Please do not cite. The complete paper is available at: http://faculty.chicagogsb.edu/erik.hurst/research/

Deconstructing Lifecycle Expenditure

Mark Aguiar
University of Rochester
and
Erik Hurst
University of Chicago

Prepared for the 10th Annual Joint Conference of the Retirement Research Consortium
"Determinants of Retirement Security"
August 7-8, 2008
Washington, D.C.

The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium (RRC). The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the RRC. We thank Jesse Shapiro for early conversations, which encouraged us to write this paper, as well as Eric French and Randy Wright for detailed comments. We also thank seminar participants at the University of Rochester, the NBER Macro Perspectives Summer Institute Session, the Institute for Fiscal Studies, the Federal Reserve Board of Governors, Wisconsin, Harvard, Yale, MIT, Chicago, CREI, Stanford, UCLA, MRRC, the University of Albany, and the PIER/IGIER conference on inequality in macroeconomics.

Abstract

In this paper we revisit two well-known facts regarding lifecycle expenditures. The first is the familiar "hump" shaped lifecycle profile of nondurable expenditures. The second is that cross-household consumption inequality increases steadily throughout the life cycle. We document that the behavior of total nondurables masks surprising heterogeneity in the lifecycle profile of individual sub-components. We find that three categories account for nearly the entire decline in mean expenditure post-middle age: food, transportation, and clothing/personal care. All other nondurable categories we study, including housing services, utilities, entertainment, domestic services, charitable giving, gambling, etc., show no decline over the life cycle. Similarly, nearly all of the increase in cross-sectional inequality is driven by these same three categories. Excluding food, clothing, and transportation from our measure of non durable expenditures reduces the increase in consumption inequality by a factor of 8, and removes nearly all of the increase post middle age. We provide evidence that the categories driving life cycle consumption are either inputs into market work (clothing and transportation) or are amenable to home production (food). Changes in the opportunity cost of time will cause movements in expenditure on such goods even if there is no change to lifetime resources. The patterns documented in the paper suggest that prior inferences from consumption data regarding the extent of uninsurable risk faced by households are sensitive to the inclusion of work related expenses and the home production of food. The disaggregated consumption data also pose a challenge to models that emphasize inter-temporal substitution or movements in income, including standard models of precautionary savings, myopia, and limited commitment, to explain the lifecycle profile of expenditures.

1. Introduction

This paper reconsiders two prominent features of life cycle consumption expenditures. The first is the fact that expenditures are "hump" shaped over the life cycle, peaking in middle age and then declining steadily throughout the second half of the life cycle.[1] The second fact is that cross-sectional consumption inequality increases steadily as individuals age.[2] Both facts have had tremendous influence on economists' inferences about household preferences, the income process that households face, and the extent to which public and private insurance markets limit household exposure to risk.

In this paper we revisit these two familiar facts by disaggregating nondurable expenditures into more detailed consumption categories. We show that there is substantial heterogeneity across consumption goods with respect to both the life cycle profile of mean expenditures and the evolution of the cross household variance in expenditures.

Specifically, we first replicate the standard finding that composite nondurable expenditures, controlling for family composition, peak in middle age at a level roughly 30 percent higher than expenditures at 25 or 65. Similarly, we document that the cross-sectional variance in log nondurable expenditure doubles between ages 25 and 75. However, we then document that the decline in nondurable expenditure post-middle age is essentially driven by three categories – food away from home, nondurable transportation, and clothing/personal care. Moreover, these three categories account for nearly the entire increase in the cross-sectional variance of expenditures over the lifecycle. All the other components of our composite nondurable measure (housing services, utilities, entertainment, domestic services, charitable giving, etc.) show *no* decline in expenditures after the age of 45 and exhibit little, if any, increase in cross sectional variance over the life cycle.

---

[1] This literature documenting the hump shaped profile of expenditures is large and extends back nearly 40 years. See, for example, Thurow (1969), Heckman (1974), Carroll and Summers (1991), Attanasio and Weber (1995), Attanasio et al (1999), Angeletos et al (2001), Gourinchas and Parker (2002), and Fernandez-Villaverde and Krueger (2007). The hump shape holds for nondurable expenditures as well as total expenditures and persists after accounting for changes in family size.

[2] See Deaton and Paxson (1994), Attanasio and Jappelli (2002), Storesletten et al (2004b), Heathcote et al (2005), and Guvenen (2007).

Canonical models of consumption emphasize movements in uninsurable permanent income as key to both the "hump" shape and the increase in cross-sectional dispersion.[3] Models based solely on fluctuations in financial resources predict that categories with larger income elasticities should display greater increases in cross-sectional dispersion and more pronounced hump shapes. However, the disaggregated data show no such pattern. For example, households increase spending on relative luxuries such as entertainment, charitable giving, and gambling after middle age while simultaneously reducing spending on food, clothing, and transportation. Similarly, the cross-sectional dispersion in the former categories all show declines over the life cycle.

The data do, however, support a prominent role for expenses that are closely linked to a households' opportunity cost of time. These categories consist of clothing and transportation, which have been shown to be inputs into market labor supply, as well as food, which is amenable to home production. As the opportunity cost of time falls over the lifecycle and, consequently, some households reduce their attachment to the labor force, expenditures on work related and home produced goods should fall, even if there is no change in lifetime resources or preferences. Such work related and food expenses account for the entire decline in nondurable expenditures after middle age, coincident with the peak in market labor supply for the average household. Moreover, while inequality in composite nondurables increases throughout the life cycle, doubling between age 25 and 75, inequality in nondurable expenditure excluding food and work related expenses increases by only 25 percent, with nearly all of the increase occurring prior to the age of 46 or after the age of 68.

