



# Heterogeneity in Household Spending and Well-being Around Retirement

Patrick Moran, Martin O'Connell, Cormac O'Dea, Francesca Parodi

MRDRC WP 2021-427

UM21-13

# Heterogeneity in Household Spending and Well-being Around Retirement

**Patrick Moran**

University of Copenhagen and IFS

**Cormac O’Dea**

Yale University, IFS and NBER

**Martin O’Connell**

University of Wisconsin-Madison and IFS

**Francesca Parodi**

University of Turin, IFS, and CEP

September 2021

Michigan Retirement and Disability Research Center, University of Michigan, P.O. Box 1248.  
Ann Arbor, MI 48104, [mrdrc.isr.umich.edu](http://mrdrc.isr.umich.edu), (734) 615-0422

## Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium through the University of Michigan Retirement and Disability Research Center Award RDR18000002-03. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

## Regents of the University of Michigan

Jordan B. Acker, Huntington Woods; Michael J. Behm, Grand Blanc; Mark J. Bernstein, Ann Arbor; Paul W. Brown, Ann Arbor; Sarah Hubbard, Okemos; Denise Ilitch, Bingham Farms; Ron Weiser, Ann Arbor; Katherine E. White, Ann Arbor; Mark S. Schlissel, *ex officio*



# Heterogeneity in Household Spending and Well-being Around Retirement

## Abstract

We study heterogeneity in spending patterns around the time of retirement. Using rich consumption data from the Panel Study of Income Dynamics, and exploiting within-household spending variation, we systematically classify households into groups characterized by differences in consumption transitions at retirement. We decompose the overall spending changes into the contribution made by different subcomponents of consumption. We find that the households that increase their spending shift budget away from food and toward transportation, recreation, and trips. In contrast, those households for which spending falls reduce the budget share spent on transportation and food away from home, while increasing the share allocated to food at home and housing expenditures. Using a life-cycle model, we characterize the mechanisms capable of driving these observed patterns.

## Citation

Moran, Patrick, Martin O'Connell, Cormac O'Dea, and Francesca Parodi. 2021. "Heterogeneity in Household Spending and Well-being Around Retirement." Ann Arbor, MI. University of Michigan Retirement and Disability Research Center (MRDRC) Working Paper; MRDRC WP 2021-427. <https://mrdrc.isr.umich.edu/publications/papers/pdf/wp427.pdf>

## Authors' acknowledgements

We are grateful to Richard Gong and Aneesha Parvathaneni for excellent research assistance and to Lucas Finamor for helpful discussions.



## 1. Introduction

For most households, retirement is the life event associated with the largest changes in the state of their personal finances. As households transition into retirement, the nature of their consumption decision changes — instead of dividing earnings between consumption and saving, retirees begin to choose how much of any wealth they have accumulated to consume. Furthermore, any mistakes in financial planning earlier in life will be very hard to mitigate at this stage. Understanding how households' living standards transition as they move into retirement is crucial for the wide array of public policies that affect the elderly's well-being. This includes policies relevant to wealth accumulation during working life (e.g., rules around private pensions and saving incentives), as well policies whose effects are most visible in retirement (e.g., Social Security).

A substantial literature has studied how living standards change as households move into retirement. This literature has, most often, used data on household spending as a measure of consumption and an indicator of living standards. Early papers in the literature documented that wealth at retirement is insufficient to maintain preretirement spending levels in retirement (Hamermesh et al. 1984) and that spending falls on retirement (Banks et al. 1998, Bernheim et al. 2001).<sup>1</sup> This fact has been characterized as the “retirement consumption puzzle” — retirement being an anticipated event, the

---

<sup>1</sup> Banks et al. (1998) and Bernheim et al. (2001) documented falls in consumption on retirement in a comprehensive measure of spending in the U.K. and food expenditure in the U.S., respectively, that both sets of authors characterize as inconsistent with the most basic formulation of life-cycle models.

fact that households do not save enough to keep consumption (measured using spending) constant through retirement potentially indicates optimization failures, information deficiencies, or other frictions that suggest a role for policy. A large subsequent literature has studied spending transitions on retirement, and in most cases found falls in average spending of at least some forms. While Bernheim et al. (2001) and Hurd and Rohwedder (2013) show that the tendency for spending to fall on retirement is greatest for those with low wealth, the rest of the literature has mainly focused on cohort averages in spending transitions.

This paper aims to study the importance of *heterogeneity* in spending patterns during households' transition to retirement. We start by showing, using data on a broadly-defined measure of spending from the Panel Study of Income Dynamics (PSID), that households' change in spending at retirement varies systematically with observable household characteristics such as wealth, lifetime earnings, education, and marital status. However, the vast majority of heterogeneity cannot be explained by observable characteristics alone: The heterogeneity *across* groups of households with different characteristics is relatively small compared to *within* group heterogeneity.

Motivated by these facts, and using methods recently developed by Bonhomme and Manresa (2015), we allocate households into discrete groups based on their nondurable spending changes on retirement. The major benefit of this approach is that the groups are defined *only* by within-household spending variation, rather than prior assumptions about which observable characteristics may or may not correlate with consumption behavior.

We classify households into three groups characterized by systematically different spending behavior patterns around retirement. The first group has stable spending dynamics during the transition to retirement, the second group experiences a substantial increase in spending upon retirement, and the third group displays a sharp fall in spending. More than half of households fall into the first group with the smooth transition, while roughly 20% to 25% of households fall into each of the remaining groups. We find that households with a stable transition are substantially better off, both in terms of a permanent income measure that we calculate and accumulated wealth, than households in the other two groups. We show that, for each of the two groups for which spending changes discontinuously, the dynamics of nondurable spending, once controls for life-cycle and macro effects are accounted for, are well-characterized by a one-off permanent discontinuous change at retirement. Moreover, we highlight that households experiencing discontinuous changes (increases or decreases) in spending around retirement are similar in terms of socioeconomic status and permanent income, but households showing the positive change in spending accumulated substantially more wealth during working life.

A number of papers point out that the spending changes at retirement that were documented by the early literature may not translate into changes in living standards. Aguiar and Hurst (2005) show that lower measured spending on retirement may not imply lower consumption, if (as the authors show is the case in their data) retirees spend more time engaged in home production.<sup>2</sup> Laitner and Silverman (2012) argue

---

<sup>2</sup> Recent work by Stephens Jr. and Toohey (2018) shows, using a broad set of food surveys that nutrient caloric intake does fall on retirement, notwithstanding the capacity of households to

that even if consumption falls systematically on retirement, marginal utility might not fall if consumption and leisure (which, of course, increases on retirement) enter the utility function nonseparable. A number of papers point out that falls in spending at retirement in various U.S. data sets are driven by work-related expenses (which likely do not directly impact living standards) and food (which could not affect consumption if accompanied by changes in home production) (see Hurst 2008, Aguilera et al. 2011, and Aguiar and Hurst 2013).<sup>3</sup>

To understand what drives changes in nondurable spending for the three groups that we identify, we decompose the overall spending changes into the contribution made by each of the different subcomponents. We also quantify how the within period share of nondurable spending each group allocates to different spending subcomponents changes at retirement. This enables us to provide evidence of whether spending changes are likely to translate into changes in living standards, and importantly, for this first time, document heterogeneity in these patterns that is driven by unobservables. We find that the group that increases spending on retirement decides to increase spending on all subcomponents, but they shift their budgets away from food and toward transportation, recreation, and trips. In contrast, the group that cuts

---

spend more time at home production. Been et al. (2021) show that, while time at home production does increase, it does not do so sufficiently to fully smooth out observed consumption falls.

<sup>3</sup> A number of papers have documented spending falls in retirement outside of the U.S., including Schwerdt (2005; Germany), Miniaci et al. (2010; Italy), Luengo-Prado and Sevilla (2013; Spain), Stephens Jr. and Unayama (2012; Japan), Moreau and Stancanelli (2015; France), Agarwal et al. (2015; Singapore) and Olafsson and Pagel (2018; Iceland), although the extent to which the falls are broad-based or limited to certain categories differs across settings.

spending at retirement reduces spending across all the subcomponents, but they reduce the share of their budget spent on transportation and food away from the home, while increasing the budget share allocated to food at home and housing expenditures.

To interpret this evidence, we outline a life-cycle model that clarifies the mechanisms through which consumption spending may change discontinuously at retirement. For households with a large increase in total spending, we find that complementarities between leisure and spending (at least for some subcomponents of spending, such as travel) play an important role in driving the spending increase that we observe. In contrast, for households that exhibit falls in total spending, our results suggest that reduced work-related expenses and additional home production after retirement are important in driving spending falls. For both of these groups it is also possible that systematic mistakes in financial planning play a role.

Overall, this paper makes three contributions to the literature. First, we document substantial heterogeneity in spending dynamics at retirement and apply newly developed methods to group households based on unobservable heterogeneity. Second, we investigate how households in the different groups vary based on a wide range of observable characteristics and how their spending patterns evolve across a range of different spending categories, using the broad measures of spending data in the modern PSID. Third, we evaluate the potential mechanisms that may explain a spending decrease in for some households and a sharp spending increase for other households. Overall, our results demonstrate that accounting for rich heterogeneity is key to understanding how and why household spending and living standards transition at retirement.

The rest of this paper is structured as follows. Section 2 gives details on the data that we use. Section 3 presents the empirical approach. Section 4 shows our results on spending dynamics. Section 5 discusses the mechanisms that could be driving these results. Section 6 concludes.

## **2. Data**

We study heterogeneity in spending patterns upon retirement using the 1999 to 2019 PSID. The PSID is well suited to our analysis because it contains detailed information on demographics, income, assets, and spending at the household-level. The PSID began collecting data in 1968 with a core sample of roughly 5,000 U.S. households. Since then, the PSID has followed both the original households and their split-offs (i.e., children and their families). In addition, a refresher sample of immigrant families was added to the PSID in 1997/1999, to ensure that the sample would continue to closely resemble the national population. PSID respondents were surveyed annually until 1996 and biennially starting in 1997.

Crucially, for our purposes, the PSID began collecting detailed information on the multiple subcomponents of spending in 1999, whereas previously it only collected information on food expenditure.<sup>4</sup> Further, in 1999, the PSID began collecting data on asset holdings in each wave of the survey, whereas previously asset holdings were recorded only intermittently. Since our analysis requires detailed information on both

---

<sup>4</sup> As a result of this data limitation, the early literature on the retirement consumption puzzle focused on food expenditure (both at home and away from home) as the primary form of spending. See Aguiar and Hurst (2005) for a discussion of the limitations of studying food expenditure.

spending and assets, we take advantage of the 1999 to 2019 waves of the PSID, bringing new data to the study of consumption transitions around retirement.

## *2.1 Sample selection*

For our baseline specification, we focus on a sample of households observed retiring (which we define precisely below) between the ages of 50 and 80, with nonmissing data on nondurable spending, and with at least one observation both before and after retirement. We include households from both the PSID's core sample (i.e., the original respondents and their split-offs) and the immigrant refresher sample. Whenever there is a change in household composition of either the head or spouse (e.g., due to marriage, divorce, or death), we drop the wave with the change and consider the family to be a new household starting in the subsequent wave.<sup>5</sup> We make no restrictions on marital status, including both single and married households in our analysis. To limit the extent to which our results will be driven by outliers, we drop observations where family wealth exceeds \$20 million or where nondurable expenditure is in either the bottom or top percentile of expenditure. Finally, we exclude observations where the household head is older than 100.

We define the transition into retirement following Bernheim et al. (2001). Household  $i$  is considered to be working in year  $t$  if both the head and spouse (if present) work more than 1,500 hours in the year. The household is considered to be

---

<sup>5</sup> Small sample sizes preclude us from looking at consumption changes between periods that the household composition changed, for example changes in consumption over a period where there was a divorce or bereavement. Small sample sizes also precluded us from looking at consumption changes for those of cohorts affected by particular aggregate shocks (e.g. those who retired at the onset of the Great Recession).

retired in year  $t$  if either the head or spouse work fewer than 500 hours in both the current year and any subsequent year for which data is available. Some households make the transition from work to retirement over the course of several years, during which the retiring member works part-time, between 1,500 hours and 500 hours. Similar to Bernheim et al. (2001), we restrict the sample to households with transition periods of less than four years.<sup>6</sup> We focus on the six years before and six years after retirement.

Given the potentially important role of income in driving heterogeneous consumption patterns around retirement, it is valuable to have information on life-time earnings, although historically this has been very difficult to measure due to data limitations. We overcome this issue by constructing a novel measure of permanent income, exploiting the rules of the U.S. Social Security system to back-out information on life-time earnings based on the Social Security benefits that individuals receive in retirement. One important advantage of this approach is that it allows us to obtain a retroactive measure of life-time earnings, even for the large number of individuals who joined the PSID survey after the start of their working life. This approach uses the fact that retirement benefits paid by the U.S. Social Security system are a function of Averaged Indexed Monthly Earnings (AIME), which averages a measure of earnings over the highest-paid 35 years of an individual's working life. To obtain AIME, we invert this function and use data on Social Security payments and age and year of retirement for each individual.

---

<sup>6</sup> Bernheim et al. (2001) uses the data from years in which the PSID had annual rather than biannual data, restricting the sample to households with transitions of less than five years.

