

The Retirement of a Consumption Puzzle

Erik Hurst

Abstract

This paper summarizes five facts that have emerged from the recent literature on consumption behavior during retirement. Collectively, the recent literature has shown that there is no puzzle with respect to the spending patterns of most households as they transition into retirement. In particular, the literature has shown that there is substantial heterogeneity in spending changes at retirement across consumption categories. The declines in spending during retirement for the average household are limited to the categories of food and work related expenses. Spending in nearly all other categories of non-durable expenditure remains constant or increases. Moreover, even though food spending declines during retirement, actual food intake remains constant. The literature also shows that there is substantial heterogeneity across households in the change in expenditure associated with retirement. Much of this heterogeneity, however, can be explained by households involuntarily retiring due to deteriorating health. Overall, the literature shows that the standard model of lifecycle consumption augmented with home production and uncertain health shocks does well in explaining the consumption patterns of most households as they transition into retirement.

1. Introduction

According to the standard lifecycle model of consumption, forward looking agents will smooth their marginal utility of consumption across predictable income changes such as retirement.¹ However, there is a large literature, discussed in the following section, which documents household expenditure falls precipitously upon retirement. This phenomenon has been referred to as “the retirement consumption puzzle”. Given some of these findings, Attanasio (1999), in his chapter on consumption in the Handbook of Macroeconomics, concludes: “The [lifecycle] model has been fitted with success only to households in the middle of their lifecycle...additional work is needed to understand the behavior of young and elderly households. The [consumption] behavior of retirees, in particular, can be quite difficult to model.” Making an even stronger claim, Bernheim et al. (2001) state: “Contrary to the central tenets of life-cycle theory, there is little evidence that households use savings to smooth effects on consumption of predictable income discontinuities [such as retirement]” (page 854).

In this paper, I review the recent research in expenditure patterns of individuals as they transition into retirement. In doing so, I highlight five stylized facts that have emerged with respect to the behavior of consumption around retirement. First, the fact that certain types of expenditures fall sharply as households enter into retirement is rather robust across data sets within the United States, across data sets from differing countries, and across differences in methodological approach. Second, the declines in expenditures are mostly limited to two types of consumption categories: work related items (such as clothing and transportation expenditures) and food (both at home and away from home). When broader measures of consumption are

¹ For classic references, see Modigliani and Brumberg (1954) and Friedman (1957). The standard lifecycle model usually assumes household utility is separable in consumption and leisure. See Heckman (1974) for a lifecycle model with utility being non-separable between consumption and leisure.

analyzed or when expenditure categories that exclude food and work related expenses are analyzed, the measured declines in spending upon retirement are either close to zero or are increasing. As a result, the retirement consumption puzzle is a bit of a misnomer. The fact that work related expenses fall upon retirement is in no sense puzzling when viewed through the lens of standard consumption models. What is potentially puzzling is why food expenditures decline sharply at the time of retirement while the rest of the household's consumption bundle remains relatively constant.

The third stylized fact that I highlight is that actual food intake (as measured by the quantity and quality of one's diet) remains constant through retirement despite the fact that food expenditures fall sharply. Appending the standard lifecycle model with a Becker (1965) model of home production would generate both sharply falling expenditures and relatively constant actual consumption upon retirement for consumption categories that are amenable to home production. Consistent with this fact, there is ample evidence within the literature documenting that retirees spend much more time on food production (preparing meals and shopping for groceries) than their non-retired counterparts.

The fourth stylized fact with respect to changing consumption of retirees is that there is substantial heterogeneity across individuals in the population with respect to changing expenditures in retirement. Specifically, the literature shows that the declines in expenditures are greatest for households that have little accumulated wealth prior to retirement. Moreover, the models of home production are unable to explain the expenditure declines for this segment of the population. In other words, even though the canonical model of consumption augmented with home production can explain consumption patterns for most households as they transition

into retirement, there is a segment of the population (perhaps as much as 25 percent) that experiences real consumption declines upon retirement.

The fifth stylized fact is that these households that experienced real consumption declines upon retirement often had experienced involuntary retirement. Moreover, these involuntary retirements are often due to health shocks. Again, standard lifecycle models would predict declining spending upon retirement if retirement status were accompanied by an unexpected change in lifetime resources. In particular, households that experience severe health shocks prior to their planned retirement date are often forced to retire early. The early retirement reduces their lifetime resources (as they are working less years) and the health shock often is accompanied by increased out of pocket medical expenses. Both effects, according to any standard lifecycle model of consumption, would result in a decline in non-medical related consumption.

Overall, the aggregation of the evidence from the existing literature suggests that the claims of Attanasio (1999) and Bernheim et al. (2001), recounted above, may have been premature. The declines in expenditures, aside from work related expenses, primarily occur in food and the declines are largest for those who involuntarily retired. Now, this is not to discount the possibility that some households are myopic with respect to their consumption decisions (or have time inconsistent preferences). It is just that these households are only a relatively small fraction of the total population. As a result, one should conclude that standard models of lifecycle consumption augmented with home production and uncertain health shocks does well in explaining the consumption patterns of most households as they transition into retirement.

One final point needs to be stressed. The fact that households smooth their consumption as they transition into retirement does not imply that they have saved adequately for retirement.

It is possible that households who planned insufficiently would not learn about their saving shortfall until well after they have retired (perhaps when they receive a future health shock). The goal of this chapter is to show that the perceived wisdom that the average household cannot managed to sustain their well being as they transition into retirement is misguided not that households have the ability to sustain their consumption throughout their entire retirement.

The remainder of the paper is organized around summarizing the existing literature to provide support for the five stylized facts discussed above. The last section of the paper concludes and offers some additional perspective by drawing on the literature about whether or not households save adequately for retirement.