To gain more insight into the importance of work status and work related expenses we perform a number of additional exercises. We document that the decline in expenditure on food away from home after middle age is associated with a decline in the frequency with which individuals frequent fast food establishments or cafeterias, with no indication that individuals reduce their patronage of restaurants with table service. This is consistent with food away from home being a work-related expense, particularly as it relates to movements over the life cycle. Similarly, time diaries indicate that there is a

[3] See the discussion in Section 5 for details.

large decline in time spent commuting to work after the age of 50. However, non-work traveling increases slightly over the lifecycle. To the extent that transportation expenditures are proportional to transportation time, these results imply that the decline in transportation expenses is due entirely to a decline in work related transportation. We also document that controlling for labor supply mitigates the decline in spending on clothes, transportation, and food away from home.

The patterns documented in this paper argue for a reassessment of the mapping of consumption to uninsurable permanent income. In particular, the patterns of "core" nondurable expenditures, excluding work related/home production expenses, suggests that the increase in cross-household consumption inequality increases much less than suggested by total nondurables and is essentially constant for households between the ages of 45 and 65. This suggests that households face less uninsurable income risk – particularly later in life - than suggested by total consumption expenditures. In this sense, this paper complements recent studies that conclude the canonical consumption models have overestimated the extent of uninsurable income risk later in the life cycle.[4]

2. Data and Empirical Methodology

A. Data

To examine the life cycle profile of expenditure and the life cycle evolution of the cross sectional dispersion, we use data from the Consumer Expenditure Survey (CEX). Specifically, we use the NBER CEX extracts, which includes all waves from 1980 through 2003. We restrict the sample to households who report expenditures in all four quarters of the survey and sum the four responses to calculate an annual expenditure measure. We also restrict the sample to households that record a non-zero annual expenditure on six key sub-components of the consumption basket: food, entertainment, transportation, clothing and personal care, utilities, and housing/rent. This latter condition is not overly restrictive, resulting in the exclusion of less than ten percent of the households. Lastly, we focus our analysis on households where the head is between the ages of 25 and 75 (inclusive). After imposing these restrictions, our analysis sample

[4] Examples from diverse fields and using different methodology include Keane and Wolpin (1997), Cunha, Heckman, and Navarro (2005), Guvenen (2007), and Huggett, Ventura, and Yaron (2007).

contains 53,412 households. Appendix A contains additional details about the construction of the dataset and sample selection.

When examining the life cycle profile of mean expenditures and cross sectional dispersion, we limit our analysis to nondurables excluding health and education expenditures. Our measure of nondurables consists of expenditure on food (both home and away), alcohol, tobacco, clothes and personal care, utilities, domestic services, nondurable transportation, airfare, nondurable entertainment, gambling, business services and charitable giving.[5] We also examine a broader measure of nondurables which includes housing services where housing services are calculated as either rent paid (for renters) or the self-reported rental equivalent of the respondent's house (for home owners). We exclude expenditures on education and health care from the analysis, as the utility (or returns) from consuming these goods vary significantly over the life cycle. Likewise, we exclude all durables aside from housing given the difficulty in creating annual service flow measures for these expenditures.

Our measure of nondurable expenditure plus housing services comprises roughly 70 percent of household annual monetary outlays. The remaining portion of annual outlays can be attributed to expenditures on durables such as automobiles, home furnishing, and large entertainment durables (16 percent); health expenditures (6 percent); education expenditures (2 percent); and other expenditures which are difficult to classify (5 percent).[6]

B. Estimating the Life Cycle Profile of Expenditure

When examining life cycle profiles of mean expenditure and cross sectional dispersion, we adjust all expenditures for cohort and family composition effects. The CEX is a cross-sectional survey and therefore age variation within a single wave

[5] Examples of expenditures that are included in each of the expenditure categories can be found in the data appendix and the corresponding documentation to the NBER CEX files.

[6] These other categories include, among others, life insurance premiums, interest paid on consumer credit, college dormitory fees, money allocated to burial plots, union dues, books, lodging expenses away from home, legal services, etc. Some of these categories were excluded because of the classification system introduced by Sabelhaus and Harris when creating the NBER CEX files. For example, the category of "books" includes money spent on books for leisure reading and books purchased for course work. Likewise, the category of "other lodging expenditures" includes both college dormitory expenses as well as vacation rentals. For consistency, we excluded from our analysis any category that combined some health or education component. However, in the NBER working paper version of this paper, we examined these categories in greater detail. None of our results are changed if we included these measures in our nondurable expenditure measure. This is not surprising given that they comprise only a small fraction of total household expenditures.

represents a mixture of life cycle and cohort effects. Moreover, expenditures are measured at the household level and not the individual level. Household size has a hump shape over the life cycle, primarily resulting from the fact that children enter and then leave the household. Likewise, children of differing ages have different consumption needs. Additionally, marriage and death probabilities change over the life cycle. We identify life cycle from cohort variation by using the multiple cross-sections in our sample, and use cross-sectional differences in family composition to identify family composition effects.

Formally, to estimate the life cycle profile of expenditures, we estimate the following regression:

$$\ln(C_{it}^{k}) = \beta_0 + \beta_{age} Age_{it} + \beta_c Cohort_{it} + \beta_{fs} Family_{it} + \varepsilon_{it}^{k} \qquad (1)$$

where C_{it}^{k} is expenditure of household *i* during year *t* on consumption category *k*, Age_{it} is a vector of 50 one-year age dummies (for ages 26-75) referring to the age of the household head[7], $Cohort_{it}$ is a vector including eleven five-year age of birth cohort dummies, and $Family_{it}$ is a vector of family structure dummies that include a marital-status dummy, 10 household size dummies, and controls for both the number and age of household children aged 21 or under. Specifically, we control for the number and age of household children by including dummy variables for the number of children in the following age categories: 0-2, 3-5, 6-13, 14-17, and 18-21. Moreover, for the latter two categories, we create separate indicators for male and female children. Our detailed family composition controls allow us to control flexibly for the potential that children of different ages and sex have different consumption needs or preferences. Moreover, the fact these family composition effects are allowed to differ across expenditure categories accommodates varying degrees of returns to scale across goods.[8]

[7] For married households, we use the husband's age. See the data appendix for additional details of how we identify household head in multi-adult households.