We study the dynamics of a broad measure of nondurable spending, which we define similarly to Blundell et al. (2016), who also use the modern PSID (i.e., the PSID from 1999). Unlike a number of early studies, which use food spending alone to evaluate spending changes at retirement, we look at changes in a broader spending measure, as well its subcomponents. Our nondurable spending measure includes food at home, food away from home, food stamps, transportation expenditure, gasoline, home insurance, and utilities.<sup>7</sup> We exclude from our baseline analysis spending on rent, health care, child care, and education, given that (i) measurement error in self-reported house prices may drive unrealistic fluctuations in imputed rent, (ii) increased health care spending is more likely to indicate the realization of an adverse shock than higher consumption levels, and (iii) child care and educational spending are likely to be driven by changes in household structure, such as children entering and exiting the household. We evaluate sensitivity to these choices in Appendix C.<sup>8</sup> The above measure of nondurable spending is available for the entirety of our sample from 1999 to 2019. That said, additional spending information is available in the PSID beginning in 2005, when a number of new expenditure categories were added to the survey (namely, clothing, trips, and recreation).<sup>9</sup> Our baseline analysis uses the measure of consumption available from 1999 to 2019; however, we provide additional results in Section C.4

---

<sup>7</sup> Transportation expenditure includes car insurance, car repair, parking, bus fare, taxi fare, etc., but does not include durables such as car purchases. Utilities expenditure includes electric, heating, water, and miscellaneous utilities.

<sup>8</sup> See Banks et al. (2019) for an investigation of spending trends at older ages with a focus on spending on healthcare.

<sup>9</sup> More specifically, PSID respondents are asked to report spending on (i) “clothing and apparel, including footwear, outerwear, and products such as watches or jewelry”; (ii) “trips and vacations, including transportation, accommodations, and recreational expenses on trips”; and (iii) “recreation and entertainment, including tickets to movies, sporting events, and performing arts and hobbies including exercise, bicycles, trailers, camping, photography, and reading materials.”

showing how the additional spending categories evolve around the transition into retirement. We convert all components of spending (as well as income, wealth, and other values) to 2015 dollars.

## 2.2 Descriptive statistics

**Table 1: Household consumption change at retirement**

	<b>p25</b>	<b>p50</b>	<b>p75</b>
<b>Permanent Income Tertile 1</b>	-21.0	-2.9	19.9
<b>Permanent Income Tertile 2</b>	-21.0	-2.9	18.4
<b>Permanent Income Tertile 3</b>	-17.5	-2.5	16.3
<b>Wealth Tertile 1</b>	-24.3	-4.8	25.0
<b>Wealth Tertile 2</b>	-20.7	-2.7	19.8
<b>Wealth Tertile 3</b>	-18.1	-1.6	17.3
<b>Income Tertile 1</b>	-24.8	-4.2	21.5
<b>Income Tertile 2</b>	-19.8	-2.5	20.5
<b>Income Tertile 3</b>	-19.1	-3.5	16.8
<b>Married</b>	-20.2	-3.3	18.5
<b>Single</b>	-26.4	-2.5	24.4
<b>Home owners</b>	-20.2	-3.2	17.7
<b>Nonowners</b>	-30.1	-4.8	27.0
<b>White</b>	-18.3	-2.3	18.4
<b>Nonwhite</b>	-27.8	-5.8	24.5
<b>Involuntary Retired</b>	-24.7	-3.3	21.2
<b>Voluntary Retired</b>	-20.6	-3.2	19.1
<b>No High School</b>	-23.4	-3.3	18.4
<b>High School Graduate</b>	-22.2	-3.7	14.7
<b>College</b>	-19.1	-2.6	20.6
<b>Urban</b>	-21.1	-3.3	18.4
<b>Rural</b>	-21.9	-3.3	19.9

**Notes:** This table shows how household nondurable spending changes at the 25th, 50th, and 75th percentiles conditional on preretirement household characteristics. Spending change is defined as the percentage change in spending averaged over three time periods directly preceding retirement and spending averaged over three time periods following retirement. AIME, wealth, and income tertiles are determined from respective means over three time periods preceding retirement.

Table 1 shows the change in nondurable spending for households in different subsamples of the data based on observable characteristics such as lifetime earnings, wealth at retirement, and marital status. We report the 25th, 50th, and 75th percentiles of consumption changes at retirement within subgroups based on these characteristics. Across all the subgroups we consider, we find that the median change in spending is broadly similar, and is between -1.6% and -5.8% at retirement. Further, we find a large degree of dispersion within each group, with households cutting spending by roughly 20% at the 25th percentile and increasing spending by around 20% at the 75th percentile. This shows that there is a large degree of heterogeneity in consumption changes at retirement, and that the observable characteristics in the PSID are unable to capture the majority of this heterogeneity. Table 1 highlights the importance of studying heterogeneity across unobservable, in addition to observable, characteristics as a substantial degree of heterogeneity remains even when we group households based on observable characteristics. While the existing literature on spending dynamics at retirement has advanced in recent years by looking at heterogeneity across observables such as income or wealth, we are not aware of any studies evaluating heterogeneity based on unobservable characteristics. This is largely due to the difficulty of performing such analysis in a systematic manner. The approach developed in the next section is designed to overcome this difficulty.

### **3. Empirical approach**

As documented in the previous section, there is substantial heterogeneity in spending patterns around retirement. While the existing literature has traditionally explored heterogeneity in spending behavior by comparing households with different

observable characteristics, we use an approach that allows us to study heterogeneity in spending patterns without making assumptions about which characteristics (observable or unobservable) may be driving such heterogeneity. Our approach entails first classifying households based on heterogeneity in spending behavior on retirement, and then evaluating which observable characteristics are associated with these differences in behavior.

### 3.1 *Baseline empirical specification*

In order to achieve this goal, we exploit within-household variation in behavior during the transition into retirement, applying the methods recently developed by Bonhomme and Manresa (2015). This method allows the formation of groups that collect households similar in terms of some empirical regularity (in our case, spending changes on retirement). The group membership of each household in the sample and the group-specific effect of retirement can be jointly estimated. In particular, we estimate the following equation which allows for heterogeneous effects of retirement:

$$\ln(C_{i,t}) = \alpha_i + \beta_{g_i} \mathbb{1}_{Retired_{i,t}} + \theta X_{i,t} + \varepsilon_{i,t} \quad (1)$$

where  $C_{i,t}$  is nondurable spending,  $\alpha_i$  is a household fixed effect,  $\mathbb{1}_{Retired_{i,t}}$  is an indicator variable equal to 1 if household  $i$  is retired in year  $t$ , and  $X_{i,t}$  is a vector of controls. We allow for group-specific heterogeneity in the retirement coefficients, denoted by  $\beta_{g_i}$ , where  $g_i$  represents group membership for household  $i$ . The retirement coefficients  $\beta_{g_i}$  are identified by within-household variation in spending at retirement (conditional on controls  $X_{i,t}$ ) for each group. Similarly, group membership  $g_i$  is identified by within-household variation in spending at retirement, assigning each household to the group containing households exhibiting the most similar variation in spending changes on

retirement. As group membership is unknown ex-ante, we jointly estimate group membership ( $g_i$ ) and the coefficients ( $\alpha_i, \beta_{g_i}, \theta$ ) using the iterative procedure developed by Bonhomme and Manresa (2015), which we describe in the next paragraph.

The parameters in Equation (1) are estimated using an iterative algorithm developed by Bonhomme and Manresa (2015), which alternates between estimating the coefficients ( $\alpha_i, \beta_{g_i}, \theta$ ) conditional on group membership ( $g_i$ ) and estimating group membership conditional on the coefficients. In the first iteration of the algorithm, an initial guess of group membership is made, which is then used to estimate the coefficients. Based on the estimated coefficients, group assignment for each household is re-estimated to minimize a least squares criterion. This iterative procedure is repeated until convergence. As the above algorithm is sensitive to the initial guess of  $g_i$ , we repeat the algorithm  $S$  times, each with a different initial guess of group membership. The final estimates are selected based on the same least squares criterion.<sup>10</sup> We report further details on this methodology in Appendix B.

We choose control variables to account for the potential effects of the life-cycle and aggregate economic shocks.  $X_{i,t}$  includes family size, an age polynomial, the unemployment rate, and log stock prices. We do not allow group-specific heterogeneity in the effect of the control variables ( $\theta$ ).<sup>11</sup> We set the number of groups  $G = 3$ .

---

<sup>10</sup> We follow Bonhomme and Manresa (2015) by setting the number of repetitions to be  $S = 1,000$ .

<sup>11</sup> While it is possible to allow for heterogeneity in  $\theta$ , allowing for such heterogeneity could result in group assignment being determined in large part by  $\theta_{g_i}$  rather than spending dynamics around retirement. Given our goal of grouping households based on spending dynamics, we maintain homogeneous coefficient on  $X_{i,t}$ .

### 3.2 *Estimates of grouped heterogeneity*

We report results from estimation of Equation (1) in Table 2. Our procedure identifies three groups of households characterized by different patterns of spending behavior around retirement: for households in Group 1, spending is stable through the transition to retirement, households in Group 2 experience a substantial increase in spending upon retirement, and households in Group 3 present a sharp fall in spending. Roughly half of the households fall into the group with a stable consumption transition as they move into retirement, while a quarter of households fall into each of the other two groups.

**Table 2: Group-specific estimates and group membership**

	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Effect of Retirement (<math>\beta_g</math>)</b>	0.03 (0.01)	0.41 (0.02)	-0.39 (0.02)
<b>Share Households</b>	52.6%	22.8%	24.7%
<b>Weighted Share Households</b>	55.6%	21.6%	22.8%

**Notes:** Weighted share uses mean household PSID Core/Immigrant Longitudinal Family Weight.

Our finding that the three groups exhibit sharply different consumption patterns raises the question of what factors drive these differences. To address this question, we first look at how various observable characteristics compare across the three groups. The results are shown in Table 3. Three main patterns emerge.

First, households with a stable transition into retirement (Group 1) are substantially better-off than households in the other two groups. These households have substantially higher permanent income, as measured by AIME, at both the mean and the median than the other groups, as well as higher family wealth at the time of

retirement.<sup>12</sup> A larger share of households in Group 1 are white (relative to nonwhite), married (relative to single), or homeowners (relative to renters) compared to the other groups. In addition, households in Group 1 are slightly less likely to be hand-to-mouth consumers, i.e., households with close to zero liquid assets at the time of retirement.<sup>13</sup>

**Table 3: Descriptive Statistics by Group Membership**

<b>Variables</b>	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Permanent income</b>			
... p25	30,962	30,666	26,094
... p50	54,904	48,822	49,174
... p75	94,503	76,941	73,720
... Mean	64,985	54,780	55,855
<b>Wealth at retirement</b>			
... p25	65,443	25,128	9,958
... p50	231,841	136,540	94,091
... p75	628,399	474,409	370,314
... Mean	474,293	386,209	299,999
<b>Wealth at retirement/Permanent income</b>			
... p25	2.1	1.5	1.1
... p50	5.6	4.7	3.7
... p75	13.1	16.2	10.7
... Mean	10.3	10.6	8.0
<b>White</b>	0.75	0.62	0.56
<b>Metro population</b>	0.69	0.71	0.71
<b>Married status of head</b>	0.78	0.68	0.61
<b>House ownership status</b>	0.87	0.75	0.70
<b>Hand-to-mouth status</b>	0.46	0.55	0.61
<b>Involuntarily retired</b>	0.12	0.11	0.17
<b>Observations</b>	473	205	222

<sup>12</sup> Our measure of wealth includes checking accounts, saving accounts, directly held stocks, housing wealth net of mortgage balances, and Defined Contribution pension wealth, but does not include Defined Benefit pension wealth or the value of future Social Security benefits.

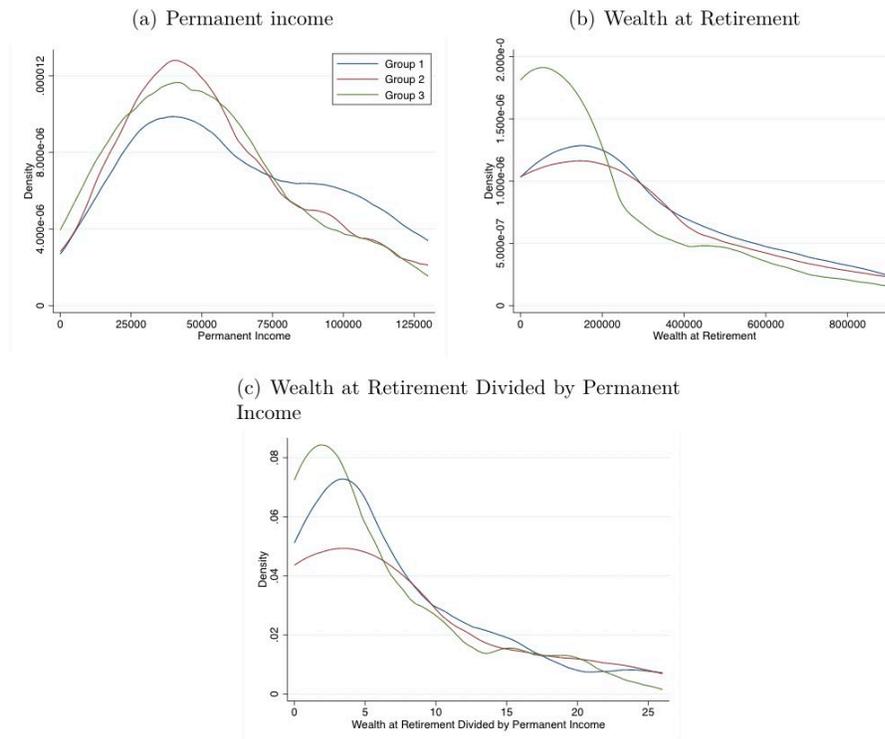
<sup>13</sup> Hand-to-mouth consumers are defined as households with nonzero illiquid wealth, but liquid assets less than two weeks of income, following Kaplan and Violante (2014).

