

Changes in Nutrient Intake at Retirement*

Melvin Stephens Jr.

Desmond Toohey

University of Michigan

University of Delaware

and NBER

This Version: May 17, 2018

*Stephens: Department of Economics, University of Michigan, 611 Tappan St., Ann Arbor, MI, 48109, e-mail: mstep@umich.edu and National Bureau of Economic Research, Cambridge, MA. Toohey: Department of Economics, University of Delaware, 420 Purnell Hall, Newark, DE, 19716, e-mail: dtoohey@udel.edu. Peter Hudomiet and He Yang provided excellent research assistance. We thank Peter Ganong and seminar participants at the ASSA meetings, Maryland AREC, Michigan Labor Lunch, the MRRC Researcher Workshop, UC-Berkeley, and Wharton for helpful comments and suggestions. We thank Mark Aguiar and Erik Hurst for generously making their data and programs available. Research reported in this publication was supported by the National Institute On Aging of the National Institutes of Health under Award Number R03AG046503 and the Michigan Center on the Demography of Aging. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Changes in Nutrient Intake at Retirement

Abstract

While the literature finding a decrease in food expenditures at retirement suggests households do not adequately save for retirement, subsequent evidence that nutrient intake is unaffected by retirement has tempered these concerns. We further examine nutrient intake changes at retirement both by analyzing a much wider range of datasets, including longitudinal data, and by improving upon the empirical methodology used in earlier work. Our analysis yields four main results. First, unlike prior work, we find that caloric and nutrient intake fall at retirement in numerous cross-sectional datasets. We can reconcile these contrasting results as being due to well-documented differences and improvements in methodologies used to measure food intake. Second, using longitudinal data, we also find that intake falls at retirement. Third, we show that a food consumption index used in prior work to capture the relationship between permanent income and foods eaten can severely underestimate the impact of retirement on consumption. We show that a minor methodological revision circumvents this bias and that the revised consumption index falls at retirement. Finally, while unemployment reduces the consumption index, we find, in contrast to prior work, that the impact of retirement on the consumption index is larger. Overall, we consistently find that retirement reduces food intake.

1 Introduction

Are households adequately saving for retirement? The “retirement-consumption puzzle,” the finding in numerous studies that expenditure falls at retirement, in particular when retirement is expected, suggests that households are ill-prepared to exit the labor force.¹ Complementary results indicate, however, that this evidence alone does not refute the Life-Cycle/Permanent Income Hypothesis (LCPIH). Some studies show that the expenditure decline is found primarily in food and work-related expenditures.² In addition, the influential work of Aguiar and Hurst (2005) finds that, although food expenditure decreases at retirement, measures of nutrient intake do not change at retirement.

In this paper, we dramatically expand the investigation of nutrient intake changes at retirement by using cross-sectional datasets that span over forty years and, for the first time in this literature, longitudinal data. We also show that the novel approach used in prior work for linking an individual’s basket of food intake to permanent income can severely attenuate the estimated impact of retirement on consumption. We propose an alternative approach that circumvents this bias. Our analysis across these numerous data sources and methodologies finds that food consumption falls at retirement, indicating that adjustments to home production are not enough to fully smooth consumption at retirement.

The issue of retirement savings adequacy generates much interest among economists for numerous reasons including the dramatic decline in the aggregate savings rate (e.g., Parker 1999), changing household retirement wealth risk as employers shifted from defined benefit to defined contribution pension plans (e.g., Poterba 2014), and the lack of household retirement planning and financial sophistication (e.g., Lusardi and Mitchell 2007). Although assessing savings adequacy can be difficult due to heterogeneity in income, preferences, risk attitudes, etc. across households, a

¹Examples of such studies include Banks, Blundell, and Tanner (1998), Bernheim, Skinner, and Weinberg (2001), Lundberg, Startz, and Stillman (2003), Smith (2006), Haider and Stephens (2007), and Fisher, Johnson, Marchand, Smeeding, and Torrey (2008). Hurd and Rohwedder (2003, 2006, 2008) and Ameriks, Caplin, and Leahy (2007) find that households anticipate a decline in expenditures at retirement which is somewhat larger than the realized decline at retirement. Blau (2008) finds that the average observed retirement consumption decrease is substantially larger than would be predicted by a model with an uncertain date of retirement. Studies using data from Italy (Minicaci, Monfardini, and Weber 2010) and Japan (Stephens and Unayama 2012), countries in which workers receive large lump sum payments at the time of retirement, find smaller consumption changes at retirement.

²These studies include Fisher, Johnson, Marchand, Smeeding, and Torrey (2008), Battistin, Brugiavini, Rettore, and Weber (2009), Aguila, Attanasio, and Meghir (2011), Li, Shi, and Wu (2014), and Dong and Yang (2014).

number of techniques based on the Life-Cycle/Permanent Income Hypothesis (LCPIH) have been utilized to answer this question. In addition to examining expenditure changes at retirement, authors compare pre-retirement income to their post-retirement annuitized income (e.g., Moore and Mitchell 1997; Gustman and Steinmeier 1998) and compare simulated optimal savings levels with observed household wealth (e.g., Engen, Gale, and Uccello 1999; Scholz, Seshadri, and Khitatrakun 2006). Results based on the latter approach find that roughly 25 percent of households have inadequate retirement savings while the other studies find that consumption falls by nine to twelve percent at retirement. Complementary work by Olafsson and Pagel (2018), which uses detailed data on both household spending and balance sheets, shows workers increasing their savings upon entering retirement rather than dissaving as one would expect from rational LCPIH agents with adequate pre-retirement savings.

Aguiar and Hurst (2005) question whether expenditures are the appropriate measure of consumption to analyze at retirement. The LCPIH implies that households should smooth the marginal utility of consumption over time, not consumption expenditures *per se*. While non-durable consumption expenditures might reasonably approximate the marginal utility of consumption in some contexts, the large change in the amount of time available to invest in consumption-related activities upon retiring may dramatically alter the way in which consumption expenditures generate utility. For example, more time available for grocery shopping and meal preparation may allow households to generate the same level of (marginal) utility after retirement even with lower food expenditures. Using cross-sectional data from the early to mid-1990s, Aguiar and Hurst find that at retirement households: decrease food expenditure, keep food intake constant, spend more time in food preparation, and spend less time eating out. They interpret these findings as showing strong support of the LCPIH augmented with household production.³

We expand upon the Aguiar and Hurst analysis in multiple dimensions and, in each analysis, we find that food consumption falls at retirement. We first use cross-sectional data spanning 1971-2012 to examine changes in nutrient intake at retirement. One concern from using OLS regressions to

³In contrast, Been, Rohwedder, and Hurd (2017) jointly analyze expenditure, consumption, and time use in data from the Health and Retirement Study (HRS) and find little evidence of households replacing decreased market expenditures with home-produced goods, even in market consumption categories that are well-substituted by home production.

investigate this issue is that unobservable differences across households may bias estimates of the impact of retirement on food intake. For example, more patient individuals may both save more and retire later which would confound these estimates. Following the prior literature, we use age as an instrument, exploiting the non-linear age-retirement profile induced by retirement incentives found in the Social Security system. We find a significant decrease in caloric intake at retirement both in our OLS and two-stage least squares (2SLS) regressions. We also find a significant decrease at retirement for a number of nutrients including vitamin E and protein.

We are able to reconcile the differences between our findings and those of Aguiar and Hurst (2005). Aguiar and Hurst examine the change in food intake at retirement using the 1989-91 and 1994-96 waves of the Continuing Survey of Food Intake of Individuals (CSFII). As we discuss in detail below, dramatic methodological changes occurred between these two waves of the CSFII stemming from important methodological and implementation concerns with the contractor who fielded the 1989-91 CSFII. The immediate result was a sizable change in reported food intake between the 1989-91 and 1994-96 waves of the CSFII as well as in the estimated impact of retirement on food intake. Using numerous nationwide food intake surveys that did not suffer from these methodological concerns as well as restricting our analysis to surveys collected after the methodological changes were implemented yields a large and significant drop in food intake at retirement. Analysis of this data using quantile regressions and quantile instrumental variable regressions (Chernozhukov and Hansen 2005) indicates that the decrease is fairly constant across the intake distribution.

We then turn to longitudinal data with repeated measures of caloric intake. Following Hall's (1978) seminal work on the LCPIH, the vast majority of the subsequent literature has tested the model's implication that households should smooth (the marginal utility of) consumption over time. The first difference framework accounts for unobserved time-invariant household characteristics that may be associated with consumption changes. Following this tradition, we find that caloric intake falls at retirement. While we again use age as an instrument for retirement, the first difference estimator implicitly accounts for a linear change in caloric intake with age such that our results account for linear physiological changes in food intake with age.

We also expand on Aguiar and Hurst's consumption index which relates individuals' entire food

consumption baskets to permanent income. They project the foods eaten by prime-age workers onto their permanent incomes. The estimated relationships between the foods eaten and permanent income are then used as weights to produce an index measuring the implied permanent incomes of workers based on their food consumption. Aguiar and Hurst find that this consumption index does not fall at retirement.

We show that their methodology for producing the consumption index will attenuate the estimated impact of retirement on this index. Aguiar and Hurst regress *predicted* permanent income on food intake categories to generate weights for their index. However, their predicted permanent income measure understates the true variability in permanent income across individuals. As a result, which we show formally, the index weights created in this way will understate the true variability in the index and, in turn, attenuate the estimated impact of retirement on consumption. We show that using actual income in place of predicted permanent income, an approach which is valid under the Life-Cycle/Permanent Income Hypothesis, circumvents this concern. Our approach yields a statistically significant and economically meaningful decrease in consumption at retirement.

Finally, we re-examine the impact of unemployment on consumption. Aguiar and Hurst find that unemployment reduces consumption expenditures and, most notably, find that unemployment reduces the consumption index even though retirement has no effect on the consumption index. While we find a similar impact of unemployment on consumption. when we use our revised consumption index, we find the exact opposite of Aguiar and Hurst: the impact of retirement on consumption is larger than the impact of unemployment on consumption.

Overall, we find consistent evidence across numerous datasets and multiple empirical methodologies that nutrient intake falls at retirement. While home production may allow households to mitigate the loss in welfare that would occur in response to a *ceteris paribus* drop in expenditures at retirement, we find that, nonetheless, nutrient intake declines at retirement. Our results are consistent with the literatures that find that many households do not have adequate retirement savings (see the survey of Poterba 2014) and both lack financial literacy and fail to plan for retirement (see the survey of Lusardi and Mitchell 2014).

2 Theoretical Motivation

We use the Life-Cycle/Permanent Income Hypothesis (LCPIH) as a framework for understanding how changes in food intake inform the question of whether households are adequately saving for retirement. Households are assumed to maximize their utility over the remainder of their lives, V_t , with

$$V_t = \max_{\{C_t, L_t\}} E_t \left[\sum_{k=t}^T \left(\frac{1}{1+\delta} \right)^{k-t} U(C_k, L_k; \boldsymbol{\theta}_k) \right] \quad (1)$$

subject to the asset equation

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

where C_t , L_t , w_t , and A_t are (composite) consumption, leisure, wage, and assets in year t , respectively, $U(\cdot)$ is the period specific utility function, $\boldsymbol{\theta}_t$ are variables (which may or may not be time-varying) thought to affect utility such as age and family size, \bar{L} is the maximum annual amount of time available for leisure, r is a constant interest rate, and δ is the subjective discount rate.

Using standard optimization techniques, it is straightforward to show that the marginal utility of consumption in year t is equal to the marginal utility of assets in period t , λ_t , or

$$U_C(C_t, L_t; \boldsymbol{\theta}_t) = \lambda_t \quad (3)$$

If $U(\cdot)$ is additively separable in consumption and leisure, then $U_C(C_t, L_t; \boldsymbol{\theta}_t)$ will not depend on L_t . Under such an assumption, along with standard utility functions (e.g., constant relative risk aversion), the marginal utility of consumption can be proxied by consumption expenditures alone if appropriate measures of $\boldsymbol{\theta}_t$ are available. However, if $U(\cdot)$ is not separable in consumption and leisure, then $U_C(C_k, L_k; \boldsymbol{\theta}_t)$ will depend upon hours of work decisions, including whether or not one retires, as well as consumption expenditures.

One mechanism for such a non-separability, as noted by Aguiar and Hurst (2005), is the production of household goods. Becker (1965) argues that consumption is based on “final goods” which are, in turn, functions of market goods and time spent in conjunction with the market goods (e.g.,

cooking meals, watching TV, etc.). The importance of time allocation is, perhaps, particularly acute when examining retirement given the sharp reduction in time spent working. Households may still enjoy the same level of food consumption after retirement even if their food expenditures fall, given that they have more time available to spend in the household production of food. E.g., Aguiar and Hurst (2005) find more time spent shopping, more time preparing meals, and less time eating out for the retired relative to the non-retired. Thus, the marginal utility of actual consumption may not change at retirement even if expenditures decline when workers leave the labor force. Rather, a measure of actual food consumption, such as food intake, is arguably much more appropriate to use for analyzing the consumption response at retirement.

Alternatively, the resulting Euler Equation which determines the household's optimal allocation of consumption between periods t and $t + 1$ is

$$U_C(C_t, L_t; \boldsymbol{\theta}_t) = \left(\frac{1+r}{1+\delta} \right) E_t[U_C(C_{t+1}, L_{t+1}; \boldsymbol{\theta}_{t+1})] \quad (4)$$

Again, if $U(\cdot)$ is additively separable in consumption and leisure then we should not observe changes in expenditures at retirement within households as marginal utility is a function of only expenditures. However, non-separabilities between consumption and leisure induced by home production invalidate the use of within-household changes in expenditures as a test of whether or not households are adequately prepared for retirement. Changes in food intake, on the other hand, might provide a better metric by which to estimate retirement consumption changes.

3 Changes in Intake at Retirement: Cross-Sectional Data

3.1 Data

Our cross-sectional analysis uses numerous nationally-representative datasets that contain information on food intake. All these surveys collect food intake using the 24-hour food recall method in which detailed data is recorded on all foods consumed during the prior day. Some of the surveys collect up to two days of additional individual intake data using recall or food consumption diaries. We briefly discuss the relevant features of these datasets below while more detail is found in the

Data Appendix.

The first set of studies that we use were collected by the U.S. Department of Agriculture (USDA). We use the 1989-91 and 1994-96 Continuing Survey of Food Intake of Individuals (CSFII), which are the datasets used by Aguiar and Hurst (2005, hereafter AH). The 1989-91 CSFII collected three days of food intake per respondent while the 1994-96 collected intake for two days for reasons we shall discuss below. We also use the 1977-78 and 1987-88 waves of the Nationwide Food Consumption Survey (NFCS).⁴ The NFCS recorded detailed food intake from all household members over three days. The NFCS imposed a higher burden on households than the CSFII as it also asked households to provide details on food items consumed by the household over a seven day period. Both the CSFII and NFCS contain measures of food expenditures at home and away from home as well as demographic, health, and labor market information.

We also use data collected by the U.S. Department of Health and Human Services (HHS). The National Health and Nutrition Examination Survey (NHANES) is well-known as a survey in which participants undergo a comprehensive health exam, but it also obtains food intake information using 24-hour recalls. The first three waves, NHANES I (1971-75), NHANES II (1976-80), and NHANES III (1988-94), each contain a single day of food intake information in addition to health, labor market, and demographic data. However, these surveys did not obtain food expenditures. Moreover, household income in the NHANES surveys is a categorical variable, not a continuous one.

NHANES returned to the field in 1999 and continued to collect one day of food intake. Beginning in 2002, HHS and USDA agreed to merge the CSFII with the NHANES and have fielded this survey on an on-going basis in all subsequent years. Starting in 2003, the Continuous NHANES has collected two days of food intake from respondents. Since 2007 the survey has also collected food expenditure information, both at home and away from home. We use Continuous NHANES data from 1999-2012.

⁴The first national survey to collect individual-level food intake was the 1965-66 Household Food Consumption Survey (HFCS). However, this earlier wave of the survey did not collect labor force information except for an indicator as to whether or not the female household head worked thirty-five or more hours. As such, we make limited use of the 1965-66 HFCS.