2. Documenting Stylized Facts About Changing Expenditures at Retirement

A large volume of research has emerged during the last decade examining household expenditures at the time of retirement. In this section, I discuss this recent work.² In doing so, I draw particular attention to a fact that emerges when aggregating results across papers: the extent to which expenditures decline in retirement varies with the measure of consumption examined. In particular, essentially the declines in expenditures at the time of retirement documented within the literature occur within two consumption categories: work related expenses (clothing and transportation costs) and food (meals at home and meals away from home).

The fact that work related expenses decline in retirement is not at all surprising. Any model that has some expenditures which are strong complements with working (such as business attire) will predict those expenditures will fall as households exit the labor force. However, to the extent that food is a large share of households' budgets and is often considered a necessity

² Although, this survey focuses on recent innovations in retirement spending, evidence about the existence of a retirement consumption puzzle extends back over two decades. The classic reference in this early literature is Hamermesh (1984). Hamermesh used data from the 1973 and 1975 Retirement History Survey to show that expenditures of retirees fall sharply in the first few years after retirement.

with a relatively low income elasticity, food spending declines in retirement could be seen as a puzzle. As I discuss below, what the literature has documented as a “retirement consumption puzzle” is a misnomer. In actuality, the literature has primarily documented a “retirement food consumption puzzle”. Moreover, the true puzzle is why food expenditures fall sharply *despite* the fact that the remaining portion of the households non-work related/non-food expenditures remain roughly constant.

One recent paper that focuses on the differential lifecycle spending patterns for different consumption categories is Aguiar and Hurst (2007b). Unlike many other studies, which are discussed below, Aguiar and Hurst (2007b) use the Consumer Expenditure Survey (CEX) as opposed to other micro data sets within the United States like the Panel Study of Income Dynamics (PSID), the Retirement History Survey (RHS), or the Health and Retirement Survey (HRS). The CEX differs in two ways from these other surveys. First, it has broader measures of consumption than the other surveys. Second, the CEX is a cross-sectional survey with only a short (4 quarter) panel component. The PSID, HRS, and RHS follow the same individual over much longer periods of time. To get around this potential drawback, Aguiar and Hurst use the large amount of cross sections from the CEX to create a synthetic panel by following a given cohort over time. In doing so, they compare the spending patterns of all households of a given cohort between the ages of 60 and 62 to the spending patterns of all households of the same cohort between the ages of 66 and 68 (six years later).

Aguiar and Hurst (2007b) examine the lifecycle profile of spending in 11 categories of non-durable consumption. The categories include: total food, alcohol and tobacco, non-durable transportation, clothing and personal care, domestic services, entertainment services, utilities, charitable giving, net gambling receipts, all other non-durable expenditures (including business

services) and housing services. They exclude education and health expenditures from their measure of non-durable expenditures. As is often done in the literature, Aguiar and Hurst compute the service flow of housing for homeowners by using the self reported answer to the question of what the homeowner would charge (net of utilities) to someone who wished to rent their housing structure today. For renters, the service flow of housing is their monthly out of pocket expenditures on rent. All expenditure are adjusted for differences in household size.

The main finding of Aguiar and Hurst (2007b) is to show that between the early 60s and the late 60s total non-durable spending plus housing services falls by roughly 5 percent. However, there is substantial variation across categories of consumption. In particular, spending on total food, clothing, and non-durable transportation falls, respectively, by 7 percent, 18 percent, and 15 percent. Conversely, spending on housing services, utilities, charitable giving, net gambling receipts, and entertainment remain constant or rise during the retirement years. For example, between the early 60s and the late 60s, entertainment spending increases by 5 percent and charitable giving increases by 4 percent. Value weighted, the categories that decline at retirement and the categories that do not decline at retirement both represent fifty percent of household total non-durable expenditures. After excluding food, clothing, and non-durable transportation from their measure of consumption, Aguiar and Hurst find that expenditures actually increase by 6 percent between the early 60s and the late 60s.

These results are hard to reconcile with households being ill prepared for retirement. Why is it that households would forego food (a necessity) while simultaneously increasing their spending on entertainment (going to the movies, golf games, vacations) and charitable giving? In their paper, Aguiar and Hurst argue that Becker's 1965 model of consumption commodities can explain well the lifecycle patterns of different consumption categories. Spending on goods

that are complements to time (like entertainment) should increase in retirement while spending on goods that are substitutes to time (like food production) should fall during retirement.

The results of Aguiar and Hurst are consistent with the results documented by Fisher et al. (2006) who also use CEX data from a similar time period but who employed a different empirical strategy. Fisher et al. compared the spending patterns of *non-retired* households between the ages of 60 and 64 to the spending patterns of *retired* households between the ages of 65 and 69 five years later. In doing so, they are assuming that retirement status prior to and after the age of 65 is completely exogenous to factors that determine household consumption. This assumption is likely not valid given that households will also optimize over their choice of retirement age. Moreover, Fisher et al. only look at total spending (with and without housing services) and food spending.

Subject to these caveats, Fisher et al. find that most of the action of the decline in total expenditures at the time of retirement occurs within the food categories (food at home and food away from home). Specifically, for their third cohort, food at home and food away from home fell upon retirement by roughly 8.3 percent and 15.9 percent, respectively.³ The corresponding change over these age ranges for their broad measure of consumption excluding housing services and their broad measure of consumption including housing services were -3.1 percent and -1.2 percent. Again, the declines in food spending associated with retirement for CEX households were much larger than the declines in total spending.⁴

³ Fisher et al. (2006) create three separate 5 year birth year cohorts within the CEX data. We report the results for the third cohort, which in terms of magnitude was in the middle of the first and second cohorts. The first cohort (the cohort that retired during the mid 1990s) looked different than the second and third cohorts. While the second and third cohorts experienced substantial declines in food spending upon retirement, their first cohort experienced little change in food spending upon retirement.