Second, households who increase spending at retirement (Group 2) are only slightly better-off, in terms of their incomes, than their counterparts who cut spending (Group 3). These two groups differ only modestly in terms of their permanent income or demographic characteristics; the distribution of permanent income for Group 2 is shifted slightly to the right compared to Group 3 and the share of white households, married households, and homeowners is slightly higher in Group 2 than in Group 3, but differences are small. This suggests that differences between these groups are not easily accounted for by socioeconomic or demographic characteristics.

Third, while households in Group 2 are relatively similar to Group 3 in terms of demographic characteristics and income, these households accumulate substantially more wealth relative to their lifetime earnings. The median household in Group 2 has a wealth to permanent income ratio of 4.7 at the time of retirement, while the median household in Group 3 has a ratio of only 3.7. The fact that households in Group 2 accumulate more wealth to income than households in Group 3 helps explain how households in Group 2 are able to afford the sharp increase in consumption that we observe following retirement.

Figure 1 shows how the distributions of permanent income and wealth at retirement vary across the three groups: Panel (a) shows the distributions of permanent income, Panel (b) shows the distribution of wealth at retirement, and Panel (c) shows the distribution of the ratio of these two variables. This highlights that there is substantial overlap in the distributions of these observables in each of the three groups.

**Figure 1: Distributions of income and wealth by group membership**



**Note:** Figure shows kernel density estimates of the distributions of permanent income, wealth at retirement and wealth at retirement/permanent income for each of the three household groups.

To further investigate the association between observable characteristics and group membership, we estimate a multinomial logistic regression that relates group membership to the characteristics in Table 3. We find that the most significant predictors of group membership are homeownership, race, and marital status. Households that own their own home are 15.6 percentage points more likely to be in the stable transition group and households that have a white household head are 11.8 percentage points more likely to be in that group than are those with a nonwhite household head. In contrast, single households are 6.2 percentage points more likely to

be in the group with a sharp fall in spending after retirement than are married households.

Overall, however, we find that observable characteristics are relatively poor predictors of group membership. For the vast majority of households, observable characteristics are unable to generate an accurate prediction of group membership. This highlights the role of unobservable heterogeneity in driving differences in consumption patterns across households.

#### 4. Spending dynamics

In this section, we document spending dynamics around retirement for nondurable expenditure and its various subcomponents. We do this using an event study design. Let  $t$  denote calendar time (years),  $K_{i,t}$  the number of years since household  $i$  retired and  $g$  denote groups of households (e.g., the three groups identified by the Bonhomme and Manresa estimator). As we observe households every two years,  $K_{i,t} \in \Omega \equiv (-6, -4, -2, 0, 2, 4, 6)$ . We estimate:

$$x_{i,t} = \alpha_i + \sum_{h \in \Omega} \beta_{g,h} \mathbb{1}[K_{i,t} = h] + \theta X_{i,t} + \varepsilon_{i,t}, \quad (2)$$

where  $x_{i,t}$  denotes a measure of spending by household  $i$  in year  $t$ ,  $\alpha_i$  is a household fixed effect,  $\beta_{g,h}$  are group-specific dynamic retirement effects,  $X_{i,t}$  are time varying controls capturing life-cycle and macroeconomic effects (an age polynomial, family size, the state unemployment rate, and log stock price), and  $\varepsilon_{i,t}$  is an idiosyncratic term. This specification allows for anticipatory effects from six years prior to retirement and post

event dynamics in the six years following retirement.<sup>14</sup> For each consumption measure we estimate Equation (2), imposing common spending retirement dynamics across all households (in which case there is only one group,  $g$ ) and allowing dynamics to differ across the three groups we identify using the Bonhomme and Manresa algorithm.

#### *4.1 Total nondurable spending*

We first estimate Equation (2) for total nondurable expenditure, setting  $x_{i,t} = \ln(C_{i,t})$ . Figure 2 plots the profile of dynamic retirement effects (i.e., the  $\hat{\beta}_{g,h}$ 's). Panel (a) gives results across all households, showing that, on average, the profile of nondurable expenditure (net of the influence of age and other time varying effects) is flat until two years after retirement, and then falls by 0.05 log points between two and six years post-retirement. However, this average effect masks heterogeneity across households.

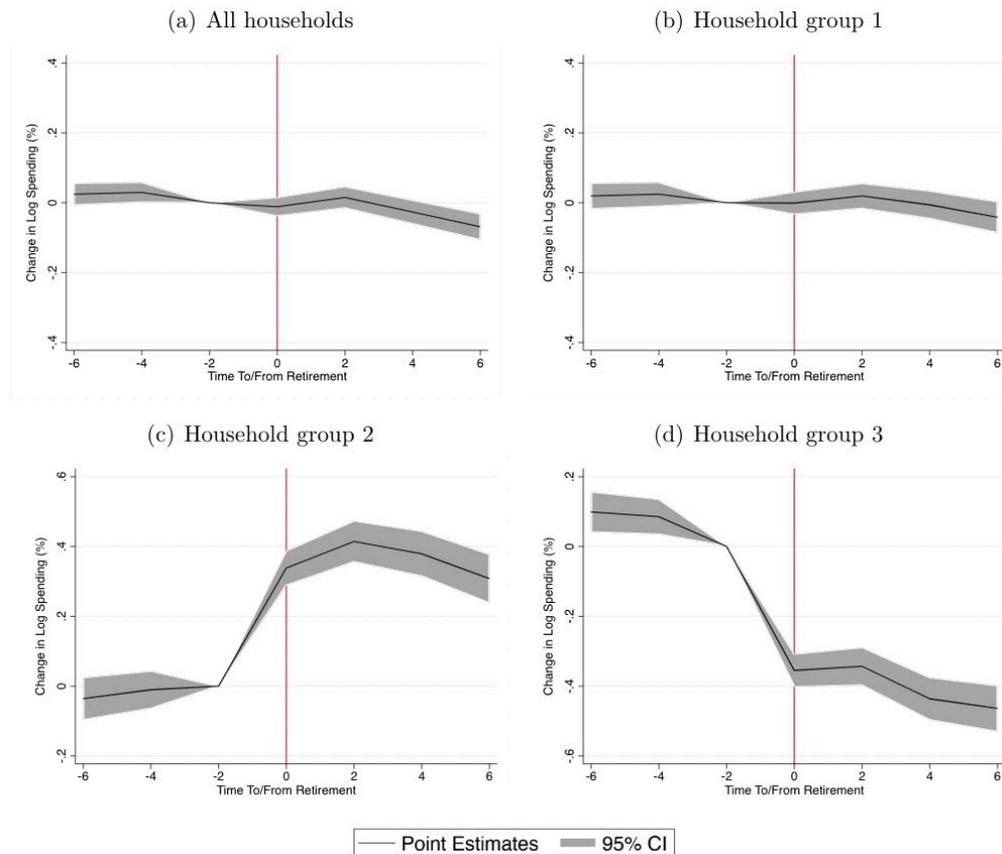
Panels (b) to (d) plot the estimated retirement effects separately for the three groups of households identified by the Bonhomme and Manresa estimator. Group 1's consumption profile is flat over the entire period, six years before to six years after retirement. In contrast, Group 2 exhibits a flat spending profile pre- and post-retirement, but a persistent increase in spending of 0.4 log points in the period of retirement. For Group 3 the pattern is reversed: In the year of retirement there is a persistent fall in nondurable spending of 0.4 log points. Together panels (b) to (d) suggest the three

---

<sup>14</sup> An identification problem, analogous to the age-period-cohort identification problem, arises if there are household fixed effects, fully dynamic retirement effects and fully flexible age effects. At least one restriction is needed; in our case we restrict the age effects to be a polynomial. As argued below, the data provide support for additional restrictions on the dynamic retirement effects, which further resolves the identification problem.

groups' consumption profiles through retirement are well captured by a persistent jump in consumption (or the absence of it, for Group 1) in the year of retirement.

**Figure 2: Retirement dynamics for total nondurable spending**



**Note:** Figures plot the estimated  $\hat{\beta}_{g,h}s$  from equation (2) when dependent variable is log nondurable spending. Panel (a) is for households. Panels (b) to (d) are for the three groups we identify in Section 3. Confidence intervals are calculated assuming allocation to groups is deterministic.

#### 4.2 Spending subcomponents

In this section, we decompose the changes in spending upon retirement depicted in Figure 2(b) to (d) into spending changes in four (exclusive and exhaustive)

subcomponents of nondurable expenditure: spending on food at home, on food away from home, on transport expenditure, and on house expenditures (excluding rent).

To do this, we estimate a version of Equation (2) for each spending component, and for total nondurable expenditure, of the form

$$x_{i,t} = \alpha_i + \beta_{g_i}^c \mathbb{1}[K_{i,t} \geq h] + \theta X_{i,t} + \varepsilon_{i,t}. \quad (3)$$

This replaces the dynamic retirement effects with a persistent static effect; i.e.,  $\beta_{g,h} = 0$  for  $h < 0$  and  $\beta_{g,h} = \beta_g$  for  $h = (0, 2, 4, 6)$ .<sup>16</sup> To distinguish the estimated retirement effects for different subcomponents of spending, we include a superscript  $c$ , (i.e.,  $\widehat{\beta}_g^c$ ) where  $c = \{\text{food at home, food away from home, transport, house expenditures}\}$ .

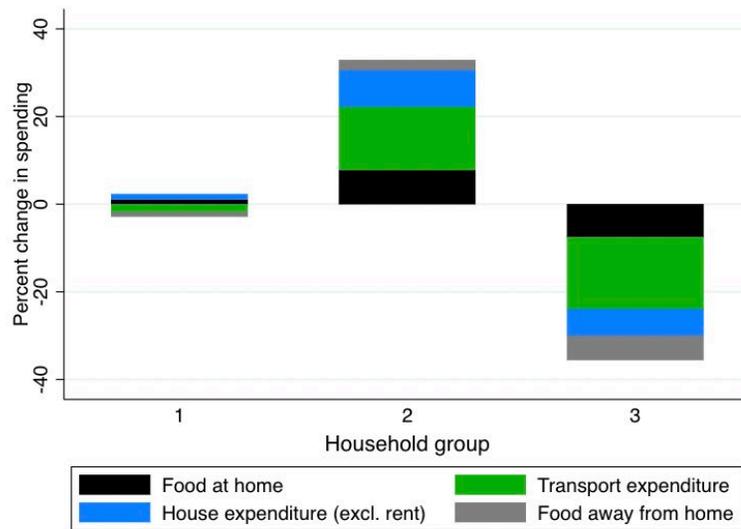
We use the estimates of Equation (3) to, for each household group, decompose the total percentage change in nondurable spending into the contribution made by each spending category. To do this, we estimate the equation in levels rather than logs, and define the contribution of spending category  $c$  as  $\Delta \widehat{x}_g^c = \widehat{\beta}_g^c / \mathbb{E}[\widetilde{C}_1 | g, h > 0]$ , where  $\mathbb{E}[\widetilde{C}_1 | g, h > 0]$  is predicted total nondurable spending from group  $g$  post-retirement omitting the contribution of the retirement dummy.  $\mathbb{E}[\widetilde{C}_1 | g, h > 0]$  serves as an estimate of counterfactual nondurable spending had households not retired. This approach is similar to that taken by Kleven et al. (2019) and allows for the inclusion of observations where spending on a subcomponent is zero (which would not be possible if we used log spending as the dependent variable).

---

<sup>16</sup> We present the dynamic results in Appendix C.4. They show that, as in Figure 2, the retirement consumption dynamics are well approximated with a static effect.

Figure 3 illustrates this decomposition. For Group 1 spending on all four categories is stable through retirement. For Group 2 approximately half of the 33% increase in nondurable expenditure at retirement is due to higher transport expenditure. The remaining half is split roughly evenly between higher spending on food at home and household expenditure, with higher spending on food away from the home contributing only marginally. For Group 3, the 36% fall in nondurable spending at retirement is driven by lower transport spending (-16%), and by lower spending in the other three categories (which each contribute approximately equally).

**Figure 3: Decomposition of retirement spending changes**



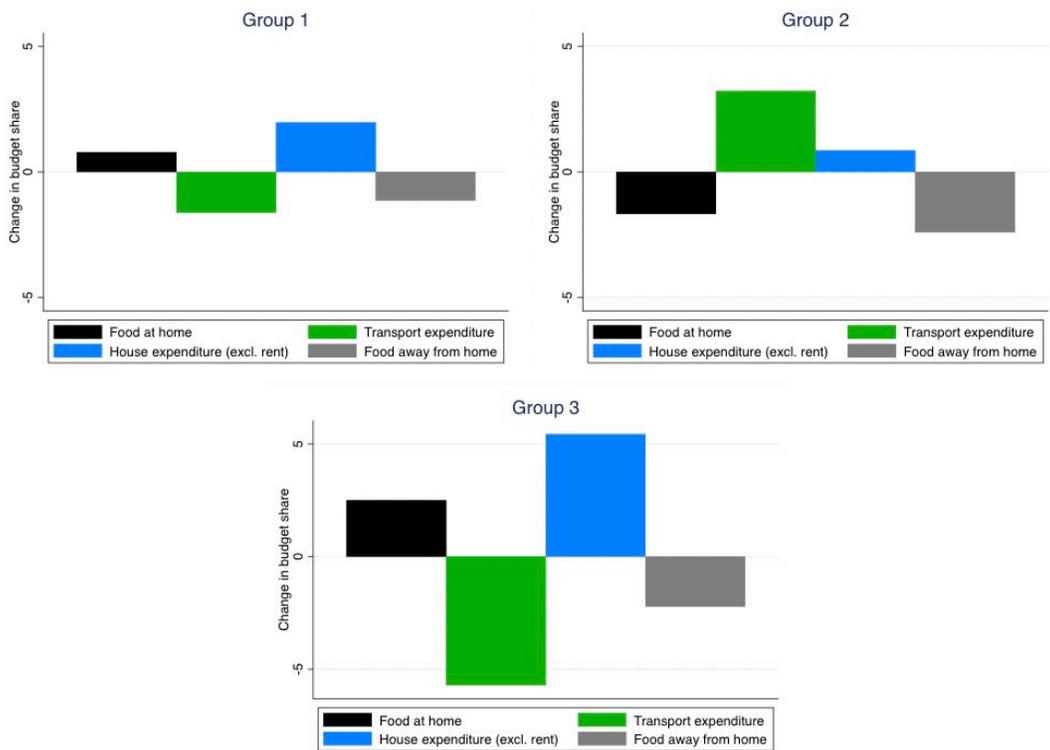
**Note:** For each of the three groups we identify in Section 3, figure shows a decomposition of changes in total nondurable spending at retirement into the contribution from four spending subcomponents.

