaves the general conclusions of the paper unchanged. Likewise, errors in expectations of the income growth process during their working years by the low wealth households cannot explain the results. Lastly, given the amount of commitment devices available to these households to save for retirement (pensions, 401K's, housing equity), such households cannot be 'sophisticated' hyperbolic consumers (see Laibson, 1997; Angeletos et. al., 2001). After performing a variety of robustness tests, I conclude that between 10 and 20 percent of the pre-retired households that I examined followed myopic rule of thumb consumption plans during their working years which left them ill-prepared for retirement.

These households who were ill-prepared for retirement were more than twice as likely to end up on public assistance during retirement. Specifically, almost 14% of the households in the bottom twenty percent of the wealth residual distribution ended up on SSI during their subsequent retirement. Finally, using information collected in the *PSID* during the early 1970s, I find that these households with low wealth entering retirement, conditioned on lifecycle factors, were self-

aware of their myopic tendencies early in their life cycle. 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. These questions are definitely noisy measures of household behavior. However, households who entered retirement in the late 1980s with lower than normal savings were, two decades earlier, 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.

In the classic fable by Aesop, a sharp distinction is drawn between ants, who saved during their summer (working years) to sustain consumption during their winter (retirement), and grasshoppers, who did not save for their future period of low earnings. The results of this paper show that a non-trivial segment of the population (about 10 to 20 percent) behaves as 'economic grasshoppers'. The consumption of such households closely tracks their income during their working years leaving them ill prepared to sustain consumption during retirement. Like Aesop, I conclude that the population is comprised of two types: "economic ants" who follow some sort of optimizing consumption rule and "economic grasshoppers" who do not plan for predictable future periods of low income. While most of the population is comprised of these economic ants, the remaining twenty percent of the population who follow myopic consumption rules cannot be ignored. For example, it is these households who would have the largest consumption response to a temporary tax cut or who would suffer the greatest welfare loss from the elimination of the social security system.

The remainder of the paper is set up as follows. In the next section, I discuss some related literature which addresses variation in pre-retirement wealth across households. In doing so, I also discuss existing research on myopic consumption behavior. In section III, I discuss how I identify households who had accumulated 'lower than normal' wealth entering retirement and how their behavior differs from other households when retirement finally occurs. Section IV sets up the main empirical test for the paper – the excess sensitivity of consumption to predictable

income changes during the pre-retired households working years. Section V presents the results and section VI explores alternate theories of consumption which may or may not reconcile the paper's results. The next to last section prevents survey evidence that these households who had saved lower than normal entering retirement were aware of their low saving tendencies nearly two decades prior to measurement of their pre-retirement wealth. The last section offers some conclusions and discussion.

II. A Review of Some Relevant Literature

It is well documented that there is very large dispersion in the accumulated wealth of families approaching retirement, conditional on total lifetime earnings (Gustman and Juster, 1996; Smith, 1997; Venti and Wise, 1998 and 2000; Hurst, Luoh and Stafford, 1998; Lusardi, 1999). By examining total lifetime earnings, shocks to health and earnings, and differences in portfolio allocation, Venti and Wise (2000) conclude that "the bulk of the dispersion (in pre-retirement wealth across households) must be attributed to differences in the amount that households choose to save". There are several papers which analyze the adequacy of saving by comparing simulated optimal saving behavior from a calibrated lifecycle model to actual household data (Bernheim, 1992 and 1997; Bernheim and Scholz, 1993; Gale, 1998; and Warshawsky and Ameriks, 1998).[1] The conclusion of almost all these papers is that a large fraction of households have saving levels entering retirement that will leave them unprepared to sustain consumption during retirement. Even Engen, Gale, and Uccello (1999), who 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, conclude that somewhere between 5 and 25 percent of the pre-retired population are under-saving for retirement.

[1] For other papers assessing the adequacy of savings across pre-retired households see Moore and Mitchell (1997), Gustman and Steinmeier (1998), and Hurd and Zissimopoulos (2002).

Many papers have documented that consumption for the average household does decline during retirement (Hamermesh, 1984; Banks, Blundell and Tanner, 1999). The authors of these papers suggest that their findings are consistent with the households' realization that they had not saved enough for retirement. By examining the relationship between accumulated wealth and the shape of the consumption profile, Bernheim, Skinner, and Weinberg (2001) conclude that lifecycle models of saving cannot explain all the variation in wealth across households. For example, they find no relation between accumulated wealth prior to retirement and consumption growth rates either prior to or after retirement. Such results imply that differences in time discount rates across households cannot explain the observed variations in accumulated wealth across pre-retired households. Similarly, they rule out differences in risk tolerance, exposure to uncertainty, relative tastes for work and leisure at advanced ages, and income replacement rates as potential explanations for the differences in saving across pre-retired households, conditional on lifetime income and demographics.

The approach taken in this paper differs from that of Bernheim, Skinner and Weinberg in many ways. Most importantly, while Bernheim et. al. conclude that the evidence that they present is an indictment of the Permanent Income Hypothesis, they do attempt to test for differences in consumption rules across different types of households. Engen, Gale, and Uccello (1999) point out that "The results in the Bernheim, Skinner, and Weinberg paper actually show that most households experience an increase in consumption after retirement, again after controlling for unanticipated events that may cause retirement and affect life income simultaneously, and for other factors". The approach in this paper is to test for differences in consumption rules among pre-retired households. I find that for most pre-retired households, the Permanent Income Hypothesis cannot be rejected. However, 10-20% of the population does appear to follow non-optimizing, myopic consumption rules which lead to very little savings as they enter retirement. I conclude that one theory cannot explain all the variation in pre-

retirement saving across households, although the standard lifecycle model does well for most households.

The success of the Permanent Income Hypothesis more generally is mixed.[2] Using micro data, some authors find that consumption growth responds to predictable income growth in a representative sample of households during their non-retired working years (for example, see Hall and Mishkin, 1982). However, this result is not universal (for example, see Attanasio and Browning, 1995). Using micro data from the Panel Study of Income Dynamics, Hall and Mishkin (1982) conclude that 1/5 of all consumption is just set to a fraction of current income instead of following the more complicated optimal consumption rule. One could interpret the results of Hall and Miskin as saying that 1 out of 5 households follow some simple Keynesian rule of thumb. Using aggregate data, Campbell and Mankiw (1989) conclude that half of all consumers follow simple consumption rules of thumb. Many authors have criticized the findings of Hall and Miskin and Campbell and Makiw because they fail to account for capital market imperfections (see Zeldes, 1989; Deaton, 1992). In this paper, I use micro data to explicitly isolate a group of households who appear to be rule of thumb in their consumption decisions. Additionally, I directly show that liquidity constraints are not driving the results. A major conclusion of this paper is that no one theory of consumption represents all households. The data suggests that the majority of households appear to follow permanent income consumption rules. However, there is a non-trivial minority of households who appear to be at least partially myopic with regard to their consumption decisions.

III. Segmenting Pre-Retired Households By Wealth Residuals

The purpose of this section is to isolate pre-retired households with lower than 'normal' wealth. For the tests that follow, I want to isolate households who, upon retirement, had similar

[2] There are many authors who have tested the Permanent Income Hypothesis using micro data sets. See Browning and Lusardi (1996) and the cites within for references.

opportunities to save over their lifetime. As discussed above, there is a large literature which documents that some households save little given their income and demographic trajectories. Using the 1989 Panel Study of Income Dynamics (*PSID*), where there are multiple decades of past income and demographic data for given pre-retired households, I can segment pre-retired households by whether they have higher or lower wealth than other households who experienced similar economic histories. The empirical approach in this part of the paper is to estimate a reduced form equation of household wealth on multiple controls for the household's past, current and expected future income, employment, pension status, health status, and demographics. The residuals from such a regression will be used to isolate households who save less than predicted by their observables.