⁴ Fisher et al (2006) only broke out food as a separate consumption category. They did not separately analyze other categories such as work related expenses or entertainment.

The main take away from Aguiar and Hurst (2007b) and Fisher et al. (2006) is that the decline in food spending and work related expenditures are the primary factors explaining the declining expenditures at the time of retirement within the United States. The changes in expenditures within other categories are either close to zero or increasing as households transition to retirement. In particular, spending on luxury goods, like entertainment services, actually increase as households transition into retirement.

One paper which provides a different view of the spending patterns of retirees using CEX data is Laitner and Silverman (2005). Using a different methodology than Aguiar and Hurst (2007b) and Fisher et al. (2006), they find that total expenditures drop by 16 percent upon retirement. Although they do not disaggregate the consumption decline into separate categories, their decline in total spending is much larger than the declines found by Aguiar and Hurst and Fisher et al. using nearly identical datasets. The reason for the difference is that their structural model which estimates the change in spending at retirement allows for an age effect in addition to the separate retirement effect. As discussed above, the underlying data show that after accounting for cohort effects, total non-durable expenditures decline slightly for households of a given cohort in the CEX as age through their 60s.

What specifically does drives the large decline in spending documented by Laitner and Silverman at the time of retirement? The answer is that they have a large offsetting positive age effect for households during their 60s which offsets the negative retirement effect. In other words, a household's desired consumption level would be to have an increasing consumption profile during their 60s. But, as almost all households retire during their 60s, consumption becomes suppressed. The net result of the positive age effect and the negative retirement effect is to have household consumption be roughly constant during their 60s. The reason, I believe,

that they estimate such a strong age effect is that households with higher permanent income retire later than households with lower permanent income. If this is true, the expenditure levels, on average, for people of a given cohort who are working at age 70 will be higher than the expenditure levels, on average, for people of that same cohort who were working at age 60. As a result, their model estimates a positive age effect for non-retired individuals. Given that these two parameters are fitting a flat aggregate expenditure profile, the large positive age effect forces a large negative retirement effect. The assumption made by Laitner and Silverman (2005) that the timing of retirement is exogenous to factors that determine expenditure levels is likely driving their strong negative estimate of expenditure declines at the time of retirement. Regardless, the underlying data suggest that total non-durable expenditure profiles (inclusive of food, clothing and non-durable transportation) fall slightly during the 60s and total non-durable expenditure profiles (exclusive of food, clothing, and non-durable transportation) increase slightly during the 60s.

The result of differential expenditure declines by consumption category at the time of retirement is not limited to data from the United States. In one of the early seminal papers in this literature, Banks et al. (1998) using data from Britain's Family Expenditure Survey (FES) derive similar conclusions to Aguiar and Hurst (2007) and Fisher et al. (2006). Creating pseudo panels, Banks et al. examine the evolution of total non-durable spending as household transition into retirement. They conclude that: "Whereas the anticipated fall in consumption growth is around 2 percent, actual consumption growth at retirement falls by as much as 3 percent." In other words, they find evidence of a retirement consumption puzzle in that expenditures are declining more at retirement than would be predicted by a standard lifecycle model. Moreover, they document that the declines in food expenditures and the declines in work related

expenditures (including canteen and restaurant meals, transport, and adult clothing) were much larger than the decline in total non-durable expenditures. Again, a key fact that emerges from the Banks et al. study is that food expenditures at retirement decline much more sharply than does the expenditures on other non-food/non-work related categories.

Similarly, using data from the Italian Survey of Family Budgets (ISFB), Miniaci et al. (2003) analyze consumption declines by consumption category for Italian households. Again, given their lack of panel information for a given household, they analyze synthetic panels following different cohorts over time. Their results for Italy are consistent with the results for the United States and the results for Britain. Miniaci et al. also analyze a much broader set of consumption categories. The only decline in expenditures for retired Italian households occurred in either work related categories (clothing and transportation) or food (food at home and food away from home). All other components of non-durable consumption either remained constant or actually increased through retirement years (households in their 60s). These remaining categories include health expenditures, fuel expenditures, and other housing expenditures. Again, their results show that to the extent that non-durable consumption falls in retirement, it is mostly (if not completely) driven by work related expenditures and food expenditures.

Battistin et al. (2006) also studied the retirement consumption puzzle among Italians using the Bank of Italy Survey on Household Income and Wealth (SHIW). Their work improves upon the work of Miniaci et al. because they use a regression discontinuity approach to instrument for retirement status. This approach offers traction given that different Italians are eligible for the state provided pension at different ages (and as a result, the incentive to retire at a given age differs among the different groups). Using this identification, they find that non-durable consumption falls by roughly nine percent as households transition to retirement.

However, they also examine the decline in a variety of different consumption categories. The greatest declines in spending were in meals away from home, clothing, and transportation. Moreover, they state: “The conclusion that we draw from this exercise is that our estimated retirement consumption drop could well be due entirely to a reduction of work-related expenses and a substitution away from market goods to home production of food” (page 17).

While a decline in work related expenses can be consistent with the standard lifecycle model of consumption with work specific expenditures (such as formal dress and work related transportation), the decline in food expenditures is harder to explain. Given that food is a necessity, and therefore has a small income elasticity, many authors have argued that analyzing food expenditures provides a strong test of consumption smoothing during retirement. The prevailing view was that if retired households do not smooth food expenditures then it is unlikely that they will smooth spending on other components of their consumption bundle. As a result, the majority of the other papers that documented a retirement consumption puzzle within the United States used exclusively (or included) data on food expenditures. Moreover, it is the synthesis of this work on food that has encouraged researchers to allege the existence of a retirement consumption puzzle.