While Figure 3 shows which spending categories drive changes in overall nondurable spending, it does not tell us what categories households reallocate their total spending toward and away from at retirement. To address this question we re-

estimate Equation (3) for each spending category, with the dependent variable equal to the category budget share. In this case, the estimate  $\widehat{\beta}_g^c$  tells us the average change in the share of their total nondurables spending households in Group  $g$  allocate to spending on  $c$  when they retire.

Figure 4 plots our estimates of  $\widehat{\beta}_g^c$  for the three household groups and four spending categories. On average, at retirement households in Group 1 reallocate spending away from transport expenditure and food away from the home and toward food at home and household expenditures. Group 3 (the group with a fall in total expenditure at retirement) exhibits changes in budgets shares that are in the same direction, on average, as Group 1, but are much larger in magnitude. In particular, on average, they reduce their budget shares allocated to transport by 5.7 percentage points and food away from home by 2.2 percentage points, and raise their budgets shares for food at home by 2.5 percentage points and household expenditure by 5.4 percentage points. In contrast, households in Group 2 (those with a rise in total spending at retirement), on average, raise their budget shares allocated to transport (by 3.2 percentage points) and household expenditure (by 0.9 percentage points) and reduce the share of their budgets allocated to food at home (by 1.7 percentage points) and food away from the home (2.4 percentage points).

**Figure 4: Changes in budget shares in retirement**

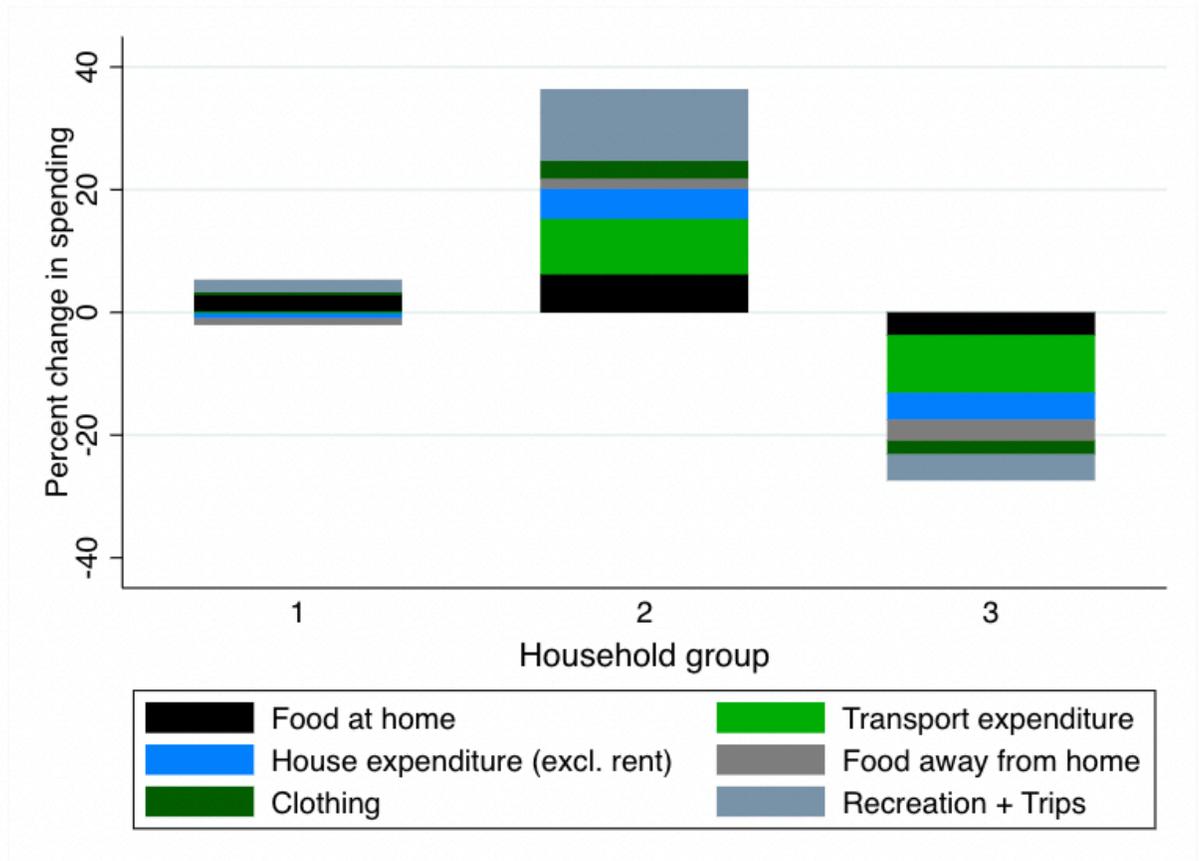


**Note:** Figures show the change in the share of their total nondurable spending that households in the three groups we identify in Section 3 allocated to each spending subcomponent at retirement.

From 2005 on, the PSID includes information on spending on two additional categories of goods: clothing and recreation and trips. For the subsample of households we observe retiring in the period 2005 to 2019, we replicate Figures 3 and 4, but include spending on these two additional components. In doing this, we retain the same grouping of households as used in the preceding analysis. Figure 5 shows that households in Group 2 (the group that raises total spending at retirement) significantly increase their spending on recreation and trips on retiring, whereas households in Group 1 exhibit little change on average and those in Group 3 exhibit a modest decline in recreation and trips spending. Figure 6 shows that, as well as increasing the

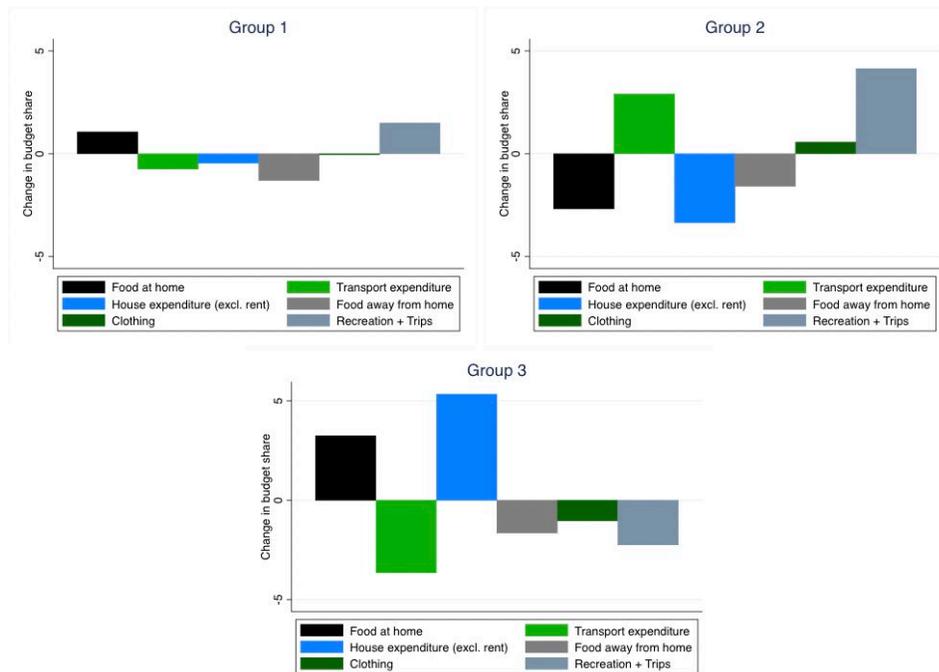
resources allocated to transport expenditure, Group 2 households also significantly increase their budget share for trips and recreation.

**Figure 5: Decomposition of retirement spending changes, additional components**



**Note:** Figure replicates the analysis in Figure 3 but using on the post 2005 PSID, which contains additional spending subcomponents.

**Figure 6: Changes in budget shares at retirement, additional spending components**



**Note:** Figure replicates the analysis in Figure 4 but using only the post 2005 PSID, which contains additional spending subcomponents.

## 5. Mechanisms

To help place our results into a broader context, we outline mechanisms that could rationalize the consumption transitions observed in the data. In discussing these mechanisms, we point to evidence in our analysis that is suggestive of the extent to which each channel is at play. We emphasize, throughout, that these mechanisms are not mutually exclusive. Further, given the role of unobserved heterogeneity uncovered in Section 3, we stress that each of these mechanisms may matter to differing degrees for the three household groups that we study.

## 5.1 Framework

That consumption falls in retirement has been labeled a ‘puzzle’ follows from the fact that the most parsimonious life-cycle model predicts that consumption changes on retirement should not differ systematically from consumption changes in other periods. Hence, in this section, we outline a life-cycle model and state the assumptions needed to rule out discontinuous changes in consumption at retirement. We then discuss the deviations from that simple model that would generate falls/jumps in consumption at retirement.

Suppose agents live for up to  $T$  periods, working for the first  $R$  periods, at which point they retire. Their probability of survival to period  $t$  is given by  $s_t$ . Their (potentially stochastic) income during working life is given by  $y_t^e$ . When working, they incur some work-related expenses ( $w_t$ ) — this is expenditure that does not provide any utility. Income in retirement is given by  $y^R$ . Utility is defined over consumption and leisure. Agents choose their consumption of each of a vector of goods,  $\mathbf{q}_t$ , available at prices  $\mathbf{p}_t$ . Total period expenditure on consumption goods is  $c_t$ . Leisure is a function of whether an agent is retired or not. Borrowing is not allowed. The interest rate is given by  $r$ .

The agent's problem can then be expressed as:

$$\max_{\{c_t\}} \sum_{t=0}^T \beta^t s_t E[U(\phi(c_t, p_t, l_t), l_t)] \quad (4)$$

$$\text{where } \phi(c_t, p_t, l_t) = \max_{q_t} u(q_t, l_t) \quad \text{s.t. } p_t' q_t = c_t$$

$$\left. \begin{array}{l} \\ \\ \text{s.t.} \end{array} \right\} \begin{array}{l} a_{t+1} = (a_t + y_t^e - we - c_t)(1 + r) \quad \text{if } t < R \\ \\ a_{t+1} = (a_t + y^r - c_t)(1 + r) \quad \text{if } t \geq R \\ \\ a_t \geq 0 \quad \forall t \\ \\ l = \mathbb{1}[t < R] \end{array}$$

Features of the technology as well as preferences could differ across households (though in the program specified in (4) we suppress  $i$  subscripts).

If we additionally assume that:

- A1. There are no work-related expenses and there is a constant, across time, price vector for consumption goods ( $\mathbf{p}$ );
- A2. Preferences over consumption and leisure are additively separable:

$$\circ U(\phi(c_t, p_t, l_t), l_t) = v(c_t, p_t) + \psi(l_t) \quad (5);$$

- A3. Agents are perfectly informed about all deterministic model features and have rational expectations over stochastic model features;

then optimal consumption can, as long as households have nonzero wealth, be characterized by the following Euler equation:

$$v'_c(c_t, p) = \beta s_t (1 + r) E[v'_c(c_{t+1}, p)], \quad (6)$$

which holds that agents will choose consumption spending to equate, over time, the expected discounted marginal utility of consumption. Such a solution can admit consumption profiles that increase or decrease over time (depending on the extent of discounting ( $\beta$ ), survival probabilities, and the interest rate), but rules out any discontinuous changes in consumption. The following three subsections discuss each of assumptions A1 to A3 and what relaxing them implies for consumption behavior around retirement.

### 5.1.1 Distinguishing consumption from expenditure

An influential vein of research has argued that, while measured expenditure ( $w_e + c_t$ ) falls on retirement, the extent of this fall overstates the fall in consumption. There are two distinct types of changes in expenditure not reflected in consumption spending.

The first channel is work expenses (for instance, commuting costs and spending on clothes for work) which fall to zero on retirement. For Group 3 (those who exhibit a drop in nondurable spending at retirement), there is a large fall in both transportation expenditure and the budget share allocated to transportation. This is also true, though to a lesser extent, of Group 1 (who exhibit a stable profile of nondurable spending through retirement). This reduction in spending (and the share of spending allocated to) transport expenditure is consistent with the disappearance of work-related travel expenses for these groups on retirement.

The second channel through which consumption per unit of expenditure might increase is if the price of consumption falls when agents retire. This could happen as the additional time that retirees have can substitute for money in the production of consumption. Aguiar and Hurst (2005) show for a sample of households in the U.S. that

while food expenditure falls on retirement, time spent shopping for food and preparing meals increases on retirement. Measuring food intake directly, the authors find that neither the quantity nor the quality falls on retirement, though recent work by Stephens Jr. and Toohey (2018) using different data does find falls in caloric intake.