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. The date of analysis 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 between 1984 and 1999. Consequently, cross-sectional studies of wealth using *PSID* data are limited to the years of 1984, 1989, 1994 and 1999, where the 1999 data is the most recent *PSID* data made available. Given the nature of the tests that I perform below, I need to follow the differing groups of 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, I chose 1989 as the year in which I segment pre-retired households by their conditional wealth levels.[3] Additionally, as discussed below, the pension questions used in the wealth prediction equation were asked only in 1989. Given the sample design of the *PSID*, nearly all of

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

the 1989 households where the head was between the age of 50 and 65 had at least one family member in the *PSID* since the survey's inception in 1968. As a result, there are almost twenty previous years of data on income, employment and demographics for each pre-retired household in 1989.

When estimating the wealth regressions, the sample was restricted to non-retired households who were between the ages of 50 and 65 in 1989. Additionally, the sample was restricted to only those households who had positive *PSID* measured wealth in 1989. Given that these households were well into their lifecycle, the requirement of positive wealth holdings was not overly restrictive. Less than four percent 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 all 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 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). In terms of data quality, the *PSID* matches up very well against similarly defined wealth measures from the Survey of Consumer Finances throughout most of the wealth distribution (Juster and Stafford, 1999; Mathematic Draft Report, 2002).

The *PSID* wealth supplements, in general, do have one major drawback when used to assess retirement savings. Up through the late 1990s, the *PSID* did not ask explicitly ask households questions concerning their private pension wealth or about expected social security retirement

benefits.[4] In 1989, however, this problem can be partially overcome. The 1989 *PSID* households were asked questions about their expectations of the percentage of their pre-retired yearly labor earnings that would be replaced by all their pension plans (including social security) during retirement. Using this information, the extent to which the measured *PSID* retirement wealth is underestimated can be assessed.

When predicting wealth for the sample of pre-retired *PSID* households, controls suggested by economic theory are included. Both the level and the lifecycle trajectory of earnings should predict household wealth holdings near retirement. To control for differences in income across households, I include the following controls: a quadratic in 1988 family labor income, a quadratic in past average family labor income between 1970 and 1987, a quadratic in recent average family labor income between 1980 and 1987, the change in family labor income between 1980 and 1988, and a dummy for whether the household had zero family labor income in 1988.[5] Additionally, a series of controls were included to proxy for whether the households received recent unexpected income shocks between 1980 and 1988. These variables include: a dummy variable equal to 1 if the head was unemployed during the 1989 survey date, a dummy variable equal to 1 if the wife was unemployed during the 1989 survey date, a dummy variable equal to 1 if the head was unemployed anytime between 1980 and 1988, a dummy variable equal to 1 if the wife was unemployed anytime between 1980 and 1988, and, separately, the number and duration of spells of unemployment experienced by the head and the wife between 1980 and 1988. In some specifications, I included information on the past unemployment shocks experienced by the households in the 1970s. These variables provided very little additional explanatory power. The results in this paper were in no way changed if household unemployment spells from the 1970s were also used in the 1989 wealth prediction equation.

[4] Starting in 2001, the *PSID* will start collecting pension data as part of its wealth supplements.

[5] 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.

Past unexpected demographic or health shocks could also affect household wealth. To account for these factors, the following additional controls are included: the level and changes in both the head's and the wife's health between 1980 and 1989, 1989 marital status, past shocks to marital status between 1980 and 1989, family size in 1989, changes in family size between 1980 and 1989, and the number of children born to the household. Additionally, race and education controls, as well as a series of occupation, industry and region dummies, were included to the regression.

Finally, two measures of pension wealth were also included in the regression. As noted above, in the 1989 survey, the *PSID* asked its respondents about their expectations of how much of their pre-retirement earnings would be replaced by their pension plans (including social security) during retirement.[6] Even though their have been many studies showing that households are unable to accurately account the details of their pension plans, their expectations about these plans is what the PIH model predicts should determine pre-retirement savings. Second, as part of the employment component of the survey, households were asked whether they were covered by a pension or retirement plan at their pace of work, whether they have contributed directly to this plan by having money deducted from their pay, and, on average, about what amount of their pay have they contributed to this plan over the last five years. Separate controls were included in the first stage regression to account for the answers to these questions. It should be noted that by controlling for may variables that are a function of a households time preference or planning abilities (i.e., education, occupation, slope of their income profile), I am biasing my groups towards being similar along these dimensions. The fact that households wealth differs in addition to these controls suggest that I am isolate groups of households who have different saving

[6] The *PSID* question concerning pension replacement rates is as follow: "We're interested in how much of earnings will be replaced by pensions. Thinking of your (and your wife's) total pension benefits when you (both) retire, including Social Security, how with they compare with your (and your wife's) pre-retirement earnings – I mean, about what percent of your pre-retirement earnings will they be?"

propensities above and beyond the extent that these saving propensities are correlated with observables.

The residuals from the cross sectional regression of log household wealth on the controls discussed above provide a measure of whether the household has saved more or less than otherwise similar households. The adjusted R-squared from this regression 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 their saving residuals is in some sense arbitrary. To begin, I segment households with residuals in the lowest twenty percent of the residual distribution as being households who save much less than predicted by their observables. 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). I initially chose the 20^{th} percentile cutoff given 1) Hall and Miskin (1982), discussed above, find that about twenty percent of the population appears to be rule of thumb, and 2) Engen, Gale, and Uccello (1999), also discussed above, find that about 20% of households dramatically under save for retirement. However, in the sections that follow, I explore the robustness of my results when the cutoff is redefined as the 10^{th}, 30^{th}, 40^{th} or 50^{th} 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 10^{th}, 20^{th}, 30^{th}, 40^{th} and the 50^{th} percentiles of the wealth residual distributions are, respectively, -1.32, -0.73, -0.36, -0.12, and 0.12. As we will see in the following sections, the twenty percent cutoff is well justified.

B. Comparison of the Two Pre-Retirement Wealth Residual Sub-Samples

Given the way that the samples were split, the two samples should be identical along all observables included in the regression. This procedure assures that, to the extent that economic circumstances can be measured in survey data, the two samples of pre-retired samples had similar

opportunities to save. 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. Column I reports the mean of variables for the pre-retired households who have wealth residuals in the top 80 percent of the wealth distribution, Column II reports the mean of variables for the pre-retired households who have wealth residuals in the bottom 20 percent of the wealth distribution, and Column III reports the *p*-value of a *t*-test on the difference between the means for the two samples. Households with wealth residuals in the top 80 percent of the wealth distribution have over ten times as much median wealth as their counterparts in the bottom of the wealth distribution ($8,300 vs. $83,000).[7] The magnitude of this difference persisted throughout the distributions. Furthermore, there are large differences in portfolio compositions between the two groups of households. For the high wealth residual households, almost 80% own a home, 30% have direct ownership of stocks, and nearly 20% own their own business. The comparable numbers for the low wealth residual households are 43%, 12% and 3%.

Given the methodology used to get the wealth residuals, it is not surprising that the mean household in each group looks identical along all observables. The average income for both groups of households earned during the 1980s was about $36,000. Despite the similarity in earned income, the wealth differences are striking. Going back to the 1970s, the level of earned income is also very similar between the two groups ($15,000 for the low residual pre-retired households vs. $17,000 for all other pre-retired households). Health shocks, income shocks, and all demographics – aside from current marital status – are nearly identical between the two groups of households

The only measure of consumption, aside from housing expenditures, that the *PSID* directly asks its respondents about is their food consumption.[8] Specifically, in all years between 1970 and

[7] Unless otherwise specified, all dollar values reported in the paper are in 1989 dollars.
[8] 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

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. As seen in Table 1, the food consumption between the two types of households during the 1980s and during the mid to late 1970s was nearly identical. Tautologically, if income levels are similar between two groups and one group saves less than the other, the consumption of that group should be higher. However, given that food consumption represents less than 14% of the median households consumption bundle (Skinner, 1987) and given that the yearly saving difference between the two groups implied by the pre-retirement wealth difference is small, we would expect the actual difference in yearly food consumption between the two groups to be small. Therefore, the means reported for the food consumption numbers during the 1970s and the 1980s for the two groups are very much consistent with their reported income and wealth statistics.