3.2 Empirical Methodology

Based on equation (3), we follow AH by estimating cross-sectional specifications

$$y_i = \gamma R_i + \mathbf{X}_i \boldsymbol{\beta} + \epsilon_i \quad (5)$$

where y_i is a measure of food intake, such as calories, expenditure, or macronutrients, \mathbf{X}_i is a vector of individual and household level characteristics, and R_i is an indicator for whether or not the respondent is retired.⁵ While we closely follow the empirical specifications used by AH, we must make minor adjustments due to differences in the data collected across surveys. \mathbf{X}_i includes indicators for black, four education groups, nine household size groups, five categories of self-reported health, and calendar year effects that differ by survey.⁶

One concern is that the error term, ϵ_i , may be correlated with retirement status, R_i , which could occur if, for example, relatively more impatient individuals tended to accumulate less in assets and retire earlier which could lead to a spurious negative correlation between retirement and food intake. To circumvent this concern, Aguiar and Hurst use the non-linear relationship between age and retirement status, which is driven primarily by the early and normal Social Security retirement ages in the United States, as an instrumental variable for retirement. The corresponding first stage equation is

$$R_i = \boldsymbol{\psi} \mathbf{age}_i + \mathbf{X}_i \boldsymbol{\delta} + u_i \quad (6)$$

where \mathbf{age}_i is a vector of age indicators. As long as this non-linear age pattern is not systematically related to the unobserved differences in food intake across households then the instrument will be valid. Since \mathbf{X}_i does not include any other age terms, this approach assumes that age only affects food intake through changes in retirement status, e.g., there are no physiological changes in food intake needs across the ages used in the analysis, 57-71.

As with most of the literature examining retirement consumption changes, we focus our analysis

⁵We define individuals as retired who self-report recent activity both as retired and not working. I.e., if someone states that they are retired but still report that they are working, we do not define them as being retired.

⁶AH also include indicators for residing in an urban area, four Census regions, and seven health conditions, variables which are not available in all surveys. However, in results not shown here but available from the authors, excluding these variables has little effect on the findings. We thank Mark Aguiar and Erik Hurst for providing their data and programs with which we were able to exactly replicate their findings using the CSFII.

on the retirement of the male household head (Banks, Blundell, and Tanner 1998; Bernheim, Skinner, and Weinberg 2001; Haider and Stephens 2007). Given that men historically have been the primary household earners, restricting to men generally limits the analysis to households in which retirement changes are likely to have a large impact on income and expenditure. To define male household heads, we adopt the convention used by the Panel Study of Income Dynamics wherein married/partnered households the male is deemed to be the household head while single-headed households can have either a male or a female head.⁷ Table A1 presents summary statistics for the cross-sectional studies that we use in our analysis. The age, race, and household size composition of the samples remain fairly constant across the samples. Predictably, education levels are higher for the more recent surveys than in the earlier ones.

Table 1 presents the OLS estimates of changes in log household income and log food expenditure at retirement using the subset of surveys which contain food expenditure information: NFCS, CSFII, and NHANES (2007-2012).⁸ Each coefficient in the table is from a separate regression using equation (5) and is an estimate of the parameter γ on the retired indicator. Our 2SLS estimates shown in the second and fourth columns using age indicators as instruments show a significant impact of retirement on income and food expenditure. In addition, the OLS estimate for the impact of retirement on food expenditure for men is near the nine percent to twelve percent range found in the aforementioned longitudinal studies, although our 2SLS estimate is somewhat larger than the findings in the earlier literature.⁹

3.3 The Impact of Retirement on Caloric Intake

Before turning to our estimates of the impact of retirement on caloric intake, we briefly examine the age-retirement profiles which are the variation used to identify the impact of expected retirement on intake. Figure 1 plots the coefficients on the complete set of age dummies from estimating the first stage equation (6) where the excluded age is age 57. When pooling across all the cross-sectional studies, as shown in the top left panel of the Figure, retirement rates are relatively flat

⁷Since we encounter very few married/partnered households with same-gendered heads in our samples, we drop such households from our analysis.

⁸Since household income is a categorical variable in all of the NHANES surveys, we restrict our income regressions to the NFCS and CSFII.

⁹As we noted above, we follow Aguiar and Hurst by not including age among the regressors in \mathbf{X}_i .

from age 57 to 61 but then rise sharply at subsequent ages beginning with a pronounced jump at age 62. The first stage F statistic for the excluded instruments, shown in brackets at the top of each panel, is 56.4 for the pooled sample. The remaining panels of the Figure show broadly consistent age-retirement patterns for each of the individual studies.¹⁰

Our estimates of the impact of retirement on caloric intake are shown in Table 2. Each coefficient shown in the table is from a separate regression and is an estimate of the parameter on the retired indicator in (5) with log calories consumed as the dependent variable. The first column presents the finding from Table 2 of Aguiar and Hurst (2005) for the impact of retirement on caloric intake using both waves of the CSFII. This 2SLS estimate is generated using age indicators as instruments for retirement in a regression including controls for race, sex, education, household size, year, self-reported health, Census region, urbanicity, and seven individual health conditions.¹¹ As AH highlight in their work, the estimated effect of retirement is economically small and statistically not different from zero.

Column (2) of Table 2 reports OLS and 2SLS estimates of the effect of retirement on log caloric intake using the same data source as AH (the two waves of the CSFII) but modifying the specification to include only variables that are available across all the cross-sectional datasets.¹² In addition, we follow the prior literature by, as discussed above, restricting the sample in column (2) to men only.¹³ These changes to the specification produce a 2SLS estimate in column (2) that is nearly identical to the AH result in column (1). Although not shown here, subsets of these changes similarly do not have meaningful effects on the finding. In contrast to the 2SLS estimate, the OLS coefficient in the adjusted replication of column (2) suggests that there is a measurable, if modestly-sized, decrease in caloric intake associated with retirement in the CSFII data.

¹⁰The F statistics for the excluded instruments shown in the Figure all exceed conventional levels.

¹¹The health conditions controls are indicators for self-reporting any history of diabetes, high blood pressure, heart disease, cancer, osteoporosis, high cholesterol, and stroke.

¹²For a consistent specification, we exclude the indicators for urban, region, and the individual health conditions from the original AH set of regressors. We also make a slight change to the coding of the retirement variable, only coding individuals as retired if they self-report their labor force status as retired (as in AH) *and* they report that they are not currently working. This change affects only 17 of the 985 retired individuals in the AH sample of column (1) and allows us to have a consistent definition of retirement in other surveys that ask about retirement only when respondents are not working.

¹³We also include observations from the oversample of low-income households in CSFII 1989-91, which are excluded from column (1), and weight the regressions using the CSFII's sampling weights. Restricting to male household heads but including the low-income sample produces an estimation sample that is approximately one-fifth smaller than that used by AH.

We next apply this same analysis to a pooled sample of all the cross-sectional datasets described in Section 3.1.¹⁴ As we find in column (3) of Table 2, the OLS estimate of the impact of retirement on consumption in this pooled sample is quite similar to what we find for the CSFII sample in column (2). However, the magnitude of the 2SLS estimate, on the other hand, dramatically jumps from two log points to nearly 18 when pooling all the studies. At -0.175, the 2SLS estimate of the effect of retirement on log caloric consumption is highly statistically significant, economically meaningful, and approaches the magnitude of the effect of retirement on food expenditure in Table 1. Taken together, the estimates of columns (1), (2), and (3) show that the different results obtained from using the larger group of datasets is not due to differences in specification but due to differences in the relationship between retirement and food consumption observed across the datasets.¹⁵

Table 3 further highlights the differences across datasets by presenting the results for each of the individual cross-sectional datasets, ordered by when they were first in the field. When applying OLS, we find a decrease in intake at retirement in every dataset with the majority of these estimates being statistically significant. When using 2SLS, we find a large and significant decrease in caloric intake at retirement for five of the eight individual datasets. For one study, the CSFII 1994-96 (column (7) of Table 3), we find a negative and economically meaningful but statistically insignificant effect of retirement on intake. However, the NFCS 1987-88 and CSFII 1989-91 (columns (4) and (6) of Table 3, respectively) stand out in that we find *positive* but small and insignificant estimates when using these two studies. Thus, the 2SLS estimates are similar across the datasets with the exception of the NFCS 1987-88 and the two waves of the CSFII. As AH make exclusive use of food intake data from both waves of the CSFII, we next turn to examining

¹⁴One issue with pooling across multiple datasets is potential changes in the age-retirement profile induced by changes in the Social Security normal retirement age. For those reaching normal retirement age across all of our sample periods, this age was 65 for those born in 1937 and earlier, increased by two months each year for those born in 1938-1942, and then was age 66 for those born in 1943-1954. Early retirement age has remained at age 62 for all of these cohorts. Thus, changes in the normal retirement age likely have little impact on our results when pooling across multiple datasets and even less effect when we focus on individual datasets in Table 3.

¹⁵Although we have excluded women from our analysis, we have also estimated the impact of retirement on caloric intake including women in the sample. To do so, we also add an indicator for gender to equation (5). Our CSFII only adjusted replication, analogous to the 2SLS estimate in column (2) of Table 2, yields a point estimate of -0.004 with a standard error of 0.046 when including women. Thus, we continue to find a result very similar to that of AH when limiting the analysis to the CSFII only. When we pool all studies, analogous to column (3) of Table 2, we find an estimated impact of retirement on caloric intake of -0.157 (0.029) when including women. Thus, the dramatic change in the estimated impact of retirement on caloric intake when moving from only using the CSFII to including a broader range of surveys is unaffected by including women in the sample.

differences in the survey methodologies used across these surveys.

3.4 Reconciling Estimates Across Cross-Sectional Datasets

Both the 1987-88 NFCS and 1989-91 CSFII were conducted within close temporal proximity (1987-1991) and, as such, use variation from very similar birth cohorts (since the analysis is restricted to men aged 57-71). One possibility is that time period and/or cohort effects may explain why the estimates when using these two datasets differ substantially from those found when using the remaining datasets. E.g., the strong economic growth during the late 1980s may have allowed for higher-than-expected retirement benefits for those exiting the labor force during this period and lowered the income drop at retirement. However, the estimate from NHANES III (column (6) of Table 3), which was fielded contemporaneously (1988-1994), shows a large and significant decline in intake at retirement, which runs counter to an explanation based on time periods and cohorts. Instead, we believe that differences in data quality across these surveys can explain the discrepancies in estimated retirement effects.

Concerns about the quality of the NFCS 1987-88 were raised in a report by the US General Accounting Office (GAO) (US GAO, 1991). The GAO found many reasons to question the quality of the NFCS data including the household response rate of 34 percent, a lengthy interview averaging three hours with very small monetary compensation for participation, the contractor's failure to follow the planned sampling design, and high turnover among the contractor's staff along with a failure to maintain minimum staff training levels. The NFCS surveys had the feature that households were asked to record all food that was consumed or disposed of for a seven day period *after which* they then recorded their individual food intakes for three consecutive days. The concerns raised by the GAO led the US Department of Agriculture's Human Nutrition Information Service (USDA HNIS) to convene an expert panel to examine the issue of non-response in the NFCS 1987-88. The panel "concluded that it is not possible, with absolute certainty, to demonstrate either the presence or absence of nonresponse bias in the 1987-88 NFCS data" but went on to state that it "does not recommend use of the data from the 1987-88 NFCS." (USDA HNIS 1991a, p.13).¹⁶

¹⁶The low household response rate of 34 percent for the 1987-88 NFCS was part of a longer trend. As shown in Appendix Table A2, the household response rate previously had fallen from 85 percent for the 1965-66 HFCS to 61 percent for the 1977-78 NFCS, both of which had a similar design as the 1987-88 survey. To some extent, this trend

The USDA subsequently implemented a major redesign in its methodology for collecting information on individual food intake in concert with a number of other federal agencies (USDA 1994). Among the changes included a reduction from three days of individual intake to two days, collecting all days using 24-hour recall (whereas previously the first day was 24-hour recall followed by a two-day food diary left with the household), and a restructured 24-hour recall methodology to encompass a wider array of recall strategies used by respondents (USDA-ARS 1997). However, these changes were first implemented in the CSFII 1994-96. The methods employed to collect individual intake in the 1989-91 CSFII were the same as those used in the 1987-88 NFCS.

The graphs in Figure 2 provide some insight into how these issues affected reported dietary intake. The darker, solid line in each panel of the Figure plots average caloric intake for men by age for each of the cross-sectional datasets as well as the 1965-66 HFCS.¹⁷ The lighter, dashed line in each panel is the average caloric intake for men across all the studies pictured. Average male caloric intake in both the 1987-88 NFCS and the 1989-91 CSFII lies substantially below average caloric intake across all studies. Relative to these two surveys and consistent with the implementation of better dietary recall methods, the Figure shows that male caloric intake rises sharply in the 1994-96 CSFII. The survey re-design also had a dramatic effect on the individual response rate to the first day of food intake which, as shown in Appendix Table A2, jumped from 58 percent for the 1989-91 CSFII to 80 percent for the 1994-96 CSFII.¹⁸

parallels the rise in earnings non-response in the Current Population Survey which rose from 11 percent in 1968 to 18 percent in 1978 to 27 percent in 1982 (Lillard, Smith, and Welch 1986). While these CPS nonresponse rates are conditional on participating in the survey, they are consistent with a broader change in the willingness to provide information to surveys during this period.

¹⁷As we discussed above, the 1965-66 HFCS does not include relevant information on labor force participation and therefore it cannot be used in our regression analysis.

¹⁸The household response rate jumped from 67 percent for the 1989-91 CSFII to 85 percent for the 1994-96 CSFII. The fact that the CSFII did not contain the seven day household food use component likely contributed to the higher CSFIII response rates relative to the 1987-88 NFCS. Aguiar and Hurst (2005, p. 922, footnote 4) report the response rates for “eligible households” as 80 and 85 percent, respectively. The discrepancy between our reported response rates and those reported by Aguiar and Hurst is that Aguiar and Hurst fail to account for, as is done in the official computation of the response rates, the responses during the initial screening of households to determine which households are eligible. As shown in Text Table 10 of Tippet et al. (1995), the screening rate, which is the fraction of occupied households that responded to the initial survey contact to determine the household roster, is 84.5% in the 1989-91 CSFII. Multiplying this number by the eligible household response rate reported by Tippet et al. of 79.6% (corresponding to the 80% figure reported by Aguiar and Hurst) yields our household response rate of 67%. The screening rate for the CSFII 1994-96, as reported in Table 4-2 of USDA-ARS (1997), is 98.5%. As such, adjusting for the initial screening has a negligible effect on the household response rate for the 1994-96 CSFII which is why the 85% household response reported by Aguiar and Hurst matches the response rate that we report. Finally, note that in Appendix Table A2 we report unconditional individual response rates, i.e, those which subsume response to the screening interview as well as household participation in the survey. The individual response rates reported

Figure 2 also shows that NHANES III has higher levels of caloric intake than the two studies fielded contemporaneously, the 1987-88 NFCS and the 1989-91 CSFII. Recall that in Table 3 we find a significant retirement intake decline in NHANES III but no such decline at retirement in either of the contemporaneous studies. The dietary recall methods for NHANES were dramatically revised between NHANES II and NHANES III which likely contributed to the increase in caloric intake between these two NHANES waves as shown in Figure 2.¹⁹ In fact, the contractor that both helped develop the improved CSFII dietary intake methodology and fielded the 1994-96 CSFII is the same one that fielded NHANES III. These NHANES revisions, pre-dating those that were implemented in the 1994-96 CSFII, help explain why NHANES III has higher caloric intake levels than either the 1987-88 NFCS or the 1989-91 CSFII.²⁰

The changes in survey methodology not only affected the level of caloric intake but also the slope of the age-caloric intake profile. This slope is directly relevant for the 2SLS estimates shown in Table 3 since age is the instrument used in the analysis. The panels of Figure 2 also report the slope of the age-caloric intake profile for ages 57-71, the same ages as are used in Table 3. At these ages, the slope of the age-caloric intake profile is -3.8 in the 1987-88 NFCS and -5.8 in the 1989-91 CSFII. These are, by far, the smallest such slopes shown in the Figure. In fact, the slope of -19.5 for the 1994-96 CSFIII is nearly four times larger than what is found in the 1989-91 CSFII.

