

Deconstructing Lifecycle Expenditure

Mark Aguiar and Erik Hurst



Project #: UM08-06

Deconstructing Lifecycle Expenditure

Mark Aguiar
University of Rochester

Erik Hurst
University of Chicago

January 2008

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
<http://www.mrrc.isr.umich.edu/>
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98362-5-04). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Regents of the University of Michigan

Julia Donovan Darrow, Ann Arbor; Laurence B. Deitch, Bingham Farms; Olivia P. Maynard, Goodrich; Rebecca McGowan, Ann Arbor; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; S. Martin Taylor, Grosse Pointe Farms; Katherine E. White, Ann Arbor; Mary Sue Coleman, ex officio

Deconstructing Lifecycle Expenditure

Mark Aguiar and Erik Hurst

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. We document that the behavior of total nondurables masks surprising heterogeneity in the lifecycle profile of individual sub-components. We find, for example, that while food expenditures decline after middle age, expenditures on entertainment continue to increase throughout the lifecycle. These patterns pose a challenge to familiar lifecycle models that emphasize inter-temporal substitution or movements in income, including standard models of precautionary savings, myopia, and limited commitment. Secondly, we document that the increase in the cross-sectional dispersion of expenditure over the lifecycle is not greater for luxuries. In particular, the dispersion in entertainment expenditure declines relative to food expenditures as households become older, casting further doubt on theories that emphasize (exclusively) shocks to permanent income. We propose and test a Beckerian model that emphasizes intra-temporal substitution between time and expenditures as the opportunity cost of time varies over the lifecycle. We find this alternative model successfully explains the joint behavior of food and entertainment expenditures in the latter half of the lifecycle. The model, however, is less successful in explaining expenditure patterns during the early half of the lifecycle.

Authors' Acknowledgements

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

1. Introduction

The well known hump-shaped profile of lifecycle expenditures has been extensively studied within economics.¹ Specifically, after accounting for changes in family size, consumption expenditure increases through middle age and then declines sharply thereafter. This holds for nondurable expenditure as well as total expenditure. For example, conditional on family size and cohort fixed effects, non-durable expenditure excluding education and health increases by roughly 30 percent between the ages of 25 and 45 and then falls by nearly the same amount between 45 and 70.²

In this paper, we revisit this familiar fact by decomposing nondurable expenditures into more detailed consumption categories. In doing so, we show that there is a tremendous amount of heterogeneity across the lifecycle profiles of individual consumption categories. Essentially the entire decline in nondurable expenditure late in the lifecycle is driven by three categories – food, nondurable transportation, and clothing/personal care. Expenditure on these categories are positively correlated with market work. Food is amenable to home production (see Aguiar and Hurst 2005, 2007) while transportation and clothing are inputs into market work.³ The remaining categories of nondurable expenditures, constituting roughly half of total nondurable expenditures, do not decline over the back half of the lifecycle. These categories include entertainment, housing services, charitable giving, and utilities. Moreover, expenditures on several of these categories, most notably entertainment, increase over the latter half of the lifecycle. Any explanation of the lifecycle profile of expenditures needs to match the fact that food

¹ This literature extends back nearly 40 years. See Thurow (1969).

² Authors' calculation (Figure 1 below). These results are consistent with the findings in the literature. See, for example, 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).

³ Banks et al (1998) and Battistin et al (2006), for examples classify clothing and nondurable transportation as work related expenses.

expenditures (a necessity) falls during the back half of the lifecycle while expenditure on entertainment (a luxury) rises during the back half of the lifecycle.

In addition, we revisit the stylized facts about the increasing cross sectional dispersion in expenditures over the lifecycle, as documented by Deaton and Paxson (1994). While the cross sectional dispersion for composite nondurable expenditures increases over the lifecycle, we also show the increase in dispersion is limited to primarily three categories: clothing, alcohol/tobacco, and a residual category which we call “other non-durables”. A composite measure of non-durables that excludes these categories shows little increase in cross sectional dispersion over the lifecycle. In fact, the cross-sectional dispersion of entertainment expenditures actually declines over the lifecycle, undermining the notion that dispersion is driven by idiosyncratic permanent income shocks.

To explain the facts documented in this paper, we propose a Beckerian model of consumption commodities that emphasizes the *intra*-temporal substitution between time and expenditures in consumption as the opportunity cost of time varies over the lifecycle (Becker 1965). It is well known that such a model can qualitatively generate a hump in lifecycle expenditures (see, for example, Ghez and Becker 1975 and Aguiar and Hurst 2007). The model has other testable implications as well, particularly regarding the joint allocation of time and expenditures over the lifecycle.

In the final part of the paper, we calibrate the model and assess its quantitative ability to match the joint behavior of food and entertainment expenditures over the lifecycle. To do so, we use data on both expenditures and time allocation. We show that entertainment expenditures and time allocated to entertainment are positively correlated

over the lifecycle, suggesting complementarity between time and goods. Conversely, food expenditures and time allocated to food preparation are negatively correlated, suggesting substitutability between time and goods. To gain a precise estimate of these intra-temporal elasticities, we use changes in behavior of both time and expenditures around retirement. We then predict the lifecycle behavior of entertainment expenditures conditional on the observed lifecycle behavior of food expenditures as well as the observed time allocated to food and entertainment.

The model suggests that entertainment expenditures should be relatively stable between ages 43 and 60. Specifically, the model predicts declines in expenditure of 3 percent between age 43 and 51 and an additional 3 percent between age 52 and 60. The data imply respective changes of +3 and -3 percent. The fact that the model matches the divergence of food and entertainment expenditures in the latter half of the lifecycle suggests that this striking feature of the data is consistent with the Beckerian model of consumption commodities.

The Beckerian model is less successful in explaining the first half of the lifecycle. In particular, the model predicts that entertainment expenditure should increase by 1 percent between age 25 and 33 and be unchanged between age 34 and 42. In the data, the respective changes are increases of 47 and 35 percent. One way to interpret this failure is through the data on time allocation. The model suggests that agents should delay time spent on entertainment until the complementary expenditure is high, that is delay entertainment time until middle age. The time freed up should instead be allocated to home production, where the margin of substitution between time and goods is high. This is not the pattern observed in the data. Relative to their 30s and 40s (and to expenditures

on entertainment), agents in their 20s allocate an abundance of time to entertainment. Note that the low level of expenditure while young may be due to liquidity constraints and/or precautionary savings. However, these forces cannot explain why the young allocate so much time to entertainment rather than food production – there is no equivalent constraint on time allocation. This allocation of time may instead reflect the high returns to building social capital for the young and the low returns to home production before the accumulation of a stock of home durables. We return to these extensions in the concluding section.

The remainder of this paper is organized as follows. Section 2 documents the lifecycle profiles for composite nondurable expenditure and its key sub-components. Section 3 introduces the Beckerian model. Section 4 calibrates the model and tests its ability to match the joint lifecycle profile of food and entertainment expenditures. Section 5 reviews the related literature, focusing on the canonical explanations for the lifecycle profile of nondurable expenditures. Section 6 concludes.

2. Lifecycle Expenditure

2A: Total Nondurable Expenditure Over the Lifecycle

In this section we document that the familiar “hump-shaped” profile of nondurable expenditure over the lifecycle reflects the aggregation of heterogeneous profiles for different types of goods.⁴ To begin, we review the facts about total nondurable expenditure over the lifecycle, and then disaggregate total nondurables into various sub-components.

⁴ Studies that document a humped shape profile for nondurables conditional on family size include, among others, Attanasio et al. (1999) and Fernandez-Villaverde and Krueger (2007).

Our data is from the Consumer Expenditure Survey (CEX). We use the NBER CEX extracts, which includes all waves from 1984 through 2003. We restrict the sample to households that 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 contains 53,412 households. Appendix A contains additional details about the construction of the dataset and sample selection.