[8] Because of concerns about a potential correlation between family composition and permanent income, we performed many robustness exercises where we control for family composition in different ways. See the robustness appendix posted online at www.markaguiar.com/aguiarhurst/lifecycle/robustness_appendix.pdf for details of the robustness exercises. In summary, our results are robust to all of these alternative specifications.

The coefficients on the age dummies, β_{age}, represent the impact of the life cycle conditional on cohort and family size fixed effects, both of which we allow to vary across expenditure categories. Each of these age coefficients should be interpreted as log deviations from the spending of 25 year olds.

As is well known, co-linearity prevents the inclusion of a vector of time dummies in our estimation of (1). To account for changes in the relative price of each consumption category, we deflate all categories into constant dollars using the relevant CPI product-level deflator, if available. Otherwise, we use the relevant PCE deflator from the National Income Accounts.[9] All data in the paper are expressed in 2000 dollars. Any movements in expenditure patterns over time that are not captured by the five-year cohort dummies or by the price deflators will be interpreted as variation over the life cycle. In Appendix B, we discuss in detail the robustness exercises we performed to assess whether our results are sensitive to our choice of including cohort effects as opposed to time effects. The main conclusions of our paper hold whether we choose to estimate life cycle profiles including cohort effects or year effects.

In Section 4, we augment our benchmark estimates by controlling explicitly for relative prices and total expenditure. We postpone discussion of that methodology until Section 4.

C. Estimating the Life Cycle Profile of Cross Sectional Expenditure Dispersion

To estimate the life cycle profile of the cross sectional expenditure dispersion, we start by computing $(\sigma^2)_{it}^k$, the variance of ε_{it}^k (the residuals from (1)) for each age and cohort. We then estimate the following equation:

$$(\sigma^2)_{it}^k = \alpha_0^k + \alpha_{age}^k Age_{it} + \alpha_{cohort}^k Cohort_{it} + \eta_{it}^k \ . \qquad (2)$$

The vector of age coefficients, α_{age}^k, for each consumption category, *k*, provides our estimates for the evolution of cross sectional variance in expenditures over the life cycle. This method is essentially the same as the one used by Deaton and Paxon (1994). The age coefficients should be interpreted as being deviations from the variance of log

[9] We have verified our results are robust to using a common aggregate deflator (CPI-U) for all categories..

expenditure observed for 25 year olds. As with mean expenditures, collinearity prevents the identification of time trends separate from age and cohort effects.[10] We, however, show that the main results of our paper with respect to the changing cross sectional variance of expenditures over the life cycle are invariant to the inclusion of time effects (as opposed to cohort effects) in Appendix B.

3. Life Cycle Expenditure Patterns by Category

In this section, we document the existence of substantial heterogeneity across different consumption categories with respect to both the life cycle profile of expenditure and the evolution of the cross sectional variance of expenditure over the life cycle.

For context, we first show the trends in life cycle expenditure and life cycle dispersion for our composite nondurable measures. These results are shown in Figures 1a (life cycle expenditure profile) and Figure 1b (life cycle profile of cross sectional expenditure variance). The solid line in each figure represents the results using nondurable expenditures without housing services. The dotted line represents the results using nondurable expenditures with housing services.

Figure 1a replicates the well-documented profile of nondurable expenditures over the life cycle, with nondurable expenditures excluding housing services peaking in middle age at roughly 25 percent (that is, 0.25 log points) higher than the level of 25 year old expenditure, and then declining by nearly 30 percent over the latter half of the life cycle. Nondurable expenditures inclusive of housing services rises faster early in the life cycle, but then does not decline as significantly later in the life cycle. The gap between the two series represents the life cycle behavior in housing services, which we will discuss on its own in the next sub-section.[11] These results are consistent with a large literature documenting the hump shaped profile of nondurable expenditures over the life cycle.[12]

[10] See Heathcote et al (2005) for a detailed sensitivity analysis regarding cohort versus time fixed effects in identifying the evolution of inequality over the life cycle.

[11] Yang (2008) documents that the life cycle profile of housing services is different from the life cycle profile of composite nondurable expenditures. She then writes down a model of where housing consumption is costly to downsize to explain the differences between housing and nondurable expenditures. As we show below, the life cycle profile of housing looks like the life cycle profile of most non durable consumption categories.

[12] See, for example, Attanasio et al (1999) and Fernandez-Villaverde and Krueger (2007) and the cites within.

Figure 1b shows the increase over the life cycle of the cross-sectional variance of log nondurable expenditures relative to the variance observed for 25 year olds. The variance for nondurable expenditures with and without housing expenditures for 25 year olds is 0.16 and 0.17, respectively. Between the ages of 25 and 75, the cross sectional variance of nondurable expenditures increase by roughly 20 points, regardless of whether or not housing services are included in the measure of nondurable expenditures. These magnitudes are similar to the results reported by Guvenen (2007) and are consistent with the findings of others that the cross sectional variance of expenditure increases by at least 100 percent over the life cycle. Additionally, most of the increase comes later in the life cycle (after the age of 40) leading some researchers to conclude that there is a prominent role for permanent income shocks later in the life cycle.

A. Life Cycle Profiles of Disaggregated Expenditures

In Figures 2a and 2b, we plot the life cycle expenditure profiles for the sub-components of our composite nondurable. Specifically, we document the life cycle spending patterns separately for housing services, utilities, nondurable entertainment, nondurable transportation, food consumed at home, food consumed away from home, alcohol and tobacco, domestic services, clothing and personal care, and a residual "other" category. The "other" nondurable category includes airfare spending, charitable giving, and net gambling receipts. All category specific life cycle spending profiles are estimated using (1) and, as a result, are adjusted for cohort and family composition. Table 1 reports the share of spending out of total nondurable expenditures (both with and without housing) for each of the consumption subcategories.