We find evidence consistent with the substitution from food outside the home to home-prepared food in each of our three groups. Spending on food away from home falls for each of Groups 1 and 3, while for Group 2, despite large increases in spending overall, food away from home increases only modestly, and the budget share falls. The results for food at home mirror these results: Spending on food at home increases for Groups 1 and 3, and while it falls for Group 2, the at-home budget share rises for this group. These patterns are indicative of a shift from prepared food toward home-cooked food.

#### 5.1.2 Nonseparability between consumption and leisure in the utility function

Well-informed, optimizing agents will smooth expected marginal utility across the life cycle. If preferences over consumption are additively separable from other variables that enter the utility function, including leisure, then smoothing of marginal utility will imply consumption-smoothing. If, on the other hand, the marginal utility of consumption depends on leisure (or any other variable) then discontinuous changes in leisure (or other relevant variables) will trigger discontinuous changes in consumption.

Program (4) allows for the possibility of two kinds of interactions in preferences between consumption and leisure. If the utility function is given by  $U(\phi(c_t, p_t, l_t), l_t)$ , preferences over consumption and leisure are weakly separable. In this case leisure impacts the marginal utility of total consumption, but not the ratio of marginal utilities

between any two specific consumption goods. Retirement, by changing leisure, can directly impact the consumer's intertemporal allocation of resources, and therefore can lead to a discontinuous jump in spending. However, its impact on within period consumption allocations is only indirect, through any income effect related to a change in period total consumption. Alternatively, if the utility function is  $U(\phi(c_t, p_t, l_t), l_t)$  leisure can also directly impact within period budget allocations by changing the marginal rate of substitution between two different goods.

If leisure and consumption (or some components of consumption) are complements for a particular household, the increase in leisure time available in retirement should lead to consumption (or consumption of the component complementary with leisure) discontinuously increasing. Conversely, if leisure and consumption (or a specific component of consumption) are substitutes, consumption will discontinuously decrease at retirement.

Households in Group 2, in contrast to those in the other two groups, exhibit a large increase in transport expenditure, as well as the budget share allocated to it, at retirement. In our analysis with the post-2005 PSID data, which collects data on an even more detailed set of consumption characteristics, we also find that the households in this group increase, both in levels and in terms of budget shares, their consumption of recreation and trips. The increase in overall spending for this group, as well as the tilting in spending toward travel expenditure is strongly suggestive of a complementarity between travel spending and leisure for this group.

### 5.1.3 Retirement frictions

If retirement is accompanied by a shock to households' information sets, then it may induce households to reset their consumption plans, leading to a discontinuous change in consumption (and marginal utility). We can distinguish between two types of shocks. First, households could be hit by adverse realizations of shocks from distributions which, before retirement, they had knowledge of. Households who receive adverse shocks (to employment possibilities or health, for example) could make up a disproportionate share of Group 3 (whose consumption falls) and those who receive favorable shocks could make up a disproportionate share of Group 2 (whose consumption increases). Second, it could be that (some) households are simply not able to make appropriate retirement plans; that is, they are unable to solve the program outlined in (4). We look for suggestive evidence of each of these phenomena.

We find mild evidence in support of shocks to households' information sets at retirement. For instance, a larger share of Group 3 households is forced to retire involuntarily; however, the magnitudes are relatively small. As shown in Table 3, involuntary retirement characterizes 17% of households in Group 3, compared to 11% to 12% in the other groups. Therefore, it seems that involuntary retirement does not explain the whole story. We further investigate the role of shocks by looking at differences in health across the three different groups of households. The results are contained in Appendix E. Overall, we find that health deteriorates around the time of retirement for all three groups, as measured by (i) self-reported health, (ii) the number of severe ailments, and (iii) disability. The health deterioration for Group 3 is slightly worse than for the other two groups, which may in part contribute to the spending decline experienced by this group. However, we find that Group 1 and Group 2 have very

similar health declines around the time of retirement, indicating that the increase in spending observed for Group 2 is not driven by these households having better than expected health during old age.

Second, an assumption underlying Program (4) is that agents are able to solve the maximization problem. But there is evidence that, at least in some settings, agents use heuristics and approximate rules of thumb as the basis for their decision-making and that they systematically form incorrect beliefs.<sup>17</sup> Relatedly, there is evidence that people may make time-inconsistent decisions, which in the context of savings, can lead them to continually postpone saving.<sup>18</sup> Both rule-of-thumb decision-making and time inconsistency can lead people to save inadequately (or, in the former case, also excessively) for retirement, relative to if they solve Program (4). This, in turn, can lead to discontinuous jumps in consumption (and marginal utility) at retirement. In our context, one way to potentially explain discontinuous changes in spending at retirement is that households may be using rule-of-thumb decision-making, which may cause Group 2 to over-save for retirement and Group 3 to under-save.

To investigate this channel, we exploit extra supplemental data on personality traits, which is available in the one-time “Well-being and Daily Life Supplement” administered to PSID respondents in 2016. We focus on personality traits that may be

---

<sup>17</sup> See, for example, Heimer et al. (2019) and O’Dea and Sturrock (2021), who study the role of subjective beliefs over life expectancy in the context of models of consumption and saving in retirement, and Delavande and Rohwedder (2011) and Luttmer and Samwick (2018) who show that there is substantial variation in beliefs over future Social Security entitlements.

<sup>18</sup> For a review of the empirical evidence see DellaVigna (2009). See also Hurst (2003) for evidence that a share of households in the PSID who arrive at retirement with low wealth behave myopically through their working life.

related to individuals' willingness to plan ahead regarding their personal finances. The results are reported in Table 4. We find that the three groups of households do differ systematically along these characteristics. In particular, households in Groups 2 and 3 are less likely to report being good at doing thorough jobs and at planning than their counterparts in Group 1.

Overall, we find evidence consistent with planning or information frictions playing some role in driving discontinuous changes in spending at retirement, but these cannot explain the entire story. If retirement frictions alone were at play, the changes in budget shares on retirement for the two groups would be driven by the income effect associated with lower/higher period total consumption. However, we find that the direction of budget share changes for Groups 2 and 3 are nonsymmetrical. While Group 2 exhibits a large spending increase and Group 3 experiences a large fall, both groups increase the share of their budget allocated to household expenditure and reduce the share allocated to food away from the home. These results demonstrate that while informational and/or behavioral frictions appear important, they are unable to account for all of the spending changes that we uncover, which suggests an additional role of nonseparabilities between consumption and leisure.

**Table 4: Differences in personality traits, by households groups**

	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Does Thorough Job</b>	0.09* (0.05)	-0.23** (0.11)	-0.15 (0.10)
<b>Worries</b>	0.07 (0.06)	-0.09 (0.12)	-0.19* (0.10)
<b>Lazy</b>	-0.07 (0.05)	0.07 (0.11)	0.25** (0.10)
<b>Unsure Life</b>	-0.02 (0.06)	0.01 (0.11)	0.04 (0.11)
<b>Plan Ahead</b>	0.07 (0.06)	-0.01 (0.11)	-0.22** (0.10)
<b>Share of Households N</b>	56.8% 262	18.2% 84	24.9% 115

In summary, our results point to a number of mechanisms driving discontinuous changes in spending at retirement. In particular, they suggest that the increase in spending we document for Group 2 is likely to be driven in some part by a complementarity between particular forms of consumption and leisure. For Group 3, those households that exhibit large falls in total spending at retirement, our results suggest that reduced work-related expenses, additional home production after retirement, and systematic inability to plan for retirement all play a role.

## **6. Conclusion**

Substantial spending changes at retirement could be indicative of a failure of households to optimally allocate their wealth across the life cycle if, for example, they reflect information deficiencies among those saving for retirement or other failures to plan effectively. On the other hand, they could be benign expressions of households effectively spreading their resources over their life cycle. The extent to which each explanation drives observed spending changes (and for which households) determines

what role, if any, these patterns suggest for policy. Exploiting rich longitudinal data on spending in the PSID, we document both average changes in spending among a sample of U.S. retirees, and investigate variation in changes across households. We find, based on a broad measure of consumption, that falls in average spending on retirement are modest. A focus on this average alone obscures the fact that there is substantial dispersion in spending changes on retirement, with some households exhibiting large falls and some increasing their spending on retirement.

Our results demonstrate that accounting for rich heterogeneity is key to understanding how and why living standards change at retirement. Substantial falls in spending on retirement are observed for only a relatively small share of the population. Among those households, adverse shocks and an inability to effectively plan for retirement play some role, and these factors can lead to consequential (and perhaps unanticipated) falls in living standards on retirement. There is evidence, though, even among these households whose spending falls, that part of the explanation lies in switches from market to home production of food, and a fall in work-related expenses, neither of which indicate falls in living standards. Further, for a majority of households in our sample, household spending transitions on retirement are best modeled as stable (or even increasing), with evidence of complementarity between leisure time and some forms of consumption driving reallocation of budget and an overall increase in spending for some households.

A promising direction for future work is to formally characterize the welfare implications of the spending dynamics we document by estimating a dynamic model of

intratemporal choice over consumption groups and the intertemporal decision of how much to consume and save.

## References

- Agarwal, S., J. Pan, and W. Qian (2015). The composition effect of consumption around retirement: Evidence from Singapore. *American Economic Review* 105(5), 426–31.
- Aguiar, M. and E. Hurst (2005). Consumption versus expenditure. *Journal of Political Economy* 113(5), 919–948.
- Aguiar, M. and E. Hurst (2013). Deconstructing life cycle expenditure. *Journal of Political Economy* 121(3), 437–492.
- Aguila, E., O. Attanasio, and C. Meghir (2011). Changes in consumption at retirement: evidence from panel data. *Review of Economics and Statistics* 93(3), 1094–1099.
- Ameriks, J., Caplin A., & Leahy, J. (2007) “Retirement Consumption: Insights from a Survey”. *The Review of Economics and Statistics*; 89 (2): 265–274
- Banks, J., R. Blundell, P. Levell, and J. P. Smith (2019). Life-cycle consumption patterns at older ages in the United States and the United Kingdom: Can medical expenditures explain the difference? *American Economic Journal: Economic Policy* 11(3), 27–54.
- Banks, J., R. Blundell, and S. Tanner (1998). Is there a retirement-savings puzzle? *American Economic Review*, 769–788.
- Been, J., S. Rohwedder, and M. Hurd (2021). Households’ joint consumption spending and home production responses to retirement in the US. *Review of Economics of the Household*, 1–27.
- Bernheim, B. D., J. Skinner, and S. Weinberg (2001). What accounts for the variation in retirement wealth among US households? *American Economic Review* 91(4), 832–857.

- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Bonhomme, S. and E. Manresa (2015). Grouped patterns of heterogeneity in panel data. *Econometrica* 83(3), 1147–1184.
- Delavande, A. and S. Rohwedder (2011). Individuals' uncertainty about future social security benefits and portfolio choice. *Journal of Applied Econometrics* 26(3), 498–519.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2), 315–72.
- Hamermesh, D. S. et al. (1984). Consumption during retirement: The missing link in the life cycle. *The Review of Economics and Statistics* 66(1), 1–7.
- Heimer, R. Z., K. O. R. Myrseth, and R. S. Schoenle (2019). Yolo: Mortality beliefs and household finance puzzles. *The Journal of Finance* 74(6), 2957–2996.
- Hurd, M. D. and S. Rohwedder (2013). Heterogeneity in spending change at retirement. *The Journal of the Economics of Ageing* 1, 60–71.
- Hurst, E. (2003). Grasshoppers, ants, and pre-retirement wealth: A test of permanent income.
- Hurst, E. (2008). The retirement of a consumption puzzle. Technical report, National Bureau of Economic Research.
- Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica* 82(4), 1199–1239.
- Kleven, H., C. Landais, and J. E. Søgaaard (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics* 11(4), 181–209.

- Laitner, J. and D. Silverman (2012). Consumption, retirement and social security: Evaluating the efficiency of reform that encourages longer careers. *Journal of Public Economics* 96(7-8), 615–634.
- Luengo-Prado, M. J. and A. Sevilla (2013). Time to cook: expenditure at retirement in Spain. *The Economic Journal* 123(569), 764–789.
- Luttmer, E. F. and A. A. Samwick (2018). The welfare cost of perceived policy uncertainty: evidence from social security. *American Economic Review* 108(2), 275–307.
- Miniaci, R., C. Monfardini, and G. Weber (2010). How does consumption change upon retirement? *Empirical Economics* 38(2), 257–280.
- Moreau, N. and E. Stancanelli (2015). Household consumption at retirement: A regression discontinuity study on French data. *Annals of Economics and Statistics/Annales d'Economie et de Statistique* (117/118), 253–276.
- O'Dea, C. and D. Sturrock (2021). Survival pessimism and demand for annuities. *Review of Economics and Statistics*.
- Olafsson, A. and M. Pagel (2018). The retirement-consumption puzzle: New evidence from personal finances. Technical report, National Bureau of Economic Research.
- Schwerdt, G. (2005). Why does consumption fall at retirement? Evidence from Germany. *Economics Letters* 89(3), 300–305.
- Stephens Jr, M. and D. Toohey (2018). Changes in nutrient intake at retirement. Technical report, National Bureau of Economic Research.
- Stephens Jr, M. and T. Unayama (2012). The impact of retirement on household consumption in Japan. *Journal of the Japanese and International Economies* 26(1), 62–83.

# Appendices

## Appendix A: Data

Table 5 gives details on our sample selection, and illustrates the number of observations dropped at each step in our sample selection process.