We adjust all expenditures for cohort and family size effects. The CEX is a cross-sectional survey and therefore age variation within a single wave represents a mixture of lifecycle and cohort effects. Moreover, expenditures are measured at the household level and not the individual level. Household size has a hump shape over the lifecycle, primarily resulting from the fact that children enter and then leave the household. Additionally, marriage and death probabilities change over the lifecycle. We identify lifecycle from cohort variation by using the multiple cross-sections in our sample, and use cross-sectional differences in family size to identify family size effects. Specifically, we estimate the following regression:

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

where C_{it}^k is expenditure (in 2000 dollars) 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), $Cohort_{it}$ 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 and 10 household size dummies. The coefficients on the age dummies, β_{age} , represent the impact of the lifecycle conditional on cohort and family size fixed effects, both of which we allow to vary across expenditure categories. The fact that family size effects are allowed to differ across expenditure categories accommodates varying degrees of returns to scale across goods.⁵ All expenditures are adjusted for family size in this manner.

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 product-level PCE deflators from the National Income Accounts. 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 lifecycle. However, we do not think that time effects are driving the results, as the patterns depicted below are relatively stable across all cohorts in the sample.

Figure 1 plots the familiar lifecycle profile of core nondurable expenditures and total nondurable expenditures. Core nondurables consist of expenditure on food (both

⁵ An alternative approach is to use “adult equivalence” scales, such as those developed by the OECD (see for example <http://www.oecd.org/dataoecd/61/52/35411111.pdf>). The difficulty with these scales is they are designed for total expenditure, and the same scale may not be suitable for all sub-components. For example, food and housing services likely have different returns to scale and therefore should have different normalizations. Our approach allows household size adjustments to vary across goods. One drawback of our methodology is that household size may be correlated with such variables as permanent income. Depending on the sign of this correlation, we may be over or under adjusting expenditure for household size. We have experimented with several alternative adjustments to account for changing family size over the lifecycle. In all specifications, entertainment expenditures were non-decreasing after middle age while food expenditures declined after middle age.

home and away), alcohol, tobacco, clothes and personal care, utilities, domestic services, nondurable transportation (including air fare), nondurable entertainment, gambling, contributions to non-profits, business services and expenses related to life insurance, publications, and lodging away from home. Total nondurables are core nondurables plus housing services, calculated as either rent paid or the self-reported rental equivalent of the respondent's house. We exclude expenditures on education and health care from the analysis, as the utility (or returns) from consuming these goods vary significantly over the lifecycle.

The figure represents log-deviations from households whose head is 25 years old. The dots represent individual data points and the lines represent the 3 period moving average of the respective series. Figure 1 replicates the well-documented profile of nondurable expenditures over the lifecycle, with core nondurable expenditure peaking in middle age at 30 percent higher than the level of 25 year old expenditure, and then declining by the same amount over the latter half of the lifecycle. Total nondurable expenditure rises faster early in the lifecycle, but then does not decline as significantly later in the lifecycle. The gap between the two series represents the lifecycle behavior in housing services, which we will discuss on its own in the next sub-section.

2B. Disaggregating Nondurable Expenditure Over the Lifecycle

The lifecycle profile of nondurable expenditures depicted in Figure 1 has been the subject of numerous studies. The standard approach to lifecycle consumption follows the canonical models of Modigliani and Brumberg (1954) and Friedman (1957). These permanent income or “consumption smoothing” models imply that marginal utility of

consumption is a martingale (up to a discounting term). Almost all studies equate consumption with expenditures, making the lifecycle “hump” at first glance a challenge to the canonical models. Perhaps the two explanations that have gained the most advocates are the precautionary savings model (with labor income risk and impatient consumers) and poor planning models, where the latter includes models in which agents do not plan or cannot commit to a plan. These popular models are discussed in section 5.

If the observed pattern of expenditure is due to inter-temporal substitution or movements in financial resources, then luxury goods should respond more than necessities. To gain insight into the validity of such mechanism we highlight the lifecycle expenditure on food and nondurable entertainment. We highlight these two categories for a reason. Food is the canonical necessary good. Entertainment, on the other hand, has a relatively high income elasticity, and therefore a relatively high inter-temporal elasticity of substitution as well.⁶ If the declines in expenditure in the second half of the lifecycle are due to poor planning, time inconsistency, or impatience, entertainment spending will decline more than food expenditure. Indeed, as we show next, the opposite occurs.

Figure 2a plots the level of these two expenditure categories adjusted for family size and cohort effects, as described by equation (1). Food consists of food purchased for consumption at home plus food consumed away from home. Non-durable 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

⁶ We take as a premise that food has a lower income elasticity than nondurable entertainment expenditure. In our sample, regressions of food expenditure on household income (or total expenditure instrumented with household income) yield an estimated income elasticity less than one, while the elasticity for entertainment expenditure is consistently much greater than one across various specifications. However, as discussed below, these types of regression mix income effects and substitution (“price of time”) effects, and therefore do not accurately isolate the income elasticity.

sets and does not include reading material and magazine subscriptions. The average annual expenditure on food is \$6,061 in year 2000 dollars, while entertainment totals \$1,429. This represents 29 and 7 percent of core non-durable expenditures, respectively. We anchor each line at the average spending of 25 year olds across all cohorts and family sizes, and then use the coefficient on the age dummies to plot the remainder of the lifecycle.

Figure 2a indicates the food follows the general shape of aggregated nondurable expenditures – expenditures increase in the first half of the lifecycle, peak in the early 40s, and then steadily decline in the latter half of the lifecycle. Entertainment expenditures exhibit a different pattern. Like food, entertainment increases until the early 40s. However, expenditures on entertainment do not fall over the lifecycle. Instead, spending on entertainment continues to increase through retirement.

Figure 2b plots the same series in log deviations from the average expenditure of households with a 25 year old head. The figure indicates that food expenditure increases roughly 30 percent by the early 40s, and then declines by that amount between age 40 and age 75. Entertainment, on the other hand, increases 70 percent by age 45, and then increases another 10 percent through the early 70s.

These patterns are at odds with the predictions of most standard theories put forth to explain the lifecycle profile of expenditures. Moreover, plausible models of poor planning or extreme impatience would not predict an increasing profile of entertainment expenditures over the back side of the lifecycle. As a result, an alternative framework is needed to explain these patterns, a point we revisit more formally in the next section.

The various panels of Figure 3 plot the lifecycle profile of other components of nondurable expenditure. We group the goods by their expenditure profile over the lifecycle. Panel A depicts nondurable transportation expenditures, such as car maintenance, gas, tolls, air fare, etc. Aside from food, nondurable transportation is the only other category of non durable expenditure that displays a prominent “hump” in expenditure like that found in total nondurable expenditures. As we will show later, the decline in expenditures on non-durable transportation correlates strongly with the measures of time spent working. This is consistent with non durable transportation being a complementary expenditure with market work.

Panel B depicts alcohol and tobacco, clothing and personal care, and a residual non-durable expenditure category that includes business services, expenses related to life insurance, publications, and lodging away from home (including spending on children’s school lodging). These categories start out at a high level of expenditure and then fall steadily throughout the lifecycle. The decline in expenditure starting in middle age is particularly pronounced for these categories. Specifically, the log deviation between 45 year olds and 68 year olds is -1.3 for alcohol and tobacco expenditures, -0.7 for other non-durable expenditures, and -0.5 for clothing and personal care expenditures. These categories decline over the back side of the lifecycle at a much greater rate than the composite non-durable consumption measure.

Panel C collects categories that, like entertainment, do not decline over the latter half of the lifecycle. These categories are utilities, housing services, and domestic services. The latter category does decline slightly in the middle of the lifecycle, perhaps

reflecting lifecycle child care needs, before increasing late in life, which likely represents health-care related assistance.

Figures 2 and 3, taken together, document substantial heterogeneity in the lifecycle profile of expenditure of various goods. We collect the key patterns in Table 1. We break the goods into two categories, corresponding to those that fall over the latter half of the lifecycle and those that do not fall. Each category totals roughly half of our measure of total non-durable expenditures.

There are two additional facts that can be discerned from Table 1. First, the categories that experience the most marked declines in expenditure between the ages of 45 and 60 are also the same goods that experience the most marked decline in expenditures during the retirement years. The literature studying the decline in expenditure associated with retirement, the so-called “retirement consumption puzzle”, has typically been pursued independently of the literature on lifecycle consumption. (See Hurst (2007) for a survey of the retirement consumption literature). Nevertheless, the standard explanations for the fall in spending at the time of retirement overlap with those proposed for the lifecycle, including poor planning (see, for example, Bernheim et al. (2001)), time inconsistent preferences (see, for example, Angeletos et al. (2001)), and non-separability in utility between consumption and leisure (see, for example, Laitner and Silverman (2005)).