In light of these data quality issues, including the recommendation of the USDA-HNIS's expert panel, it seems quite clear that we should not use the 1987-88 NFCS in our analysis. Given that

by Aguiar and Hurst (2005, p. 922, footnote 4) are conditional on the household's participation in the survey.

¹⁹These revisions included moving from a paper and pencil to a computerized system, changing from a manual to an automated coding system, and implementing more systematic methods for probing respondents to recall food items (McDowell et al. 1994). As shown in Appendix Table A2, this change only had modest effects on the NHANES dietary intake response rate which was 74 percent for NHANES I, 73 percent for NHANES II, and 79 percent for NHANES III. Continuous NHANES response rates range between 75 and 80 percent except for 2011-2012 which is 70 percent.

²⁰The NHANES surveys have dramatically lower rates of proxy reporting for dietary intake, i.e., one household member reporting the intake of another household member, than either the NFCS or the 1989-91 CSFII. This difference is explained, in part, by the fact that the NHANES 24 hour recall is collected when respondents undergo the physical exam portion of the survey which requires the respondent to be present. In addition, while proxy reporting was allowed in the 1989-91 CSFII (USDA HNIS 1991b), it was not in the 1994-96 CSFII except for those "who could not report for themselves due to physical or mental limitations" (USDA-ARS 1997, p.54). However, within each survey, we find that the relationship between proxy reporting and calorie levels in the retirement-age sample cannot explain the differences in calorie levels across surveys. The effect of proxy reporting on caloric intake is statistically insignificant for most surveys after including our set of controls. In NFCS87-88, proxy reports are slightly higher than self-reports and in Continuing NHANES, proxy reports are slightly lower. In any case, the estimates of equation (5) across surveys are not appreciably changed by either including a control for proxy reporting or restricting the sample to observations without proxy reporting. These results are available from the authors.

the 1987-88 NFCS and 1989-91 CSFII used the same individual intake survey methodology and the documented impact of the methodological changes on individual intake shown in Figure 2, we would argue that the 1989-91 CSFII should be excluded as well. Returning to Table 2, we report estimates when dropping these low-response-rate studies in column (4). Our 2SLS estimate of the impact of retirement on caloric intake becomes -0.206 with a standard error of 0.034, yielding a statistically significant decline in caloric intake at retirement.

Alternatively, we can focus only on the surveys implemented after the major improvements to survey methodology. As shown in Table 3, the 2SLS estimate in the Continuous NHANES survey is -0.205 and is highly statistically significant. If we combine the 1994-96 CSFII with the Continuous NHANES data as in Table 2 column (5), our 2SLS estimate of the impact of retirement on caloric intake is -0.194 with a standard error of 0.048, again yielding a statistically significant decline in caloric intake at retirement.²¹ Overall, our estimates using cross-sectional data show a significant decrease in caloric intake at retirement.

3.5 Quantile and Quantile IV Estimates

Because prior studies have found heterogeneity in the consumption response at retirement, we also examine the decrease in caloric intake across the intake distribution. Bernheim, Skinner, and Weinberg (2001) find relatively larger decreases in expenditure for those in the lower wealth and income quartiles. Scholz, Seshadri, and Khitatrakun (2006), find that a majority of households have wealth levels which match or exceed those predicted by a stochastic life-cycle model although 20 percent are saving less than their target level. Aguiar and Hurst (2005) find much larger decreases in caloric intake for the 10 percent of their sample they identify as being low wealth. Since we do not have a consistent measure of wealth across our numerous surveys, we instead examine heterogeneity in the expenditure and intake responses at retirement by estimating quantile and quantile instrumental variables regressions.²²

Figure 3 shows quantile and quantile IV estimates for caloric intake in the sample of pooled

²¹If Continuous NHANES and the 1994-96 CSFII are combined with NHANES III, which also made use of a new methodological dietary intake survey instrument as noted above, our 2SLS estimate of the impact of retirement on caloric intake is -0.197 with a standard error of 0.044.

²²We implement the quantile instrumental variable estimator developed by Chernozukhov and Hansen (2005). As before, we use age indicators as instruments for retirement with age dummies.

studies excluding those with low response rates, as in column 4 of Table 2.²³ The two panels show the quantile estimates along with the associated 95 percent confidence intervals for each fifth quantile. The horizontal dashed line represents the corresponding OLS estimate in the left panel and the corresponding 2SLS estimate in the right panel.

In both panels, the caloric intake estimates are relatively constant across the quantiles. Across the distribution, the standard quantile estimates are negative and quite similar to the OLS estimate. The quantile IV estimates are again larger in magnitude but very similar to the corresponding 2SLS estimate. Due to rather wide confidence intervals, we are unable to reject the null hypothesis of no effect on caloric intake at the lowest quantile we examine, although the quantile point estimate remains quite similar to the 2SLS estimate. Overall, the retirement caloric intake response does not vary dramatically across the distribution.

3.6 The Impact of Retirement on Nutrient Intake

In addition to the amount of calories eaten, the transition into retirement may affect the quantity of nutrients that individuals consume. Table 4 presents the estimated impact of retirement on a wide range of nutrients: vitamins A, C, and E, calcium, cholesterol, saturated fat, and protein. These nutrients were selected to match those examined by AH. Each coefficient reported in the Table is the parameter on the retirement indicator found when applying 2SLS to estimate specifications using equation (5).

The results in column (1) of Table 4 use data from pooling all the studies except for the 1987-88 NFCS and the CSFII 89-91 for reasons discussed in the previous subsection.²⁴ Each point estimate shown is for the coefficient on the retirement indicator when estimating equation (5) where the outcome is noted on the left side of each row. As discussed above, we find that caloric intake drops dramatically at retirement when pooling these samples. We also find that all the nutrients either significantly decline at retirement (vitamin E, calcium, cholesterol, saturated fat, and protein) or

²³The low-response surveys we drop in this analysis are the NFCS 1987-88 and CSFII 1989-91. We obtain similar results by restricting to the the post-revision studies used in column 5 of Table 2: CSFII 1994-96 and Continuous NHANES.

²⁴Since information on individual nutrients are not available in all of the studies, the sample sizes vary across each regression as detailed in the notes to Table 4.

yield a positive but insignificant estimate for the impact of retirement (vitamins A and C).²⁵

Of course, changes in income may lead to a different response for different nutrients. To understand variation in the response across nutrients, AH regress intake of each nutrient on a measure of permanent income using a sample of employed individuals.²⁶ They find that intake of vitamins A, C, and E, and calcium increase with permanent income, intake of cholesterol and saturated fat decrease with permanent income, and that there is no effect of permanent income on protein intake. Comparing these permanent income responses to our findings in column (1), the transition into retirement is not consistent with either an increase or a decrease in permanent income.

For their analogous set of results, AH include calories as a control in the specifications. They interpret their findings as providing insight into diet quality holding calories constant. However, since caloric intake changes dramatically at retirement in our data, calories should be treated as an endogenous regressor when it is included as a control. In the absence of a plausible instrument for calories, in column (2) of Table 4 we present results which control for calories with the caveat that these results should be treated cautiously. After controlling for calories, we find positive and significant effects vitamins A and C as well as for calcium. The remaining nutrients do not change significantly at retirement.

The large changes in caloric intake at retirement complicate the interpretation of the calorie-constant results in column (2) of Table 4. For example, we find that the level of calcium intake falls at retirement (column (1)) but that calorie-constant calcium intake increases at retirement (column (2)). Even ignoring the endogeneity of caloric intake, it is difficult to interpret the finding that calorie-constant calcium intake rises at retirement as a beneficial outcome given that total calcium intake falls at retirement. Similarly, the calorie-constant null effect of retirement on protein fails to convey the overall dramatic drop in protein intake at retirement.

In the final two columns of Table 4, we present the nutrient intake results from using both CSFII waves, the same data that is used by AH. In the CSFII sample, the results for each nutrient are invariant to the inclusion of caloric intake as a control, as shown by comparing the results in

²⁵While these surveys collected information on vitamin supplement use, this information pertains to supplement use over the past 30 days and is not recorded as part of the dietary intake data. Beginning in 2007, NHANES collected vitamin supplement data that directly corresponds to the 24 hour recall period. However, the 24 hour supplement use data is contained in separate files from the dietary intake data and is not used here.

²⁶See Table 2 of Aguiar and Hurst (2005) for more details on this procedure.

columns (3) and (4).²⁷ This finding is consistent with the result at the top of column (3) which shows that caloric intake does not change at retirement in the CSFII sample. Given this lack of caloric intake response in the CSFII, the treatment of caloric intake as exogenous by AH makes sense for their analysis. However, since we find that caloric intake falls at retirement, in our data we believe that the results in column (1), which do not control for calories, are the relevant set of findings for understanding whether intake falls at retirement. As such, we view the results in Table 4 as evidence that nutrient intake falls at retirement.

4 Changes in Intake at Retirement: Longitudinal Data

4.1 Data

As discussed in the theoretical section, one concern with the cross-sectional analysis is that the included covariates may not adequately capture heterogeneity in preferences and lifetime resources across households. Ideally, the 2SLS estimates would yield consistent estimates by exploiting variation in retirement driven by incentives that are orthogonal to these unobserved characteristics. However, that analysis relies on the strong identifying assumption that nutrient intake differences across ages are driven by retirement propensities and not other time-varying characteristics such as physiological changes. To address these concerns, we also present results using longitudinal food intake data. Our first difference estimates presented below not only allow us to account for fixed sources of heterogeneity across households, but also subsume a linear rate of change in the level of food intake in the intercept term such that the 2SLS estimates, as discussed below, are identified by non-linearities in the age-retirement profile. This combination of features of the longitudinal analysis further mitigates concerns that the estimated caloric intake response to retirement is biased.

The Multiple Risk Factor Intervention Trial for the Prevention of Coronary Heart Disease (MRFIT) sample of nearly 13,000 men aged 35 to 58 at initial contact was interviewed on multiple occasions beginning in 1973 and extending over the next six to eight years (Multiple Risk Factor Intervention Trial Research Group 1982). The survey obtained demographic, labor market, health

²⁷Our results in column (4) are quite similar to those reported by AH in their Table 2.

exam information, and food intake using 24-hour recall during the initial visit as well as during seven annual follow-up visits. However, not all of this information was collected at each visit which influences the analysis that we are able to perform below.

All participants were asked to provide 24-hour food recall information at the initial visit as well as at the first, second, third, and sixth annual visits. However, direct information on current labor force status, including whether or not the participant is currently retired, was only collected at two points in time: the initial interview and the sixth annual visit. Therefore, our main longitudinal analysis will focus on the change in food intake between the initial visit and the sixth annual visit.

MRFIT was an experiment designed to investigate whether a multiple intervention treatment affected coronary heart disease. The multiple interventions, which were provided as a bundle, included treatment for hypertension, smoking cessation advice, and counseling on lowering cholesterol.²⁸ Since the treatment group received advice on modifying their diet, we focus our primary analysis on the control group, which is roughly half of the sample. To use a sample with ages that are roughly comparable to those in the cross-sectional sample, we restrict our analysis to those who are ages 47 to 58 at the initial visit, which means they are ages 53 to 64 at the sixth annual visit. This restriction cuts our sample roughly in half.

The income information collected in MRFIT during the initial visit and the sixth annual visit is categorical with the values for the nine categories remaining constant over time.²⁹ In our analysis, we follow the convention used in the MRFIT data which is to assign the midpoint of each interval as the income value for all but the lowest and highest categories.³⁰ Since we do not use limited dependent variable methods to analyze the income data, we view our income results as descriptive in that they demonstrate that income falls at retirement in the MRFIT sample.

²⁸More broadly, the third intervention “sought to encourage the development of lifelong shopping, cooking, and eating patterns” (Multiple Risk Factor Intervention Trial Research Group 1982, p.1466). The treatment group was also given recommendations regarding limits on fat intake as a share of total calories.

²⁹The categories were less than \$4,200, \$4,200 to \$7,199, \$7,200 to \$9,999, \$10,000 to \$11,999, \$12,000 to \$14,999, \$15,000 to \$17,999, \$18,000 to \$24,999, \$25,000 to \$34,999, and \$35,000 and over.

³⁰The lowest category is assigned to be \$4,200 while the highest category is assigned \$45,000.

4.2 Empirical Methodology

Based on the Euler Equation (4), we estimate longitudinal specifications

$$\Delta y_{it} = \gamma R_{it} + \Delta \mathbf{X}_{it} \beta + e_{it} \quad (7)$$

where Δy_{it} is the change in caloric intake between the initial visit and the sixth annual visit, $\Delta \mathbf{X}_{it}$ is the change in time-varying individual and household level characteristics over this same time span, and R_{it} is an indicator for whether the respondent retired between the initial visit and the sixth annual visit. By using a first-difference specification, any endogeneity that is attributable to permanent, fixed differences across individuals (e.g., impatient individuals retire sooner) will be eliminated from the analysis and OLS will yield consistent estimates. Given the limited set of time-varying covariates available in MRFIT, $\Delta \mathbf{X}_{it}$ includes changes in the indicator for being married and, in the caloric intake regressions, changes in indicators for high blood pressure, heart disease, stroke, diabetes, and cancer.

If retirement between visits is endogenous, e.g., due to unanticipated changes in health status between surveys, then OLS estimates of equation (7) will be inconsistent. One approach is to account for time-varying factors that may be correlated with caloric intake and the decision to retire. Our specifications which include controls for changes in health conditions are designed to address this concern. However, to account for unobserved factors that may affect the retirement decision, we again instrument for retirement using age indicators.

One subtle but important benefit to note with the first-difference equation (7) is how it accounts for changes in the dependent variable that occur with age. Specifically, if the outcome depends linearly on age in the level equation, then this age dependence will be part of the constant term after applying the first differencing. Thus, both our OLS and 2SLS estimates of the effect of retirement on caloric intake using equation (7) are not simply identified by a linear age-calorie relationship which is a potential concern with the cross-sectional analysis.

4.3 2SLS Estimates

We present our results examining the impact of retirement on outcomes in the MRFIT data in Table 5. We limit our samples to individuals employed at the initial survey so that we are comparing outcomes for those who retire against those who do not. Before turning to the longitudinal analysis, we first present the results of applying the cross-sectional methods used in the previous section to the second wave of the MRFIT data. We do so to compare the findings across the data sources, especially since our MRFIT sample is somewhat younger at the second wave (ages 53 to 64) than the cross-sectional samples (ages 57 to 71).

Panel A of Table 5 presents the results of the cross-sectional analysis using MRFIT data. The first row of the Panel shows that income falls at retirement both when using OLS and 2SLS. The estimates are quite similar in magnitude to what we find for our sample in Table 1.³¹ The second row shows the impact on caloric intake. The OLS estimate is slightly smaller than those found in Table 3. The 2SLS estimates are significant and somewhat larger than the analogous cross-sectional estimates shown in Table 3.³² Thus, even though our longitudinal sample is slightly younger than our cross-sectional sample, the results from applying cross-sectional methods are similar in both the cross-sectional and longitudinal samples.

Our findings from applying the longitudinal methods to the MRFIT data are shown in Panel B of Table 5. In the first row of the Panel, we find that the decrease in income at retirement when estimating the first-difference specification is quite similar to our results from using the cross-sectional approach. Our OLS estimates of the change in caloric intake at retirement in the first-difference specification are also similar to our findings from using the cross-sectional specification and again are insignificant. The 2SLS estimates of the effect of retirement on caloric intake are marginally significant. Overall, our results from using the longitudinal data are quite similar to findings from using cross-sectional data and further indicate that caloric intake falls at retirement.

³¹We also include a indicator for married, black, and four education categories.

³²The corresponding first stage estimates are presented in Appendix Table A3.