Table 1 illustrates that the low wealth residual households and the other wealth residual households are nearly identical along all observables other than wealth and portfolio composition. These households had similar demographic, health and income trajectories, yet one group entered retirement with large amounts of wealth and the other group entered with little wealth. I cannot rule out that the low wealth residual group consumed slightly more food than the other pre-retired group. Additionally, the household could have spent the majority of the added consumption on non-food items.

C. Differences in Subsequent Retirement Behavior By Wealth Residual Groups

It is quite possible that even though the two households appear to have different wealth levels entering retirement, their subsequent retirement behavior could be quite similar. This could happen if the low wealth group planned on receiving and had actually received a wealth windfall upon retirement (i.e., expecting a large bequest). In this section, we compare the

---

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 measure of *PSID* consumption.

retirement behavior of the two groups to see if the low wealth households actually experienced a greater consumption decline upon retirement. If that is the case, it suggests that these low wealth residual households – like the proverbial grasshopper - were ill prepared to sustain consumption during retirement.

There are potentially two interrelated drawbacks to comparing the subsequent retirement behavior of the grasshoppers (low wealth residuals) and the ants (all other wealth residuals). 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 early 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. However, there is no information to suggest that the households in the two pre-retirement wealth residual groups, on average, 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. Furthermore, accounting for death prior to retirement, similar percentages of both groups were observed actually retiring before 1999 (43% of ants and 38% of grasshoppers).

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. As before, the measure of consumption is total food consumption. Panel B of Table 2 shows level of consumption averaged over the three years preceding retirement, the level of consumption averaged over the three years after retiring, and the mean and median percentage decline in consumption during retirement. 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 decline over all households.

Panel B of Table 2 shows that the mean consumption prior to retirement for both grasshoppers and ants was quite similar. 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 retired showed a similar pattern: low saving pre-retired households experienced nearly a 20% consumption decline compared to a 11% decline for the other households. The fact that the average household experiences a consumption decline during retirement is consistent with almost all existing empirical work (see, Bernheim, Skinner and Weinberg, 2001). The interesting result from Table 2 is that, as would be predicted, the low wealth residual households experienced a much more severe consumption decline during retirement.

The results in Table 2 are based on very few grasshopper 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 because there are so few low wealth residual households actually retiring or is it because there is no actual difference in behavior between the two groups. To explore the robustness of these results, I performed exactly the same analysis using 1984 as the year in which I segmented the households by pre-retirement wealth residuals.[9] Doing so allowed me to observe many more households who actually retired prior to 1999. 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. These average and median declines are statistically different from each other at standard levels of significance. From this, I conclude that there were large differences in consumption declines between the two groups.

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

Additionally, Table 3 provides evidence that the low wealth residual households were more likely to take up Supplemental Security Income (SSI) once retired. SSI is a government assistance program that targets, among others, households over the age of 65 with limited assets and limited income. Column I shows the percentage of households with pre-retirement wealth residuals in the top 80% of the residual distribution, Column II shows the low residual households and, as above, Column III shows the *p*-value of the difference in means. Households in the low wealth residual group, who subsequently retired, were statistically more likely to end up on welfare in either 1997 or 1999 (13.7% vs. 6.4%).[10] So, not only were the low wealth residual households more likely to experience a consumption decline in retirement, they were also more likely to receive government assistance.

IV. Identifying Differences in Consumption Rules Across Wealth Residual Groups

One interpretation of the above results is that retirement comes as a 'surprise' to households who have accumulated little pre-retirement wealth, adjusted for lifecycle characteristics. Like the proverbial grasshopper, such households consumed their earnings during their working years leaving them ill-prepared to sustain consumption during retirement. It should be acknowledged that there are many other theories that could explain the decline in consumption at retirement (i.e., substitutability between consumption and leisure; reduction in expenses associated with work; and more time to shop and seek consumption bargains). However, in this section, we formally explore whether the drop in consumption at retirement is due to the inability of the low residual households to smooth expected income shocks. In particular, I test whether these low wealth residual households consistently did not smooth expected income changes throughout their working years. In the subsequent sections, I will explore the alternate theories as to why consumption could decline so sharply for these low wealth residual households.

[10] The *PSID* went to biannual interviewing starting in 1997. As a result, there is no 1998 survey.

The dramatic decline in consumption during retirement for these low wealth residual households seems at odds with the Permanent Income Hypothesis (PIH). According to the PIH with perfect capital markets, 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. Specifically, a household who has a utility function where consumption and leisure are separable and who expects their income to fall during retirement should have saved during their working years so that their discounted marginal utility of consumption is equated through the retirement period.

In this section, I explore whether the two groups of pre-retired households discussed above follow standard permanent income consumption rules during their working lives. Below, the standard Permanent Income Hypothesis (PIH) model of household consumption and saving decisions is outlined (Modigliani and Brumberg, 1954; Friedman, 1957). The model will be used to specify the empirical tests of whether the consumption behavior between the two pre-retired groups discussed in the previous section differs during their working years.

A. Empirically Testing for Permanent Income Consumption Behavior

In section III, we identified two potentially different types of consumers in the population – those that save little relative to others given their lifecycle characteristics and all other households. As in Zeldes (1989), assume that households in each of the two groups solve the following maximization problem:

$$\max_{C_{ikt}} \quad u(C_{ikt}, \Theta_{ikt}) + E_t \sum_{s=t+1}^{T} \frac{1}{1+\delta_k}^{(s-t)} u(C_{iks}, \Theta_{iks}) \tag{1}$$

s.t. $$X_{ik,t+1} = (1 + r_{ik,t+1})(X_{ikt} - C_{ikt}) + Y_{ik,t+1}$$

$$Y_{ikt} = P_{ikt} V_{ikt}$$

$$P_{ikt} = g_{ik} P_{ik,t-1} N_{ikt}$$

$$u(C_{ikt},\Theta_{ikt}) = \frac{C_{ikt}^{1-\rho_k}}{1-\rho_k} \exp(\Theta_{ikt}),\ \rho_k > 1;$$

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, which could differ across different types of households. δ and ρ are assumed to be constant across all households of a given type, but the parameters can differ across types. 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)):[11]

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

[11] The solution to this model is discussed in Appendix 1. Included in the Appendix is a discussion of estimating the consumption Euler equation when consumption is measured with error.

where $\dot{Z}_{ik,t+1} = \Delta \ln Z_{ik,t+1}$, for any variable Z, $\varepsilon_{ik,t+1}$ 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:

$$\Theta_{ikt} = b_{0k} age_{ikt} + b_{1k} age_{ikt}^2 + b_{2k} \ln(famsize_{ikt}) + \eta_k + \tau_{ik} + \mu_t + \xi_{ik,t+1} \qquad (3)$$

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 fixed effect which is constant over time across types (η_k), a family fixed component which is constant over time for any family within a type (τ_{ik}), 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 three (ξ_{ikt}).[12, 13] Substituting (3) into (2), one gets:

$$\dot{C}_{ik,t+1} = \lambda_{0k} + \lambda_{1k} \ln(1 + r_{ik,t+1}) + \lambda_{2k} \dot{famsize}_{ik,t+1} + \lambda_{3k} age_{ikt} + \mu_{t+1} - \mu_t + \varepsilon^*_{ik,t+1} \qquad (4)$$

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 $\gamma_k(\delta_k - \omega_{ik}^2/2 + b_{0k} + b_{1k})$, where γ_k is the intertemporal elasticity of substitution ($1/\rho_k$). The coefficient λ_{1k} in (4) is equal to γ_k.[14]

As outlined in the previous section, (4) 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 (4) to differ accordingly between the two groups of households:

---

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

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

[14] Given that changes in family size are planned in advance, $\Delta ln\, famsize_{it}$ 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.

$$\dot{C}_{ik,t+1} = \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi D_{Year} + \varepsilon^*_{ik,t+1} \quad (5)$$

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 by type. $\boldsymbol{D}_{Year}$ 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:

$$\dot{C}_{ik,t+1} = \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi D_{Year} + \beta_1 \dot{Y}^{Predict}_{ik,t+1} + \beta_2 D_{<20} \dot{Y}^{Predict}_{ik,t+1} + \varepsilon^*_{ik,t+1} \quad (6)$$

where $\dot{Y}^{Predict}_{ik,t+1}$ is the predictable component of income growth rate between t and $t+1$ estimated simultaneously with (6). If households are not sufficiently 'impatient', the Permanent Income 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.