During the retirement years, food expenditures decline while entertainment expenditures increase. This is inconsistent with plausible stories of poor planning and time inconsistent preferences. More generally, half of the components of total non durable expenditures actually rise during the retirement years. The results in this paper

cast doubt on the existence of a retirement consumption “puzzle.”⁷ As we show below, these patterns at the time of retirement are consistent with some goods being complements with time (like entertainment) and others being substitutes with time (like food via home production or clothing and transportation via work related expenses).

The second fact that Table 1 highlights is that the timing of the declines for the “falling” categories is closely tied to the lifecycle profile of market work hours. For household heads, employment rates and market work hours begin to decline in the mid 40s and begin to fall off sharply starting in the early 50s. For reference, Figure A1 in the appendix plots the lifecycle profile of employment rates for household heads (solid line) and the hours per week spent working (unconditional on employment) household heads (dashed line) for our cross sectional sample of CEX respondents. Given the lack of panel data for almost all consumption categories and the fact that work hours are strongly correlated with permanent income in the cross section, it is hard to identify the correlation between work hours and spending conditional on income at the household level using cross sectional data. To isolate the effect of variation in work hours, we therefore look at the correlation of average market work hours at each age with average expenditures on a category. The lifecycle profiles of non-durable transportation, food, and clothing/personal care have a correlation with the lifecycle profile of work hours of 0.68, 0.52, and 0.90, respectively. The corresponding correlations for entertainment and utilities are -0.51 and -0.62. These results are consistent with the hypothesis that transportation and clothing are complements to market work, while food expenditures are substitutes for time spent in home production. Conversely, the evidence suggests that

⁷ 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).

expenditures on entertainment and utilities complement time spent away from work, a claim we revisit below.

2C. Cross Sectional Dispersion in Expenditure over the Lifecycle

Along with the “hump” in mean expenditure over the lifecycle, a second influential finding concerns the evolution of cross-sectional expenditure inequality over the lifecycle. The influential paper of Deaton and Paxson (1994) documented that the cross sectional variance of log consumption expenditures increases over the lifecycle. In the standard model, this is a violation of insurance and implies uninsurable shocks to permanent income that accumulate over the lifecycle (see also Storesletten et al (2004b), Heathcote et al (2005), and Guvenen (2007)).

A related issue is whether the shape of lifecycle expenditure profile varies by income and education. Several papers have argued that less educated households are more likely to be poor planners or to exhibit time inconsistent preferences (see, for example, Lusardi and Mitchell (2007) and Laibson et al. (2007)). Moreover, a prominent result of Carroll and Summers (1995) is that lifecycle expenditure tracks income profiles across educational attainment. Lastly, Bernheim et al (2004) document that food expenditures drop relatively more at retirement for low wealth and low income households.

In this section we revisit the evolution of expenditure dispersion over the lifecycle. We first characterize the changing cross-sectional distribution of expenditures over the lifecycle, and then explore mean expenditures conditional on educational attainment. Our first measure of dispersion is the standard deviation of log expenditure at each age after controlling for cohort and family status. To be precise, we analyze the

residuals from the regression of log expenditure on age, cohort, and family status dummies (equation (1)). We compute the sample standard deviation of the residuals at each age for each cohort. We then regress the standard deviations on age and cohort dummies in order to remove any cross-cohort variation in dispersion. As with mean expenditures, collinearity prevents the identification of time trends separate from age and cohort effects.⁸

In Figure 4a we plot the cross sectional standard deviation of log expenditure for core nondurables, food, and entertainment. The pattern for core nondurables is roughly similar to that in Deaton and Paxson's study. In particular, dispersion is relatively stable until age 40, and then increases steadily over the remainder of the lifecycle. If the dispersion represents shocks to income, then entertainment expenditures should show a greater increase in dispersion than do food expenditures. However, this is not the case. In fact, the cross sectional dispersion in entertainment *declines* over the lifecycle, while the dispersion of food expenditures is relatively stable.

The relative stability of food and the declines in entertainment expenditure dispersion over the lifecycle begs the question of the source of the increasing dispersion for total nondurables shown in Figure 4a and documented in Deaton and Paxson (1994). One source is that there are several categories of nondurable expenditures for which dispersion increases sharply over the lifecycle. We summarize the dispersion in expenditures for the main sub-components of total non-durables in Table 2. Three categories stand out as exhibiting sharp increases in dispersion later in the lifecycle – clothing and personal care items, alcohol and tobacco, and the residual “other non-

⁸ 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.

“durable” category (which includes books and publications, lodging away from home, and business services). Recall from Table 1 and Figure 3b that these three categories experience the sharpest declines over the back side of the lifecycle. Figure 4b plots the cross sectional dispersion of log expenditures on clothing and personal care, alcohol and tobacco, other nondurables, and core nondurables excluding these three categories. We see that differences across households in spending on clothing/personal care, other non durables, and alcohol/tobacco increase by 0.4, 0.6 and 0.6 log points between age 25 and age 68, respectively. Excluding these categories from nondurables reduces the increase in lifecycle dispersion to below 0.15 log points.⁹

To provide a more complete picture of expenditure inequality over the lifecycle, we document the evolution of key percentiles of the cross sectional distribution for food and entertainment. Specifically, we estimate quantile regressions using the same covariates as in equation (1). We do this for the 10th, 50th, and 90th percentiles. We then plot the coefficients on the age dummies (with age 25 the omitted group) in figure 5 for food (panel A) and entertainment (panel B).¹⁰ Note that this normalization allows us to depict the *change* in each percentile over the lifecycle. For context, the notes to the figures report the respective percentiles for 25 year olds, pooling the sample across cohorts.

⁹ In Appendix Figure A2 we plot the standard deviation of work hours over the lifecycle (relative to 25 year olds). Notice that the dispersion in work hours increase dramatically between the ages of 50 and 65. As a result, it should not be surprising that the dispersion in work related expenses (such as clothing and non durable transportation) increase sharply between the ages of 50 and 65.

¹⁰ Note that this exercise is not equivalent to tracking a particular household (say, the 10th percentile household in total lifetime expenditures) over the lifecycle. Rather, it characterizes the distribution of expenditure on food and entertainment at each age (after conditioning on cohort and household size), regardless of whether, say, the 10th percentile household in food or entertainment expenditure at age 25 is the same as the 10th percentile household at age 45.

Figure 5 reinforces the conclusions of figure 4a. Specifically, in panel A we see that the percentiles of the food expenditure distribution all move in lock-step over the lifecycle, with each percentile following the same hump-shaped profile. That is, there is no narrowing or widening of the food expenditure distribution. This holds as well for the 25th and 75th percentiles, which we omit from the figure for clarity. On the other hand, in panel B the 10th percentile of the entertainment expenditure distribution experiences the fastest relative growth over the lifecycle. Similarly, the (omitted) 25th percentile is increasing relative to the median, but not as much as the 10th percentile. That is, the bottom of the entertainment expenditure is catching up to the middle of the distribution, indicating a narrowing of the entertainment expenditure distribution over the lifecycle.

The fact that the bottom of the expenditure distribution for entertainment is increasing relative to the rest of the sample raises questions regarding the hypothesis that poorer households are particularly unprepared for declines in income later in the lifecycle. We shed more light on this question by exploring whether less educated households reduce their entertainment expenditures during the latter half of the lifecycle.

Figures 6a and 6b document the lifecycle pattern of food and entertainment expenditure for more and less educated households, respectively. Less educated refers to households whose head has completed 12 years or less of education. More educated households have heads with at least some college education. Figures 6a and 6b document that the profile of food expenditure does not rise as sharply early in the lifecycle for less educated households, consistent with the fact that income does not rise as steeply for these households, and falls more sharply later in the lifecycle, consistent with the evidence of Bernheim et al (2001). However, the profile of entertainment expenditures

by educational attainment contradicts claims that the sharp drop in food and total nondurable expenditures for less educated households is due to poor planning or declines in current income. In fact, entertainment rises relatively more later in the lifecycle for less educated households than it does for more educated households, consistent with our previous results on the narrowing of the cross-sectional gap in entertainment expenditures over the lifecycle.