5 Changes in a Food Intake Index at Retirement

The results from the previous two sections show that caloric and nutrient intake fall at retirement both in cross-sectional and longitudinal data. Aguiar and Hurst also examine the impact of retirement on a more comprehensive measure of food intake: a consumption index which relates permanent income to individual food items eaten by the respondent. Consistent with their earlier results, they find that this measure of intake does not fall at retirement. Using the same CSFII data as AH, we show that their finding is not robust: a seemingly innocuous change in the method relating permanent income to consumption yields strikingly different results.³³ Intuitively, by first predicting permanent income, their method inadvertently removes some of the variation in permanent income across households which attenuates the coefficients of their consumption index and, in turn, attenuates the relationship between the consumption index and retirement.³⁴

5.1 Baseline Results

AH formally link an index of food intake to permanent income by extending the first order condition found in (3) by assuming that consumption, C , is an aggregator function of J home-produced goods, λ_t can be approximated as a linear projection on permanent income, age, and additional controls, and assuming a CRRA utility function.³⁵ The resulting empirical specification is

$$\begin{aligned} \ln(y^{perm}) = & \beta_0 + \alpha_1 c_1 + \dots + \alpha_J c_J + \beta_S \ln S_t + \beta_{\theta} \theta_t \\ & + \beta_{age} age_t + \beta_{age^2} age_t^2 + \beta_{hours} hours_t + \beta_{hours^2} hours_t^2 + \varepsilon_t \end{aligned} \quad (8)$$

where the y^{perm} is permanent income, c_j variables include information on actual food intake from the 24 hour recalls including calories, eleven nutrients, and nearly 80 food category indicators in

³³Although we have highlighted our concerns with the CSFII data in Section 3.4, we can only apply the AH consumption index methodology to datasets which have continuous income measures. Unfortunately, the only other available datasets with such measures are the NFCS surveys which both were fielded prior to the survey methodology changes discussed above (The Continuing NHANES has categorical income measures). Thus, we limit our consumption index analysis in sections 5 and 6 to the CSFII.

³⁴Although we have focused on the male only sample in the previous sections, we use the same sample as AH in this section to emphasize that the differences between their results and what we find in this section are driven entirely by methodological considerations.

³⁵See the working paper version of their paper for more details (Aguiar and Hurst 2003).

order to capture a broad spectrum of food items, S_t is total household spending on market goods in year t , and θ_t are variables (which may or may not be time-varying) thought to affect utility such as self-reported health and family size and $hours_t$ is weekly hours worked.

AH estimate equation (8) using a sample of full-time workers between the ages of 25 and 55 for whom, by assumption, the first order condition (3) holds. They then use these estimates to produce an index, $\ln \widehat{C}$, which relates food intake to permanent income where $\ln \widehat{C} = \widehat{\alpha}_1 c_1 + \dots + \widehat{\alpha}_{JCJ} + \widehat{\beta}_S \ln S_t + \widehat{\beta}_\theta \theta_t$.³⁶ These estimated parameters can be used to link observed food intake to the permanent income of all individuals ages 57 to 71 including retired individuals for whom current income is likely a poor proxy for lifetime resources. Finally, $\ln \widehat{C}$ is used as the outcome when estimating the cross-sectional equation (5) to test whether the consumption index changes at retirement on the sample of 57 to 71 year olds.

Since permanent income is unobserved, AH choose to construct a measure of permanent income to use as the dependent variable when estimating equation (8). Thus, they first regress the log of observed income, $\ln(y^{observed})$, on a number of observable characteristics thought to be correlated with permanent income including occupation, education, their interactions, as well as sex and race interactions. They use these estimates to predict permanent income for each household, $\ln(\widehat{y^{perm}})$, and then use this predicted value as the dependent variable when estimating equation (8). The results of this last specification are used to construct the consumption index $\ln \widehat{C}_{AH}$.

Alternatively, we can estimate equation (8) by replacing $\ln(y^{perm})$ with the log of observed income, $\ln(y^{observed})$ which is, in turn, the sum of permanent and transitory income.³⁷ Since consumption decisions should be unaffected by transitory income for unconstrained households under the LCPIH, the regressors in (8) are correlated with the permanent component of income but not with the transitory component of income. Thus, using $\ln(y^{observed})$ as the dependent

³⁶AH include total household spending both in equation (8) and when computing $\ln \widehat{C}$ to account for differences in prices paid across households. Also, in the text of their article where AH describe the construction of their index, they do *not* include $\widehat{\beta}_\theta \theta_t$. However, they do include $\widehat{\beta}_\theta \theta_t$ in the programs they kindly provided with which we were able to exactly replicate their published results. As such, we include $\widehat{\beta}_\theta \theta_t$ as part of $\ln \widehat{C}$. We discuss below how dropping $\widehat{\beta}_\theta \theta_t$ from the index affects the analysis.

³⁷In our context, the log of observed income equals the sum of the log of permanent income and the log of transitory income. E.g., in his analysis of buffer stock consumers, Carroll (1997) specifies income in year t as $Y_t = P_t V_t$ where P_t is permanent income (which, in logs, evolves as a random walk with drift) and V_t is a multiplicative transitory income shock with mean one. Taking logs of both sides yields $\ln(Y_t) = \ln(P_t) + \ln(V_t)$, the specification of observed income that we consider here.

variable will generate consistent estimates of the parameters in (8) for the same reasons that AH are able to consistently estimate (8) when predicted permanent income is the dependent variable. These estimates can then be used to construct the consumption index, $\ln \widehat{C}_{ST}$.

Both of these approaches for constructing the consumption index are imperfect. A concern with the construction of $\ln \widehat{C}_{ST}$ is that the transitory component of income may be correlated with the consumption measures. However, since the LCPIH predicts that only the annuity value of transitory income will be spent on current consumption and the average individual in the full-time worker sample is relatively young, the impact of transitory income on $\ln \widehat{C}_{ST}$ should be rather small. On the other hand, the $\ln \widehat{C}_{AH}$ consumption index only uses variation in permanent income that is associated with the characteristics that appear in the permanent income regression. Moreover, it further relies on the extent to which these characteristics are, in turn, correlated with the consumption measures. By understating the true variation in permanent income, this procedure attenuates the estimates in the consumption index, $\ln \widehat{C}_{AH}$, and also attenuates the estimated impact of retirement on consumption.³⁸

The results of estimating equation (5) using both methods of constructing $\ln \widehat{C}$ are shown in the first row of Table 6.³⁹ The result in the first column replicates the estimate of -0.006 found by AH. However, using the alternative approach proposed above to construct $\ln \widehat{C}_{ST}$, we find that the impact of retirement on consumption is -0.078 as shown in the second column. Because the dependent variable in these regressions is generated, we estimate standard errors using 1000 repetitions of two bootstrap procedures. The first, reported in parentheses, is a block bootstrap clustered at the level of primary sampling unit stratified by the two CSFII surveys. The second, reported in brackets, is unclustered and unstratified.⁴⁰ The result using $\ln \widehat{C}_{ST}$ is marginally significant at conventional levels.⁴¹ Further, it is economically meaningful and markedly larger than the $\ln \widehat{C}_{AH}$ estimate.⁴²

³⁸We provide a formal discussion of this issue below.

³⁹As discussed in the notes to Table 6, we drop five observations with zero income from the original AH sample in order to use the exact same sample when creating $\ln \widehat{C}_{AH}$ and $\ln \widehat{C}_{ST}$.

⁴⁰The clustered, stratified bootstrap is intended to account for the sampling design of the CSFII surveys and is analogous to the clustering we employ in our other analyses. The alternate bootstrap without clustering or stratification is the procedure employed by Aguiar and Hurst (2005).

⁴¹The result using $\ln \widehat{C}_{ST}$ has a larger standard error than the estimate using $\ln \widehat{C}_{AH}$. However, that is consistent with the fact that the variability in $\ln \widehat{C}_{AH}$ is attenuated relative to that found in $\ln \widehat{C}_{ST}$ as is shown below.

⁴²As we noted above, the text of the Aguiar and Hurst article suggests that $\widehat{\beta}_\theta \theta_t$ is excluded from their index while

5.2 Understanding the Retirement Impact Across Food Indices

This subsection provides a formal discussion of why first predicting permanent income as opposed to directly using observed income when estimating equation (8) attenuates the impact of retirement on consumption. We can illustrate the difference in the estimated retirement consumption response between these two approaches beginning with a stylized version of the estimating equations. Let

$$\ln(y^{perm}) = \pi_0 + \pi_1 \mathcal{C} + u \quad (9)$$

be an equation that is analogous to (8) but relates permanent income to a single measure of household consumption, \mathcal{C} . The OLS regression of *observed* income on \mathcal{C} , where \mathcal{C} is correlated with permanent but not transitory income, yields

$$\hat{\pi}_1 = \frac{Cov(\ln(y^{observed}), \mathcal{C})}{Var(\mathcal{C})} = \frac{Cov(\ln(y^{perm}), \mathcal{C})}{Var(\mathcal{C})}$$

The resulting consumption index is $\ln \hat{C} = \hat{\pi}_0 + \hat{\pi}_1 \mathcal{C}$.

To generate a consumption index as in AH, we first estimate

$$\ln(y^{perm}) = \delta_0 + \delta_1 W + v \quad (10)$$

where W is a single regressor that is associated with permanent income but uncorrelated with transitory income (e.g., educational attainment). The OLS regression of *observed* income on W yields

$$\hat{\delta}_1 = \frac{Cov(\ln(y^{observed}), W)}{Var(W)} = \frac{Cov(\ln(y^{perm}), W)}{Var(W)}.$$

Comparable to the method used by AH, the predicted measure of permanent income, $\ln(\widehat{y^{perm}}) = \hat{\delta}_0 + \hat{\delta}_1 W$, can be used as the outcome when estimating (9) which results in the OLS estimator for

the programs they provided to us include $\widehat{\beta}_\theta \theta_t$ in the index. We have also estimated the impact of retirement on the consumption indices by excluding these covariates. We find that the impact of retirement on $\ln \widehat{C}_{ST}$ becomes somewhat less negative, changing from -0.078 to -0.056. The impact on $\ln \widehat{C}_{AH}$ is also shifted in a positive direction, changing from -0.006 to 0.009. Thus, we still find a substantially more negative estimate of the impact of retirement on consumption when using our consumption index.

the coefficient on \mathcal{C} being

$$\tilde{\pi}_1 = \frac{Cov(\ln(\widehat{y^{perm}}), \mathcal{C})}{Var(\mathcal{C})} = \frac{Cov(\widehat{\delta}_1 W, \mathcal{C})}{Var(\mathcal{C})} = \widehat{\delta}_1 \frac{Cov(W, \mathcal{C})}{Var(\mathcal{C})} = \widehat{\delta}_1 \cdot \widehat{\omega}_1$$

where $\widehat{\omega}_1$ is the OLS estimator for the coefficient on \mathcal{C} in the equation

$$W = \omega_0 + \omega_1 \mathcal{C} + \varepsilon \quad (11)$$

Thus, the resulting consumption index is $\ln \widetilde{C} = \widetilde{\pi}_0 + \widetilde{\pi}_1 \mathcal{C} = \widetilde{\pi}_0 + \widehat{\delta}_1 \widehat{\omega}_1 \mathcal{C}$.

Specifying the equation relating the consumption index to the binary retirement indicator, R , as

$$\ln C = \lambda_0 + \lambda_1 R + e \quad (12)$$

the IV estimators for the effect of retirement on consumption corresponding to the two approaches for constructing the consumption indexes, where A (e.g., age) is the single instrumental variable, are

$$\widehat{\lambda}_1 = \frac{Cov(\widehat{\pi}_1 \mathcal{C}, A)}{Cov(R, A)} = \widehat{\pi}_1 \cdot \frac{Cov(\mathcal{C}, A)}{Cov(R, A)} \quad \text{and} \quad \widetilde{\lambda}_1 = \frac{Cov(\widetilde{\pi}_1 \mathcal{C}, A)}{Cov(R, A)} = \widetilde{\pi}_1 \cdot \frac{Cov(\mathcal{C}, A)}{Cov(R, A)}$$

The relative estimated impact of retirement on consumption across the two methods for creating the consumption index reduces to

$$\frac{\widehat{\lambda}_1}{\widetilde{\lambda}_1} = \frac{\widehat{\pi}_1 \cdot \frac{Cov(\mathcal{C}, A)}{Cov(R, A)}}{\widetilde{\pi}_1 \cdot \frac{Cov(\mathcal{C}, A)}{Cov(R, A)}} = \frac{\widehat{\pi}_1}{\widehat{\delta}_1 \widehat{\omega}_1} = \frac{\rho_{perm, \mathcal{C}}}{\rho_{perm, W} \cdot \rho_{W, \mathcal{C}}}$$

where $\rho_{perm, \mathcal{C}}$ is the correlation between permanent income and \mathcal{C} , $\rho_{perm, W}$ is the correlation between permanent income and W , and $\rho_{W, \mathcal{C}}$ is the correlation between W and \mathcal{C} .⁴³

To the extent that changes in permanent income are perfectly reflected by changes in consumption, as AH assume when using the first order condition to relate permanent income and consumption, we anticipate that $\rho_{perm, \mathcal{C}} = 1$. At the same time, we have no reason to believe that *both* changes in permanent income are perfectly correlated with changes in the observable charac-

⁴³This simplification makes repeated use of the fact that the slope coefficient from the two variable regression of Y on X equals $\rho_{X, Y} \cdot (SD_Y / SD_X)$ where $\rho_{X, Y}$ is the correlation between X and Y and SD_n is the standard deviation of n .

teristic, W , and changes in the observable characteristic are perfectly correlated with changes in consumption. Thus, $\rho_{perm,C} \geq \rho_{perm,W} \cdot \rho_{W,C}$ and, as such, $\widehat{\lambda}_1 \geq \widetilde{\lambda}_1$ where the inequality is almost certainly non-binding.⁴⁴ Thus, the method for constructing the consumption index implemented by AH almost certainly will yield an estimate of the impact of retirement on consumption that is smaller in magnitude than one generated by the valid alternative proposed above.

We next extend the analysis by replacing the single permanent income instrument, W , with j instruments, W_1, W_2, \dots, W_j . Using the resulting predicted permanent income measure, $\ln(\widehat{y^{perm}})' = \widehat{\delta}'_0 + \widehat{\delta}'_1 W_1 + \dots + \widehat{\delta}'_j W_j = \widehat{\delta}'_0 + \sum_{m=1}^j \widehat{\delta}'_m W_m$, to estimate (9) yields the consumption index $\ln \widetilde{C}' = \widetilde{\pi}'_0 + \widetilde{\pi}'_1 C$ where

$$\widetilde{\pi}'_1 = \frac{Cov(\ln(\widehat{y^{perm}})', C)}{Var(C)} = \frac{Cov(\sum_{m=1}^j \widehat{\delta}'_m W_m, C)}{Var(C)} = \sum_{m=1}^j \widehat{\delta}'_m \frac{Cov(W_m, C)}{Var(C)}$$

Whereas $\widetilde{\pi}_1$ is the product of the OLS estimators resulting from the use of a single income instrument, W , the estimator $\widetilde{\pi}'_1$ is the sum of j products, one for each instrument W_m . Assuming that each W_m is positively correlated with permanent income and is also positively correlated with C , we will find that $\widetilde{\pi}'_1 > \widetilde{\pi}_1$. It follows that, under these conditions, $\widetilde{\lambda}'_1 > \widetilde{\lambda}_1$ and the estimated impact of retirement on consumption will be increase when using multiple instruments for permanent income.

However, each term in $\widetilde{\pi}'_1$ accounts for the relationship between permanent income and consumption only through the correlations of each with the instruments W_m . The analogous estimator, $\widehat{\pi}_1$, from using observed income to construct the consumption index, directly captures the relationship between permanent income and consumption. Again assuming that each W_m is positively correlated with permanent income and is also positively correlated with C , we will find that $\widehat{\pi}_1 > \widetilde{\pi}'_1$.⁴⁵ As such, we expect to find that $\widehat{\lambda}_1 > \widetilde{\lambda}'_1$ whereby the estimated impact of retirement on consumption will be larger when using observed income to construct the consumption index relative to using multiple instruments for permanent income.