For expositional purposes, we group the categories by whether or not they decline after middle age. In particular, Figure 2a depicts categories that show no decline over the life cycle, while Figure 2b collects those categories that exhibit declines after middle age. This categorization underscores that not all categories share the prominent "hump" seen in composite nondurables. As reported in Table 1, the non-decreasing categories constitute 54 percent of nondurable expenditures including housing services.

We begin our discussion of Figure 2a with nondurable entertainment spending. Nondurable entertainment consists of such expenditures as cable subscriptions, movie and theatre tickets, country club dues, pet services, etc. It does not include durable

expenditures such as television sets and does not include reading material and magazine subscriptions. The average annual expenditure on entertainment totals $1,260 in year 2000 dollars and accounts for 7 percent of non-durable expenditure excluding housing services (Table 1). Entertainment expenditures increase until the early 40s (by roughly 70 percent). However, expenditures on entertainment do not fall after middle age. Instead, spending on entertainment remains roughly constant between 45 and 59 and then increases by an additional 8 percent between the ages of 60 and 68.

As seen in Figure 2a, housing services, utilities, domestic services, and other nondurables exhibit similar life cycle profiles to that displayed by entertainment. All these categories increase significantly between the age of 25 and the age of 44, continue to increase up through the age of 59, and then increase sharply after the age of 60.

The continuous increases in categories depicted in Figure 2a begs the question of what categories drive the decline in composite nondurable consumption spending after middle age. Figure 2b answers this question. Specifically, food at home spending increases 24 percent between 25 and 44, declines by 7 percent between 45 and 59, and then declines another 4 percent between 60 and 68. The middle age declines in expenditures are even larger for transportation (22 percent between 45 and 59), clothing and personal care (39 percent between 45 and 59), and food away from home (62 percent between 45 and 59). These four categories comprise 42 percent of nondurable expenditures with housing services.

The final "declining" category is alcohol and tobacco, which is not included in Figure 2 but is included in Table 1. This category behaves in a manner distinct from the other categories depicted in Figures 2a and 2b. Alcohol and tobacco expenditure falls continuously over the entire life cycle. Moreover, the decline in expenditure is very large: Spending on alcohol and tobacco falls by 1.22 log points between 25 and 44, another 1.65 log points between 45 and 59, and another 1.03 log points between 60 and 68. Even though alcohol and tobacco comprises only 4 percent of composite nondurables, its large decline contributes significantly to the overall decline in nondurable spending after middle age.

Table 1 summarizes the patterns shown in Figures 2a and 2b. It should be noted that expenditures on all subcategories displayed in Figures 2a and 2b increase over the

front half of the life cycle. The difference between the two groups of categories occurs after the mid-40s. Models that predict declines in spending on all consumption goods after middle age (like standard models incorporating household impatience or poor planning) are inconsistent with the disaggregated spending data.

The data reported in Table 1 and Figure 2 also suggest a re-interpretation of the so-called "retirement consumption puzzle," which refers to the decline in expenditures observed around retirement. In particular, the declines at retirement should be placed in the context of the broader trend of declining life cycle expenditures after middle age.[13] The categories that exhibit declining expenditures during the peak retirement years (60-68) are the same categories that exhibit declining expenditures over the second half of the life cycle (after the age of 45). Specifically, declines in expenditure around retirement are observed only for food, clothing/personal care, transportation, and alcohol and tobacco.[14] There is no evidence that entertainment, housing services, utilities, domestic services, charitable giving, or airline travel declines during the peak retirement years. Taken together, the results cast doubt on the existence of a generic decline in expenditure for all consumption categories around the time of retirement.

3b. Lifecycle Profile of Cross Sectional Dispersion for Disaggregated Categories

The life cycle profiles of mean expenditures by category are mirrored in the life cycle profiles of cross-sectional dispersion. As a general rule, the categories for which expenditure does not decline past middle age are the same categories for which there is no substantial increase in cross-sectional dispersion. Conversely, the categories for which expenditures decline over the latter half of the life cycle are the same categories for which the cross sectional dispersion increases dramatically over the life cycle. These results are shown in Figures 3a and 3b and are summarized in Table 2. As with Figure 1b, the age coefficients represent the change in the cross-sectional variance of log expenditure at each age relative to that of 25 year olds (see equation (2)). For reference, Table 2 also includes the variance of log consumption at age 25 for each consumption category. As seen in the table, spending on food at home, housing services, and clothing

[13] See Hurst (2007) for a survey of the retirement consumption literature

[14] The fact that declines in expenditures at the time of retirement are limited to food, clothing, and non-durable transportation has also been emphasized by Battistin et al (2006) and Hurst (2007).

exhibit rather low cross sectional variances, while alcohol and tobacco, domestic services, and other nondurables exhibit substantial cross sectional variance.

In Figure 3a, we plot the life cycle profile of the cross sectional variance of log expenditure for goods that do not experience any increase in the cross sectional variance over the life cycle. The goods that display no increase in variance are essentially the same goods that do not decline over the back side of the life cycle. The one difference between the categories in Figures 2a and 3a is food at home, which is not included in the former but is included in the latter. Expenditures on food at home do not exhibit an increasing cross sectional variance over the entire life cycle, although the variance does increase slightly after the age of 45. The other categories for which inequality does not increase over the life cycle are housing services, utilities, entertainment, and other nondurable expenditures.

Figure 3b reveals which categories drive the increasing cross sectional variance of log expenditures over the life cycle. These categories include nondurable transportation, clothing and personal care, food away from home, and domestic services. From the figure and the upper panel of Table 2, we see that at the lower end, the variance of log transportation expenditure increases by 54 percent (from 0.70 to 1.08). At the upper end, the variance of log expenditures on food away from home more than doubles, increasing from 1.54 to 3.45.