3. A Beckerian Model of Consumption

In this section, we introduce a simple lifecycle model of consumption that builds on Becker (1965) in its emphasis on time as an input into consumption. We propose this model as an alternative to the standard models that emphasize income fluctuations and inter-temporal substitution as an explanation for the lifecycle profile of expenditures. Such a model seems plausible *ex-ante* given that the primary categories that have expenditures declining over the lifecycle are either amenable to home production (food) or are work related (clothing and non-durable transportation). In this section, we provide additional structure that will allow us to assess whether the movements in expenditure we observe are quantitatively consistent with the Beckerian model.

The model gives a prominent role to *intra*-temporal substitution between time and expenditures. The model yields several testable implications beyond a qualitative “hump” in lifecycle expenditures. In the next section, we assess the model’s ability to quantitatively match the joint behavior of different consumption categories over the lifecycle.

Agents consume a number of commodities indexed by $n=1, 2, \dots, N$. Period utility is given by $u(c^1, \dots, c^N)$, which is additively separable across time but not restricted to be separable across commodities within a period. Agents live for T periods and maximize the expected discounted sum of life time utility, with a discount factor β .

We follow Becker (1965) by representing the consumption commodities that enter utility as the outputs of production functions that take time and market goods as inputs. Specifically, each commodity n is formed by combining time inputs h^n and market goods x^n using a technology given by a production function f^n . That is, $c^n = f^n(h^n, x^n)$.¹¹ To give concrete examples, a commodity may be watching a television show, which combines a durable (the television), a cable subscription, and time. Similarly, another commodity may be a meal that takes groceries and time spent cooking as inputs. Note that in the former example time and market goods are complements, while in the latter example time and market goods may be substitutes (given the option to purchase food prepared by others). As we shall see, the degree of substitutability between time and market inputs in production is a key feature that distinguishes various commodities.

For our purposes, the remainder of the agents' environment is not crucial. We therefore follow the canonical partial-equilibrium models and assume that agents self insure by borrowing and lending at a constant interest rate r . Assets must be greater than some lower bound, \underline{a} . Agents face a competitive labor market with a spot wage w , which follows a Markov process whose transition probabilities may vary over the lifecycle. Time spent in market work is denoted L , and the total time endowment for a period is

¹¹ We assume these production functions satisfy the Inada conditions. We also rule out "joint production." That is, a time or market good used to produce commodity n cannot be simultaneously used to produce commodity $n' \neq n$.

normalized to one. The price of market input bundle n is given by p^n , which we take to be constant over the lifecycle.

The problem of an agent of age t , with assets a and facing a wage rate w , can be expressed in recursive form as:

$$V(a, w, t) = \max u(c^1, \dots, c^N) + \beta E' V(a', w', t+1),$$

subject to

$$\begin{aligned} c^n &= f^n(h^n, x^n), \quad n = 1, \dots, N \\ \sum_n h^n + L &= 1 \\ a' &= (1+r)a + wL - \sum_n p^n x^n \\ L &\geq 0, \quad a' \geq \underline{a}. \end{aligned}$$

Let μ be the multiplier on the time budget constraint, λ the multiplier on the income budget constraint, θ the multiplier on non-negativity of labor time, and γ the multiplier on the borrowing constraint. The home technology constraints can be substituted directly into the agent's utility function. Note that the optimal solution is the same as that of the relaxed problem in which all the constraints are expressed as inequality constraints. Therefore, we can write the Lagrangian in such a way that all multipliers are non-negative. The first order necessary conditions are (where subscripts denote the corresponding partial derivatives):

$$\begin{aligned} x^n &: u_n f_x^n = \lambda p^n, \quad \forall n \\ h^n &: u_n f_h^n = \mu, \quad \forall n \\ L &: \lambda w + \theta = \mu \\ a' &: \beta E' V_{a'}(a', w', t+1) = \lambda - \gamma. \end{aligned}$$

The envelope condition implies $V_a(a, w, t) = \lambda(1+r)$.

Note that if agents supply non-zero market labor, then the ratio μ/λ equals the market wage, w . More generally, the ratio μ/λ represents the price of time (in units of the numeraire). To simplify future expressions, we define $\omega \equiv \mu/\lambda$. Given that empirically many respondents are not actively working, and the fact that a spot labor market in which agents face linear wage schedules may not be the best characterization of empirical labor markets, we do not emphasize the market wage as the relevant price of time. We shall describe below an alternative estimate of ω .

We focus on the intra-period tradeoff between time and goods. Specifically, divide the first order condition for h^n by that of x^n to obtain:

$$\frac{f_h^n}{f_x^n} = \frac{\omega}{p^n}. \quad (2)$$

This condition states that the marginal rate of transformation between time and goods in the production of commodity n will be equated to the relative price of time. This condition is a static first order condition that holds regardless of the nature of the utility function, the completeness of asset markets, and the nature of risk facing the agents.

A key element of our analysis is the response of expenditures relative to time inputs as the price of time varies, $\frac{d \ln(x^n/h^n)}{d \ln(\omega)}$. Let σ^n denote the elasticity of substitution

between time and market inputs into the production of commodity n , which we assume to differ across commodities but remain constant as we vary inputs for a given commodity.

That is, $\sigma^n = \frac{d \ln(x^n/h^n)}{d \ln(f_h^n/f_x^n)}$. Taking logs and differentiating equation (2), we have

$\frac{d \ln \left(\frac{x^n}{h^n} \right)}{d \ln(\omega)} = \sigma^n$. The greater σ^n , the more expenditures will respond (relative to time inputs) to changes in the price of time over the lifecycle, given permanent income.

Given a price of time, the static first order condition pins down expenditures *relative* to time inputs. For our simulation exercises this is sufficient, given that we bring in independent observations on time allocation over the lifecycle. We delay discussion of this data until the next section. Nevertheless, to provide intuition we discuss how the level of expenditure inputs for different commodities varies with the price of time and with financial resources. For clarity, we make additional simplifying assumptions. Specifically, we assume u is additively separable across commodities and f^n are constant returns to scale. To repeat, we make these additional assumptions to gain intuition for the profile of lifecycle expenditures. We do not use them when calibrating and testing the model in the next section.

We start by differentiating the first order condition for x^n with respect to ω , holding λ constant:

$$\frac{c^n u_{nn}}{u_n} \frac{d \ln c^n}{d \ln \omega} + \frac{x^n f_{xx}^n}{f_x^n} \frac{d \ln x^n}{d \ln \omega} + \frac{h^n f_{xh}^n}{f_x^n} \frac{d \ln h^n}{d \ln \omega} = 0,$$

where all derivatives are taken holding λ fixed. Here we appeal to separability in utility by assuming $u_{nn'} = 0$ for $n' \neq n$. Using the constant returns relations $f_{xx}^n x^n = -f_{hx}^n h^n$ and

$\sigma^n = \frac{f_h^n f_x^n}{f_{xh}^n c^n}$, along with the fact that $\left. \frac{d \ln h^n}{d \ln \omega} \right|_{d\lambda=0} = \left. \frac{d \ln x^n}{d \ln \omega} \right|_{d\lambda=0} - \sigma^n$, we have a result of

Ghez and Becker (1975):

$$\left. \frac{d \ln x^n}{d \ln \omega} \right|_{d\lambda=0} = s_h^n \left(\sigma^n - \frac{u_n}{c^n u_{nn}} \right), \quad (3)$$

where the notation $d\lambda=0$ reminds the reader we are holding resources constant and varying the price of time. We denote the cost share of time in the production of consumption good n , $\frac{h^n f_h^n}{c^n}$, as s_h^n .

The expression states that expenditures for commodity n increase with the price of time if the intra-temporal elasticity of substitution between time and goods is greater than the inter-temporal elasticity of substitution in consumption. The size of the increase (or decrease) depends as well on the share of time in production. In particular, if the share of time in consumption is zero, then the movements in the price of time have no effect on expenditures. This is the implicit assumption of the vast majority of the literature on consumption.

The intuition of equation (3) is as follows. As the price of time increases, agents will substitute away from time and toward market inputs to achieve a given level of consumption. This is movement along a production isoquant and is parameterized by σ^n . However, the fact that time is costlier in the current period relative to other periods suggests shifting consumption to a period in which the total cost of consumption (time plus market goods) is less. The willingness to do this is given by the inter-temporal elasticity of substitution in consumption. For a fixed ratio of inputs, this reduces market expenditures. This is a parallel movement across isoquants (or levels of consumption).