[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.

In order for OLS to yield an unbiased coefficient on 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. By definition, these lagged variables are orthogonal to the error term, $\varepsilon^*_{i,t+1}$. According to the permanent income hypothesis with perfect capital markets, no patient household should have a consumption response to a predictable change in income.

One concern, however, with this empirical approach is that the dummy indicating whether a household has lower than normal pre-retirement wealth. The sample splitting on wealth is done in 1989. The consumption Euler equations were estimated prior to 1989. As a result, there are no guarantees that the wealth residuals are orthogonal to $\varepsilon^*_{i,t+1}$ in (6). Despite this, estimating (6) with OLS will still yield unbiased estimates of β_1 and β_2 as long as the predictable changes in income are identified for both groups. The concern, then, is whether the instruments used to identify predictable income changes actually yields the predictable component of income changes for both groups.

I allow the predictable income growth process, as a function of lagged income growth, to differ for each group. The benefit of this is that I do not impose that the income generating process is the same for each group. If the income generating process is different across, my procedure will lead to unbiased predictable income growth components. However, my method has one pronounced drawback. It is possible that households with low wealth residuals perpetually received poor income draws throughout their lifetime. These households were perpetually unlucky. In each period, they expected high income growth and, as a result, saved little. Once the subsequent period came, they realized bad income growth, adjusted their consumption slightly downward, and expected high income growth in the following period. Such households who were persistently unlucky could end up in retirement with little wealth. The

lagged income process I observed for these households would not represent the household's expectation about their future income. In each period, they expected higher future income growth and realized lower income growth. Their expectation, on average, and their realizations, on average, were not equal.

There are four ways in which I defend myself against this criticism. First, by definition, these households are identical along all observables including actual income shocks (unemployment spells) and expected income shocks (education - occupation - industry interactions). Given the way that I split the sample, there is no reason to believe that the expected income processes differ between the two groups of pre-retired households.

Additionally, if this 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. In other words, these households will be consistently revising their consumption downward with each 'unlucky' income draw. If these households are truly unlucky, their consumption growth will be lower than the other group of pre-retired households. If this underestimate is constant over time, the differences in expectation will be captured in the α_1 coefficient (the dummy indicating the low residual group).

Even if the downward bias is not constant over time, we can test for whether households are persistently unlucky. For the low residual households, their lagged income process is not a valid representation of their expectations. In this case, we may observe the household responding to predictable income shocks even when the household is an optimizing permanent income household. For example, given the lagged income structure, suppose a household's income growth is predicted to be low. This low wealth household, by the definition of being unlucky, expected their income to be higher. As a result, the low realization of income would actually cause the household to adjust their consumption downward. A low realization in our predictable income measure would be associated with a fall in consumption. We would expect to see a positive relationship between income declines and consumption declines. However, by the same

logic, if we predict the household to have an income increase, this household would have expected an even higher income increase (the prediction equation would persistently underestimate the expected income growth). Likewise, we should see the household's consumption fall with predicted income increases (a negative relationship between predicted income increases and consumption increases). In other words, consumption growth will be persistently biased downward. We can test this response directly.

Finally, the most intellectually appealing robustness test I perform with regards to this issue is to assume that the income generating process is the same for both households. This income generating process is based off the full sample of pre-retired households. Basically, on average, what do lags about income growth predict for future income growth. In this case, no distinction is being made across household types. I am assuming all households – both the lucky and the unlucky ones –face the same expected income path. As I discuss below, there is no difference in results between when I impose separate or the same income processes on the two groups of households. There is no evidence whatsoever to indicate that the low wealth residual households were unlucky. Given this, OLS estimation of (6) will yield unbiased estimates of β_1 and β_2.

B. Empirically Testing for the Existence of Binding Liquidity Constraints

The above test does not isolate whether the correlation between changes in consumption and predictable changes in income result from the existence of 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. 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:

$$\begin{aligned}\dot{C}_{ik,t+1} = {} & \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} \\ & + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi D_{Year} + \beta_3 D_{Up} \dot{Y}^{\text{Predict}}_{ik,t+1} \\ & + \beta_4 D_{Down} \dot{Y}^{\text{Predict}}_{ik,t+1} + \beta_5 D_{<20} D_{Up} \dot{Y}^{\text{Predict}}_{ik,t+1} + \beta_6 D_{<20} D_{Down} \dot{Y}^{\text{Predict}}_{ik,t+1} + \varepsilon^{*}_{ik,t+1}\end{aligned} \tag{7}$$

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.

V. Euler Equation Estimation

In this section, I estimate equations (6) and (7) on data from 1975–1987 for the sample of *PSID* households who were 'pre-retired' in 1989. As discussed above, these households were between the ages of 50 and 65 in 1989, had positive 1989 wealth and were not retired in 1989. The majority of these pre-retired households in the 1989 *PSID* were respondents in the *PSID* since its inception in 1968. I start the analysis in 1975 because in calculating predictable income growth (discussed formally below), I use four lags of income growth as instruments. The estimation of the consumption Euler equations end in 1987 because the *PSID* did not ask consumption questions in 1988 and 1989.

A. Data

The *PSID* reports the age of the household head and the family size of the household in each year of the supplement. 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. 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 (5) and (6) will be biased against finding significant coefficients on β_1 and β_2.

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. I choose to date all food consumption reported during the early summer of year *t* as applying to year *t*. This is similar to the choices made by many authors who use *PSID* food consumption to estimate consumption Euler equations (for example, see Zeldes 1989).

Formally, a household's after-tax interest rate can be expressed as $r_{i,t+1} = r_{t+1}^{tbill} (1 - \tau_{i,t+1}) - \pi^{e}_{t+1}$, where r_{t+1}^{tbill} 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^{e}_{t+1} 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. I assume 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 $t+1$.[16] All values were converted into 1989 dollars. My 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). The main results of the paper were unchanged if I computed the forecastable component of inflation using lagged inflation rates.

As indicated in earlier sections, household consumption response to predictable changes in household labor income can shed light on whether the household is a permanent income consumer, liquidity constrained or 'rule of thumb' household. Unlike Shea (1995), I do not have a direct measure of the household's expected total labor income changes. However, 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. Therefore, when estimating (6) and (7), I instrument for the predictable component of household labor income growth using 4 lags of household labor income growth. Given my methodology, 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% of households (F-statistic = 10.7, p-value <

[16] The Livingston Survey of Inflation Forecasts is maintained by the Philadelphia Federal Reserve. I also used 1) the Survey of Professional Forecaster's inflation forecast and 2) the actual CPI inflation rate in place of the Livingston forecast. The results were nearly identical.

0.01) and for households who have 1989 wealth residuals in the bottom 20% of households (F-statistic = 4.2, p-value < 0.01). Given the sample selection procedure, I estimated (6) and (7) restricting the income process of the two groups to be the same. This should account for differences in bad luck between the low wealth residual sample and the sample of all other households. Below, I discuss how the results between these two methodologies were nearly identical.

B. Base Results

Panel A of Table 4 shows the results from estimating equation (6) over the time period 1975-1987 for the sample of 1989 pre-retired households. If both the low residual and high residual groups followed 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).

The results were unchanged if the income process of the two groups were restricted to be the same (i.e., use the same instruments to get predictable income for both groups). The estimate for β_1 and β_2 (not shown in a Table) were, respectively, -0.18 and 0.49. However, the standard errors, as expected, were larger.

One explanation for the positive coefficient on predicted income changes for the low wealth residual group is that this group faced binding liquidity constraints. As a result, the households could not borrow to smooth their marginal utility of consumption over time. As was discussed earlier, it is possible to test whether the failure of the permanent income hypothesis is due to the existence of binding liquidity constraints. Panel B of Table 4 estimates equation (7). 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. The second panel of Table 4 shows that the low residual households respond equally to predictable income increases and to predictable income declines.