For example, Bernheim et al. (2001) use panel data on households from the Panel Study of Income Dynamics (PSID) to examine changes in household spending at the time of retirement. Their measure of consumption includes food expenditures at home, food expenditures away from home, and the imputed or actual rental value of one’s residence. They show results for their composite measure of consumption and separately for food at home and food away from home. Results for a measure of consumption that only includes the imputed or

actual rental value of one's residence at the time of retirement are not shown. They find that, on average, their composite expenditure measure falls by 14 percent.

A variety of other studies confirm that food expenditures drop sharply upon retirement. Hurst (2006) uses a different methodology and a different time period from the PSID and finds similar results with respect to food spending. Following a given household through retirement, food spending, on average, declines by 12 percent at the median. Likewise, Haider and Stephens (2007) use panel data from the Retirement History Survey (RHS) and find that households that retire when expected experienced a 10 percent decline in food expenditures, on average. However, Haider and Stephens also analyze data from the Health and Retirement Survey (HRS) and find no decline in food spending among the recently retired. This latter result is interesting in the sense that it is the only study that found that food expenditures do not decline sharply with the incidence of retirement. Fisher et al. (2006) suggest that either period effects or cohort effects from the late 1990s may explain the lack of findings in the HRS data analyzed Haider and Stephens (2007).⁵

Hurd and Rohwedder (2003, 2005) and Ameriks et al. (2007) take a different approach to analyzing changes in spending at the time of retirement by using retroactive survey data. Instead of using the data sets described above where households were asked about their spending patterns during the last month or during the last quarter, the survey data used by Hurd and Rohwedder (2003, 2006) and Ameriks et al. (2007) asked household to retrospectively assess how much their expenditures fell upon retirement. For example, consider an individual in their survey who is currently 69 years old but who retired when he was 63. In this example, the individual would be asked to recount his change in spending from six years earlier.

⁵ Lundberg et al. (2003) documents declines in food expenditures within the PSID for married households.

Hurd and Rohwedder use data from the HRS and a supplemental survey to the HRS, the Consumption and Activities Mail Survey (CAMS). The CAMS survey asked current retirees to report how their total spending changed with retirement. They do so in two steps. First, they reported the direction of the change in spending at the time of retirement (increase, decrease, stay the same). Second, the household was asked to report the percentage change in spending if they reported that their spending increased or decreased. Using a very different methodology than the earlier surveys, they find that, on average, total spending fell by roughly 14% at the time of retirement. The median decline in spending, however, was zero. This corresponds almost exactly to the median results on total spending changes reported by Fisher et al. discussed above.

Ameriks et al. (2007) use data from two separate surveys of TIAA-CREF participants: The Survey of Participant Finances (SPF) and the Survey of Financial Attitudes and Behavior (FAB). Similar to the CAMS data, households were asked to assess the direction of their change in spending at the time of retirement and the amount of the change. The TIAA-CREF samples differs from the CAMS sample in the sense that TIAA-CREF respondents are much more educated and much wealthier than the households in CAMS. Within the samples from the TIAA-CREF surveys, retirees again, at the median, experienced no decline in total spending at the time of retirement.

One thing that distinguishes the work of Hurd and Rohwedder (2003 and 2006) and Ameriks et al. (2007) from the other work on the retirement consumption puzzle is that their surveys ask pre-retired households about their expected declines in spending upon retirement. Specifically, pre-retired households in both surveys were asked to report whether they expected their total spending to increase, decrease or stay the same upon retirement. If their spending were to change, the respondents were asked to provide the percentage amount of the change.

Collectively, their work answers the question of whether pre-retired households expect their expenditures to fall upon retirement. Hurd and Rowhder show that nearly seventy percent of pre-retired respondents in CAMS actually expected their expenditures to fall in retirement. Ameriks et al. report that nearly sixty percent of pre-retired households in the TIAA-CREF data expect to decrease their expenditures upon retirement. This research sheds light on the possible mechanisms as to why spending decreases upon retirement. Whatever the reason that results in expenditures falling upon retirement, that reason - for most households - is forecastable well in advance of their actual date of retirement.

What does this synthesis of all of the above research allow one to conclude about the retirement consumption puzzle? There is evidence from micro data across many countries showing that household expenditures drop precipitously at the incidence of retirement. However, collectively analyzing these studies shows that most of the declines are found in work related expenditures and in food expenditures. Broader measures of consumption always show less of a decline than the narrow categories of food or work related expenses. Moreover, although it is rarely documented directly, it appears that measures of total expenditures excluding food and work related items, remain relatively constant as households transition in to retirement. Furthermore, the declines in spending at retirement are predictable by the household prior to their actual retirement.

From a standard lifecycle perspective, it makes sense that expenditures that are complements with working (i.e., professional clothing) should fall when households exit the labor force. However, the standard lifecycle model of consumption, without augmentation, would have a difficult time explaining why food expenditures fall and the rest of the consumption bundle remains relatively constant. If that is truly the case, the retirement

consumption puzzle should be more appropriately named the “retirement food consumption” puzzle.

3. Explaining the “Retirement Food Consumption Puzzle”

How is it that food expenditures fall sharply at the time of retirement relative to pre-retirement trends while the rest of the consumption bundle (aside from work related expenditures) remain relatively constant as households transition into retirement relative to pre-retirement trends? Aguiar and Hurst (2005) propose an explanation. Standard tests of the permanent income hypothesis (PIH) using data on nondurables typically equate consumption with expenditure. However, as noted by Becker (1965), consumption is the output of “home production” which uses as inputs both market expenditures and time. To the extent possible, individuals will substitute away from market expenditures toward time spent in home production, including more intensive searching for bargains, as the relative price of time falls. Retirees have a decreased opportunity cost of time relative to their pre-retired counterparts and, as a result, should be able to engage in non-market production to reduce the cost of their consumption bundle while keeping their actual consumption intake relatively constant. Such a model, if true, would be most applicable to explaining the behavior of food expenditures during retirement given that food is amenable to home production.