In general, the net effect is theoretically ambiguous. However, consider food consumption. This commodity has a relatively low inter-temporal elasticity of substitution and a relatively high intra-temporal elasticity of substitution. It should

therefore positively covary with the price of time. Given that the price of time peaks in middle age and then declines through retirement, this is consistent with figure 2. Similarly, time and market goods are difficult to substitute in the production of entertainment (low σ) and entertainment has a relatively high inter-temporal elasticity of substitution, we should expect expenditures on entertainment to rise as the price of time falls. Again, this is consistent with the data assuming a reduction in the price of time later in the lifecycle. We shall return to these insights in the next section.

Now consider the response of expenditures as financial resources vary and the price of time is held constant. Differentiating the same first order condition while holding ω constant, we have:

$$\frac{c^n u_{mn}}{u_n} \frac{d \ln c^n}{d \ln \lambda} + \frac{x^n f_{xx}^n}{f_x^n} \frac{d \ln x^n}{d \ln \lambda} + \frac{h^n f_{xh}^n}{f_x^n} \frac{d \ln h^n}{d \ln \lambda} = 1.$$

The fact that ω is fixed implies that $d \ln h^n = d \ln x^n$, and together with the CRS fact that

$$x f_{xx} = h f_{xh}, \text{ we have } \left. \frac{d \ln c^n}{d \ln \lambda} \right|_{d\omega=0} = \frac{u_n}{c^n u_{mn}}. \text{ Viewed in this light, a negative income shock}$$

(an increase in λ) results in reduced consumption, with commodities with the greatest inter-temporal elasticity of substitution exhibiting the greatest decrease. That is, if agents find themselves short of resources, food and other necessities should be the expenditure category that declines the least, while entertainment should decline the most.

The canonical model focuses on the fact that expenditures respond to movements in life time resources, abstracting from movements in the price of time. This is not an issue if time is a small share of consumption inputs. However, if the Beckerian forces are empirically prominent, such analyses conflate price and income effects. There is a parallel to the analysis of labor supply, where it has long been recognized that

movements in wages represent both income and substitution effects. The same identification problem arises with consumption expenditures in the Beckerian framework.

The model provides a useful context to revisit the Deaton and Paxson dispersion plots presented in figure 4. If dispersion is driven by income shocks (that is, relative movements in the multiplier on resources, λ), then goods with the largest inter-temporal

elasticity of substitution $\left(\frac{-u_n}{c^n u_{nn}} \right)$ will display the greatest dispersion. This follows from

$$\frac{d \ln c^n}{d \ln \lambda} \Big|_{d\omega=0} = \frac{u_n}{c^n u_{nn}}.$$

This is the sense in which the cross-sectional standard deviation of entertainment expenditures should demonstrate a larger increase over the lifecycle if the dispersion is driven by relative movements in income. The fact that this is not the case provides strong evidence against the standard interpretation.

However, there are also large movements in the price of time over the lifecycle that may differ across households. Without taking a stand on several parameters for which we do not have strong priors, the model does not have clear empirical implications for the increase in food dispersion relative to entertainment in response to price of time shocks. The response of expenditure to movements in the price of time depends on the share of time in the production of the commodity as well as the difference between the intra-temporal and the inter-temporal elasticities of substitution (equation 3). It is safe to assume that the share of time is higher for entertainment. However, while the difference between the elasticities is likely to be positive for food and negative for entertainment, there is no clear prior on the relative *magnitudes* of the difference. Nevertheless, it is possible that food expenditures are relatively sensitive (in absolute value) to movements in the price of time. If this is the case, then the fact that the dispersion of food increases

as much (or more) than the dispersion in entertainment expenditures over the lifecycle suggests that cross-sectional movements in the price of time are at least as important (if not the dominant factor) as movements in permanent income. In any event, the data make clear that there must be a countervailing force that offsets movements in income, and the cross-sectional differences in the opportunity cost of time are a promising candidate.

4. Rationalizing Food and Entertainment Expenditures

The preceding analysis demonstrates that models of poor planning or the standard permanent income model that abstracts from time as an input into consumption cannot rationalize a decline in food expenditure that occurs simultaneously with an increase in entertainment. That is, a drop in available resources should generate a fall in entertainment expenditures that exceeds the fall in food expenditures, given plausible income elasticities. Such behavior, however, is potentially consistent with a Beckerian model that allows for time and market inputs to have different degrees of substitutability across commodities and large movements in the price of time over the lifecycle. We now turn to the question of whether an appropriately calibrated Beckerian model can rationalize the joint lifecycle behavior of food and entertainment expenditures.

The calibration exercises focuses on food and entertainment for two reasons. The first is, as mentioned above, that these two goods are particularly informative regarding whether expenditure is responding to a shock to resources or a shock to the cost of time. The second reason is that the time input into food and entertainment can be readily measured. The time input into other consumption categories, such as clothing or tobacco, is more difficult to delineate.

4A: Calibration

The Beckerian approach places the intra-temporal elasticity of substitution, and lifecycle movements in the price of time, at the forefront of the analysis. We, therefore, need to obtain empirical measures of these elasticities. We do so by exploiting equation (2), which equates the marginal rate of transformation between time and market inputs to the relative price of time.

More precisely, we use the implication that $\frac{d \ln \left(x^n / h^n \right)}{d \ln (\omega)} = \sigma^n$. Given an elasticity

of substitution, the model has a tight prediction for the response of market expenditures relative to time inputs to a movement in the price of time (holding resources constant).

Empirical testing of this prediction requires data on both market inputs and time inputs. The data on market inputs is from the CEX, as presented in section 2. The data on time inputs is from the American Time Use Survey (ATUS). This survey collects time diaries from a large, nationally representative sample of individuals each year, beginning in 2003. Each diary records every activity (as coded into over 400 categories) in a 24 hour period. Appendix B contains additional details on the ATUS.

Figure 7 plots the time allocated to food production (cooking and clean-up), total non-market production (cleaning, maintenance, shopping, etc. plus food production), and time allocated to entertainment (watching television, watching movies, socializing, exercise and sporting events, hobbies, etc.). Both series are conditional on household size and marital status (keep in mind, however, that time diaries are collected for individuals and not for households). Given the single cross-section, we do not control for

cohort effects. Specifically, we regress time allocated to each task in levels on age and family status controls. We construct the conditional mean time allocated at each age by adding the coefficient on the relevant age dummy to the unconditional mean for respondents aged 25. The figure depicts the log difference of this constructed conditional mean at each age minus the log mean time allocated at age 25.

The figure indicates that time spent on total home production and time devoted to food production both increase by roughly 30 percent between the age of 25 and the middle 40s, and then increase by another 40 and 30 percent respectively, as individuals reach the retirement years. Time spent on entertainment does not increase substantially early in the lifecycle, but then increases by roughly 35 percent between age 50 and age 70.

To test the model's predictions, we need estimates of the elasticities of substitution. In previous work (Aguiar and Hurst 2007), we have estimated the elasticity for food production to be between 1.5 and 2.2. In that paper, we also discuss how these estimates are consistent with other studies.¹² For the following analysis we take σ to be 1.5 for food consumption. We clarify below which results are sensitive to this choice.

Given this elasticity, we can calculate the movement in the opportunity cost of time implied by observed movements in food expenditures and time allocated to food production. Table 3 reports the relevant calculations. The first row reports the changes in food expenditure (using the CEX data and the methodology of Section 3) for various sub-sections of the lifecycle, while the second row reports the corresponding changes in the time allocated to food preparation and clean up ("food production" using the time use data). As in previous tables, the column headings indicate the respective age span, with

¹² See, for example, the estimates discussed in Rupert, Rogerson, and Wright (1995).

the end points representing three-year averages. For example, row 1, column 1 of table 3 reports that the average food expenditure for those aged 33 through 35 is 17 percent higher than those aged 25 through 28. The corresponding increase in time allocated to food production is 13 percent (second row).