⁴⁴For this result to hold, notice that it need not be the case that $\rho_{perm,C} = 1$ but only that the correlation between permanent income and the consumption measure is larger than the product $\rho_{perm,W} \cdot \rho_{W,C}$.

⁴⁵This result will hold except for the unlikely scenario in which the portion of permanent income that is unexplained by the W_m is negatively correlated with the consumption measure.

We can also expand the analysis to allow for k consumption measures, $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$,

$$\ln(y^{perm}) = \pi_0 + \pi_1 \mathcal{C}_1 + \dots + \pi_k \mathcal{C}_k + u \quad (9')$$

Assuming that each consumption measure is correlated with permanent but not transitory income, the OLS estimators for the slope coefficients of (9') when using observed income as the dependent variable are

$$\hat{\pi}_j'' = \frac{\sum_{i=1}^N \hat{r}_{ij} \cdot \ln(y_i^{perm})}{\sum_{i=1}^N \hat{r}_{ij}^2}, \quad j = 1, \dots, k$$

where \hat{r}_{ij} is the residual from regressing \mathcal{C}_j on the remaining consumption measures.⁴⁶ Each estimator, $\hat{\pi}_j''$, captures the relationship between permanent income and consumption measure \mathcal{C}_j . The resulting consumption index is $\ln \hat{C}'' = \hat{\pi}_0'' + \sum_{j=1}^k \hat{\pi}_j'' \mathcal{C}_j$.

If we were to instead use the predicted permanent income measure as the outcome, then in the case with a single W predicting permanent income, the resulting OLS estimators for equation (9') become

$$\tilde{\pi}_j'' = \frac{\sum_{i=1}^N \hat{r}_{ij} \cdot \hat{\delta}_1 W_i}{\sum_{i=1}^N \hat{r}_{ij}^2} = \hat{\delta}_1 \cdot \frac{\sum_{i=1}^N \hat{r}_{ij} \cdot W_i}{\sum_{i=1}^N \hat{r}_{ij}^2} = \hat{\delta}_1 \cdot \hat{\omega}_j'', \quad j = 1, \dots, k$$

where the $\hat{\omega}_j''$'s are the OLS estimators for the coefficients in

$$W = \omega_0 + \omega_1 \mathcal{C}_1 + \dots + \omega_k \mathcal{C}_k + \varepsilon \quad (11')$$

In this case, each estimator, $\tilde{\pi}_j''$, captures the relationship between permanent income and consumption measure \mathcal{C}_n to the extent that it works through the permanent income instrument W . Thus, the resulting consumption index is $\ln \tilde{C}'' = \tilde{\pi}_0'' + \sum_{j=1}^k \tilde{\pi}_j'' \mathcal{C}_j = \tilde{\pi}_0'' + \sum_{j=1}^k \hat{\delta}_1 \hat{\omega}_j'' \mathcal{C}_j$.

Using these indexes comprised of multiple consumption measures as the outcome when estimating equation (12), the resulting IV estimators for the coefficient on retirement are

$$\hat{\lambda}_1'' = \frac{Cov(\ln \hat{C}'', A)}{Cov(R, A)} = \sum_{j=1}^k \hat{\pi}_j'' \frac{Cov(\mathcal{C}_j, A)}{Cov(R, A)} \quad \text{and} \quad \tilde{\lambda}_1'' = \frac{Cov(\ln \tilde{C}'', A)}{Cov(R, A)} = \sum_{j=1}^k \hat{\delta}_1 \hat{\omega}_j'' \frac{Cov(\mathcal{C}_j, A)}{Cov(R, A)}$$

⁴⁶This expression for the OLS estimators when using multiple regression can be found in Wooldridge (2016).

These estimators are different weighted averages of the impact of retirement on each individual measure that enters the consumption index. When using observed income to construct the consumption index, the weights are the relationship between permanent income and consumption measure \mathcal{C}_n . However, when first predicting permanent income, the weights only capture the relationship between permanent income and each consumption measure through their individual relationship with the instrument W . If the direct relationship is stronger than the product of the indirect relationship, then the impact of retirement on consumption using $\ln \widehat{C}_{ST}$ will be larger than that found by using $\ln \widehat{C}_{AH}$.

The remaining rows in Table 6 show how varying the number of permanent income instruments and the number of consumption measures affects the impact of retirement on the the consumption indexes $\ln \widehat{C}_{AH}$ and $\ln \widehat{C}_{ST}$. The second row of the Table shows the results when we repeat the methodology as above except only using a single years of schooling measure to predict permanent income and only including total food expenditure, S , in the consumption index.⁴⁷ Consistent with our derivations, when we first predict permanent income the estimated drop in consumption at retirement of -0.010 (column 1) is nearly ten times smaller in magnitude than our estimate of -0.097 from using observed income. This result confirms that predicting permanent income first removes much of the important variation in permanent income. In the third row, we include the entire set of permanent income instruments used by AH but maintain the single regressor, S , to measure \mathcal{C} . We find, as our analysis anticipates, that the drop in consumption at retirement increases by a factor of nearly five to 0.048 (relative to the second row) due to dramatically increasing the variation in permanent income used in the analysis. Because the construction of $\ln \widehat{C}_{ST}$ does not use the permanent income instruments (W), the estimate in row 3 as the same as in row 2.

The final row in Table 6 uses only educational attainment to predict permanent income but uses the complete set of measures in the consumption index. Using the $\ln \widehat{C}_{AH}$ index we find that the change in consumption at retirement is 0.001 while using the $\ln \widehat{C}_{ST}$ index we again report that the change is -0.078.⁴⁸ As these results are quite comparable to the full specification shown

⁴⁷We include all the remaining covariates in the other regressions. Our formulas in this section are applicable after first transforming these variables to remove the effects of the remaining covariates analogous to an application of the Frisch-Waugh-Lovell Theorem.

⁴⁸Because $\ln \widehat{C}_{ST}$ does not use W , the estimated impact in row 4 is exactly the same as that in row 1.

in the first row of the Table, it suggests that the differences in the estimated impact of retirement on consumption are driven entirely by the weighting of the impact of retirement on each individual consumption measure shown in our expressions for $\widehat{\lambda}_1''$ and $\widetilde{\lambda}_1''$. As we argued above, the weights resulting from the use of $\ln \widehat{C}_{AH}$ will likely understate the impact of retirement on consumption relative to those that arise from using $\ln \widehat{C}_{ST}$.

Overall, our analysis of the estimated impact of retirement on consumption using the constructed consumption index indicates that the method used by AH inadvertently attenuates the estimates. As we have shown empirically, the extent of the attenuation can be quite dramatic as the point estimate of -0.078 that results from using $\ln \widehat{C}_{ST}$ suggests an economically significant reduction in consumption at retirement.

6 The Impact of Unemployment on Intake

The prior literature finds a significant reduction in expenditures for workers who become unemployed (Dynarski and Gruber 1997; Stephens 2001; Saporta-Eksten 2014). Although the unemployed are a heterogeneous group, including those with both short and long-term unemployment spells, the defining characteristic is the intention to find future employment. Thus, unemployment is a transitory state which stands in contrast to retirement. However, the fact that both the unemployed and the retired undergo large earnings losses as well as being non-employed provides a useful point of comparison for understanding how intake responds to changes in earnings.

Indeed, using the two waves of the CSFII, Aguiar and Hurst (2005) find comparable expenditure losses and increases in food preparation time for the retired and the unemployed. However, while they find no effect of retirement on their consumption index, they find a significant drop in the consumption index for the unemployed, suggesting that the index is able to measure consumption responses to income shocks. They interpret these contrasting results as further evidence that intake does not fall at retirement.

To increase the number of unemployed respondent observations in their analysis of the impact of unemployment, AH include the 1989-91 CSFII low-income “oversample” in their analysis. This separate sample was collected by restricting data collection to households with income at or below

130 percent of the federal poverty line (USDA HNIS 1991b). In addition, the sampling procedures targeted the low income sample to be 50% the size of the “basic” sample, which means that one-third of the combined sample is from the low income oversample.

However, AH do not use sampling weights when combining the samples for their unemployment analysis and therefore fail to get a consistent population estimate of the impact of unemployment on intake. In addition, although the 1994-96 CSFII does not have a separate oversample, it “included an oversampling of the low-income population,” (p. 3, USDA-ARS 1997) and, thus, analysis of these data also requires sampling weights. Therefore, we re-examine the effect of unemployment on intake after applying sampling weights and extending the analysis to include the alternative intake index discussed in Section 5.

Table 7 reports estimates from regressions of the the form of

$$y_i = \gamma U_i + \mathbf{X}_i \boldsymbol{\beta} + \epsilon_i \quad (13)$$

which is analogous to equation (5) except it includes the binary unemployment indicator U_i in place of the retirement indicator R_i . Each estimate shown in Table 7 is the coefficient on U_i from a separate regression of (13). We use the same sample that AH use in their unemployment analysis, restricting to prime-age workers who are either full-time employed or unemployed. To investigate the impact of weighting, we show results for the analysis both without (column (1)) and with (column (2)) the use of sampling weights.

We replicate the main AH unemployment findings as shown in column (1) of Table 7. They find a significant decrease in both food expenditure (-0.187) and their food intake index (-0.053) due to unemployment.⁴⁹ As shown in the last row of the column, we also find a substantially larger unemployment decrease when using our alternative food intake index (-0.133), consistent with our analysis in Section 5.

However, although we find a comparable impact on food expenditure when applying sample weights, we find strikingly different estimates for the impact of unemployment on the food intake index. We estimate the impact of unemployment on the AH food intake index to be -0.013 which is

⁴⁹We are able to replicate the unemployment findings reported in the first and last rows of Table 6 in Aguiar and Hurst (2005).

not only much smaller than the unweighted estimate of -0.053, but is also statistically insignificant. We find a much larger estimate of -0.045 when using our alternative food intake index, although this result is also statistically insignificant.

Importantly, the estimated impact of unemployment on the food intake index (-0.045) is smaller in magnitude than our estimated impact of retirement on the food intake (-0.078).⁵⁰ These findings stand in contrast to AH who find no effect of retirement on the food intake index but do find an effect due to unemployment. As our results are based on the use of sampling weights and the alternative food index, we believe that they provide a more appropriate comparison for the estimated retirement and unemployment effects. These estimates indicate a larger consumption response to retirement than to unemployment and provide additional support for the finding that food intake falls at retirement.

7 Conclusion

Evidence accumulated across multiple datasets and countries shows that expenditures fall, on average, at retirement. Important work by Aguiar and Hurst, however, finds that food intake does not fall at retirement. In this paper, we dramatically expand the investigation of nutrient intake changes at retirement by using cross-sectional datasets that span over forty years and by using longitudinal data. We show how the methodology used to create consumption indices in prior work will greatly attenuate the estimated impact of retirement on consumption and how to construct these indices in a way that circumvents this concern. Across multiple methods, we find that nutrient intake falls at retirement.

Our results run counter to those of Aguiar and Hurst (2005) who find that food intake does not fall at retirement. As we document, major revisions to the survey methodology used to collect dietary intake data had substantial effects on both survey response rates and the level of reported intake in the data used in their analysis. Either using studies where these concerns are less pertinent

⁵⁰While our analysis of the impact of retirement on food intake in Table 6 does not use sampling weights, consistent with AH's methodology, we estimate the impact of retirement on the alternative food intake index is also -0.078 when we use sampling weights. The retirement coefficient is likely invariant to the use of sampling weights, whereas the unemployment estimate is not, since the AH retirement analysis does not include the low income oversample which has the large disparity in sampling weights as noted above.

or limiting the analysis to datasets fielded after the improved survey methods were implemented yields consistent evidence that nutrient intake falls at retirement. We provide further evidence by using longitudinal data to estimate specifications analogous to Euler Equations and also find that nutrient intake falls at retirement. While the use of age as an instrument for retirement raises potential concerns, our longitudinal first-difference estimates, which implicitly control for a linear trend in age, yield an estimate which is quite similar in magnitude to our cross-sectional findings. Finally, we show that a consumption index relating the type of food consumed to permanent income also falls at retirement and that this decline is larger than that found for the unemployed.

Another possible concern is that food intake may fall at retirement if individuals expend less energy and, hence, have lower intake requirements. Our 2SLS estimates across both methodological approaches could then be viewed as providing evidence on the magnitude of the planned reduction in intake at retirement. Such a planned drop in intake would not refute the Life-Cycle/Permanent Income Hypothesis. On the other hand, our finding of a significant intake reduction at retirement, in contrast to the AH findings, does not confirm this hypothesis either.

Do our findings contradict evidence of an increase in household production at retirement? We do not believe so. Prior evidence of increased time spent shopping and lower prices paid after retirement indicates that households shift the composition of consumption upon leaving the labor force. However, shifts in time allocation may not be part of an optimal life-cycle plan but rather a response to a fall in resources at retirement brought about by a failure to adequately save. A growing literature has found that not only do a large share of adults lack basic financial literacy but that financial literacy is strongly associated with retirement planning (e.g., see the survey by Lusardi and Mitchell (2014)). Another strand of literature sharpens the foundations of the observed (lack of) planning behavior, as in the boundedly rational agents of Gabaix (2016), who weigh the costs and benefits of being attentive to the retirement planning problem and rationally undersave relative to canonical benchmarks. In a related context, Ganong and Noel (2017) show that consumption expenditure falls at another largely predictable income drop: the exhaustion of unemployment benefits. They find that, out of a broad set of models, the observed patterns are only consistent with an economy partially populated by hand-to-mouth consumers who do not

smooth their consumption through the income drop.⁵¹ This is evidence of similar behavior to that which is reflected by our results. Future research linking financial literacy, retirement planning, and changes in consumption and intake at retirement could provide insight on the extent to which retirement preparedness (or lack thereof) can explain our findings.

References

- Aguiar, Mark and Erik Hurst. (2003) "Consumption vs. Expenditure," National Bureau of Economic Research Working Paper No. 10307.
- Aguiar, Mark and Erik Hurst. (2005) "Consumption vs. Expenditure," *Journal of Political Economy*, 133(5):919-48.
- Aguila, Emma, Orazio Attanasio, and Costas Meghir (2011) "Changes in Consumption at Retirement: Evidence from Panel Data," *The Review of Economics and Statistics*, 93(3):1094-1099.
- Ameriks, John, Andrew Caplin, and John Leahy (2007) "Retirement Consumption: Insights from a Survey," *The Review of Economics and Statistics*, 89(2): 265-274.
- Banks, James, Richard Blundell, and Sarah Tanner (1998) "Is There a Retirement-Savings Puzzle?" *American Economic Review*, 88(4):769-88.
- Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber (2009) "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach," *American Economic Review*, 99(5):2209-2226.
- Becker, Gary S. (1965) "A Theory of the Allocation of Time," *Economic Journal*, 75:299, 493-517.
- Been, Jim, Susann Rohwedder, and Michael Hurd (2017) "Households' consumption spending and home production responses to retirement: Do retirees substitute consumption spending?" *Unpublished Working Paper*.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg (2001). "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review*, 91:4,

⁵¹Fadlon and Laibson (2017) show that, even in the presence of non-optimizing households, the government can use policy levers to ensure that household Euler equations hold on average, effectively substituting for the lack of optimization by some households. However, this result and those that follow require the government to behave in the manner of a rational, benevolent social planner. It is unclear the extent to which existing retirement policy reflects such design.

832-57.