Using a variety of data sources, Aguiar and Hurst explore how actual food consumption changes in retirement. To begin, they exploit a novel dataset - the Continuing Survey of Food Intake of Individuals (CSFII), conducted by the U.S. Department of Agriculture – which tracks the dollar value, the quantity, and the quality of food consumed within U.S. households. Using a variety of statistical tests, they find no actual deterioration of a household’s diet as they

transition into retirement.⁶ Actual food consumption does not decline despite the declining expenditure.

How do retirees maintain their food consumption despite their declining food expenditures? Aguiar and Hurst draw on the literature of time allocation set forth in Becker (1965). If Becker's theory of consumption commodities is correct, the mechanism by which households could reduce their food expenditures while keeping their food consumption constant would be through an increased allocation of time towards food production. Using detailed time diaries from National Human Activity Pattern Survey and from the American Time Use Survey, Aguiar and Hurst (2005) show that retirees dramatically increase their time spent on food production relative to otherwise similar non-retired households. The fact that retirees allocate more time to non-market production than their non-retired counterparts was also shown by Hurd and Rohwedder (2003 and 2005) using the CAMS supplement to the HRS and was shown by Schwerdt (2005) using data from the German Socio-Economic Panel.

In separate work, Aguiar and Hurst (2007a) examine the mechanism by which retirees reduce their spending on food. Is it that retirees are shopping more frequently and, as a result, are paying less for their exact same food consumption bundle? Or, are they actually switching their consumption bundle from relative expensive pre-made groceries (like using the grocery store's salad bar to purchase a pre-made salad) to relatively cheaper raw ingredients which they can combine themselves into a meal (like buying all the vegetables separately and chopping them up themselves to make the salad). Using household specific data from the ACNielsen company, which tracks the purchases of the household at the UPC level and links those purchases to detailed information about the purchaser, Aguiar and Hurst (2007a) shows that holding constant

⁶ Actually, Aguiar and Hurst (2005) find that measured food consumption increases slightly as households transition into retirement. This would be consistent with a modest substitution effect resulting from the fact that the price of "producing" a unit of food has declined after retirement.

the exact good (as measured by UPC code), retirees pay lower prices for their grocery bundle than slightly younger non-retired households. They conclude that roughly 20 percent of the declining expenditures on food for older households can be attributed to increased shopping intensity resulting in lower prices paid for the same good. The remaining 80 percent, they find, is due to increased amounts of home production.

Broadly, their results suggest that retired households should experience a slight decline in non-food items simply resulting from the increased shopping intensity of retired households. This is consistent with the facts in Aguiar and Hurst (2005) which shows retired households spend 60 percent more time shopping for non-food goods than their non-retired counterparts.

4. The Heterogeneity of Expenditure Declines across Individuals

From a decade of recent work, we have learned three things about the retirement expenditure puzzle: 1) declines in expenditure, on average, are expected upon retirement, 2) the bulk of the decline in expenditures at retirement are concentrated among work related expenditures and food, and 3) the decline in food expenditures can be explained by an increase home production of food by retirees in the sense that the time allocated to food production goes up dramatically in retirement and actual food intake does not change in any meaningful way as households retire. Aside from the results discussed above, the literature has also demonstrated one additional fact about changes in expenditure among retirees: there is a tremendous amount of heterogeneity in the change in expenditure experienced by retirees.

In one of the seminal papers in this literature, Bernheim et al. (2001) use annual panel data from the PSID to document the heterogeneity in expenditure changes at the time of retirement. As discussed above, the panel data allows the authors to follow a given household as

they transition through retirement. One of the most innovative parts of their research is that they examined food consumption declines for individuals with differing amount of retirement resources.⁷ They characterized household resources along two dimensions: 1) accumulated total assets prior to retirement relative to pre-retirement non-asset income and 2) post-retirement non-asset income relative to pre-retirement non-asset income. Their prior is that households with higher accumulated assets prior to retirement or higher income replacement rates post retirement should be better able to maintain consumption during retirement.

The results of Bernheim et al. are striking. First, they show that that essentially all households based on pre-retirement wealth and post-retirement income replacement rates experienced a decline in (primarily food) expenditure during retirement. However, they also show that the declines in expenditure were greatest for households with the lowest amount of retirement resources. For example, households within the lowest pre-retirement wealth quartile (irrespective of post-retirement income replacement quartile), experienced a 31.2 percent decline in expenditures up to four years after retirement. The comparable expenditure declines for households within the second, third, and top pre-retirement wealth quartiles (irrespective of post-retirement income replacement quartile) were 13.8 percent, 13.9 percent, and 8.9 percent, respectively.⁸ In other words, the declines in expenditures for the wealthiest households (top pre-retirement wealth quartile) were similar to the declines in expenditures for households in the second and third pre-retirement wealth quartiles. Those households in the bottom pre-retirement wealth quartile, however, experienced a much larger decline in expenditures upon retirement.

⁷ As discussed above, Bernheim et al. use a composite consumption measure which is based on food consumed at home, food consumed away from home, and the implicit or actual rental cost of housing.

⁸ These statistics come from using the statistics reported in Table 2a and Appendix Table A1 of Bernheim et al. (2001).