It should be noted that there were no statistical differences in the estimated preference parameters between the two groups (i.e., α_0 through α_5). For example, the coefficient on α_0 was 0.10 (standard error = 0.04) while the coefficient on α_1 was 0.03 (standard error = 0.03). Like other studies using micro data, the estimates of the intertemporal elasticity of substitution for both groups were small and noisily estimated.

From Table 4 we can conclude that households who enter retirement with extremely low wealth relative to their lifetime income, employment, demographic and health trajectories appear not to follow permanent income consumption rules during their working years. The households treat predictable income changes as a surprise. As their predictable income increases, their consumption increases. As their predictable income declines, their consumption declines. Their retirement behavior is consistent with such consumption decisions during their working years. By not planning for the eventuality of retirement, these households do not accumulate enough wealth to sustain their previous levels of consumption throughout the remainder of their life. Drawing on Aesop's fable, such households behave as 'economic grasshoppers' – consuming most of their income during the spring and summer of their lives only to suffer in their winter.

C. Alternate Specifications of Low Wealth Residual Households

The results above results 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 we may be interested in is how robust are the findings in Table 3 to alternate definitions of 'lower than normal' pre-retirement wealth. Tables 5 and 6 explore this question. In Table 4, I define the cutoff for the low wealth residuals, respectively, as the bottom 10%, 20%, 30%, 40% and 50% of the wealth distribution. The row reporting the results for the bottom 20% is the same as the results reported in Table 4.

Panel A of Table 5 shows that as the cutoff increases, the likelihood that the low residual households respond to predictable income changes declines. This is consistent with those households at the bottom of the wealth distribution being more likely to follow consumption rules other than the standard PIH rule. The response to predictable income changes is strongest among households with first-stage wealth residuals in the bottom 10% of the pre-retirement wealth distribution. For households in the bottom 10% of the wealth distribution, the additional marginal propensity to consume out of predictable income shocks is 0.97 (t-statistic = 1.8). Households in the bottom 10%, 20% or 30% of the pre-retirement wealth residuals, on average, respond to predictable income changes. Like Panel B of Table 4, Panel B of Table 5 shows that this behavior is not driven by liquidity constraints. Households in the bottom 10% and 20% of the wealth distribution respond both to predictable income increases and predictable income declines.

Table 6 takes a different approach to isolating which households respond to predictable income changes. In Table 5, we are unsure if the strong responsiveness of the bottom 10% of residuals is driving the significance of the results when the cutoff is defined as the bottom 20% or 30% pre-retirement wealth residuals. This is possible given that those households in the bottom

10% of the wealth residual distribution are also, by definition, in the bottom 20% of residuals. Suppose that households in the bottom 10% of the wealth residual distribution have a marginal propensity to consume out of predictable income changes of 1.00. Suppose further that those in the 10^{th} – 20^{th} percentile of the residual distribution have a marginal propensity to consume out of predictable income changes of 0.00. When looking at those in the bottom 20% of the wealth distribution, the marginal propensity to consume out of predictable would be roughly 0.50. All of the responsiveness, however, to predictable income changes would be driven by those households in the bottom 10% of the pre-retired wealth residual distribution.

Instead of defining the low pre-retirement wealth residual cutoff as an absolute cutoff, I allowed the responsiveness to predictable income changes to differ by different residual deciles. Specifically, I separately isolated the behavior of those with pre-retirement residuals (estimated from the first stage regression discussed in section III) in the following ranges: 0-10%, 10%-20%, 20%-30%, 30%-40%, and 40%-50%. Doing this allowed me to ascertain whether both households in the 0-10^{th} and the 10^{th}-20^{th} percentile ranges of the wealth residual distribution responded to predictable income changes. Table 5 reports the results of such a regression. While households with wealth residuals in the bottom 10% of the distribution respond the strongest to predictable income changes, households with residuals in the 10^{th}-20^{th} percentile also respond positively to predictable income changes. Both net effects are statistically different from zero. Both of these groups also statistically respond to both predictable income increases and income declines (results not shown). No other group responds to predictable income changes. 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.

VI. Alternate Explanations For Consumption Responses of Low Saving Households

The above analysis has provided four facts about the variation in wealth of pre-retired households. First, conditional on many controls for lifetime income, employment, health, pension status and demographics, some households have very little wealth compared to other households with similar lifecycle characteristics. Second, these same low wealth households suffer much larger consumption declines upon subsequent retirement compared to the other pre-retired households. Third, the consumption growth of these households responds strongly to predictable changes in income – a direct violation of the standard permanent income hypothesis with perfect capital markets. Fourth, and finally, the consumption profile of these households respond equally to both predictable income increases and predictable income declines. In this section, I discuss whether existing theories of consumption can jointly reconcile all of these findings.

The results above 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 households working years, neither would predict such a sharp decline in consumption upon retirement. For example, standard consumption habit models (Deaton, 1992; Dynan, 2000) would predict that consumption would decline less slowly during retirement. 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, but such households would still smooth their consumption across the period of retirement (Gourichas and Parker, 2000). Additionally, any differences in precautionary motives between low wealth residual households and all other households must come from differences in time preferences. Both of these groups, given the regression used to predict pre-retirement wealth, had similar levels of income growth and income volatility during their lifetime (see Table 1).

Likewise, it is possible that consumers with time inconsistent preferences could have consumption growth that responds to both predictable income increases and predictable income declines (Laibson, 1997). Such households have relatively high discount over short horizons and a relatively low discount rate 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 they choose to postpone 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.

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. Specifically, such households will tend to shift their savings toward illiquid assets (housing, business or pension). The low pre-retirement wealth residual households documented above have behavior that is inconsistent with sophisticated hyperbolic consumers. A majority of such households do not own any 'illiquid' assets (defined as the sum of housing, other real estate and business equity). Specifically, 44% of the low residual group owns no illiquid assets at all. The comparable number for the other pre-retired households was less than 10%. 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 certainly are not taking actions to commit themselves to save for the future.

I cannot rule out that the consumption behavior documented above is inconsistent with the households being "naïve hyperbolic consumers". While naïve hyperbolic consumers are modeled as optimizing utility at every moment in time, their extremely high current time preference rate makes them act equivalently to a myopic rule of thumb consumer. Their consumption decisions

today are made with little weight placed on their consumption in future states of the world. Whether we call a household a naïve hyperbolic consumer or a myopic rule of thumb consumer is semantics. The two types of households (naïve hyperbolic or rule of thumb consumer) would have similar consumption responses to a temporary tax cut or to a change in social security benefits and, as a result, both would suffer large welfare losses from their consumption rules.

A strong substitutability between leisure and consumption could reconcile some of the results outlined above. Such households would optimally plan for lower consumption in retirement and, as a result, would save less during their working years. 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. Households who know they will be working less in the future would still optimally to choose higher consumption today and lower consumption while not working. However, their total utility could be equated across the periods. When leisure is low, consumption would be high and when leisure is high, consumption would be low (Attanasio and Browning, 1995). To rule out that the substitutability between leisure and consumption is causing the failure of my permanent income tests, I included the change in total family hours worked directly into my estimation of regressions (6) and (7). This is the same procedure used by Attanasio and Browning, 1995. The results reported in Tables 4-6 were essentially unchanged. The substitutability between consumption and leisure is not causing the excess sensitivity tests documented above.

VII. The Self-Awareness of Economic Grasshoppers

In the early 1970s, the Panel Study of Income Dynamics asked all its respondents to self assess many of their socio-economic characteristics. Such questions were asked in both 1972 and 1975, with the 1972 survey asking a much broader range of questions. 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, whether they are satisfied with themselves, etc. 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 7. 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). As in sections III-V, lower than normal pre-retirement wealth is defined as being in the bottom 20% of the wealth residual distribution. 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 grasshoppers 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 consumptions 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 7. 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.