- Blau, David M. (2008) "Retirement and Consumption in a Life Cycle Model," *Journal of Labor Economics*, 26(1):35-71.
- Caggiula, Arlene W., George Christakis, Marilyn Farrand, Stephen B. Hulley, Ruth Johnson, Norman L. Lasser, Jeremiah Stamler, Graham Widdowson (1981) "The Multiple Risk Factor Intervention Trial (MRFIT): IV. Intervention on Blood Lipids," *Preventive Medicine*, 10(4):443-475.
- Carroll, Christopher D. (1997) "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112(1):1-55.
- Chernozhukov, Victor and Christian Hansen (2005) "An IV Model of Quantile Treatment Effects," *Econometrica*, 73(1):245-261.
- Dong, Yingying and Dennis Yang (2014) "Mandatory Retirement and the Consumption Puzzle: Prices Decline or Quantities Decline?" *Unpublished Manuscript*.
- Dynarski, Susan and Jonathan Gruber (1997) "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity*, 1, 229-84.
- Engen, Eric M., William G. Gale, and Cori E. Uccello (1999) "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, 1999:3, 65-187.
- Fadlon, Itzik and David Laibson (2017) "Paternalism and Pseudo-Rationality," National Bureau of Economic Research Working Paper No. 23620.
- Fisher, Jonathan D., David S. Johnson, Joseph Marchand, Timothy M. Smeeding, and Barbara Boyle Torrey (2008) "The Retirement Consumption Conundrum: Evidence from a Consumption Survey," *Economics Letters*, 99:482-485.
- Gabaix, Xavier (2017) "Behavioral Macroeconomics Via Sparse Dynamic Programming," *Unpublished Working Paper*.
- Ganong, Peter and Pascal Noel (2017) "Consumer Spending during Unemployment: Positive and Normative Implications," *Unpublished Working Paper*.
- Gustman, Alan L. and Thomas L. Steinmeier (1998) "Effects of Pensions on Saving: Analysis with Data From the Health and Retirement Study," National Bureau of Economic Research Working Paper No. 6681.
- Haider, Steven J. and Melvin Stephens Jr. (2007) "Is There a Retirement-Consumption Puzzle?"

- Evidence Using Subjective Retirement Expectations” *Review of Economics and Statistics*, 89(2): 247-264.
- Hall, Robert E. (1978) “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, 86(6):971-87.
- Hurd, Michael D. and Susann Rohwedder (2003) “The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement,” National Bureau of Economic Research Working Paper no. 9586.
- Hurd, Michael D. and Susann Rohwedder (2006) “Some Answers to the Retirement-Consumption Puzzle,” National Bureau of Economic Research Working Paper no. 12057.
- Hurd, Michael D. and Susann Rohwedder (2008) “The Retirement Consumption Puzzle: Actual Spending Change in Panel Data,” National Bureau of Economic Research Working Paper no. 13929.
- Interagency Board for Nutrition Monitoring and Related Research. (2000) *Nutrition monitoring in the United States: The directory of Federal and State nutrition monitoring and related research activities*. Bialostosky K, ed. Hyattsville, Maryland: National Center for Health Statistics.
- Li, Hongbin, Xinzheng Shi, and Binzhen Wu (2014) “The Retirement Consumption Puzzle Revisited: Evidence from the Mandatory Retirement Policy in China,” *Unpublished Manuscript*
- Lillard, Lee, James P. Smith, and Finis Welch (1986) “What Do We Really Know about Wages? The Importance of Nonreporting and Census Imputation,” *Journal of Political Economy*, 94(3):489-506.
- Lundberg, Shelly, Richard Startz, and Steven Stillman (2003). “The Retirement-Consumption Puzzle: A Marital Bargaining Approach,” *Journal of Public Economics*, 87(5-6):1199-1218.
- Lusardi, Annamaria and Olivia S. Mitchell (2007) “Baby Boomer Retirement Security: The Role of Planning, Financial Literacy, and Housing Wealth,” *Journal of Monetary Economics*, 54(1):205-224.
- Lusardi, Annamaria and Olivia S. Mitchell (2014) “The Economic Importance of Financial Literacy: Theory and Evidence,” *Journal of Economic Literature*, 52(1), 5-44.
- McDowell MA, BriefeI RR, Alaimo K et al. (1994) “Energy and Macronutrient Intakes of Persons Ages 2 Months and Over in the United States: Third National Health and Nutrition Examination

- Survey, Phase 1, 1988-91,” *Advance Data from Vital and Health Statistics, No. 255*. Hyattsville, Maryland: National Center for Health Statistics.
- Minicaci, Raffaele, Chiara Monfardini, and Guglielmo Webber (2010) “How Does Consumption Change Upon Retirement?” *Empirical Economics*, 38:257-280.
- Moore, James F. and Olivia S. Mitchell (1997) “Projected Retirement Wealth and Savings Adequacy in the Health and Retirement Study,” National Bureau of Economic Research Working Paper No. 6240.
- Multiple Risk Factor Intervention Trial Research Group (1982) “Multiple Risk Factor Intervention Trial. Risk Factor Changes and Mortality Results,” *The Journal of the American Medical Association*, 248(12):1465-77.
- Olafsson, Arna and Michaela Pagel (2018) “The Retirement-Consumption Puzzle: New Evidence from Personal Finances,” National Bureau of Economic Research Working Paper No. 24405.
- Parker, Jonathan (1999) “Spendthrift in America? On Two Decades of Decline in the United States Saving Rate,” in B. Bernanke and J. Rotemberg, eds., *NBER Macroeconomic Annual 1999*.
- Poterba, James M. (2014) “Retirement Security in an Aging Population,” *American Economic Review*, 104(5): 1-30.
- Reese, Robert B. (1977) *The 1977-78 Nationwide Food Consumption Survey* [Memorandum]. Washington, DC: US Department of Agriculture, Agricultural Research Service.
- Saporta-Eksten, Itay (2014) “Job Loss, Consumption and Unemployment Insurance,” *Unpublished Working Paper*.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006) “Are Americans Saving ‘Optimally’ for Retirement?” *Journal of Political Economy*, 114(4): 607-643.
- Smith, Sarah (2006) “The Retirement-Consumption Puzzle and Involuntary Early Retirement: Evidence from the British Household Panel Survey,” *The Economic Journal*, 116:C130-C148.
- Stephens Jr., Melvin (2001) “The Long-Run Consumption Effects of Earnings Shocks,” *The Review of Economics and Statistics*, 83(1):28-36.
- Stephens Jr., Melvin and Takashi Unayama (2012) “The Impact of Retirement on Household Consumption in Japan,” *Journal of The Japanese and International Economies*, 26:62-83.
- Tippett, Katherine S., Sharon J. Mickle, Joseph D. Goldman, et al. (1995) *Food and Nutrient*

- Intake by Individuals in the United States, 1 Day, 1989-91. Continuing Survey of Food Intakes by Individuals, 1989-91*, Nationwide Food Surveys Rep. No. 91-2.
- US Department of Agriculture (1994) *1992 Report on USDA Human Nutrition Research and Education Activities: A Report to Congress*, Washington, DC.
- US Department of Agriculture, Agricultural Research Service (1997) *Design and Operation: The Continuing Survey of Food Intakes by Individuals and the Diet and Health Knowledge Survey, 1994-96*, NFS Report No.96-1, Washington, DC.
- US Department of Agriculture, Human Nutrition Information Service (1991a) *Impact of Nonresponse on Dietary Data From the 1987-88 Nationwide Food Consumption Survey*, Washington, DC.
- US Department of Agriculture, Human Nutrition Information Service (1991b) *The Continuing Survey of Food Intakes by Individuals and the Diet and Health Knowledge Survey, 1989, Survey Operations Report*, Washington, DC.
- US General Accounting Office (1991) *Nutrition Monitoring: Mismanagement of Nutrition Survey Has Resulted in Questionable Data*, GAO/RCED-91-117, Washington, DC.
- Wooldridge, Jeffrey M., (2016) *Introductory Econometrics : A Modern Approach, 6E*. Mason, Ohio :South-Western Cengage Learning.

8 Data Appendix

8.1 NFCS 1977-78

The 1977-78 Nationwide Food Consumption Survey was a survey of American eating patterns conducted from the second quarter of 1977 through the first quarter of 1978. The survey was fielded by National Analysts, Inc. under contract with the United States Department of Agriculture (USDA). Approximately 15,000 households in the contiguous US were selected to participate in the main arm of the survey through a stratified, multistage sampling design. In addition to this representative sample, the survey also included oversamples of the elderly and low-income populations and samples from Alaska, Hawaii, and Puerto Rico. Basic interviews were conducted in-home, with the main meal planner and food preparer answering questions about household food either consumed or disposed of over the previous seven days. Households were contacted in advance requesting that they maintain records, if necessary, for the prior week's household food consumption. Household data were also collected on demographics, location, lodging, income, shopping behavior, usual food expenditure, participation in food programs, and patterns of eating at and away from home.

Labor force data were collected for the male and female heads of household. We code heads of household as retired if they did not work during the week prior to the interview and they reported being retired. The day of the initial interview, all individual household members completed a 24-hour food recall for the day prior to the interview. They were also given two-day food diaries to be completed for the interview day and the following day. In the first quarter of the survey, all household members completed individual food records. In the final three quarters, they were completed by all members aged 18 or younger and half of adults aged 19 or older. A few days after initial interviews, interviewers returned to collect diaries, check for errors or inconsistencies, and provide compensation. The main portion of the survey was intended to be self-weighting, but completion rates varied significantly across PSUs. Weights were calculated at the PSU-by-quarter level to recover the expected number of completed intake schedules.

8.2 NFCS 1987-88

The 1987-88 Nationwide Food Consumption Survey collected data on American food consumption between April 1987 and August 1988. The survey was conducted by National Analysts, Inc. under contract with the USDA. Households from the contiguous US were selected to participate through a stratified, multistage sampling design. Interviewers visited homes to conduct the main portion of the interview, in which the main meal planner and food preparer answers questions about household food either consumed or disposed of over the previous week. Households were contacted in advance requesting that they maintain records, if necessary, for the prior week's household food consumption. Household data were collected on demographics, location, lodging, income, shopping, usual food expenditure, participation in food programs, and patterns of eating at and away from home.

The survey asked what each of the household members aged 15 and older were doing most during the prior week. A follow-up question asked whether those who did not report working performed any work in the prior week. We code individuals as retired if their response to the first question was "retired" and their response to the follow-up question was that they did not work at all. Household members performed 24-hour food recall with the interviewer at the time of the in-home interview. They were then given two-day food diaries and instructed in their use for completing intake data for the interview day and the following day. Interviewers later returned one more time to collect the diaries and provide compensation. Weights were developed to match the population mean values of household size and 14 characteristics related to household food consumption.

8.3 CSFII 1989-91

The 1989-91 Continuing Survey of Food Intakes by Individuals surveyed Americans on their food consumption from April of 1989 through March of 1992. The survey was conducted by National Analysts, Inc. under contract with the USDA. The 1989-91 Diet and Health Knowledge Survey (DHKS) was administered to household main meal preparers at the same time. A stratified, multistage sampling design selected households for a main sample and an oversample of low-income households. Interviewers visited every address in person and attempted to meet with the main meal

preparer. A household questionnaire collected data on demographics, location, lodging, income, shopping, usual food expenditure, participation in food programs, and patterns of eating at and away from home.

The survey asked what each of the household members aged 15 and older were doing most during the prior week. A follow-up question asked whether those who did not report working performed any work in the prior week. We code individuals as retired if their response to the first question was “retired” and their response to the follow-up question was that they did not work at all. Household members performed 24-hour food recall with the interviewer at the time of the in-home interview. They were then given two-day food diaries and instructed in their use for completing intake data for the interview day and the following day. Interviewers later returned one more time to collect the diaries and provide compensation. Inverse probability of selection was computed for each household, then adjusted to account for nonresponse at the area segment level. This was done individually for basic and low-income samples. The samples were then combined and weights were adjusted using regression techniques to match a number of population characteristics.

8.4 CSFII 1994-96

The 1994-96 Continuing Survey of Food Intakes by Individuals initially surveyed Americans on their food consumption initially from the beginning of 1994 through the end of 1996. A later fielding in 1998 sampled only children aged nine or under. The survey was conducted by Westat, Inc. under contract with the USDA. In each household, one adult who had completed dietary intake data was surveyed as part of the 1994-96 DHKS. A stratified, multistage sampling design selected households for a main sample and an oversample of low-income households. Interviewers visited every address in person and attempted to meet with the main meal preparer. A household questionnaire collected data on demographics, location, lodging, income, shopping, usual food expenditure, participation in food programs, and patterns of eating at and away from home.

Household members aged 15 or older were asked if they worked at all in the previous week. If they did not work and did not report being temporarily absent from work, they were asked the reason for not working. We code individuals as retired if they did not work during the previous week and were described as retired when asked for a reason. Household members performed a

24-hour food recall with the interviewer at the time of the first in-home interview. A second day of 24-hour recall data was collected between three and ten days after the first day of recall while avoiding the same day of the week as the first recall day. In a small number of cases (3% in 1994), the second day of recall data was collected by telephone. Base weights were calculated as the inverse of sampling probability, where sampling probability incorporated the probability of selecting the PSU, the segment, the household, retaining the dwelling, and the probability of selecting a sample person in the household. Within sample person classes (given by grouping segments by region, MSA status, percent of the segment that was black or Hispanic, and quarter of operations), weights were increased to correspond to nonresponse. Weights were then scaled so the sum of weights matched the March CPS totals within a number of demographic, geographic, timing, and labor force cells.

8.5 NHANES I

The first National Health and Nutrition Examination Survey (NHANES I) was a detailed study of the health and dietary intake of Americans conducted by the National Center for Health Statistics (NCHS) from 1971 to 1975. The target population was the civilian, noninstitutionalized population ages 1 to 74 in the contiguous US, excluding Native American reservations. Low-income individuals, young children, women of childbearing age, and the elderly were oversampled. Households and sample persons were selected using a stratified, multistage sampling design. Data was collected using three Mobile Examination Centers (MECs) that moved between stands associated with primary sampling units over the survey period. The MECs stayed at each stand from three to six weeks and examined 300 to 600 people. The main dietary and health portion of the survey was fielded over two 2-year cycles from 1971 to 1974. A subset of respondents aged 25 to 74 completing the main portion of the survey also completed a more detailed supplemental health examination. Additional respondents examined in 1974 and 1975 completed only the detailed health examination and not the dietary survey.

Initial screening was completed and household data were collected during in-home visits. As part of the household questionnaire, data was collected on the working activity of adults aged 17 and over. Specifically, respondents were coded as working, keeping house, or doing “something else” over the prior three months. An additional question was asked about those who were doing

something else, allowing for these respondents to be reported as retired. Follow-up questions probed whether respondents participated at all in the labor force during the prior week. We code individuals as retired if they were reported as being retired and did not perform any work, have a job, or look for work during the prior week. If household members were selected for the full examination, appointments were made for these sample persons to visit the MECs. In addition to completing a detailed health examination, sample persons also recorded food intake during their MEC visits. A small number of respondents provided their intake data when interviewed at home. Both a 24-hour dietary recall and a food frequency questionnaire were completed by most respondents. Sampling weights were constructed to account for nonresponse and adjust the sample to match the US population in age-sex-race groups at the survey midpoint.

8.6 NHANES II

The second National Health and Nutrition Examination Survey (NHANES II) was a detailed study of the health and dietary intake of Americans conducted by the NCHS from 1976 to 1980. The design of the survey was broadly similar to that of NHANES I. Dietary data was collected for all sample persons and a detailed health questionnaire and examination was completed for those over the age of 12. The target population was the civilian, noninstitutionalized population of all 50 states between six months and 74 years of age. Children and households living in poverty were oversampled. Households and sample persons were selected using a stratified, multistage sampling design. Data was collected using three MECs that moved between stands associated with primary sampling units over the survey period. The MECs stayed at each stand from four to six weeks and examined 300 to 600 people.

Initial screening was completed and household data were collected during in-home visits. As part of the household questionnaire, data was collected on the working activity of adults aged 17 and over. Specifically, respondents were coded as working, keeping house, or doing something else over the prior twelve months. An additional question was asked about those who were doing something else, allowing for these respondents to be reported as retired. Follow-up questions probed whether respondents participated at all in the labor force during the prior two weeks. We code individuals as retired if they were reported as being retired and did not perform any work, have a

job, or look for work during the prior two weeks. If household members were selected for the full examination, appointments were made for these sample persons to visit the MECs. In addition to completing a detailed health examination, sample persons also recorded food intake during their MEC visits. Both a 24-hour dietary recall and a food frequency questionnaire were completed by most respondents. Sampling weights were constructed to account for nonresponse and adjust the sample to match the US population in age-sex-income groups at the survey midpoint.