While the declines in food expenditures for the households in the top three wealth quartiles can be explained by changing home production and shopping activities, such a modification to the lifecycle model has a hard time matching the magnitudes of the decline in expenditures for households in the bottom quartile of the wealth distribution. To this end, Aguiar and Hurst (2005) find evidence supporting the conjecture of Bernheim et al. While the average household in the data examined by Aguiar and Hurst did not experience any decline in actual food intake associated with retirement, households with very little accumulated wealth (less than \$1,000 of non-pension assets) did experience some decline in the quantity and quality of food intake associated with retirement. Aguiar and Hurst conclude that: “Average households are modeled well by the PIH in the sense that they smooth consumption across predictable income shocks such as retirement. However, there may be a segment of the population with very low (pre-retirement) wealth that experiences a measured consumption decline upon retirement.” (pg. 939)

Other researchers have confirmed the general finding of Bernheim et al. (2001) showing a large heterogeneity in expenditure decline associated at the time of retirement. For example, Hurst (2006) uses PSID data and regresses pre-retirement wealth on a full vector of income and demographic variables. Then, he splits households into a sample with low pre-retirement wealth residuals (bottom 20 percent) and all other households. He shows that the food expenditure declines associated with retirement are twice as large for those households with low pre-retirement wealth residuals compared to other households (20 percent declines vs. 10 percent declines).

Using survey evidence on retrospective consumption changes at the time of retirement from the CAMS supplement to the HRS, Hurd and Rohwedder (2003) document an extremely

large amount of heterogeneity in expenditure changes at the time of retirement. Specifically, they report that only slightly over half (53.1 percent) of households that are currently retired reported experiencing a decline in total expenditure at the time of retirement. Of the remaining, 11.5 percent reported experiencing an *increase* in total expenditures at the time of retirement while 35.5 percent reported that retirement was associated with no change in total expenditures. As in Bernheim et al., the actual decline in expenditure at the time of retirement increased as net worth declined. Households in the lowest wealth quartile experienced a 22 percent decline in actual expenditure while households in the second, third, and top wealth quartile experienced 17, 13, and 7 percent declines, respectively.

Ameriks et al. (2007), using their survey of TIAA-CREF participants find results similar to those reported by Hurd and Rohwedder. Specifically, 47 percent of retired households reported experiencing a decline in total expenditures at the time of retirement while 22 percent experienced an increase in expenditures at the time of retirement. As in the other studies, the decline in expenditure was largest for those with low wealth.⁹ These results are encouraging given the fact that the TIAA-CREF sample analyzed by Ameriks et al. is much more educated and much more likely to be high income than the nationally representative sample of CAMS participants analyzed by Hurd and Rohwedder.

Given these results, the focus of changes in expenditures in retirement should be limited to the minority of households who enter retirement with very low wealth and, as a result, experience very dramatic declines in expenditures at the time of retirement relative to other households with higher amounts of wealth.

⁹ Using the German Socio-Economic Panel, Schwerdt (2005) also found similar evidence. He found that households with low retirement income replacement rates experienced much larger expenditure declines than households with high retirement income replacement rates.

5. The Role of Unanticipated Retirement in Explaining the Heterogeneity

One of concern that motivated the identification of the “retirement consumption puzzle” is that retirement is often endogenous to life events that change the household’s consumption trajectory. Among the most commonly cited causes of involuntary retirement are health shocks. McClellan (1998) finds that workers who have worse health are more likely to have subsequent negative health shocks and are more likely to retire early. Hurd and Rohwedder (2005) report that 29 percent of the CAMS sample report that adverse health was “very important” or “moderately important” for their decision to enter retirement.

A health shock can affect the optimal consumption decision in multiple ways. First, households who are forced to retire earlier than expected will likely experience a sharp permanent decline in their lifetime resources. According to standard lifecycle theories, such a shock should cause a household to optimally lower their level of consumption, all else equal. As a result, one should expect to see declining consumption growth as households transition into their retirement.¹⁰ Also, health shocks should cause a reallocation of the consumption bundle, all else equal, towards health expenditures away from other consumption categories. If the measure of consumption excludes health expenditures, one may observe declining expenditures in retirement. Third, health shocks often affect consumption needs. For example, someone stricken with a severe illness that affects their ability to work may also have decreased appetite causing them to spend less on food during a given period. Lastly, the health shock could alter the household’s expected length of life. Again, according to standard consumption theories, an abrupt change in the planning horizon will alter the household’s consumption path.

¹⁰ This is the view expressed by Banks et al. (1998) which states: “We argue that the only way to reconcile fully the fall in consumption [at retirement] with the life-cycle hypothesis is with the systematic arrival of unexpected adverse information” (page 769).

A relevant question is to what extent do health shocks (or unexpected retirements more broadly) explain the heterogeneity in expenditure declines at the time of retirement, particularly among those with low pre-retirement wealth.¹¹ Haider and Stephens (2007) tackle the question of unexpected retirements directly. Using their data from the RHS, Haider and Stephens instrument for the time of a household's retirement with that household's own expectation of their retirement date some years prior to their actual retirement.

Given their methodology, they can compare the overall change in food spending for all households as they transition to retirement with the overall change in food spending for only those households where the date of retirement was predicted well in advance. Their IV estimates of the decline in food expenditures at the time of retirement, where age was used as an instrument for retirement status, was roughly -15 percent. The use of age as instruments was common by many of the studies documenting consumption declines at the time of retirement (see, for example, Aguiar and Hurst 2005). Using retirement expectations as an instrument instead reduces the estimated decline in food expenditures at retirement to -10 percent. In other words, the decline in food expenditures for households where the date of retirement is not forecastable is much larger than the decline in food expenditures for households where the date of retirement is known in advance.

Smith (2006) came to similar conclusions about the importance of involuntary early retirement after studying data from the British Household Panel (BHP). In her work, Smith divides the retirees in her sample into households who retire "voluntarily" and those who retire "involuntarily". She defines those who retire involuntarily as those individuals who transition into retirement from a non-work employment state (usually unemployment or long term

¹¹ There is a known relationship between household wealth and household health. Hurd and Rohwedder (2005) shows that those who cite adverse health shocks as a reason for retirement in the CAMS and HRS had significantly worse reported health prior to retirement.

disability). Her measure of expenditures is total spending on food consumed at home (meals away from the home are not included). Although her sample sizes are small (226 voluntary retirees and 57 involuntary retirees), she still is able to find that those who retire involuntarily experience much larger consumption declines than those who retire voluntarily. In her analysis, she cannot reject that those who retire voluntarily experience any expenditure declines upon retirement. However, those who retire involuntarily experience expenditure declines of over 10 percent.