**Table 5: Sample selection and sample size**

	Number of observations	Number of unique IDs (pid)
Total number in PSID	96,288	17,641
Drop the wave in which there is a change in head/ spouse		
Number with change in head/ spouse	8,589	6,376
Remaining number	87,699	17,641
Reset the pid for households whom head/ spouse changes		
Remaining number	87,699	21,409
Drop rows for which family wealth $\geq$ \$ 20 Million		
Number with family wealth $\geq$ \$ 20 Million	34	14
Remaining number	87,665	21,407
Drop rows for which head age $\geq$ 100		
Number with head age $\geq$ 100	18	13
Remaining number	87,647	21,401
Drop pids for which data doesn't exist continuously		
Number for which data is missing for intermediate waves	66	9
Remaining number	87,581	21,392
Drop rows in which pid is out of U.S.		
Number that move out of U.S.	644	166
Remaining number	86,937	21,288
Dropping pids that do not go through retirement transition		
Number that do not go through retirement transition	80,222	20,369
Remaining number	6,715	919

Drop pids for which expenditure is missing in any wave		
Number for which expenditure is missing in any wave	68	10
Remaining number	6,647	909
Drop rows for which expenditure is $\leq p1$ or $\geq p99$		
Number for which expenditure $\leq p1$ or $\geq p99$	132	92
Remaining number	6,515	907
Drop pids for which we don't observe at least one wave pre/post retirement		
Number for which we don't observe atleast one wave pre/post retirement	19	7
Number in baseline sample	6,496	900

**Notes:** This table shows the sample selection criterion for the baseline specification.

### *A.1 Summary statistics by average indexed monthly earnings (permanent income)*

In this paper, we use Social Security entitlements to construct a measure of permanent income (Average Indexed Monthly Earnings). To illustrate how this measure varies with other economic and demographic variables, Table 6 gives the 25th, 50<sup>th</sup>, and 75th percentiles of AIME, conditional on a set of observables.

**Table 6: Household AIME**

	<b>p25</b>	<b>p50</b>	<b>p75</b>
<b>AIME Tertile 1</b>	9,707	21,489	29,705
<b>AIME Tertile 2</b>	45,159	51,718	60,523
<b>AIME Tertile 3</b>	87,904	102,966	123,385
<b>Wealth Tertile 1</b>	17,846	29,988	48,545
<b>Wealth Tertile 2</b>	32,094	53,754	76,029
<b>Wealth Tertile 3</b>	40,878	73,951	109,471
<b>Income Tertile 1</b>	16,930	30,844	48,590
<b>Income Tertile 2</b>	35,228	50,945	76,741
<b>Income Tertile 3</b>	48,290	80,659	114,850
<b>Involuntary Retired</b>	18,917	34,025	48,590
<b>Voluntary Retired</b>	30,474	53,462	90,636
<b>No High School</b>	29,038	43,001	59,205
<b>High School Graduate</b>	31,765	50,732	81,636
<b>College</b>	35,058	61,969	102,815

**Notes:** This table shows household AIME at the 25th, 50th, and 75th percentiles conditional on preretirement household characteristics.

## Appendix B: Bonhomme-Manresa estimation procedure

We estimate the parameters of Equation (1) based on the following minimization problem:

$$(\hat{\theta}, \hat{\beta}, \hat{\gamma}) = \arg \min_{(\theta, \beta, \gamma) \in \Theta \times B^G \times \Gamma_G} \sum_{i=1}^N \sum_{t=1}^T (\ln(C_{i,t}) - \alpha_i - \beta_{g_i} \mathbb{1}_{\text{Retired}_{i,t}} - \theta X_{i,t})^2 \quad (7)$$

where  $\theta \in \Theta$  is the vector of parameters common across households,  $\beta_g \in B$  for  $g \in \{1, \dots, G\}$  are the group-specific effects of retirement, and  $g_i$  for  $i \in \{1, \dots, N\}$  denotes group-membership for each household. We denote  $\beta$  as the set of all  $\beta_g$  and  $\gamma$  as the set of all  $g_i$ 's so that  $\gamma \in \Gamma_G$  represents a particular grouping of the  $N$  households within  $\Gamma_G$  that is the set of all possible groupings of  $\{1, \dots, N\}$  into at most  $G$  groups.

Bonhomme and Manresa (2015) show that it is possible to jointly estimate the regression coefficients ( $\theta$  and  $\beta$ ) and group membership ( $g_i$ ) using an algorithm that iterates between (i) estimating group membership conditional on the regression coefficients and (ii) estimating the regression coefficients conditional on group membership. More specifically, for given values of  $\theta$  and  $\beta$ , the optimal group assignment for each household is found as the solution to the following minimization problem:

$$\hat{g}_i(\theta, \beta) = \arg \min_{g \in \{1, \dots, G\}} \sum_{t=1}^T (\ln(C_{i,t}) - \alpha_i - \beta_g \mathbb{1}_{\text{Retired}_{i,t}} - \theta X_{i,t})^2 \quad (8)$$

And then the estimator of  $(\theta, \beta)$  in (7) is rewritten as:

$$(\hat{\theta}, \hat{\beta}) = \arg \min_{(\theta, \beta) \in \Theta \times B^G} \sum_{i=1}^N \sum_{t=1}^T (\ln(C_{i,t}) - \alpha_i - \beta_{\hat{g}_i(\theta, \beta)} \mathbb{1}_{\text{Retired}_{i,t}} - \theta X_{i,t})^2 \quad (9)$$

where,  $\hat{g}_i(\theta, \beta)$  is given by Equation (8). Finally, the estimator of  $g_i$  is obtained as  $\hat{g}_i(\hat{\theta}, \hat{\beta})$ .

## Appendix C: Sensitivity

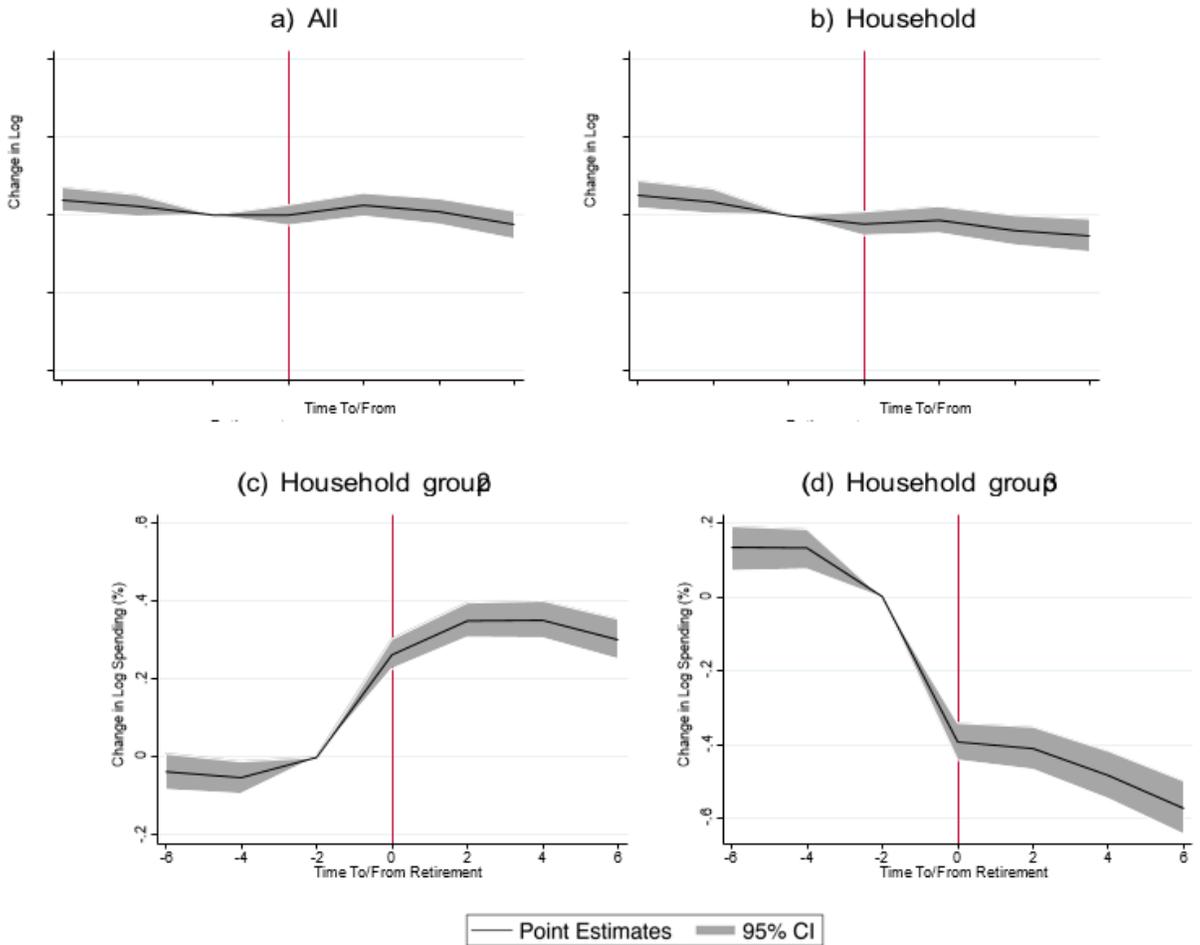
This section gives the sensitivity of our main results to i) the use of a broader measure of consumption (in Section C.1), ii) the use of a sample that excludes the involuntary retired (in Section C.2), and (iii) to a specification of our baseline event studies which rules out preretirement effects (C.3). A final subsection (C.4) shows event studies for our consumption subcomponents allowing flexible effects post-retirement.

### *C.1 Broader measure of consumption*

In our baseline analysis we excluded spending on rent, health care, child care, and education, given that (i) measurement error in self-reported house prices may drive unrealistic fluctuations in imputed rent, (ii) increased health care spending is more likely to indicate the realization of an adverse shock than higher consumption levels, and (iii) child care and educational spending are likely to be driven by changes in household structure, such as children entering and exiting the household. Here we show a selection of results for a broader measure of spending that includes health, education, childcare, and rent.

Figure 7 gives the analogue to Figure 2(a) which illustrates retirement event studies for all households and for each of our three groups. Group membership here is estimated using the broader measure, and a comparison of group membership in each of the two settings is given in Table 7.

**Figure 7: Retirement dynamics for total nondurable spending including health, education, childcare, and rent**



**Notes:** Figures plot the estimated  $\hat{\beta}_{g,h,s}$  from Equation (2) when the dependent variable is the log of a broader measure of spending used in our baseline. Panel (a) is for all households. Panel (b) to (d) are for the three groups we identify in Section 3. Confidence intervals are calculated assuming the allocation to groups is deterministic.

**Table 7: Baseline versus broader measure of consumption group matrix**

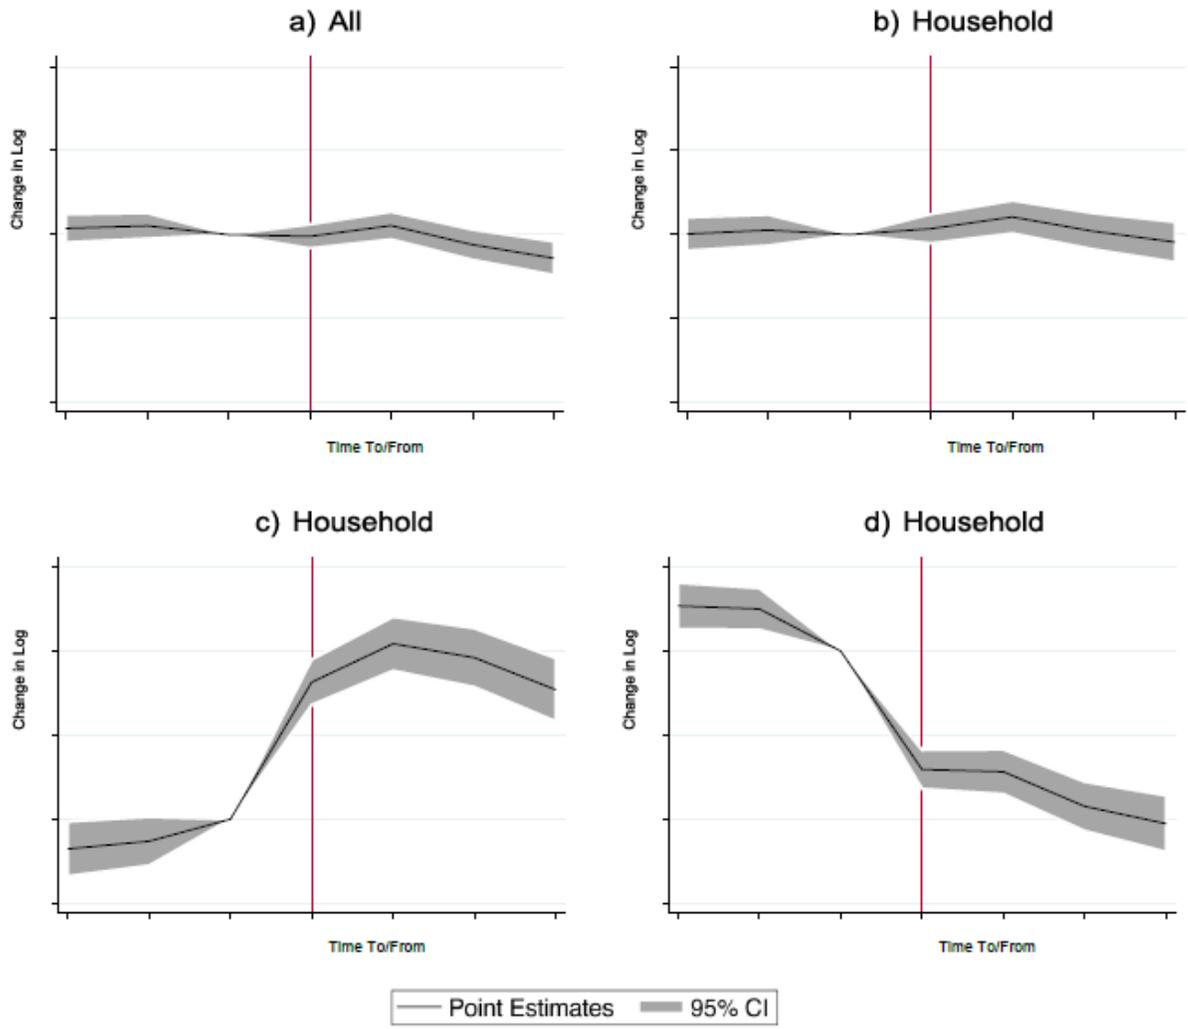
	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Group 1</b>	72.50	26.27	4.07
<b>Group 2</b>	25.50	64.19	43.89
<b>Group 3</b>	2.00	9.53	52.04

**Notes:** Overlap between groups displayed in %. Baseline is along the horizontal axis.

### *C.2 Leaving out the involuntary retired*

Figure 8 gives an analogue to Figure 2 for a sample excluding those who report that their retirement was involuntary. Table 8 compares group membership in both cases for those who are in both samples (where in the baseline group allocation is done using all households).