VIII. Welfare Loss of 'Myopic' Consumption Rules (To Be Completed)

One question that might be asked given the results in the previous section is "What is the welfare loss to these low residual household from not smoothing their consumption in response to predictable income shocks?" In this section, I use the parameters from the previous sections to

calibrate a lifecycle model under standard assumptions. Doing so, allows me to compute the welfare loss associated with not smoothing consumption during a household's working years and not smoothing consumption across retirement. Previewing the results, the welfare loss associated with not smoothing consumption during retirement is very large compared to not smoothing retirement during working years. To be completed ……

IX. Conclusions

In this paper, I document that part of the variation in pre-retirement wealth across households who have similar lifecycle income, health and demographic trajectories is due to the fact some households make myopic consumption and savings decisions throughout their lifecycle. These households are not only ill prepared for retirement, they also do not fully smooth predictable income changes during their working years.

To isolate households with similar economic circumstances but had different levels of pre-retirement savings, I regressed log 1989 wealth for non-retired households in the Panel Study of Income Dynamics (*PSID*) between the ages of 50 and 65 on measures of lifetime earnings and earnings variability, lifecycle demographics – including changes in family structure, past and current health status, and expected future retirement benefits. The benefit of using the *PSID* for this type of analysis is that there exists 20 years of prior data for the pre-retired households in the 1989 survey. As a result, lifecycle events for each pre-retired household can be carefully controlled for. Taking the residuals from this regression allows me to isolate two groups who are identical along all observable dimensions, yet have dramatically differing pre-retirement wealth levels. Households in the bottom 20% of the wealth residuals had wealth levels that were only ten percent of the average household not in the bottom 20% of the wealth residuals.

Additionally, these households with low wealth residuals suffered percentage declines in consumption during their subsequent retirement that were twice as large as other pre-retired households. Likewise, these households were twice as likely to take up government assistance –

specifically SSI – during retirement. All indications point to the fact that these low residual households were ill prepared to sustain their consumption during retirement.

Did these households end up entering retirement with so little wealth because they had different time preference rates or was it because they simply did not plan for the eventuality of retirement? In the main part of the paper, I explore the latter proposition. I find that those households with retirement wealth much lower than predicted were much more likely to have consumption plans that responded to predictable changes in income during their working years. Not only do these households seem to not smooth consumption between their working years and retirement, these households do not fully smooth smaller predictable changes in income early in their lifecycle. I find that between 10-20% of the pre-retired low wealth population seems myopic with regards to future predictable income changes. Households in the bottom ten percent of the pre-retirement wealth distribution are estimated to have a marginal propensity to consume out of predictable income changes of about 0.8. Households in with wealth residuals between the 10^{th} and the 20^{th} percentile of the distribution had a marginal propensity to consume out of predictable income changes of about 0.25. The corresponding estimated marginal propensity to consume out of predicted income changes for the rest of the pre-retired households was essentially zero.

Given that these households respond equally to predictable increases in income as they do to predictable income declines rules out that liquidity constraints are driving the results. Liquidity constraints may prevent a household from borrowing to smooth predictable income changes, but nothing prevents them from saving to smooth their consumption. Additionally, the fact that the sample split in 1989 based on pre-retirement wealth may not be orthogonal to the error term in the consumption Euler equation during the household's working years implies that care is needed when interpreting the results. I show that persistently unexpected bad income realizations cannot explain the Euler equation results. In other words, the evidence suggests that the way the sample is split is orthogonal to the instrumented predictable component of income. Finally, I show that

other theories of consumption cannot explain all the results in the paper. Specifically, the substitutability between consumption and leisure, a theory of habit consumption, theories of precautionary savings or time inconsistent preferences are ruled out as explaining the results.

These households seem to be self-aware of their propensity to not plan for the future. In the early 1970s, *PSID* respondents were asked to describe their planning and saving propensities. While these measures are noisy measures of actual behavior, they are indicative of how the respondents view themselves. Given the panel nature of the *PSID*, the pre-retired households in the late 1980s responded to these self-assessment measures during their early working years. Compared to other pre-retired households, those households who entered retirement with lower than predicted wealth were much more likely to describe themselves as 'spenders' rather than 'savers', less likely to report that they plan for the future and were less likely to report carrying out their plans conditional on having them.

Additionally, those with low pre-retirement wealth were perpetual under savers. In the early 1970s, the households were asked whether they had accumulated savings equal to one month's income. Only 30% of the low wealth 1989 pre-retired households reported having one month's worth of accumulated saving in 1972 (where the household ages ranged from 33 – 48). The comparable number for the other group of 1989 pre-retired households was over 50%. The lack of saving was something that these low wealth pre-retired households perpetuated throughout their lifetimes.

Taken together, all the results of the paper imply that Aesop, in his fable, was correct in assessing differences in savings behavior within the population. Some households behave as the proverbial ants by saving during their working years so as to sustain consumption during winter. Other households are akin to the grasshoppers, consuming most of the resources during summertime, leaving them ill-prepared for the future. The former households behave as would be predicted by the permanent income model; they smooth consumption across periods where income changes in a predictable manner. However, the latter grasshopper group appears myopic

in their decision-making. One reason we observe such heterogeneous amounts of pre-retirement wealth, conditional on lifetime income, health and demographic profiles, is because some households simply do not plan for retirement.

From a public policy perspective, it is important to identify such households in the population. It is these households that will have large consumption responses to temporary tax cuts. Additionally, reducing or eliminating social security benefits will most adversely affect such households who do not plan for retirement on their own. In future research, more work needs to be done to isolate what causes differences in consumption rules across households, and whether these consumption rules change as a household ages.

Appendix

As is well known (Hall, 1978), the solution to the model presented in equation 1 obeys the following consumption Euler equation:

$$E_t\left[\frac{(1+r_{ik,t+1})}{(1+\delta_k)}\ \exp\left(\Theta_{ik,t+1}-\Theta_{ikt}\right)\left[\frac{C_{ik,t+1}}{C_{ikt}}\right]^{-\rho_k}\right]= \quad 1. \qquad \text{(N1)}$$

If expectations are rational and prices adjust to clear markets, household *i* satisfies (N1) up to a random forecast error, $e_{ik,t+1}$. Rewriting (N1), we get:

$$\frac{(1+r_{ik,t+1})}{(1+\delta_k)}\ \exp(\Theta_{ik,t+1}-\Theta_{ikt})\left[\frac{C_{ik,t+1}}{C_{ikt}}\right]^{-\rho_k} \quad = \quad 1+e_{ik,t+1}. \qquad \text{(N2)}$$

The nonlinear equation (N2) is difficult to estimate consistently with panel data because of measurement error in consumption data. Linearization of (N2) potentially allows identification of the consumption measurement error. As in Lawrance (1991), suppose that measured consumption, C, is related to true consumption, C*, according to:

$$C^*_{ikt} = C_{ikt}\exp(v_{ikt}), \qquad \text{(N3)}$$

where v_{ikt} is the mean zero random measurement error of household *i* at time *t*. According to this specification, measurement error is distributed across households as a percentage of consumption rather than by an absolute amount. Taking logarithms of (N2) and (N3) and substituting yields the following estimable equation in terms of measured consumption:

$$\Delta \ln C_{ik,t+1} = \frac{-\ln(1+\delta_k)}{\rho_k}+\frac{\omega^2_{ik}}{2\rho_k}+\frac{\ln(1+r_{ik,t+1})}{\rho_k}+\frac{(\Theta_{ik,t+1}-\Theta_{ikt})}{\rho_k}+\varepsilon_{ik,t+1}, \qquad \text{(N4)}$$

where $\varepsilon_{ik,t+1}$ is $(e_{ik,t+1}-e^2_{ik,t+1}/2-\omega^2_{ik}/2)/\rho_k-v_{ik,t+1}+v_{ikt}$, has mean zero and the law of iterated expectations implies that it is uncorrelated with any variable known at time *t* (Hall, 1978), where ω_{ik}^2 is the variance of the forecast error ($E(e^2_{ik,t+1})$), and where consumption and the interest

rate are jointly lognormal (see Hansen and Singleton (1983)). Even if consumption and the interest rate are not jointly lognormal, (4) is valid up to a second order Taylor series expansion of $e_{i,t+1}$.

References

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.

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 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 (1994). "How Does Future Income Affect Current Consumption." *Quarterly Journal of Economics* 109(1): 111-147.