Collectively, the results of Haider and Stephens (2007) and Smith (2006) show that some of the observed heterogeneity in the declines in spending associated with retirement is due to involuntary retirement. Specifically, those that are forced to retire involuntarily experience much larger expenditure declines than households who retire when planned.

Hurd and Rohwedder (2005) shed light on the role of health shocks in particular in explaining the decline in expenditures at the time of retirement. Using their survey data, Hurd and Rohwedder examined the expenditure changes for households who self report that poor health was a very important reason for their retirement to households who self report that poor health was not important at all for their decision to retire. There is evidence that those who experienced a poor health shock that forced them to retire were more likely to report expenditure declines at the time of retirement (67.5 percent vs. 48.4 percent) and experienced larger expenditure declines at the time of retirement (24.5 percent vs. 11.4 percent). After reading their results, it is evident that adverse health shocks do explain some of the large heterogeneity in expenditure declines as households transition to retirement.

6. Changes in Retirement Well Being

One alternate way to measure whether households are unable to sustain their well-being as they transition into retirement is to ask households directly whether retirement made them worse off. The 1992 wave of the Health and Retirement Survey (HRS) asked households the following question: “Thinking about your retirement years compared to the years just before you retired, would you say the retirement years have been better, about the same, or not as good?” Focusing only a sample of all households currently retired for more than 1 year in the 1992 HRS, 41 percent report that their retirement years have been better, 29 percent report that their retirement years have been about the same, and 29 percent report their retirement years have not been as good. So, similar to the results on the heterogeneity described above, unconditional on anything, 70 percent of households report that they are at least as well off during retirement.¹²

These results do not change too much if we condition on the retired individual being less than 70 years of age or if we condition on the individual being retired for less than 5 years. For example, for household retired less than 5 years, 45 percent report being better off, 30 percent report being the same and 25 percent report being worse off.

However, the results differ substantially between households who retired voluntarily and those who did not. Specifically, HRS households were also asked the following question: “Thinking back to the time you retired, was that something you wanted to do or something that you felt you were forced into?” Roughly half of the sample reported that they retired voluntarily. For those that reported retiring voluntarily, an astonishing 95 percent of them reported that their well being increased (61 percent) or stayed the same (34 percent) after

¹² All calculations on household well being after retirement using the 1992 HRS were performed by the author for this paper. If interested, email the author for details about the calculations discussed in the text.

retirement. For those retiring involuntarily, only 50 percent reported either being better off or the same during retirement.

Of those that retired involuntarily, 72 percent report they did so because they had received an adverse health shock that prevented them from continuing to work. When households received the health shock and were forced to retire involuntarily, they self reported that their well being declined in retirement relative to the periods prior to retirement. Collectively, these results are consistent with the expenditure results described above. Overall, most people report that their well being (expenditures) remains constant or increases during their retirement years. Yet, there are a fraction of households who retire involuntarily (usually due to health shocks) that experience declining well-being (expenditures).

7. Conclusion

Up until recently, there was a view that consumption was not modeled well by standard lifecycle models as households transition into retirement. The basis of this claim was that even though retirement – for most households – is fairly predictable, consumption expenditures declined precipitously for nearly all households as they exited the labor force. Such a phenomenon had been referred to as the “retirement consumption puzzle”.

However, during the last five years, a number of papers have emerged to challenge the belief that the canonical lifecycle model of consumption is inconsistent with household behavior during retirement. Aggregating results across a variety of recent research shows that the fall in expenditures at the time of retirement for the average household is confined to only two consumption categories: 1) work related expenses and 2) food. The decline in work related expenses is completely consistent with a lifecycle model of consumption where some

consumption categories are complements with working. The real puzzle should have been cast as why food expenditures fall sharply despite the fact that the rest of the household's consumption bundle remained relatively constant through the retirement period.

But, as discussed above, even the fall in food expenditures relative to other types of expenditures is no puzzle when view through the lens of a home production model. As retirees become abundant in time, they should be willing to engage in more time intensive activities. By engaging in "food production" (preparing meals, shopping more efficiently, etc.), households could reduce the cost of their food bundle while keeping their food consumption relatively constant. There is strong support for such a model across a variety of data sources. The most significant finding shows no change in actual food consumption (as measured by the quantity and quality of their actual diets) as households transition to retirement (despite sharply falling expenditures). The reason that food falls in retirement relative to the consumption of all goods is that food, of all the consumption categories, is the most amenable to home production.

The bottom line is that, for most households, there is no retirement consumption puzzle at all. Consumption of work related items falls (this is no surprise). Food consumption remains constant (despite falling expenditures). Consumption (and expenditures) of all other goods remain constant or increase. It is hard to tell a story of why a household who is ill prepared for retirement will cut their food expenditure while increasing spending on luxuries such as entertainment and charitable giving. In other words, most households are maintaining their marginal utility of consumption as they transition into retirement across all consumption categories.

There is evidence however, that some households experience real declines in expenditures (and well being) at the time of retirement. For example, there is evidence that

households in the bottom quartile of the pre-retirement wealth distribution experience declines in food expenditures that are nearly three times as large as the median households. What causes these large declines for such households with low pre-retirement wealth? There is evidence that involuntary retirements due to health shocks can explain much of this variation.