**Figure 8: Voluntary retirement dynamics for total nondurable spending**



Notes as per Figure 2.

**Table 8: Baseline versus voluntary retirement group matrix**

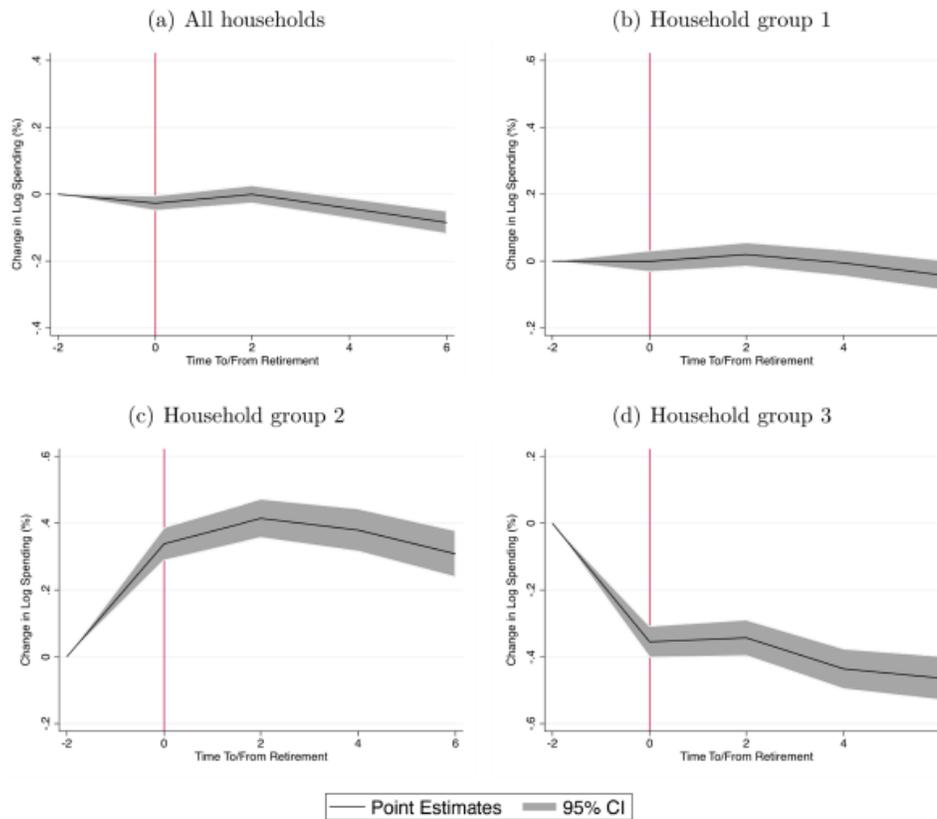
	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Group 1</b>	89.62	0.48	0.00
<b>Group 2</b>	9.84	89.90	0.55
<b>Group 3</b>	0.55	9.61	99.45

**Notes:** Overlap between groups displayed in %. Baseline is along the horizontal axis.

### *C.3 Event studies ruling out preretirement effects*

We discussed in Section 4 that in our event studies, an identification problem, analogous to the age-period-cohort identification problem, would arise if there are households fixed effects, fully dynamic retirement effects, and fully flexible age effects. We avoided this by making a parametric restriction: that the age effects be a polynomial. In this section, we show that we get very similar results for these event studies if we make an additional (over-identifying) restriction on the event study design by ruling out preretirement trends.

**Figure 9: Preretirement dynamics for total nondurable spending**



**Note:** Preretirement trends are not admitted in this specification. Otherwise notes are per Figure 2.

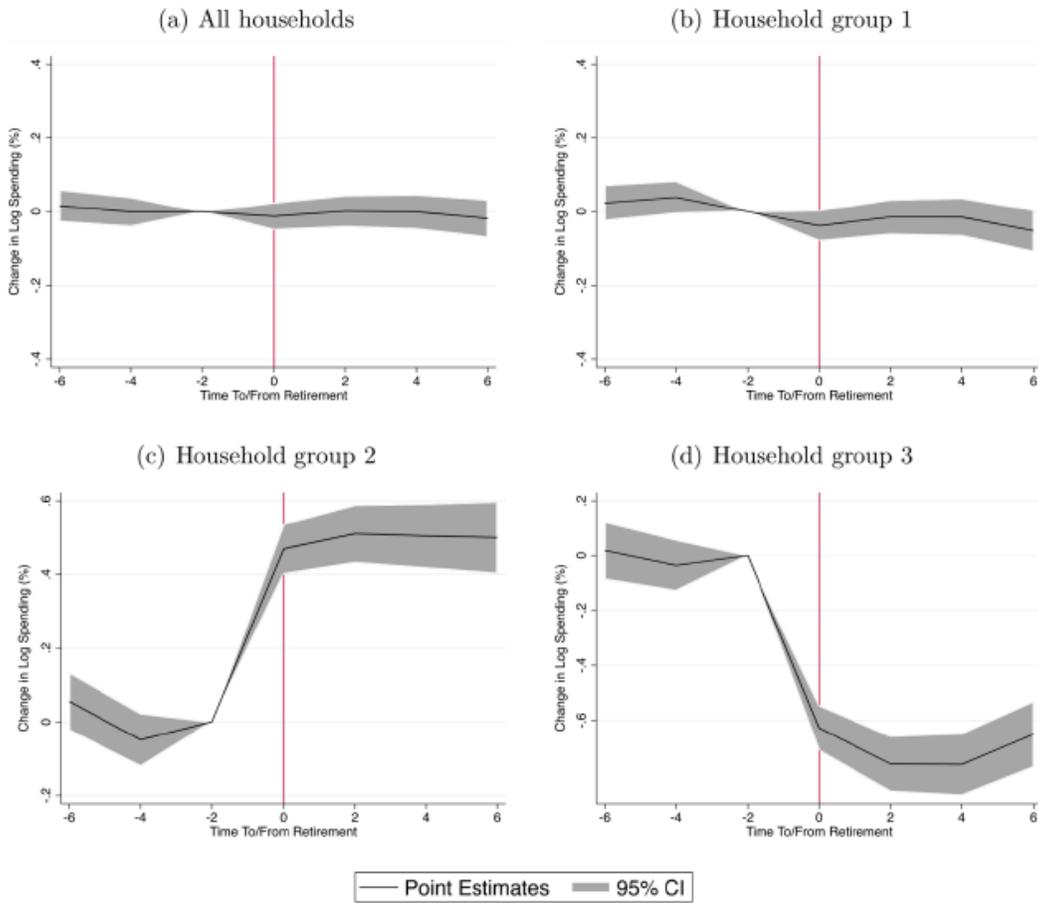
#### C.4 Dynamic graphs for subcomponents

In section C.4, we decomposed the changes in spending upon retirement depicted in Figure 2 into spending changes in four (exclusive and exhaustive) subcomponents of nondurable expenditure: spending on food at home, on food away from home, on transport expenditure, and on house expenditures (excluding rent).

In presenting those results, we replaced the dynamic retirement effects, with a persistent static effect. In this appendix, we show that for each of these categories, the retirement consumption dynamics are well approximated with a static effect. Figure 10

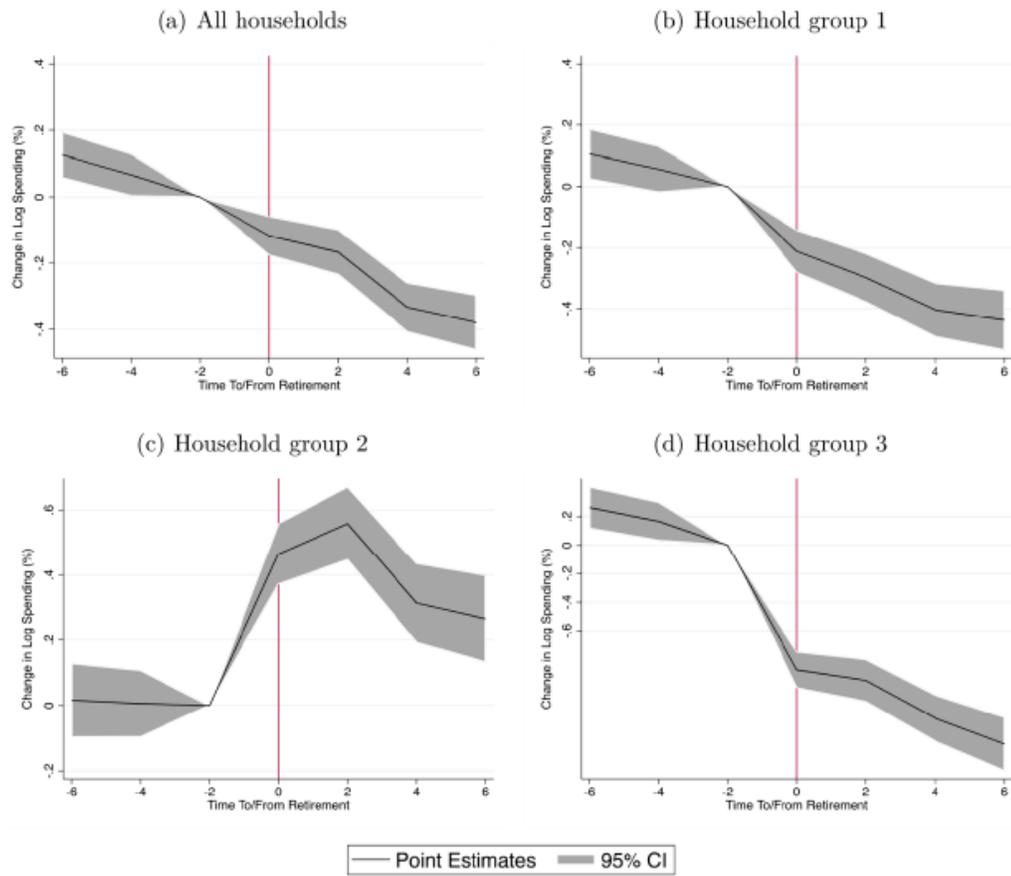
gives event studies for all households and for our three groups for spending on food in the home, Figure 11 gives them for transport expenditure, Figure 12 gives them for household expenditure, and Figure 13 gives it for food outside the house.

**Figure 10: Retirement dynamics for spending on food at home**



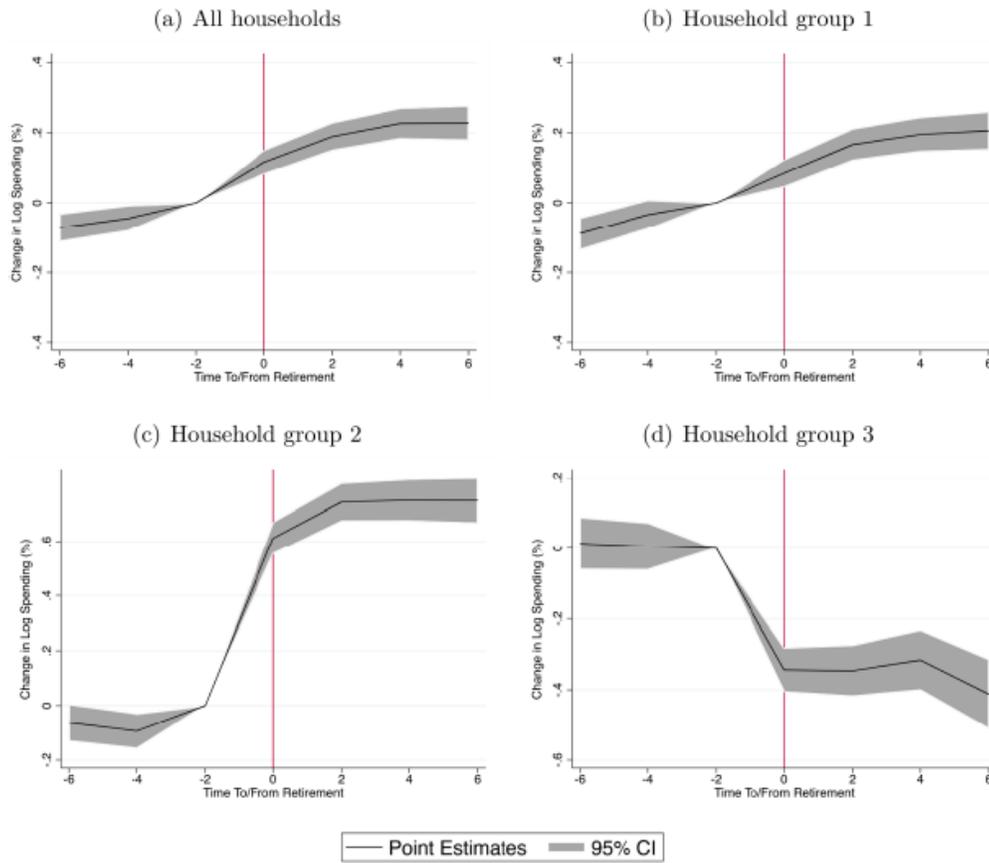
**Note:** As per Figure 2.