Taking σ for food to be 1.5, the implied change in the price of time can be calculated from the model's prediction that $\Delta \ln(\omega) \approx \frac{\Delta \ln\left(\frac{x^n}{h^n}\right)}{\sigma^n}$. For example, the ratio of expenditure to time devoted to food increases by approximately 4 percent between the ages of 25 and 34 (that is, 17 percent-13 percent). The corresponding change in the price of time is therefore 3 percent (that is, $4/1.5$). The remaining columns of row three report the implied price of time for the remainder of the lifecycle. As one would expect, the opportunity cost of time increases early in the lifecycle and then declines after middle age. Peak retirement years (age 60 through 68) involve a decline in the price of time of roughly 10 percent.¹³

Note that we do not use wages as the price of time. The model relates wages to the price of time for those who supply non-zero market labor. However, this is not the case for retirees and others who are not employed. More generally, and as pointed out in the previous section, it is not robust to alternative models of labor markets. We therefore do not advocate the use wages as the price of time.

While we have an established literature on the elasticity of substitution between time and goods for food production, no comparable estimates exist for entertainment. We calibrate this parameter from the behavior during peak retirement years, and then

¹³ For comparison, the implied price of time based on shopping intensity (Aguiar and Hurst 2007, Figure 4), increases by a little more than 5 percent between age 25 and 34, and declines by a little more than 10 percent between age 60 and age 74. The price of time in our earlier study also peaked before the age of 40 and declined roughly 30 percent from the peak in the mid 30s through age 74.

assess the model's ability to match entertainment expenditure for the remainder of the lifecycle. Specifically, as discussed above, the behavior of time and market inputs into food production suggest a decline in the price of time of 10 percent between the ages 60 and 68. Over this period of the lifecycle, the average person increases time allocated to entertainment by 20 percent (row 4 of table 3). The last row of table 3 reports that market inputs into entertainment increase by 11 percent between age 60 and 68. That is, the ratio of market to time inputs into entertainment falls by 9 percent. The implied elasticity of substitution is therefore $9/10$ or 0.9 .

As one would expect, the elasticity of substitution between time and goods in entertainment is less than that for food. Indeed, from the extension of the previous section, the model predicts entertainment expenditure to increase during retirement if the inter-temporal elasticity of substitution for entertainment is greater than 0.9 , a plausible parameterization given the prior that time and goods are less substitutable in entertainment than in food production.

4B: Predictions for the Lifecycle

With this elasticity in hand, we predict entertainment expenditures over the lifecycle. Moving back through the lifecycle, the price of time declines by 15 percent between the ages of 52 and 60. Time allocated to entertainment increases by 11 percent over this period. The intra-temporal first order condition then predicts a decrease in expenditure of 3 percent, which is identical to the actual decrease in expenditure. Similarly, the model predicts a decrease in entertainment expenditures of 3 percent between age 43 and 52. The actual change is an increase of 3 percent. Overall, the model performs relatively well for this part of the lifecycle. That is, the observed pattern

of food expenditure and time allocated to food and entertainment suggests movements in entertainment expenditure that are essentially those seen in the data.

However, the model's predictions are strongly at odds with observed entertainment expenditures early in the lifecycle. The decline of 1 percent in the price of time between ages 34 and 43, combined with no change in the time allocated to entertainment suggest approximately no change in entertainment expenditures. Instead, the data show that entertainment expenditures *increase* by 35 percent during this age span. Similarly, the model suggests a small increase in entertainment expenditures of 1 percent between the ages of 25 and 34. However, actual expenditures increase by 47 percent.

Therefore, while the model rationalizes the decline in food expenditures combined with increasing entertainment expenditures in the latter half of the lifecycle, it fails to predict the sharp increase in entertainment expenditures early in the lifecycle. Note that this failure is not due to the presence of liquidity constraints, myopia, or precautionary savings. We are exploiting a static first order condition that abstracts from the inter-temporal allocation of consumption.

This failure is also not due to our choice of 1.5 for the elasticity of substitution in food production. A higher elasticity would scale down the movements in the price of time proportionally. However, to match entertainment expenditures at retirement, the calibration would scale up the entertainment elasticity proportionally. Therefore, while a different elasticity would lead to a different magnitude for the movement in the price of

time, it would generate no change in the predicted lifecycle pattern for entertainment expenditures.¹⁴

Where the model fails is in its prediction that a small increase in the ratio of food expenditures to time spent on food production should be accompanied by an even smaller increase (or even a decrease) in the ratio of entertainment expenditures to time spent on entertainment. Instead, consumers dramatically increase their entertainment expenditures relative to entertainment time early in the lifecycle.

One way to view this phenomenon is that the model suggests agents should delay entertainment time until entertainment expenditures are high. The time allocated to entertainment absent complementary market inputs would be more profitably spent on home production, an activity which has a greater degree of substitution. Instead, agents maintain a fairly stable level of time allocated to entertainment until late in the lifecycle, despite the large movements in expenditures.

There are at least two plausible forces that may be at work, both relating to our maintained assumption that production functions are stable over the lifecycle. The first is that home production of food and entertainment both require household durables. By calibrating based on allocations late in the lifecycle, we are capturing trade offs after the stock of durables has been built up. Younger individuals have a different stock of household durables and therefore possess different technologies for translating time and

¹⁴ More formally, note that $\hat{\omega}_t = (\hat{x}_{f,t} - \hat{h}_{f,t}) / \sigma_f$, where a “hat” indicates percentage change. We calibrate σ_e such that $\sigma_e = (\hat{x}_{e,T} - \hat{h}_{e,T}) / \hat{\omega}_T$, where T is the last period. Predicted expenditures are obtained from the relation

$$(\hat{x}_{e,t} - \hat{h}_{e,t}) = \sigma_e \hat{\omega}_t. \quad \text{Substituting in, we have } (\hat{x}_{e,t} - \hat{h}_{e,t}) = (\hat{x}_{f,t} - \hat{h}_{f,t}) \frac{(\hat{x}_{e,T} - \hat{h}_{e,T})}{(\hat{x}_{f,T} - \hat{h}_{f,T})}, \text{ which is}$$

independent of our choice of σ_f .

inputs into consumption. This is consistent with the fact that the model performs well after age 40.

A second issue is that young individuals may have a larger incentive to accumulate social capital, both for personal and professional reasons. This may raise the return to time spent on entertainment relative to expenditures on entertainment early in the lifecycle. The sharp increase in expenditures relative to time inputs between ages 25 and 43 may then reflect expenditures “catching up” with the time inputs as well as a decline in the marginal return to time spent networking as agents approach middle age. Moreover, the nature of entertainment time may be changing. For example, the young may be socializing with friends while the old may be watching TV. Friends may be a substitute to market expenditures on entertainment, rather than a complement. Therefore, the substitutability of time and market inputs into entertainment may vary over the lifecycle.

5. 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 food expenditures while simultaneously increasing entertainment expenditures, given that food has a lower income elasticity than entertainment.

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.

Such precautionary savings models have been extremely influential, in part due to their ability to explain the prominent shape of lifecycle expenditure 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 Guevenen (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.¹⁵ However, if impatience is the predominant force driving the decline in expenditures over the back side 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) should decline more in the latter half of the lifecycle than necessities (such as food).

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 particular, 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

¹⁵ 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.

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).

When focusing on the disaggregated data, we documented that the primary movements of expenditures over the back side of 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. Our model emphasizes a static first order condition, without placing strong restrictions on the inter-temporal allocation of consumption. However, the tests of the model validate that intra-temporal substitution between time and goods is quantitatively important. At a minimum, our goal is to show that inferences about income risk, impatience, and planning, that exploit the lifecycle profile of expenditure while ignoring this margin of substitution will therefore be incorrect. Expenditures respond to both changes in the price of time given permanent income as well as shifts in permanent income. The large class of papers that map observed lifecycle (or, for that matter, business cycle) movements in consumption into movements into life time resources without controlling for movements in the price of time are conflating the two effects. This is not a problem if time is not an important input into consumption (or if the price of time is not moving). But, as we show in this paper, such assumption is empirically invalid.

This paper is a third in a series re-visiting consumption expenditures in light of time allocation. In Aguiar and Hurst (2005), we use food diaries to document that food expenditure declines at retirement do not imply food consumption declines. In Aguiar

and Hurst (2007), we document that shopping intensity for food, and therefore price paid for equivalent grocery items, varies systematically over the lifecycle. Both studies focused exclusively on food expenditures. In standard models, food expenditure is particularly revealing about income movements given its low income elasticity. However, food is also amenable to home production, making large movements in food expenditures perfectly consistent with constant income and constant consumption. By focusing on food, a relatively small component of non-durable expenditures, our previous work left open the question of whether, and to what extent, these Beckerian insights were relevant for other types of expenditure and to what extent that they were important for explaining the well documented profile of total non durable expenditures over the lifecycle. This paper, among other goals, serves to resolve these outstanding questions.