Standard models that focus exclusively on shocks to income predict that goods with a high income elasticity should experience the largest changes in both means and cross-sectional variances. Table 2 and Figures 3 suggest no such pattern. Relative luxuries like entertainment, gambling, charitable giving, and airfare, together with such basics as housing and utilities, show no similarity to the life cycle pattern of composite nondurables. Other than the contribution of the idiosyncratic category of alcohol and tobacco, the prominent features of the composite nondurable category are driven by food, transportation, and clothing/personal care. These latter categories are perhaps best considered as inputs into market work (or, in the case of food, amenable to home production), rather than categories with relatively large income elasticities. In the next section, we explore this premise in more detail.

4. The Importance of Food, Clothing and Transportation in Explaining Life Cycle Profiles

As seen from the discussion in Section 3, there is substantial heterogeneity across consumption categories with respect to both the life cycle profile of mean expenditures and the life cycle profiles of the cross-sectional variance. Leaving aside alcohol and tobacco spending, spending on food (particularly food away from home), clothing and personal care, and transportation seem to be driving both the decline in nondurable spending after middle age and the increase in the cross sectional variance of log nun durable spending over the life cycle.

One potential reason why these categories may behave differently over the life cycle is that food is amenable to home production and clothing and transportation spending are complements with market work. To the extent that the opportunity cost of time evolves over the life cycle, one would predict changes in spending to occur within these categories given a standard model of household optimization augmented with a home produced good and work related expenses.

Appendix Figure A-1a shows the life cycle profile of the labor supply of household heads from the Consumer Expenditure Survey. Our analysis sample for this exercise is identical to the sample used above to document the life cycle consumption profiles. We show two measures of labor supply – the fraction of heads working (solid line) and the normal hours per week worked by the head (dotted line). This latter measure is not conditioned on working. Both labor force measures start to decline for individuals in their early 50s. Referring back to Figure 1a, this is the time that nondurable expenditures begin to decline. Appendix Figure A-1b shows the variance of the two labor force measures over the life cycle. The variance of work hours start to increase sharply starting for individuals in their early fifties. Given the decline in work hours starting for individuals after middle age, it is not surprising to find that work related expenditures should start to decline. Likewise, given the increase in the variance of labor force participation starting in the early 50s, it is not surprising to see the variance of work related expenditures increase between the ages of 50 and 65.

In this section, we discuss evidence that food is amenable to home production and that clothing and transportation are complements with market work. We then divide

nondurables into three sub-aggregates – food, work related expenses, and all other nondurable expenditures – and document their respective means and variances over the life cycle. Finally, we show that controlling for work status mitigates the decline in work related spending over the latter half of the life cycle.

A. The Home Production of Food

In Aguiar and Hurst (2005, 2007b), we have explored the differences between food expenditures and food intake. Using data from the Continuing Survey of Food Intake of Individuals (CSFII) which measures food intake at the individual level using detailed food diaries (including the quality of food consumed), Aguiar and Hurst (2005) shows that food intake does not decline over the life cycle despite the decline in expenditures after middle age. On the contrary, using the detailed data on the quantity and quality of food consumed, they find that food intake actually increases after middle age.[15]

Aguiar and Hurst (2007b) estimate a model of home production and food shopping to explain the differences between food expenditures and food intake. Using data from the American Time Use Survey (ATUS), which measures the amount of time individuals spend preparing meals and shopping for food, and data from the Nielson Homescan Panel, which measures the prices that households pay for a given food good (measured at the level of the universal product code), Aguiar and Hurst (2007b) finds that after middle age individuals allocate more time to preparing meals and shopping for food, and as a result, pay lower prices for constant quality consumption goods. Their estimated model of food production and food shopping matches the decline in food at home spending and food away spending. Moreover, like the actual data on food intake, their estimated model predicts rising food intake over the life cycle.

Figure 4 sheds additional light on the margins of substitution that takes place with respect to food spending over the life cycle. Using data from the Continuing Survey of Food Intake of Individuals (CSFII), we measure an individual's propensity to eat away from home at various types of eating establishments. The primary design of the CSFII is

[15] This can be seen from Appendix Figure A1 of the NBER working paper version of Aguiar and Hurst (2005).

to measure food intake via food diaries.[16] At any time food was consumed, the respondents were asked to provide very detailed comments about what they consumed, when they consumed it, and where they purchased it. We construct a variable called "eating away from home" which takes the value of 1 if the respondent reported purchasing food at a restaurant with table service, a restaurant without table service (i.e., establishments like fast food chains), a cafeteria, or a bar/tavern. Respondents in the CSFII spend roughly 2.5 days in the sample. For the entire sample, 64 percent of individuals eat away from home at least once during their time in the sample.[17] Of those, 38 percent eat at fast food establishments, 33 percent eat at restaurants with table service, 10 percent eat at cafeterias, and 6 percent eat at bars. The percentages sum to more than 64 percent given some individuals eat at multiple establishments during their time in the sample.

Figure 4 depicts the propensity to eat at the various types of restaurants. As with the expenditure data, we adjust the propensity to eat away from home for changing family composition and all comparisons are made relative to households in their late 20s (25-29). The overall pattern is similar to expenditures on food away from home, especially as it relate to the declines after middle age. In particular, the propensity to eat away from home falls by nearly 28 percentage points for individuals in their late 60s relative to individuals in their early 20s. However, the entire decline is due to a declining propensity to eat at fast food restaurants and cafeterias. There is no decline in the propensity for individuals to eat at restaurants with table service as they age. This finding is consistent with the premise that the decline in food expenditures reflects households switching towards home production as their opportunity cost declines past middle age. The shift toward home production of food results in households purchasing fewer meals from fast food and cafeterias, which are close substitutes to home-produced food.