8.7 NHANES III

The third National Health and Nutrition Examination Survey (NHANES III) was a detailed study of the health and dietary intake of Americans conducted by the NCHS through a contract with Westat, Inc. from 1988 to 1994. Dietary data was collected for all sample persons and a detailed health questionnaire and examination was completed for those over the age of 17. The target population was the civilian, noninstitutionalized population of all 50 states older than two months. Areas with high shares of black and Mexican-American populations were oversampled, as were young children and elderly individuals. Households and sample persons were selected using a stratified, multistage sampling design. The survey was fielded in two three-year phases, allowing for individual estimates of outcomes for each phase. Data was collected using three MECs that moved between stands associated with primary sampling units over the survey period. The MECs stayed at each stand from four to six weeks and examined 300 to 600 people.

Initial screening was completed and household data were collected during in-home visits. Household questionnaires were given and examination appointments were made. As part of the household adult questionnaire, respondents were asked about their labor force status during the prior 12 months. We code respondents as retired if that was the given status and they did not work during the prior two weeks, did not have a job during the prior two weeks, and were not on layoff or looking for a job. A food frequency questionnaire was completed by individuals 17 and older in the home interview. If household members were selected for the full examination, appointments were made for these sample persons to visit the MECs. A small proportion of individuals were examined at home. Sample persons provided 24-hour food recall during the MEC visit. In the first three-year phase, a subset of these respondents aged 50 and older were also eligible to participate

in a follow-up telephone 24-hour recall data collection. Base weights were generated from the inverse probability of sampling. These were adjusted for nonresponse, generally by weighting classes defined by age group, sex, race-ethnicity, income, SMSA residence, and region. Final weights were further adjusted to match the total population in each age-sex-ethnicity group at the midpoint of the study.

8.8 Continuous NHANES

The Continuous National Health and Nutrition Examination Survey (Continuous NHANES) is a survey of the health and diet of Americans conducted by the NCHS through a contract with Westat, Inc. It was born out of two-year NHANES cycles starting in 1999. In 2003, the program was combined with the CSFII to officially become the Continuous NHANES that is ongoing in two-year cycles as of 2016. Households and sample persons are selected using a stratified, multistage sampling design for each cycle. Although each individual year is a representative sample of the population, data is released in two-year cycles because of the imprecise estimates and potential loss of confidentiality associated with single years of data. Non-Hispanic black persons have always been oversampled. Low-income white persons have been oversampled since 2000 Mexican-American persons and pregnant women were oversampled in the first four two-year cycles (1999-2006). Hispanic persons have been oversampled since 2007. Persons aged 70 and over were oversampled in the first four two-year cycles after which persons aged 80 and over have been oversampled. An oversample of adolescents aged 12 to 19 was discontinued in 2007.

A screener and basic questionnaire are completed as part of initial in-home visits. A randomizing computer algorithm selects sample persons from household rosters. Sample individuals complete a questionnaire covering demographic, dietary, socioeconomic, and health topics. A questionnaire covering consumer behavior has been included since the 2007 survey cycle. Family questionnaires are given for any family unit including a sampled person. The data do not include a flag for household head, so we count individuals as household head if their gender, age, education, marital status, and place of birth match those recorded for the household head. Data on labor force activity is recorded for respondents aged 16 and over. We code the individuals as retired if the given reason is retired and the respondent is not working, looking for work, or with a job but absent. Following

the in-home interview, selected sample persons make appointments to visit MECs for detailed examination. Intake data is recorded using a 24-hour recall in the MEC. Since the 2003 survey cycle, a second recall has been completed over the phone 3 to 10 days after examination. Weights are calculated by survey cycle and are designed to account for the probability of selection, the rate of nonresponse, and the differences between the final sample and the total population.

8.9 MRFIT

The Multiple Risk Factor Intervention Trial for the Prevention of Coronary Heart Disease (MRFIT) was a randomized controlled trial aimed at studying the effectiveness of counseling men at high risk for coronary heart disease to change to modifiable risk factors. Screening of potential participants began in the early 1970s and ultimately produced 12,866 men who appeared to be at high risk for coronary heart disease but could possibly lower their risk through modified behavior. Nearly half the sample was assigned to a treatment program, in which they were counseled to follow particular nutritional guidelines, cease smoking, and manage hypertension through weight maintenance and medication. The remaining participants, the control sample, was instructed to continue any standard care through their physicians or other providers and were not given specific counseling or instructions on behavior through the study. We restrict our main analysis to this latter (control) group.

Individuals in both the treatment and control groups participated in an initial screening as well as follow-up visits which occurred at annual intervals. At these visits, respondents completed questionnaires covering various activities and behaviors and also participated in a number of tests to record biomarker data. At the initial screening, as well as the visits that occur one, two, three, and six years after program enrollment, participants also completed 24-hour recall questionnaires covering food consumption during the prior day.

Current retirement status was asked directly at the initial screening visit and the sixth annual visit. At the initial screening, in response to a question whether the participant holds two or more jobs, the possible responses are yes, no, and retired. At the sixth annual visit, respondents were asked to describe their present job status as “working at a job full-time for pay,” “working at a job part-time for pay,” or “unemployed.” Those who reported “unemployed” were then asked an

additional question allowing themselves to report as one of “temporarily laid off,” “temporarily disabled,” “permanently disabled,” “retired,” or “other.”

We use five health conditions as controls in the analysis: high blood pressure, heart disease, stroke, diabetes, and cancer. At baseline, participants are asked if a doctor has ever told them that they have each of these conditions (plus a number of additional conditions). For the remaining annual visits, participants are asked if a doctor told them in the last twelve months that they have each condition. We construct a cumulative measure of having been ever told that you have a condition equals one if participants state that they were affirmatively told they had the condition at baseline or if they responded yes to being told they had the condition in the last twelve months either at the current annual visit or at any of the prior annual visits.

Table 1: Impact of Retirement on Income and Food Expenditure in Cross-Sectional Data

$y_i = \gamma R_i + \mathbf{X}_i\boldsymbol{\beta} + \epsilon_i$				
	Log Income		Log Food Expenditure	
	OLS	2SLS	OLS	2SLS
Retired	-0.256	-0.510	-0.081	-0.227
	(0.032)	(0.066)	(0.026)	(0.067)

Notes - The expenditure estimation sample is restricted to male heads of household ages 57 to 71 amounting to 4,382 observations from the CSFII 1989-91 and 1994-96, NFCS 1977-78 and 1987-88, and NHANES 2007-12. The NHANES data are excluded from the income regressions because only categorical income information is available, as are any respondents from the other surveys that are missing income data. The remaining 2,983 men compose the income estimation sample. The regressors included in \mathbf{X}_i are indicators for black, four education groups, nine household size groups, five self-reported health groups, and year-by-survey controls. Exact year is not reported in the public Continuous NHANES data, so two-year cycles are used instead. Estimates are generated using sampling weights. Retirement is instrumented with age dummies in 2SLS regressions. The first stage F statistics are 41.9 for the income regression and 48.1 for the expenditure regression. Huber/White/sandwich standard errors are clustered at the survey*PSU level.

Table 2: Impact of Retirement on Caloric Intake in Pooled Cross-Sectional Data

	AH Result (1)	Adjusted Replication (2)	Pooling Studies (3)	Dropping Low Response Studies (4)	Post-Revision Studies Only (5)
<i>A. OLS</i>					
Retired		-0.048 (0.021)	-0.060 (0.013)	-0.059 (0.015)	-0.047 (0.020)
<i>B. 2SLS</i>					
Retired	-0.02 (0.03)	-0.021 (0.047)	-0.175 (0.031)	-0.206 (0.034)	-0.194 (0.048)
<i>N</i>	<i>2,052</i>	<i>1,654</i>	<i>9,610</i>	<i>8,513</i>	<i>3,429</i>

Notes - The estimation sample is restricted to household heads ages 57 to 71. Each coefficient is estimated from a different regression of the form of equation (5) with log calories consumed as the dependent variable. Column (1) shows the result from Table 2 of Aguiar and Hurst (2005) which uses the CSFII 1989-91 and CSFII 1994-94 surveys. Column (2) is our replication of the Aguiar and Hurst result in column (1) using the same survey data except restricting to male heads of household, weighting the regressions with sampling weights, including the CSFII 1989-94 low-income sample, and omitting the urban, region, and individual health condition indicators used by Aguiar and Hurst. Column (3) is our estimate using all the cross-sectional datasets discussed in the text. Column (4) uses the same data as in column (3) except for the NFCS 1987-88 and CSFII 1989-91. Column (5) only uses the CSFII 1994-96 and the Continuous NHANES. The regressors included in \mathbf{X}_i are indicators for black, four education groups, nine household size groups, year-by-survey, and self-reported health status. The 2SLS results use age dummies as instruments for retirement. The first stage estimates are shown in Figure 1. The first stage F statistics are 42.4, 24.2, 56.4, 48.4, and 24.1 for columns (1) through (5), respectively. In columns (2)-(5), the estimates are generated using sampling weights and we use Huber/White/sandwich standard errors that are clustered at the survey*PSU level.

Table 3: Impact of Retirement on Caloric Intake in Individual Cross-Sectional Surveys

	NHANES I (1)	NHANES II (2)	NFCS 1977-78 (3)	NFCS 1987-88 (4)	NHANES III (5)	CSFII 1989-91 (6)	CSFII 1994-96 (7)	Continuous NHANES (8)
<i>A. OLS</i>								
Retired	-0.172 (0.048)	-0.043 (0.021)	-0.087 (0.020)	-0.090 (0.040)	-0.050 (0.030)	-0.039 (0.036)	-0.059 (0.024)	-0.044 (0.023)
<i>B. 2SLS</i>								
Retired	-0.180 (0.087)	-0.237 (0.050)	-0.203 (0.035)	0.030 (0.088)	-0.221 (0.077)	0.034 (0.073)	-0.089 (0.059)	-0.205 (0.055)
<i>N</i>	<i>570</i>	<i>1,938</i>	<i>1,181</i>	<i>490</i>	<i>1,395</i>	<i>607</i>	<i>1,047</i>	<i>2,382</i>

Notes - The estimation sample is restricted to male household heads ages 57 to 71. The regressors included in \mathbf{X}_i are indicators for black, four education groups, nine household size groups, year-by-survey, and self-reported health status. The 2SLS results use age dummies as instruments for retirement. The first stage estimates are shown in Figure 1. The first stage F statistics are 24.2, 28.1, 30.3, 13.2, 19.5, 13.0, 23.5, and 18.5 for columns (1) through (8), respectively. The estimates are generated using sampling weights. Huber/White/sandwich standard errors are clustered at the survey*PSU level.

Table 4: Impact of Retirement on Nutrient Intake in Cross-Sectional Data

Surveys Used:	All Surveys Except NFCS 87-88 and CSFII 89-91		Both Waves of CSFII	
	(1)	(2)	(3)	(4)
Log Calories	-0.206 (0.034)		-0.021 (0.047)	
Log Vitamin A	0.068 (0.082)	0.215 (0.080)	0.366 (0.130)	0.383 (0.117)
Log Vitamin C	0.122 (0.099)	0.276 (0.097)	0.328 (0.120)	0.348 (0.105)
Log Vitamin E	-0.174 (0.078)	-0.009 (0.072)	0.108 (0.090)	0.131 (0.072)
Log Calcium	-0.080 (0.045)	0.088 (0.040)	0.048 (0.076)	0.067 (0.059)
Log Cholesterol	-0.192 (0.064)	-0.026 (0.060)	-0.061 (0.073)	-0.042 (0.066)
Log Saturated Fat	-0.236 (0.059)	-0.043 (0.039)	-0.119 (0.064)	-0.094 (0.044)
Log Protein	-0.189 (0.036)	-0.038 (0.028)	-0.082 (0.051)	-0.065 (0.031)
Calories Control?	No	Yes	No	Yes

Notes - Each point estimate shown is the estimated coefficient on the retired indicator from a separate 2SLS regression of equation (5) using age dummies as instruments for retirement. The estimation sample is restricted to male household heads ages 57 to 71. Not all nutrients are observed in all surveys in columns (1) and (2), so some outcomes are estimated using restricted samples. In particular, Vitamin E and Saturated Fat are not recorded in NHANES I, NHANES II, or NFCS 1977-78. Cholesterol is also not recorded in NFCS 1977-78. Accordingly, while the sample size for most outcomes in columns (1) and (2) is 8,513, the sample size for Log Vitamin E and Log Saturated Fat is 4,824, and the sample size for Log Cholesterol is 7,328. The sample is consistently 1,654 across all outcomes in columns (3) and (4). The regressors included in \mathbf{X}_i are indicators for black, four education groups, nine household size groups, year-by-survey, and self-reported health status. The first stage F statistics for column (1) are 48.4 for most outcomes, but are 29.1 for Log Vitamin E and Log Saturated Fat and 40.7 for Log Cholesterol. For the column (2) estimates, first stage F statistics are 47.9 for most outcomes, but are 29.0 for Log Vitamin E and Log Saturated Fat and 40.3 for Log Cholesterol. The first stage F statistics are all 24.2 for column (3) and 24.6 for column (4). The estimates are generated using sampling weights. Huber/White/sandwich standard errors are clustered at the survey*PSU level.

Table 5: Impact of Retirement on Income and Caloric Intake in MRFIT

<i>A. Cross-Sectional Analysis</i>		
	$y_i = \gamma R_i + \mathbf{X}_i \boldsymbol{\beta} + \epsilon_i$	
	OLS	2SLS
Outcome:		
Log Income	-0.212 (0.025)	-0.584 (0.094)
Log Caloric Intake	-0.020 (0.025)	-0.389 (0.084)
<i>B. Longitudinal Analysis</i>		
	$\Delta y_{it} = \gamma R_{it} + \Delta \mathbf{X}_{it} \boldsymbol{\beta} + e_{it}$	
	OLS	2SLS
Outcome:		
Log Income	-0.231 (0.028)	-0.492 (0.107)
Log Caloric Intake	-0.016 (0.025)	-0.174 (0.097)

Notes - The sample is restricted to individuals from the control sample ages 47 and older at the initial visit. The income and caloric intake regressions use 2,532 and 2,609 observations, respectively. The standard errors are clustered at the clinic level (22 clinics). Retirement is instrumented with age dummies in 2SLS regressions. Panel A: The cross-sectional analysis of MRFIT uses observations in the second wave which is six years after baseline. All cross-sectional regressions include indicators for married, black, and four education categories. The health controls included in the caloric intake regressions are indicators for high blood pressure, heart disease, stroke, diabetes, and cancer. Panel B: The longitudinal analysis of MRFIT examines changes between the baseline period and the observation six years later. All first difference regressions include the change in the indicator for married. The health controls used in the caloric intake regressions are changes in indicators for high blood pressure, heart disease, stroke, diabetes, and cancer.