6. Conclusion

This paper documented that the hump in lifecycle expenditures on nondurables masks informative heterogeneity across individual expenditure categories. In particular, we highlight that food, a necessary good, declines relative to entertainment (and several other categories) in the second half of the lifecycle. Moreover, the increase in the cross-sectional dispersion of expenditure does not vary across goods with different income elasticities. This poses a challenge to theories that exclusively emphasize movements in income or inter-temporal substitution. However, the qualitative pattern is consistent with a Beckerian model in which time and good are substitutes for food, but complements in entertainment. Quantitatively, such a model does well in matching the joint allocation of expenditures and time on food and entertainment in the latter half of the lifecycle.

However, the model fails to account for the patterns observed early in the lifecycle, potentially suggesting extensions that emphasize the accumulation of home durables and social capital early in the lifecycle.

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 the price of time. The former will have a relatively small impact on necessities such as food, while the latter will have a larger impact on goods, like food, for which time and market inputs are substitutes. 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. Recessions are periods in which both income and the price of time fall. Analyses that ignore the latter will draw misleading conclusions about the importance of income uncertainty.

Data Appendix¹⁶

A. CEX Data

This paper uses data from the Consumer Expenditure Survey's quarterly interview survey. The survey unit is a household (consumer unit). Each consumer unit is interviewed once per quarter for five consecutive quarters. The first interview collects demographic data and inventories major durables. The subsequent four interviews collect recall data on expenditures over the preceding three months. We collapse the four interviews into a single annual observation per household, summing over the quarterly expenditures. In particular, we do not use the panel dimension of the four quarterly interviews.

While expenditure is reported at the household level, demographics are reported for individuals. We use demographic characteristics reported by the household head. A head is defined as the member who identifies himself or herself as the "head of household" in the survey. If there are multiple heads, we identify the head as the male (if one is present) and resolve any remaining ties by employment (employed over nonemployed), age (eldest), and marital status (married over non-married).¹⁷

We use the extracts compiled by Ed Harris and John Sabelhaus and provided by the NBER (http://www.nber.org/data/ces_cbo.html). All data, programs, and documentation for this paper can be found on the authors' website (<http://troi.cc.rochester.edu/~maguiar/deconstructing/datapage.html>). Harris and

¹⁶ All data sets and programs are posted on the authors' website (<http://troi.cc.rochester.edu/~maguiar/deconstructing/datapage.html>).

¹⁷ There are a handful of households with multiple heads who share the same sex, age, employment status, and marital status (as well as household size). However, as these are the only demographic variables used in this paper, this duplication is immaterial to identifying the demographic characteristics of the household.

Sabelhaus aggregate expenditures into 47 categories, which are listed in the documentation posted on the authors' website. The Harris and Sabelhaus dataset includes households whose first interview was conducted between the first quarter of 1980 through the second quarter of 2003. Due to changes in the survey methodology, data from the last two quarters of 1985 and 1995 are omitted.¹⁸ The data set contains a total of 167,133 households.

We restrict the Harris and Sabelhaus sample in the following ways. First, we keep households whose heads are between age 25 and 75. To obtain reliable estimates of cohort effects, we restrict attention to cohorts with at least 10 years of data. In particular, we restrict the sample to households whose head is at most 65 in 1980Q1 and at least 35 in 2003Q2. This leaves 122,962 households. Second, the household must have completed all four expenditure surveys, providing a complete picture of *annual* expenditures. There are 75,883 such households in the sample, or roughly 62 percent. Harris and Sabelhaus provide adjusted weights to use with the restricted sample. However, the restricted sample of Harris and Sabelhaus also excludes households with incomplete income reports and students. Usage of their adjusted weights necessitates excluding these households as well, leaving 58,305 households.

Our final sample restriction is that households must have strictly positive expenditure on six major expenditure categories: food, housing services, utilities, clothing and personal care, nondurable transportation, and nondurable entertainment. Roughly 92 percent of the sample satisfied this last criterion, resulting in a sample of

¹⁸ Prior to 1984, only urban consumers were surveyed. Exclusion of these years does not significantly alter the results reported in the paper.

53,412 households. This is our main sample for analysis. Summary demographics for the sample are reported in Table A1 (to be completed).

To deflate expenditures, we use the National Income and Product Account's Price Indexes for Personal Consumption Expenditures by Type of Product. We deflate each expenditure category to year 2000 dollars using the corresponding product deflator, matching the calendar year of the annual deflator with the calendar year in which the household was first surveyed.

B. Time Use Data

We use the 2003, 2004, and 2005 waves of the American Time Use Survey (ATUS) conducted by the U.S. Bureau of Labor Statistics (BLS). Participants in ATUS, which includes children over the age of 15, are drawn from the existing sample of the Current Population Survey (CPS). The individual is sampled approximately 3 months after completion of the final CPS survey. At the time of the ATUS survey, the BLS updated the respondent's employment and demographic information. The ATUS waves totaled 20,720, 13,973, and 13,038 respondents in 2003, 2004, and 2005, respectively. We restrict our sample to respondents aged 25 through 75, resulting in sample sizes of 16,860, 11,436, and 10,580, respectively. We pool these 38,876 respondents into a single cross section.

The survey uses a 24-hour recall of the previous day's activities to record time diary information. The unit of analysis is an individual, and only one individual per household is surveyed. We control for effects of marriage and family size by regressing the amount of time (in levels) for a specific activity on one-year age dummies, a dummy

for marital status, and ten family size dummy variables, and report the coefficients on the age dummies.

The ATUS reports time allocation using over 400 detailed activity codes. For our analysis we focus on three aggregates. Food production includes time spent on cooking, meal preparation, and meal clean up. Total non-market production includes food production plus time spent on housekeeping, indoor and outdoor home maintenance, vehicle maintenance, and obtaining goods and services. Entertainment consists of time spent socializing, watching television, enjoying non-tv entertainment, and pursuing hobbies.

References

Aguair, Mark and Erik Hurst (2007). "Lifecycle Prices and Production", *American Economic Review*, 97(5).

Aguair, Mark and Erik Hurst (2005). "Consumption vs. Expenditure", *Journal of Political Economy*,

Aguair, Mark and Erik Hurst (2004). "Consumption vs. Expenditure", NBER Working Paper 10307.

Aiyagari, S Rao, 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109(3), 659-84.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001), "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation", *Journal of Economic Perspectives*, 15(3), 47-68.

Attanasio, Orazio, James Banks, Costas Meghir, and Guglielmo Weber (1999). "Humps and Bumps in Lifetime Consumption." *Journal of Business and Economic Statistics*, January 1999, 17, 1, 22-35.

Attanasio, Orazio, and Guglielmo Weber (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, 103(6), 1121 – 1157.

Banks, James, Richard Blundell, and Sarah Tanner (1998). "Is There a Retirement Savings Puzzle?" *American Economic Review*, 88(4), 769 – 788.

Martin Browning & Thomas F. Crossley, 2000. "Luxuries Are Easier to Postpone: A Proof," *Journal of Political Economy*, University of Chicago Press, vol. 108(5), pages 1022-1026, October

Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber (2006). "How Large is the Retirement Consumption Drop in Italy?", mimeo.

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.

Blundell, Richard, Martin Browning, and Costas Meghir (1994), "Consumer Demand and the Life-Cycle Allocation of Household Expenditures," *Review of Economic Studies*, 61, 57-80.

Carroll, Christopher (1997). "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112(1): 1-56.

Carroll, Christopher and Lawrence Summers (1991). "Consumption Growth Parrells Income Growth: Some New Evidence" In *National Saving and Economic Performance*, edited by B. Douglas Bernheim and John Shoven. Chicago: University of Chicago Press.

Charles, Kerwin and Mel Stephens (2006). "The Level and Composition of Consumption Over the Business Cycle: The Role of 'Quasi-Fixed' Expenditures", in *Working and Poor: How Economic and Policy Changes Are Affecting Low-Wage Workers*, ed. Rebecca Blank, Sheldon Danziger, and Robert Schoeni, Russell Sage Foundation.