B. Transportation and Clothing As Work Related Expenses

[16] The CSFII is a large nationally representative survey of individuals (as opposed to households). As in Aguiar and Hurst (2005), we use the surveys waves conducted in 19xx-xx and 19xx-xx for our analysis. For our analysis, we restricted the sample to 25-75 year olds. Our total sample size used for the results in Figure 4 was 6,615 individuals. See the data appendix for a more detailed description of the CSFII data.

[17] This number is approximately consistent with data from the 2003 American Time Use Survey that shows that for a similar sample, 26 percent of individuals eat at a restaurant, fast food establishment or bar on any given day. If eating away from home is i.i.d., this implies that for a 2.5 day interval, 53 percent of individuals would report time spent in a restaurant, fast food establishment or bar.

Spending on clothing and transportation has long been viewed as complements with market work. For example, Nelson (1989) and DeWeese and Norton (1991) both use data from the Consumer Expenditure Survey to show that even after controlling for total income, total expenditures, and family composition, married couples spend much more on clothing when both spouses are working than otherwise similar households in which only the husband worked. Both papers conclude that work status directly effects household spending on clothing.

Similarly, when analyzing consumption decisions of retirees, Banks et al (1998) explicitly refer to spending on canteen and restaurant meals, transportation, and adult clothing as "work-related expenditures". A similar classification was employed recently by Battistin et al (2006) when examining consumption changes around retirement in Italy.

Lastly, there is a literature on computing welfare calculations across people that advocates netting out work related expenses from measured income. For example, Lazear and Michael (1980) argued that, "There are certain costs of employment, such as costs of transportation to work, requisite clothing expenditures for employed persons, union or professional fees, etc. These work-related costs should be netted out of income in obtaining an indicator of standard of living".[18]

Transportation expenditures reflect the need to commute to work as well as travel for other purposes. While expenditure data does not separately measure costs due to work travel from non-work travel, we use time diaries from the pooled 2003-2005 American Time Use Survey (ATUS).to gauge the relative importance of each.[19] The detailed categories of the ATUS allow us to separately identify time spent traveling to and from work from time spent traveling for other reasons (including going to grocery store, going to visit friends, going to the movies, etc.). The average individual between the ages of 25 and 75 spends 9.0 hours per week traveling, with 2.3 hours per week being

[18] See Lazear and Michael, page 203.

[19] The ATUS is a nationally representative survey which uses time diaries to measure how individuals allocate their day. For a detailed account of the ATUS, see Aguiar and Hurst 2007a. For this analysis, we restrict the sample to only households between the ages of 25 and 75. Our total sample size was 38,876 individuals. See the data appendix for additional details about the ATUS, our sample selection, and our definition of variables.

traveling associated with commuting to/from work. For those that work, work related travel represents roughly one-third of all time spent traveling.

Figure 5 shows the life cycle profile of travel time after adjusting for changing family composition. The family composition controls include a marital status dummy, dummies for household size, and a dummy for whether the household has a child under the age of 5. The life cycle profile is expressed as hours per week deviation from households aged 25-29. Consistent with the decline in transportation expenditures over the life cycle starting for households in their early 50s, the decline in transportation time also starts for individuals in their early 50s. However, as seen from Figure 5, the entire decline in travel time occurs due to a decline in traveling to and from work. Non-work travel time actually increases over the life cycle. The data from the time use surveys suggests that the decline in transportation spending over the life cycle stems from the decline in time spent commuting to work. Again, this is consistent with the fact that transportation expenditures, and particularly their fluctuations over the life cycle, have a substantial work related component.

C. Work Hours and Lifecycle Consumption Profiles

The previous discussion and the patterns depicted in figures 2 and 3 suggest a natural decomposition of nondurable expenditures. Namely, we consider three sub-aggregates: (i) work related expenses, consisting of clothing/personal care and transportation; (ii) food, which includes food consumed at and away from home; and (iii) all other expenditure categories including housing services but excluding tobacco and alcohol. We refer to this latter measure as "core nondurable" expenditures. This 3-way decomposition simplifies the analysis while retaining the key patterns of the more disaggregated data.

The mean and cross-sectional variances of these categories are depicted in Figures 6a and 6b, respectively. Table 3 summarizes the life cycle profiles for the three composite consumption goods. In Figure 6a, we see that core nondurables increases sharply up through middle age and then increases slightly thereafter. Both food and work related expenditures fall sharply after middle age.

A striking reflection of our results on consumption inequality is clear in Figure 6b. The cross sectional variance of core nondurable expenditures displays a dramatically

different life cycle pattern than does the cross sectional variance of total nondurable expenditures as analyzed by Deaton and Paxson and others, and replicated in Figure 1b above. In particular, up through the age of 68, the cross sectional dispersion in core nondurables increases by only .05 points, with nearly all of the increase coming prior to the age of 45. Given that the variance of core nondurables for 25 year olds is 0.28, the cross sectional dispersion of core nondurables increases by less than 18 percent over the life cycle. This is an order of magnitude lower than the increase in cross sectional variance for total nondurables.