Table 6: Impact of Retirement on a Food Intake Index, 2SLS Estimates

	Consumption Index:	
	$\ln \widehat{C}_{AH}$ (1)	$\ln \widehat{C}_{ST}$ (2)
A. Full Specification	-0.006 (0.024) [0.025]	-0.078 (0.040) [0.044]
B. One W , One \mathcal{C}	-0.010 (0.003) [0.004]	-0.097 (0.025) [0.027]
C. Full W , One \mathcal{C}	-0.048 (0.013) [0.014]	-0.097 (0.025) [0.027]
D. One W , Full \mathcal{C}	0.001 (0.021) [0.021]	-0.078 (0.040) [0.044]

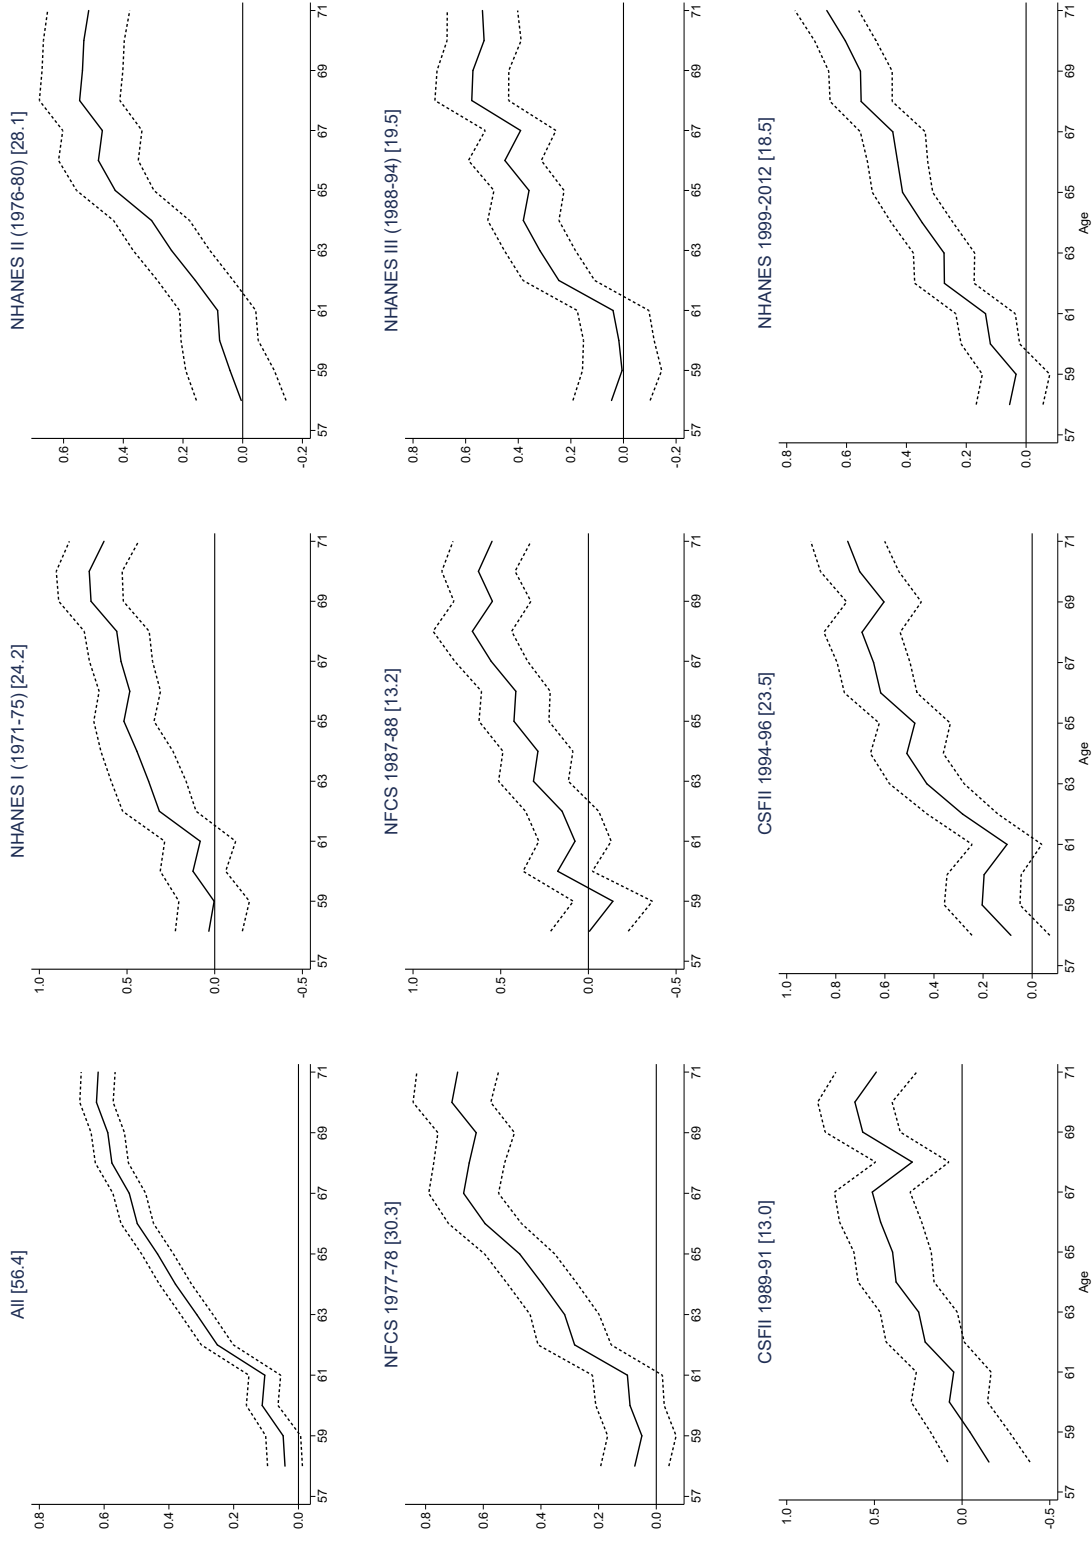
Notes - Each point estimate is an estimated coefficient on a retirement indicator from a separate 2SLS regression with a consumption index as the dependent variable and age indicators as the excluded instruments. All regressions also include indicators for male, black, four education groups, nine household size groups, year, urban residency, and Census region. The sample for these regressions is 2,052 household heads ages 57 to 71 from the CSFII 1989-91 and CSFII 1994-96, the composition of which is discussed in detail Aguiar and Hurst (2005). The construction of the dependent variables is detailed in Section 5. The regressions of permanent income on consumption categories, using equation (8), are estimated on a sample of 2,961 full-time employed household heads ages 25 to 55 which is five fewer observations than are used by Aguiar and Hurst (2005). This discrepancy is due to the dropping of five household heads with zero income such that there are identical samples used in columns (1) and (2); column (1) would be exactly the same to the third decimal place if these observations were included. Standard errors in parentheses are generated by 1000 repetitions of a bootstrap stratified by survey and clustered at the level of primary sampling unit. Standard errors in brackets are generated by 1000 repetitions of an unstratified, unclustered bootstrap.

Table 7: Impact of Unemployment on Intake, 2SLS Estimates

	Unweighted (1)	Weighted (2)
Log Expenditure	-0.187 (0.034)	-0.146 (0.063)
Log ln \widehat{C}_{AH}	-0.053 (0.011)	-0.013 (0.018)
Log ln \widehat{C}_{ST}	-0.133 (0.025)	-0.045 (0.035)


Notes - Each point estimate is an estimated coefficient on a retirement indicator from a separate 2SLS regression with a consumption index as the dependent variable and age indicators as the excluded instruments. All regressions also include indicators for sex, black, four education categories, nine household size groups, and Census region. The construction of the indices in rows two and three is described in section 5 and the notes for Table 6. The sample for all regressions is 3,874 household heads ages 25 to 55 from the CSFII 1989-91 and CSFII 1994-96 who report either full-time employment in actual hours or report being unemployed and looking for a job. The first two point estimates in column (1) replicate results reported in Aguiar and Hurst (2005). Estimates in column (2) are generated using CSFII sample weights. Huber/White/sandwich standard errors are reported in parentheses for the Log Expenditure regressions. Standard errors reported in parentheses for the index regressions are generated by 1000 repetitions of a bootstrap stratified by survey and clustered at the level of primary sampling unit.

Figure 1: First Stage Estimates



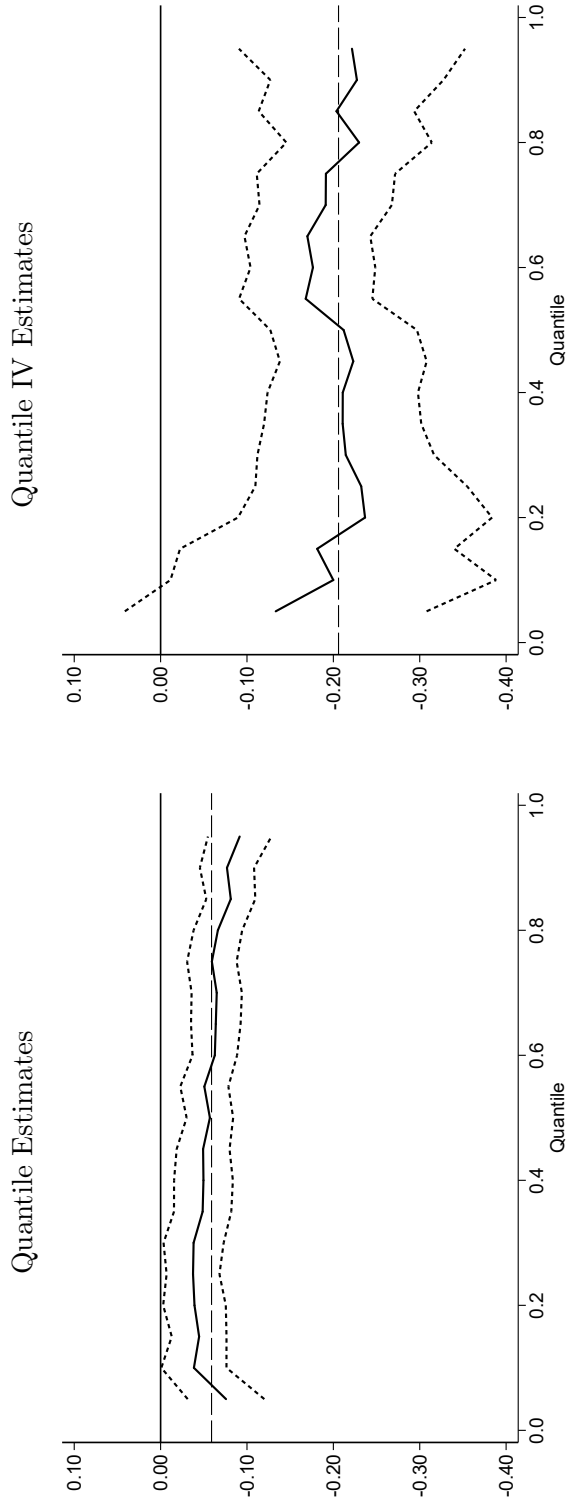
Notes - Each panel shows the estimated coefficients from first-stage regressions of the form of equation (6) for a different sample used in the 2SLS cross-sectional caloric intake regressions reported in Table 3. The solid line in each panel shows the estimated coefficients on the excluded instruments: age indicators for each age from 58 to 71, with age 57 serving as the omitted group. The dashed lines are 95 percent confidence intervals calculated from Huber/White/sandwich variance estimation clustered at the survey*PSU level. These regressions also include indicators for black, four education groups, nine household size groups, year-by-survey, and self-reported health status, all of which are included as controls in the second stage regressions. The F statistics on the excluded instruments are displayed in brackets in the title for each panel.

Figure 2: Age Profiles of Caloric Intake for Men



Notes - Means are generated using single-day data for all men appearing in each survey. The darker, solid line in each panel shows mean caloric intake by age for the indicated survey. Xs mark age ranges used in the retirement analysis. The lighter, dashed line shows mean caloric intake across all surveys and is identical across panels. We use sampling weights and group ages into five-year ranges. Reported slope for each panel is the coefficient estimate from a regression of calories on a linear age term only using respondents ages 57 to 71 in each survey.

Figure 3: Quantile Effects of Retirement on Caloric Intake



Note: These estimates are generated using male heads of household from the pooled sample of cross-sectional surveys excluding those with low response rates. This sample matches the sample used to produce the results appearing in column 4 of Table 2. The controls also match those used to produce those results.

Table A1: Summary Statistics

	Pooled Studies (1)	NHANES I (2)	NHANES II (3)	NFCS 1977-78 (4)	NFCS 1987-88 (5)	NHANES III (6)	CSFII 1989-91 (7)	CSFII 1994-96 (8)	Continuous NHANES (9)
Age	63.4	63.0	63.3	63.4	63.7	63.7	64.1	63.8	63.3
Black	0.09	0.08	0.08	0.09	0.05	0.08	0.11	0.09	0.09
Hhld Size	2.26	2.36	2.30	2.39	2.41	2.32	2.27	2.33	2.18
<i>Education:</i>									
HS Grad	0.26	0.24	0.24	0.29	0.35	0.29	0.35	0.35	0.23
Some College	0.20	0.09	0.12	0.13	0.13	0.14	0.13	0.17	0.26
College Grad	0.25	0.10	0.12	0.12	0.20	0.21	0.26	0.25	0.32
<i>Self-Reported Health:</i>									
Very Good	0.24	0.23	0.21	0.00	0.25	0.24	0.25	0.29	0.27
Good	0.34	0.32	0.30	0.48	0.33	0.33	0.34	0.34	0.32
Fair	0.18	0.23	0.23	0.23	0.20	0.19	0.19	0.15	0.17
Poor	0.06	0.08	0.12	0.09	0.04	0.07	0.05	0.05	0.05
Retired	0.40	0.36	0.34	0.41	0.37	0.45	0.49	0.55	0.38
Calories	2,127	1,925	1,990	2,064	1,926	2,180	1,892	2,059	2,239
<i>N</i>	9,610	570	1,938	1,181	490	1,395	607	1,047	2,382

Notes - The sample is restricted to male household heads from ages 57 to 71. Sample weights are used to compute the statistics shown in the table. The excluded education and health categories are high school dropout and excellent, respectively.

Table A2: Dietary Intake Survey Response Rates

Survey	Household Response Rate ^a	Individual Day 1 Response Rate ^b	Average Male Caloric Intake
HFCS 1965-66	85	81	2,536
NFCS 1977-78	61	57	2,280
NFCS 1987-88	34	31	2,119
CSFII 1989-91	67	58	2,202
CSFII 1994-96	85	80	2,445
NHANES I (1971-75)	99	74	2,308
NHANES II (1976-80)	91	73	2,371
NHANES III (1988-94)	86	78	2,654
Continuing NHANES (1999-2012)	73 to 84	70 to 80	2,593

^aFor NHANES, this column reports the response rate to the in-home interview.

^bFor NHANES, this column reports the response rate to mobile exam which includes the 24 hour recall interview.

Notes - The individual response rates shown are unconditional, i.e., they are the product of the household response rate and the conditional individual response rate to the first day food intake instrument given that the household participates in the survey. First day caloric intake averages shown in the final column are restricted to men ages 25-75 due to an upper age limit in some NHANES surveys and are calculated using sampling weights. HFCS 1965-66 response rates reported by Reese (1977). NFCS 1977-78 and NHANES I, II, and III response rates reported by Interagency Board for Nutrition Monitoring and Related Research (2000). NFCS 1987-88 response rates reported by US GAO (1991). CSFII 1989-91 response rates reported by Tippet et al. (1995). CSFII 1994-96 response rates reported by USDA-ARS (1997). Continuing NHANES response rates sourced from https://www.cdc.gov/nchs/nhanes/response_rates_cps.htm.

Table A3: First Stage Estimates for MRFIT

	Cross-Sectional Analysis	Longitudinal Analysis
Age=54	0.014 (0.014)	0.016 (0.014)
Age=55	0.027 (0.012)	0.029 (0.013)
Age=56	0.060 (0.015)	0.059 (0.015)
Age=57	0.040 (0.016)	0.043 (0.016)
Age=58	0.085 (0.028)	0.092 (0.029)
Age=59	0.115 (0.020)	0.120 (0.020)
Age=60	0.150 (0.032)	0.155 (0.031)
Age=61	0.161 (0.038)	0.165 (0.037)
Age=62	0.254 (0.042)	0.257 (0.041)
Age=63	0.312 (0.039)	0.320 (0.041)
Age=64	0.391 (0.096)	0.406 (0.095)
<i>F</i> statistic	23.3	24.7

Notes - The sample is restricted to individuals from the control sample from MRFIT who are ages 47 and older at the initial visit. The standard errors are clustered at the clinic level (22 clinics). See the notes to Table 5 for information on the additional regressors included in the analysis.