**Figure 11: Retirement dynamics for transport expenditure**



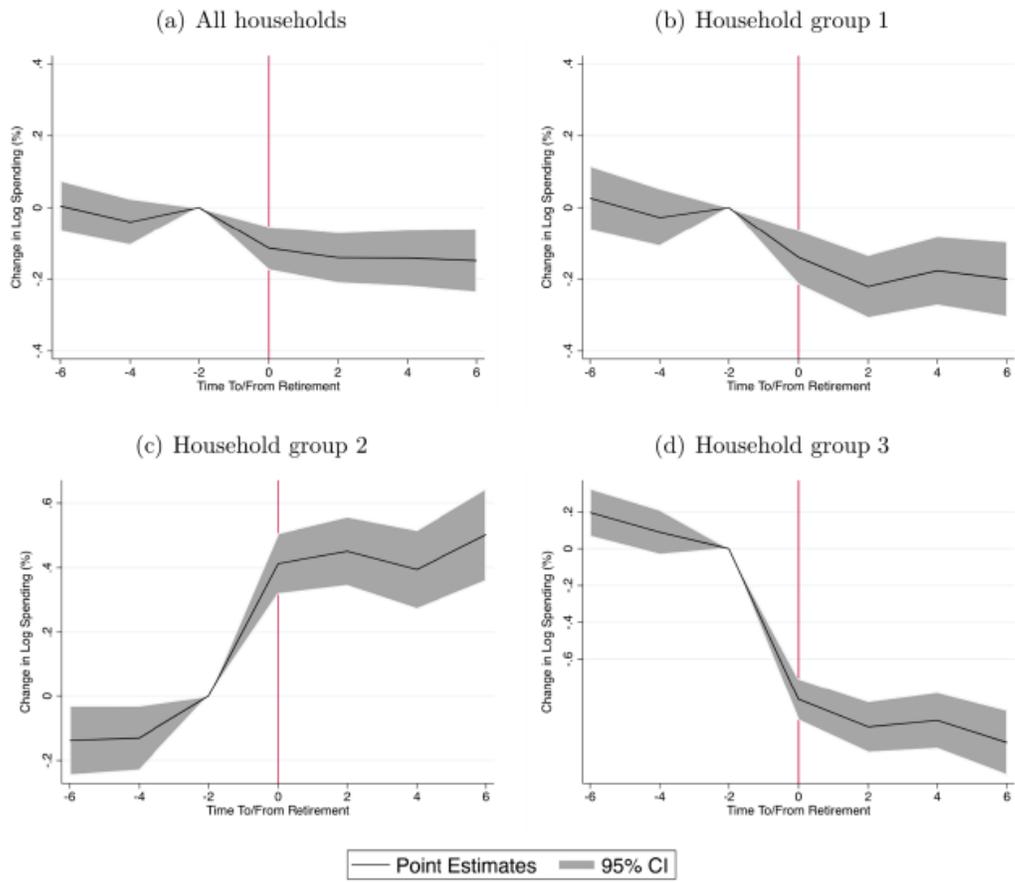
**Note:** As per Figure 2.

**Figure 12: Retirement dynamics for housing expenditure excluding rent**



**Note:** As per Figure 2.

**Figure 13: Retirement dynamics for spending on food away from home**

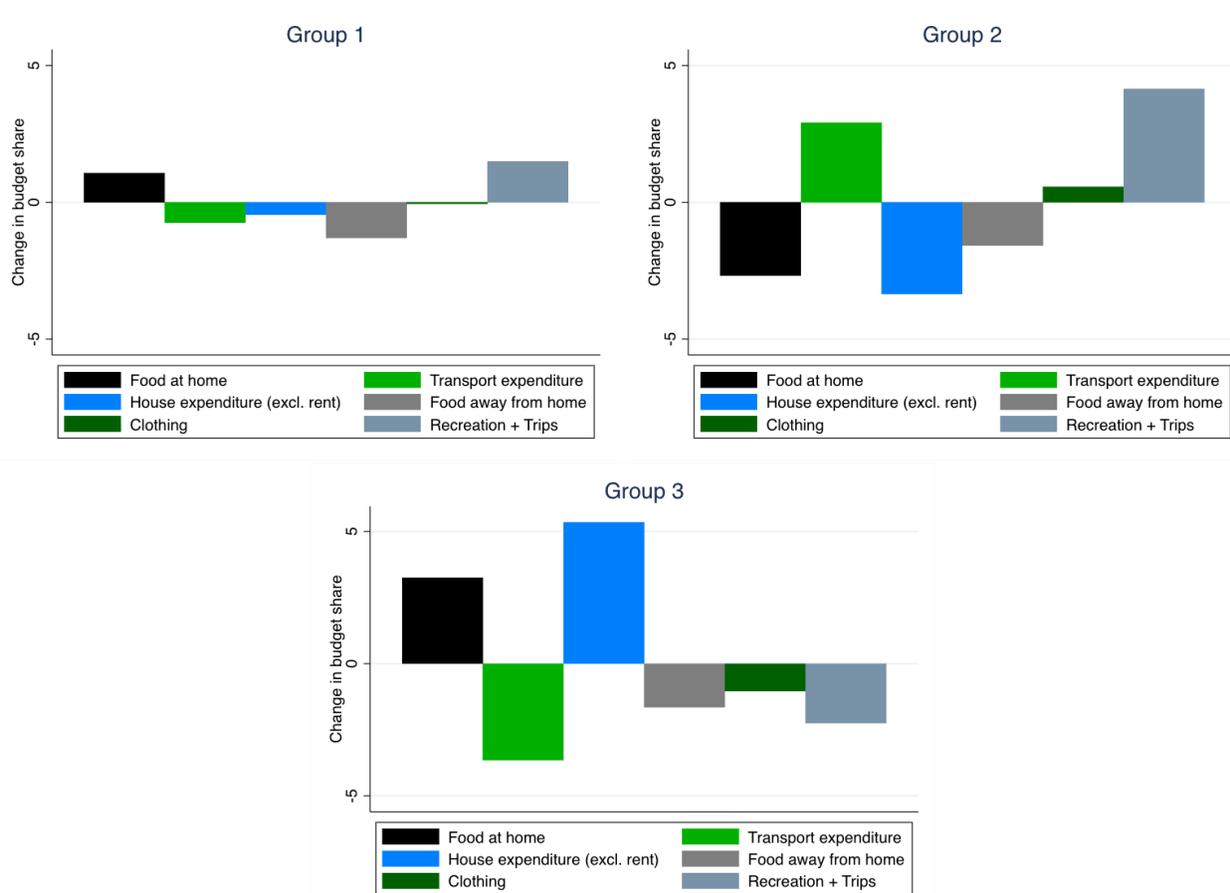


**Note:** As per Figure 2.

## Appendix D: Post-2005 consumption measures

Figure 4 in the body of the paper gives changes in budget shares on retirement for our four baseline consumption groups. Figure 14 analogously gives changes in budget shares on retirement for the years after 2005, when measures of spending on additional components are available.

**Figure 14: Changes in budget shares at retirement**



**Notes:** In these figures, we plot the regression coefficient on group retirement interaction dummies estimated in an equation with fixed effects and family size, unemployment rate, log spending, and quadratic age controls.

## Appendix E: The role of health

We explore differences in health for the three different groups of households that we obtained from our estimator. To this end, we use three different measures of health in the PSID:

- The first is a self-reported measure of health, assessed with the question, “Would you say your health in general is excellent, very good, good, fair, or poor?”
- The second is a summation of the number of health conditions the individual reports. We restrict our analysis to health conditions asked about in all survey waves between 1999 and 2019.
- The third is a measure of disability, assessed with the question, “Do you have any physical or nervous condition that limits the amount of work you can do?”

Table 9 reports each of the health measures for households in the three groups during the wave directly prior to retirement. We find that the three groups differ relatively little in terms of subjective health, ailments, and disability. The largest difference is that Group 2 has fewer ailments than the other groups, however, this effect is relatively small.

**Table 9: Health prior to retirement**

<b>Variables</b>	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Subjective health</b>	.04	-.01	-.09
<b>Ailment index</b>	-.12	-.02	-.07
<b>Disability</b>	.22	.22	.20

**Notes:** This table shows summary statistics by group for the time period preceding retirement. We report the z-score of subjective health (initially measured on a 1 to 5 index). Ailment index is the standardized score for household mean affliction by the following conditions: stroke, hypertension, diabetes, cancer, lung disease, heart attack, heart disease, psychiatric illness, arthritis, asthma, and memory impairment.

Table 10 shows the change in health at retirement for each of the three groups. Here we find that Group 3 has the largest deterioration in health around the time of retirement. For instance, self-reported health falls by 0.31 standard deviations for Group 3, compared to 0.2 standard deviations for Group 1. Similarly, Group 3 has more ailments and disability. In comparison, Group 2 appears relatively similar to Group 1. This suggests that the increase in spending for Group 2 cannot be explained by better-than-expected changes in health at the time of retirement.

**Table 10: Change in health at retirement**

	<b>Group 1</b>	<b>Group 2</b>	<b>Group 3</b>
<b>Subjective Health</b>	-0.20	-0.18	-0.31
<b>Ailment Index</b>	0.39	0.40	0.48
<b>Disability</b>	0.29	0.25	0.44

**Notes:** This table shows the difference in average health outcomes before and after retirement. Three periods preceding retirement and three following retirement are used. Ailment index is the standardized score for household mean affliction by the following conditions: stroke, hypertension, diabetes, cancer, lung disease, heart attack, heart disease, psychiatric illness, arthritis, asthma, and memory impairment.

## **Appendix F: Supplemental variables**

We complement our baseline empirical analysis with additional data on personality traits, time use, and well-being elicited from PSID respondents in two supplemental surveys: the “Disability and Use of Time Supplement” administered in 2009 and 2013 and the one-time “Well-being and Daily Life Supplement” administered in 2016. This appendix lists the survey questions we use from these supplemental surveys.

### *F.1 Disability and use of time (2009)*

- (DF1GWB1) Taking all things together, how satisfied are you with your life these days?
- (DF1GWB4) From 0 to 6 (6 being very satisfied), how satisfied are you with your current financial situation?
- (DF1GWB6) From 0 to 6, how satisfied are you with your marriage?
- (DF1GWB2) From 0 to 6, how satisfied are you with your health?

## *F.2 Disability and use of time (2013)*

- (DF2GWB1) Taking all things together, how satisfied are you with your life these days?
- (DF2GWB4) From 0 to 6, how satisfied are you with your current financial situation?
- (DF2GWB2) From 0 to 6, how satisfied are you with your health?
  
- (DF2GWB6) From 0 to 6, how satisfied are you with your marriage?
  
- (DF2PER1) You are someone who does a thorough job. Does this describe you not at all, a little, some, or a lot?
- (DF2PER5) You are someone who worries a lot. Does this describe you not at all, a little, some, or a lot?
- (DF2PER7) You are someone who tends to be lazy. Does this describe you not at all, a little, some, or a lot?
  - 1: Not at all
  - 2: A little
  - 3: Some
  - 4: A lot
  - 8: DK
  - 9: NA; refused
  
- (DF2SES1) You have usually felt pretty sure your life would work out how you wanted. Or, there have been more times when you haven't been very sure.
  - 1: Usually pretty sure life would work out
  - 2: More times when haven't been very sure
  - 8: DK
  - 9: NA; refused
  
- (DF2SES2) Are you the kind of person that plans your life ahead all the time, or do you live more from day to day?

- 1: Plans life ahead all the time
- 2: Live more from day to day
- 8: DK
- 9: NA; refused

### *F.3 Well-being and daily life supplement*

- (WB16A1) How satisfied are you with your life as a whole these days?
  - 1: Completely satisfied
  - 2: Very satisfied
  - 3: Somewhat satisfied
  - 4: A little satisfied
  - 5: Not at all satisfied
  - 9: NA; not answered
- (WB16A5D) How satisfied are you with each of the following: My financial situation.
- (WB16A5E) How satisfied are you with each of the following: My hobbies.
- (WB16A5F) How satisfied are you with each of the following: My marriage or romantic relationship.
- (WB16A5G) How satisfied are you with each of the following: My family life.
- (WB16A5H) How satisfied are you with each of the following: My friendships.
- (WB16A5I) How satisfied are you with each of the following: My health.
- (WB16A5J) How satisfied are you with each of the following: My faith.
  - 1: Completely satisfied
  - 2: Very satisfied
  - 3: Somewhat satisfied
  - 4: A little satisfied
  - 5: Not at all satisfied
  - 6: Does not apply to me
  - 9: NA; not answered