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).

Flavin, Marjorie (1981). "The Adjustment of Consumption to Changing Expectations About Future Income." *Journal of Political Economy* 89(5): 974-1009.

Friedman, Milton (1957). *A Theory of the Consumption Function*. Princeton: Princeton University Press.

Gale, William G. (1998). "The Effects of Pensions on Household Wealth: A Re-evaluation of Theory and Evidence." *Journal of Political Economy*, 106(4), 706-23.

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 Miskin (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Pane Data on Households." *Econometrica* 50(2): 461-481.

Hammermesh, Daniel (1984). "Consumption During Retirement: The Missing Link in the Life Cycle." *Review of Economics and Statistics*, 66(1): 1-7.

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 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 (1996). "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets." *Journal Business and Economics Statistics* 14(1): 81-90.

Lusardi (2001)

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.

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.

Figure 1: Distribution of Wealth Residuals
Sample: Pre-Retired Households in the PSID Aged 50-65

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
Selected Income and Demographic Variables	*1st Stage Residual Top 80 Percent*	*1st Stage Residual Bottom 20 Percent*	p-*value of difference Column I and II*
Wealth Distribution and Portfolio Composition			
25th percentile of Household Wealth	$28,900	$2,014	<0.01
Median Household Wealth	$83,150	$8,275	<0.01
75th 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
Demographics			
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
Labor Income and Labor Income Variability			
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
Health Shocks			
% 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
Retirement Pension			
% of 1989 Income Replaced During Retirement	0.54	0.52	0.66
Consumption			
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

Sample split using a first stage regression of 1989 household wealth on a vector of household observables.
Sample restricted to non-retired households aged 50-65 in the PSID in 1989.
Top 80% of residual sample includes 654 households. Bottom 20% of residuals includes 165 households.
All dollar values are in 1989 dollars.

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

Panel A: Average Age of Retirement

	I	II	III
	1st Stage Residual Top 80 Percent	*1st Stage Residual Bottom 20 Percent*	p-*value of difference Column I and II*
Average Age of Retirement	62.4	62.3	0.99
Sample Size	267	51	

* Sample: 50-65 year old, non-retired households in 1989, who were subsequently retired between 1990 and 1999.
* Wealth residuals defined as in Table 1 (see text for details).
* Retirement Age is defined as the first year that the household self-reports being retired.

Panel B: Average Decline in Consumption Upon Retirement

	I	II	III
	1st Stage Residual Top 80 Percent	*1st Stage Residual Bottom 20 Percent*	p-*value of difference Column I and II*
Average Yearly Food Consumption Prior to Retirement	$4,045	$3,700	0.39
Average Yearly Food Consumption After Retirement	$3,744	$2,913	<0.01
Average Percentage Point Decline in Consumption At Retirement	-0.03	-0.11	0.14
Median Percentage Point Decline in Consumption At Retirement	-0.11	-0.19	0.40
Sample Size	223	46	

* All dollar amounts in 1989 dollars.

* *Yearly Food Consumption Prior to Retirement* is defined as the household's food consumption averaged over the three years prior to retirement.
* *Yearly Food Consumption After Retirement* is defined as the household's food consumption averaged over the three years after retirement.
* *Average Percentage Point Decline in Consumption At Retirement* is computed for each household and then averaged over households.

Note: Sample sizes differ between panels A and B of Table A1 due to truncating top/bottom 5% of outliers in the change in consumption distribution. See text for details.

Table 3
Incidence of SSI Take-Up Among Subsequently Retired Households,
By 1989 Wealth Residual Groups

	I	II	III
	1st Stage Residual Top 80 Percent	*1st Stage Residual Bottom 20 Percent*	p-*value of difference Column I and II*
Percent of Subsequently Retired Households on SSI in either 1997 or 1999	0.064	0.137	0.065

* Sample: 50-65 year old, non-retired households in 1989, who were subsequently retired in 1997 and 1999. Wealth residuals defined as in Table 1 (see text for details).

* Number of households with 1st stage residuals in top 80 percent who were retired in either 1997 or 1999: 267.
* Number of households with 1st stage residuals in bottom 20 percent who were retired in either 1997 or 1999: 51.
* Note: PSID did not collect any data in 1998.

Table 4:
Responsiveness of Early Life Consumption to Predictable Income Changes, Segmenting By 1989 Pre-Retired Wealth Residuals

Panel A: Test for Excess Sensitivity

$$\dot{C}_{ik,t+1} = \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1}$$
$$+\alpha_5 D_{<20} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi \mathbf{D}_{\mathbf{Year}} + \beta_1 \dot{Y}^{Predict}_{ik,t+1} + \beta_2 D_{<20} \dot{Y}^{Predict}_{ik,t+1} + \varepsilon^{*}_{ik,t+1}$$

where: $D_{<20}$ is a dummy variable indicating whether the households had first stage residuals in the bottom twenty percent of the pre-retired wealth distribution.
$\mathbf{D}_{\mathbf{Year}}$ is a vector of year dummies.

	β_1	β_2
Estimated Coefficient	-0.16	**0.56**
	(0.15)	**(0.28)**

Panel B: Test for the Existence of Liquidity Constraints

$$\dot{C}_{ik,t+1} = \alpha_0 + \alpha_1 D_{<20} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{<20} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1} + \alpha_5 D_{<20} \dot{famsize}_{ik,t+1}$$
$$+ \alpha_6 age_{ikt} + \alpha_7 D_{<20} age_{ikt} + \varphi D_{Year} + \beta_3 D_{Up} \dot{Y}^{Predict}_{ik,t+1} + \beta_4 D_{Down} \dot{Y}^{Predict}_{ik,t+1} + \beta_5 D_{<20} D_{Up} \dot{Y}^{Predict}_{ik,t+1}$$
$$+ \beta_6 D_{<20} D_{Down} \dot{Y}^{Predict}_{ik,t+1} + \varepsilon^{*}_{ik,t+1}$$

where: $D_{<20}$ and $\mathbf{D}_{\mathbf{Time}}$ are defined as in Panel A.
D_{up} is a dummy variable indicating whether the household had predictable income increases.
D_{down} is a dummy variable indicating whether the household had predictable income declines.

	β_3	β_4	β_5	β_6
Estimated Coefficient	-0.11	-0.01	**0.35**	**0.31**
	(0.25)	(0.29)	**(0.14)**	**(0.17)**

Notes For Both Panel A and Panel B
* 4,668 Observations (727 households)
* Standard errors (in parentheses) adjusted for within household heterogeneity.
* Predictable component of household labor income growth instrumented for using four lags of household labor income growth.
* Income processes estimated separately for low residual and other residual households.
* Coefficients in bold are statistically significant at 5% level.

Table 5:
Responsiveness of Early Life Consumption to Predictable Income Changes, Alternate Definitions of 1989 Low Wealth Residuals

Panel A: Excess Sensitivity Test

$$\dot{C}_{ik,t+1} = \alpha_0 + \alpha_1 D_{residual} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{residual} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1}$$
$$+\alpha_5 D_{residual} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{residual} age_{ikt} + \varphi \mathbf{D}_{\mathbf{Year}} + \beta_1 \dot{Y}^{Predict}_{ik,t+1} + \beta_2 D_{residual} \dot{Y}^{Predict}_{ik,t+1} + \varepsilon^{*}_{ik,t+1}$$

where: $D_{residual}$ is a dummy variable indicating whether the households had first stage residuals in the bottom X% of the wealth distribution, where X is, separately, 10, 20, 30, 40 and 50. $\mathbf{D}_{\mathbf{Time}}$ is a vector of time dummies.