The implication is that essentially all of the increase in cross-sectional variance over the life cycle stems from work related expenses and food, with the increase in food inequality primarily reflecting food away from home. The sharp increase in inequality in expenditure on work related expenses is clear in Figure 6b. Note in particular that the variance of work related expenses increases significantly after middle age, while core nondurables shows no comparable increase. The cross-sectional variance of total nondurables increases by nearly 10 percentage points between the ages of 46 and 68 (Figure 1b), which represents nearly two-thirds of the increase in life cycle dispersion of total nondurables. All of the increase in variance between the ages of 45 and 68 in total nondurables is due to an increase in the variance of work related expenditures and food (as well as the covariances). The increase in dispersion of work related expenses reflects the fact that the cross-sectional variance in labor supply increases after middle age, as well (see Figure A-1b). This suggests that uninsurable shocks to permanent income are not a major concern between the ages of 40 and 68. There is an increase in variance after the age of 68 in core nondurables, which is unlikely to reflect earning risk but perhaps reflects late-in-life health shocks.[20]

Given the importance of work related expenses in driving changes in expenditure over the life cycle, a natural approach would be to control for work status. A difficulty with simply adding controls for employment status is the fact that labor supply is closely associated with permanent income. For example, lower wage workers tend to work

[20] The Appendix shows a series of robustness exercises for Figures 6a and 6b. First, we show that if we estimate the specifications plotted in 6a and 6b with year effects as opposed to cohort effects, the conclusions are unchanged (see Appendix B). Also, we show that the patterns are the same for low educated households as they are for high educated households (see Appendix C).

fewer hours than high wage workers (see Aguiar and Hurst (2007a) and Vandenbroucke (2007)). Absent a panel, controls for labor supply will also proxy for permanent income. However, we can accommodate this for consumption sub-categories by including nondurable expenditures as an additional control.

Specifically, we estimate the following for each of our three sub-aggregates:

$$s_{it}^k = \omega_0 + \omega_{age} Age_{it} + \omega_c Cohort_{it} + \omega_{fs} Family_{it} + \sum_k \omega_p^k \ln P_t^k + \omega_p \ln P_t + \omega_X \ln X_{it} + \omega_L L_{it} + \varepsilon_{it}^k, \quad (3)$$

where s_{it}^k is the share of consumption category *k* out of total nondurable expenditures excluding tobacco and alcohol, where *k* equals work related, food, and core nondurables expenditures. These shares sum to one. The age, cohort, and family status controls are the same as in equation (1). We include as additional controls the log price index of each of our sub-aggregates (lnP^k) as well as the overall price index (lnP). We compute the category specific price indices by using the weighted share of the price indices for the goods that comprise the consumption category. We include log expenditure on total nondurables with housing services ($ln\ X$) excluding alcohol and tobacco, which is also just the sum of work related, food, and core nondurable expenditures.

Finally, we include a vector of dummies describing the household labor supply (L). Our work status controls include: a vector of 7 dummies indicating the number of weeks worked by husbands during the previous year, a vector of 7 dummies indicating the number of weeks worked by the wife during the previous year, a vector of 9 dummies indicating the number of hours usually worked during the week by the head, and a vector of 9 dummies indicating the number of hours usually worked during the week by the wife. We found that labor supply has a nonlinear effect on consumption, and so include a very rich set of work status controls. To avoid issues of changes in household formation and its effect on labor supply, we restrict our analysis sample to include only married households. This leaves us with a sample of 32,204 households.

Given the fact that expenditures on the individual consumption categories are determined simultaneously and the fact that any measurement error in one category will lead to measurement in total nondurables, we follow the standard practice of instrumenting total expenditure with log total household family income and education

dummies. Our measure of total household family income includes labor, asset, and transfer income of both husbands and wives.

Note that equation (3) is essentially the almost ideal demand system (AIDS) of Deaton and Muelbauer (1980), conditioned on work status, family size, and age. We impose the restriction that the overall price index is given by the CPI-U, but do not impose restrictions related to consumer optimization such as symmetry and homogeneity. The inclusion of work status controls to form a conditional demand system follows the important work of Browning and Meghir (1991) and Blundell et al (1994), whose work we discuss in Section 5.

The results from the estimation are shown in Figures 7a (work related expenses), 7b (food), and 7c (core nondurables). In each figure we depict the age coefficients from (3) (dotted line), as well as the age coefficients when we omit work status controls (solid line). As always, each point represents the deviation from age 25, with the units for Figure 7 being share of total nondurable expenditures. As seen in 7a, and consistent with the results in 6a, the share of total expenditures allocated to work related expenses declines over the second half of the life cycle. Specifically, the share decreases by 5 percentage points between the ages of 50 and 68. However, after accounting for household labor supply, the decline is reduced by nearly half.

While adjusting for labor supply substantially mitigates the decline in work related expenses, it does not explain the entire decline. When we disaggregate the work related expenditures into its two components and re-estimate (3) on the disaggregated categories, we find that work status controls explain almost all the decline in transportation expenditures and less than half the decline in clothing expenditures (results not depicted). There are several reasons why clothing behaves differently than transportation. The first is that clothing is semi-durable, unlike the other categories of nondurable consumption. Therefore, there may be a wedge between the service flow of clothing and measured expenditures. Moreover, there may be changes in required clothing expenses over the life cycle depending on occupation and career trajectories. Therefore, labor supply will absorb only some of the life cycle variation in work related clothing expenses. Finally, Charles et al (2008) have argued that clothing expenditures are a status good. As the need to signal status becomes less important over the life cycle

(as individuals have already found a spouse, gotten a job, and made friends etc.), the use of clothing as a symbol of status would diminish.

We also estimated (3) restricting our sample to households between the ages of 25 and 50 and including a dummy variable indicating whether the household head is working (and omitting all other work status controls). This captures the effect of employment for individuals in prime working years. The coefficient is large and positive on the working dummy for both clothing and transportation. In other words, all else equal, households spend a larger share of their consumption on clothing and transportation when the household head is working (even after controlling for total expenditures).

Figures 7b and 7c show the similar patterns for food at home, food away from home, and core nondurables. As with work related expenses, both food at home and food away from home decline over the latter half of the life cycle. Including work status controls explains nearly the entire decline in food away from home. The share spent on food at home actually declines more when work status controls are included. This reflects that food away from home is more sensitive to work status, so in terms of shares, households substitute away from restaurants towards food at home as they stop working.