But, there are other potential explanations as well. Hurst (2006) suggest that households with low pre-retirement wealth entering retirement may be myopic with respect to their consumption decisions and, as a result, planned insufficiently for retirement. In a recent seminal paper, Scholz et al. (2006) run individual earnings, demographic and health trajectories (for an actual household) through a calibrated lifecycle consumption model (with idiosyncratic income and health shocks). They then compare the predicted household wealth on the eve of retirement from such a model to the household's actual pre-retirement wealth and find that roughly 20 percent of households are ill-prepared to sustain consumption during retirement. The remaining 80 percent of households have accumulated enough wealth to maintain their marginal utility of consumption through retirement.

In summary, there is some evidence that a small subset of households may be ill-prepared to sustain their consumption through retirement. However, the standard lifecycle model augmented with home production and idiosyncratic health shocks can explain the retirement consumption behavior for the overwhelming majority of households. Future work, however, needs to be spent learning more about these households who may, in fact, be ill prepared to maintain their consumption levels post retirement. Moreover, the data seems to suggest there may potentially be room for an improved insurance market that would allow households to maintain consumption in the event that they receive a health shock. Studying the consumption

needs of such households after a health shock, therefore, would also be a fruitful area for future research.

Lastly, this paper only focused on consumption movements around the period of retirement. As households live longer, the real question of interest will increasingly shift towards whether households can maintain their consumption well into their periods of retirement. Households may be able to smooth their consumption as they transition into retirement but may be unable to sustain that consumption level over all periods of their remaining life. Very little work has been conducted on whether households can maintain their desired consumption well into retirement. Answering this question is also an excellent area for future work.

References

- Aguiar, Mark and Erik Hurst (2007a). "Lifecycle Prices and Lifecycle Production", *American Economic Review* (forthcoming).
- Aguiar, Mark and Erik Hurst (2007b). "Re-examining Lifecycle Consumption", *University of Chicago* working paper.
- Aguiar, Mark and Erik Hurst (2005). "Consumption vs Expenditure". *Journal of Political Economy*. 113(5), 919 -948.
- Ameriks, John, Andrew Caplin, and John Leahy (2007). "Retirement Consumption: Insights from a Survey". *Review of Economics and Statistics* (forthcoming).
- Attanasio, Orazio (1999). "Consumption". *Handbook of Macroeconomics*, chapter 11, North Holland, Elsevier.
- Banks, James, Richard Blundell, and Sarah Tanner (1998). "Is There a Retirement Savings Puzzle?" *American Economic Review*, 88(4), 769 – 788.
- Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber (2006). "How Large is the Retirement Consumption Drop in Italy?", Working Paper http://www.entelugieinaudi.it/pdf/Seminari/Econometria/S_20061123.pdf.
- Becker, Gary (1965). "A Theory of The Allocation of Time." *The Economic Journal*, 75 (September), 493 – 517.
- Bernheim, Douglas, Jonathan Skinner, and Steven Weinberg (2001). "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review*, 91(4), 832 – 857.
- Fisher, Jonathan, David Johnson, Joseph Marchand, Timothy Smeeding, and Barbara Boyle Torrey (2006). "The Retirement Consumption Conundrum: Evidence From a Consumption Survey", Center for Retirement Research Working Paper 109.
- Friedman, Milton (1957). *A Theory of the Consumption Function*. Princeton: Princeton University Press.
- Hamermesh, Daniel (1984). "Consumption During Retirement: The Missing Link in the the Life Cycle" *Review of Economics and Statistics*, 66(1), 1 – 7.
- Haider, Steven and Melvin Stephens (2007). "Is There A Retirement Consumption Puzzle? Evidence Using Subjective Retirement Expectations", *Review of Economics and Statistics*, forthcoming.

- Heckman, James (1974). "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship Between Income and Consumption Over the Life Cycle". *American Economic Review*, 64(1), 188-94.
- Hurd, Michael and Susann Rohwedder (2003) "The Retirement Consumption Puzzle: Anticipated and Actual Declines in Retirement Spending", NBER Working Paper 9586.
- Hurd, Michael and Susann Rohwedder (2006) "Some Answers to the Retirement Consumption Puzzle", NBER Working Paper 9586.
- Hurst, Erik (2006). "Grasshoppers, Ants and Pre-Retirement Wealth: A Test of Permanent Income Consumers", NBER Working Paper 10098.
- Laitner, John and Dan Silverman (2005). "Estimating Life-Cycle Parameters from Consumption Behavior at Retirement", NBER Working Paper 11163.
- Lundberg, Shelly, Richard Startz, and Steven Stillman (2003). "The Retirement Consumption Puzzle: A Marital Bargaining Approach" *Journal of Public Economics*, 87, 1199-1218.
- McClellan, Mark (1998). "Health Events, Health Insurance and Labor Supply: Evidence from Health and Retirement Survey," in David Wise, ed., *Frontiers in the Economics of Aging*, 1998, Chicago, University of Chicago Press, 301-346.
- Miniaci, Raffaele, Chiara Monfardini, and Guglielmo Weber (2003). "Is There a Retirement Consumption Puzzle in Italy?" *Institute for Fiscal Studies Working Paper* WP03/14.
- Modigliani, Franco and Richard Brumberg (1954). "Utility Analysis and the Consumption Function: An Interpretation of the Cross Section Data," in Post-Keynesian Economics, Eds.: Kenneth Kurihara. New Brunswick, NJ, Rutgers University Press, pp 388-436.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006). "Are Americans Saving 'Optimally' for Retirement?", *Journal of Political Economy*, 114(4), 607-643.
- Schwerdt, Guido (2005). "Why Does Consumption Fall at Retirement? Evidence from Germany" *Economics Letters* (89), 300-305.
- Smith, Sarah (2006). "The Retirement Consumption Puzzle and Involuntary Early Retirement: Evidence from the British Household Panel Survey", *The Economic Journal*, 116 (March), C130 – C148.