Definition of $D_{residual}$	**λ_1**	**λ_2**
1. Bottom 10% of first stage residuals	-0.21 (0.15)	**0.97** **(0.53)**
2. Bottom 20% of first stage residuals	-0.24 (0.15)	**0.56** **(0.28)**
3. Bottom 30% of first stage residuals	-0.23 (0.14)	**0.46** **(0.20)**
4. Bottom 40% of first stage residuals	-0.24 (0.14)	0.27 (0.19)
5. Bottom 50% of first stage residuals	-0.19 (0.15)	0.11 (0.21)

Table 5 (continued):
Responsiveness of Early Life Consumption to Predictable Income Changes, Alternate Definitions of 1989 Low Wealth Residuals

Panel B: Test for the Existence of Liquidity Constraints

$$\dot{C}_{ik,t+1} = \alpha_0 + \alpha_1 D_{residual} + \alpha_2 \ln(1+r_{ik,t+1}) + \alpha_3 D_{residual} \ln(1+r_{ik,t+1}) + \alpha_4 \dot{famsize}_{ik,t+1}$$
$$+ \alpha_5 D_{residual} \dot{famsize}_{ik,t+1} + \alpha_6 age_{ikt} + \alpha_7 D_{residual} age_{ikt} + \varphi D_{Year} + \beta_3 D_{Up} \dot{Y}^{Predict}_{ik,t+1} +$$
$$+ \beta_4 D_{Down} \dot{Y}^{Predict}_{ik,t+1} + \beta_5 D_{residual} D_{Up} \dot{Y}^{Predict}_{ik,t+1} + \beta_6 D_{residual} D_{Down} \dot{Y}^{Predict}_{ik,t+1} + \varepsilon^{*}_{ik,t+1}$$

where: $D_{residual}$ and $\mathbf{D_{Time}}$ are defined as in Panel A.
D_{up} is a dummy variable indicating whether the household had predictable income increases.
D_{down} is a dummy variable indicating whether the household had predictable income declines.

Definition of $D_{residual}$	λ_3	λ_4	λ_5	λ_6
1. Bottom 10% of first stage residuals	-0.20	-0.09	**0.53**	**0.44**
	(0.25)	(0.29)	**(0.24)**	**(0.20)**
2. Bottom 20% of first stage residuals	-0.17	-0.01	**0.35**	**0.31**
	(0.25)	(0.29)	**(0.14)**	**(0.17)**
3. Bottom 30% of first stage residuals	-0.13	-0.09	**0.29**	0.15
	(0.25)	(0.29)	**(0.10)**	(0.19)
4. Bottom 40% of first stage residuals	-0.11	-0.14	0.20	0.13
	(0.25)	(0.29)	(0.09)	(0.20)
5. Bottom 50% of first stage residuals	-0.11	-0.10	0.11	0.15
	(0.27)	(0.30)	(0.10)	(0.21)

Notes For Both Panel A and Panel B
* 4,668 Observations (727 households)
* Standard errors (in parentheses) adjusted for within household heterogeneity.
* Predictable component of household labor income instrumented for using four lags of household labor income.
* Income processes estimated separately for low residual and other residual households.
* Coefficients in bold are statistically significant at 5% level.

Table 6:
Responsiveness of Early Life Consumption to Predictable Income Changes, Additive Definitions of 1989 Low Wealth Residual

This table reports the estimation of an equation similar to that reported in Tables 3 and 4 except that separate interactions are included for different percentiles of the pre-retirement wealth residuals. Unlike Table 4, all these separate interactions are jointly included in the regression.

where: $D_{0\text{-}10}$, $D_{10\text{-}20}$, $D_{20\text{-}30}$, $D_{30\text{-}40}$, *and* $D_{40\text{-}50}$ are, respectively, dummy variables indicating whether the households had first stage residuals between the 0-10th percentiles, between the 10th – 20th percentiles, between the 20th – 30th percentiles, between the 30th – 40th percentiles, and between the 40th – 50th percentiles of the residual distribution.

Variable	Coefficient Estimate (Standard Error)
$\dot{Y}^{P}_{t,t+1}$	-0.26 (0.16)
$D_{0-10}\,\dot{Y}^{P}_{t,t+1}$	**0.93** **(0.45)**
$D_{10-20}\,\dot{Y}^{P}_{t,t+1}$	**0.34** **(0.17)**
$D_{20-30}\,\dot{Y}^{P}_{t,t+1}$	0.14 (0.19)
$D_{30-40}\,\dot{Y}^{P}_{t,t+1}$	-0.29 (0.25)
$D_{40-50}\,\dot{Y}^{P}_{t,t+1}$	-0.18 (0.29)

Notes:

* 4,668 Observations (727 households)
* Standard errors (in parentheses) adjusted for within household heterogeneity.
* Predictable component of household labor income instrumented for using four lags of household labor income.
* Income processes estimated separately for low residual and other residual households.
* Coefficients in bold are statistically significant at 5% level.

Table 6 (continued):
Responsiveness of Early Life Consumption to Predictable Income Changes, Additive Definitions of 1989 Low Wealth Residual

Panel B: Test For Existence of Liquidity Constraints

Variable	Coefficient Estimate (Standard Error)
$D_{Up} \dot{Y}^{P}_{t,t+1}$	-0.22 (0.25)
$D_{Down} \dot{Y}^{P}_{t,t+1}$	0.04 (0.29)
$D_{Up} D_{0-10} \dot{Y}^{P}_{t,t+1}$	**0.47 (0.21)**
$D_{Down} D_{0-10} \dot{Y}^{P}_{t,t+1}$	**0.59 (0.34)**
$D_{Up} D_{10-20} \dot{Y}^{P}_{t,t+1}$	**0.26 (0.12)**
$D_{Down} D_{10-20} \dot{Y}^{P}_{t,t+1}$	0.22 (0.18)

Notes For Both Panel A and Panel B
* 4,668 Observations (727 households)
* Standard errors (in parentheses) adjusted for within household heterogeneity.
* Predictable component of household labor income instrumented for using four lags of household labor income.
* Income processes estimated separately for low residual and other residual households.
* Coefficients in bold are statistically significant at 5% level.

Table 7
Household Response to Historical Saving Questions and to Subjective Self-Assessment of Planning and Spending Behavior: By 1989 Wealth Residual Groups

1972 and 1975 Saving and Subjective Self Assessment Questions.	I *1st Stage Residual Top 80 Percent*	II *1st Stage Residual Bottom 20 Percent*	III p-*value of difference Column I and II*
1972 Variables			
% 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
1975 Variables			
% 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?"

* Subjective self assessment questions were asked only in the PSID in 1972 and 1975. The 1972 set of questions were much more extensive. As a result, the carry out plans question and the saver versus spender question were not asked in 1975.
* Sample: 50-65 year old, non-retired households in 1989. Wealth residuals defined as in Table 1 (see text for details).

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

Panel A: Average Age of Retirement for 1984 Sample

	I	II	III
	1st Stage Residual Top 80 Percent	*1st Stage Residual Bottom 20 Percent*	p-*value of difference Column I and II*
Average Age of Retirement	62.7	62.8	0.99
Sample Size	421	91	

* Sample for Panels A and B of Table A1: 50-65 year old, non-retired households in 1984, who were subsequently retired between 1985 and 1999.
* Residuals computed from a first stage regression of household wealth in 1984 on observables.
* Retirement Age is defined as the first year that the household self-reports being retired.

Panel B: Average Decline in Consumption Upon Retirement for 1984 Sample

	I	II	III
	1st Stage Residual Top 80 Percent	*1st Stage Residual Bottom 20 Percent*	p-*value of difference Column I and II*
Average Yearly Food Consumption Prior to Retirement	$4,238	$4,093	0.61
Average Yearly Food Consumption After Retirement	$3,703	$3,122	0.02
Average Percentage Point Decline in Consumption At Retirement	-0.08	-0.16	0.07
Median Percentage Point Decline in Consumption At Retirement	-0.13	-0.22	0.08
Sample Size	391	79	

* All dollar amounts in 1989 dollars.

* *Yearly Food Consumption Prior to Retirement* is defined as the household's food consumption averaged over the three years prior to retirement.
* *Yearly Food Consumption After Retirement* is defined as the household's food consumption averaged over the three years after retirement.
* *Average Percentage Point Decline in Consumption At Retirement* is computed for each household and then averaged over households.

Note: Sample sizes differ between panels A and B of Table A1 due to truncating top/bottom 5% of outliers in the change in consumption distribution. See text for details.