Figure 7 and the results in this section generally, support the premise that the categories that drive the aggregate life cycle profiles for means and variances – clothing, transportation, and food – are sensitive to work status. Equation (3) can only be estimated for shares given the inability to observe permanent income absent a panel. However, the behavior of shares can be mapped into total expenditures using a structural model, something we are pursuing in ongoing research. What is clear directly from the data is that work related goods are behind the decline in total expenditures and the increase in inequality after the age of 50.

5. Discussion and Related Literature

There is a large body of work that tries to explain the lifecycle profile of *composite* nondurable expenditures, without addressing the heterogeneity found in disaggregated consumption categories. For example, some authors have argued that the lifecycle profile represents evidence against the forward-looking consumption "smoothing" behavior implied by permanent income models, particularly since the hump

in expenditures tracks the hump in labor income (as documented by Carroll and Summers (1991)). This view interprets expenditure declines in the latter half of the lifecycle as evidence of poor planning. A related literature has developed which also emphasizes imperfect household planning based on the sharp decline in expenditures at the onset of retirement (see, for example, Bernheim, Skinner, and Weinberg 2001). Models of limited commitment to plans (such as Angeletos et al 2001) share the implication that the decline in expenditures late in life is due to insufficient resources. Standard models of poor planning or dynamic inconsistency, however, do not predict that households late in the life cycle reduce some expenditures while simultaneously increasing others.

Another literature has combined rational, forward looking agents with incomplete markets. In particular, the hump shaped profile in expenditure reflects optimal behavior if households face liquidity constraints combined with a need to self-insure against idiosyncratic income risks (see, for example, Zeldes 1989, Deaton 1991, Carroll 1997, Gourinchas and Parker 2002). Households build up a buffer stock of assets early in the lifecycle, generating the increasing expenditure profile found during the first half of the lifecycle. The decline in the latter half of the lifecycle is then attributed to impatience coming to the fore, once households accumulate a sufficient stock of precautionary savings. Similarly, the increase in cross-sectional dispersion of total expenditures has been interpreted as evidence of large, uninsurable idiosyncratic shocks to income.

Such precautionary savings models have been extremely influential, in part due to their ability to explain the prominent features of lifecycle expenditures in a rational agent, incomplete markets framework. Indeed, several important studies have used expenditure profiles to "back out" or verify measures of labor income risk over the lifecycle (see, for example, Deaton and Paxson (1994) as well as more recent papers by Storesletten et al (2004a, 2004b) and Guvenen (2007)). A related literature uses movements in consumption to infer movements in permanent income (see, for example, Blundell and Preston (1998) and Aguiar and Gopinath (2007)). Of course, the quality of these measures of income risk depends crucially on the validity of the underlying model of consumption.

Precautionary savings models also have strong predictions for the lifecycle behavior of goods with different income elasticities. The standard precautionary savings

model works off the tension between the need to accumulate assets for insurance versus impatience relative to the market interest rate. A fairly high degree of impatience is necessary to explain the sharp decline in expenditures in the latter half of the lifecycle.[21] However, if impatience is the predominant force driving the decline in expenditures over the second half of the lifecycle, then categories of consumption for which there is a high degree of inter-temporal elasticity should decline faster than those with a low degree of substitutability. Given the equivalence between inter-temporal elasticity and income elasticity (see Browning and Crossley 2000), this implies that luxury goods (such as entertainment, gambling, and charity) should decline more in the latter half of the lifecycle than necessities (such as food and transportation).

Note that both the precautionary saving models and the poor planning models place an emphasis on income fluctuations. The poor planning models emphasize deterministic trends in lifecycle labor income. The precautionary savings models emphasize income uncertainty. In particularly, the high degree of impatience in the precautionary savings model needed to explain the sharp decline in expenditures late in life must be matched with a commensurately high degree of income uncertainty early in the lifecycle. This latter component is necessary to explain why agents save and exhibit an upward profile of expenditure early in the lifecycle, despite the high subjective discount rate. The tight link between income risk and impatience relative to the interest rate is also a familiar feature in incomplete market models with infinitely lived agents (see for example, Huggett 1993 and Aiyagari 1994).

By focusing on the disaggregated data, we documented that the primary movements of expenditures later in the lifecycle are inconsistent with the models that rely exclusively on precautionary savings, myopia, or limited commitment. We should stress, however, that our work does not imply these forces are not at work at all. For example, the increases early in the life cycle may indeed reflect a precautionary savings motive. However, without accounting for the role that work related expenses play, attempts to use total expenditure to back out permanent income risk will be mis-specified.

[21] The important role impatience plays in these models is highlighted by the fact that Gourinchas and Parker (2002) are able to obtain a very precise estimate of time preference. As discussed in that paper (p. 73), this reflects that fact that the precautionary savings model's predictions are extremely sensitive to the discount rate.

The "non-separability" between consumption and work status has been documented by, among others, Browning and Meghir (1991) and Blundell et al (1994), using conditional demand system analysis. Our work complements these studies by arguing that once we strip out work related expenses, mean nondurable expenditures do not decline and there is little increase in cross-sectional inequality after middle age.

6. Conclusion

This paper documented that the hump in lifecycle expenditures and the increase in consumption inequality over the life cycle is driven primarily by work related expenses or goods amenable to home production. There is no sense in which goods with a high income elasticity display more prominent patterns over the life cycle, suggestion a force other than uninsurable permanent shocks is also at work. We provide evidence that changing labor supply is a promising candidate.

The facts documented in the paper have important implications for linking consumption movements to income shocks. In particular, it is important to separate shocks to resources from movements in labor supply (or demands on time more generally). We have documented the relevance of this point for both average expenditures over the lifecycle as well as the cross sectional dispersion of expenditures. While this paper focuses on lifecycle movements, the same issue arises in studies of the business cycle (see, for example, Greenwood, Rogerson, and Wright (1995)). Recessions are periods in which both income and employment fall. Analyses that ignore the latter will draw misleading conclusions about the importance of income uncertainty and the extent of insurance.