Deaton, Angus, and Christina Paxson (1994), "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102, 437-67.

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

Ghez, Gilbert and Gary Becker (1975). *The Allocation of Time and Goods over the Life Cycle*, Columbia University Press, New York.

Gourinchas, Pierre-Olivier and Jonathan Parker (2002), "Consumption over the Life Cycle," *Econometrica*, January 2002, 70, 1, 47-89.

Guvenen, Fatih (2007), "An Empirical Investigation of Labor Income Processes", NBER Working Paper 13394.

Heathcote, Jonathan, Kjetil Storelsetten, and Giovanni Violante (2005), "Two Views of Inequality over the Life Cycle," *Journal of the European Economic Association*, April-May 2005, 3, 2-3, 765-775

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.

Hurst, Erik (2007). "Understanding Consumption in Retirement: Recent Developments", Pension Research Council Working Paper WP2007-26.

Krueger, Dirk, and Jesus Fernandez-Villaverde (forthcoming), "Consumption over the Life Cycle: Some Facts from CEX Data," *Review of Economics and Statistics*.

Laibson, David, Andrea Repetto, and Jeremy Tobacman (2007). "Estimating Discount Functions with Consumption Choices over the Lifecycle". NBER Working Paper 13314.

Laitner, John and Dan Silverman (2005). "Estimating Life-Cycle Parameters from Consumption Behavior at Retirement", NBER Working Paper 11163.

Lusardi, Annamaria and Olivia Mitchell (2007). "Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth", *Journal of Monetary Economics*, forthcoming.

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.

Storelsletten, Kjetil, Chris Telmer, and Amir Yaron (2004a), "Cyclical Dynamics in Idiosyncratic Labor-Market Risk," *Journal of Political Economy*, 112(3), 695-717.

Storelsletten, Kjetil, Chris Telmer, and Amir Yaron (2004b), "Consumption and Risk Sharing Over the Lifecycle," *Journal of Monetary Economics*, 51(3), 609-633.

Thurow, Lester (1969). "The Optimum Lifetime Distribution of Consumption Expenditures", *American Economic Review*, 59(3), 282-97.

Zeldes, Stephen (1989). "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2): 305-346.

Table 1: Summary of Change in Expenditure over the Lifecycle by Consumption Category

<i>Disaggregated Consumption Group</i>	<i>Share of Core Non Durable Expenditure</i>	<i>Share of Non Durable Expenditure with Housing</i>	<i>Log Change Between 25 and 45</i>	<i>Log Change Between 45 and 60</i>	<i>Log Change Between 60 and 68</i>
Group 1					
Food	0.31	0.21	0.30	-0.11	-0.07
Transportation	0.19	0.13	0.35	-0.13	-0.15
Clothing and Personal Care	0.10	0.07	0.26	-0.30	-0.18
Alcohol and Tobacco	0.04	0.03	-0.62	-0.78	-0.55
Other Non-Durable	0.09	0.06	0.16	-0.40	-0.26
Group 1 Total/Weighted Mean	0.73	0.50	0.24	-0.21	-0.15
Group 2					
Entertainment	0.06	0.04	0.82	0.04	0.05
Housing Services		0.31	0.62	0.15	0.07
Utilities	0.16	0.11	0.73	0.25	0.05
Domestic Services	0.03	0.02	0.81	0.27	0.14
Charitable Giving	0.03	0.02	0.77	0.18	0.04
Group 2 Total/Weighted Mean	0.27	0.50	0.67	0.17	0.06
Core Non Durable Expenditure	1.00	0.69	0.33	-0.10	-0.09
Non Durable Expenditure w/Housing		1.00	0.42	-0.03	-0.05

Note: See text for definition of each category and the appendix for data sources. Total/Weighted Mean refers to sum for shares and the weighted mean for the other columns, using the share of nondurable expenditure with housing as weights. For each category, we regress log expenditures in 2000 dollars on four-year cohort dummies, marital status and family size dummies, as well as one year age dummies. The log changes are the difference in the coefficients on the age dummies at the respective ages in each column. For each age, we average over three years centered on the age indicated in the column head (e.g., the change between 45 and 60 is the average of 59-61 minus the average of 44-46). The exceptions are age 25, which represents the average of 25 through 27, and age 68, which is the average of 66 through 68.

Table 2: Dispersion of Expenditure over the Lifecycle by Consumption Category

<i>Disaggregated Consumption Group</i>	<i>Cross-sectional Standard Deviation of log Expenditure at 25</i>	<i>Change Between 25 and 45</i>	<i>Change Between 45 and 68</i>	<i>Change Between 25 and 68</i>
Other Non-Durable	1.00	0.22	0.35	0.56
Alcohol and Tobacco	1.32	0.21	0.34	0.56
Clothing and Personal Care	0.78	0.07	0.32	0.39
Charitable Giving and Gambling	1.71	0.00	0.11	0.11
Transportation	0.83	-0.10	0.12	0.02
Food	0.44	-0.02	0.05	0.03
Housing Services	0.61	-0.06	-0.11	-0.17
Domestic Services	1.46	-0.21	0.01	-0.20
Entertainment	1.12	-0.12	-0.13	-0.25
Utilities	0.88	-0.50	-0.13	-0.63
Weighted Mean	0.75	-0.07	0.03	-0.04
Core Non Durable Expenditure	0.39	0.07	0.10	0.17
Non Durable Expenditure w/Housing	0.38	0.07	0.09	0.16

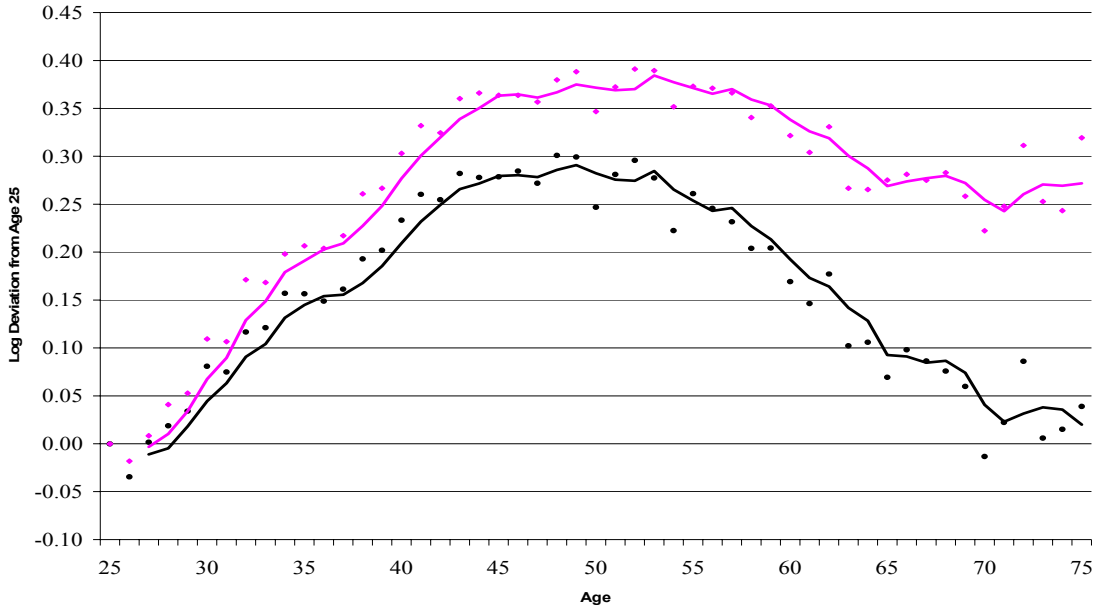
Note: See text for definition of each category and the appendix for data sources. Weighted mean uses the share of nondurable expenditure with housing from Table 1 as weights. For each category, we regress log expenditures in 2000 dollars on four-year cohort dummies, marital status and family size dummies, as well as one year age dummies. At each age and for each cohort, we compute the standard deviation of the residuals from this regression. The first column is the average standard deviation for age 25, pooling all cohorts. For the remaining columns, we regress the standard deviation on age and cohort dummies, and use the age dummies to report the lifecycle profile. For each age, we average over three years, with 25 representing ages 25-27, 45 representing ages 44-46, and 68 representing ages 66-68.

Table 3: Predictions of the Beckerian Model

<i>Category</i>	<i>Percent Change Between 25 and 34</i>	<i>Percent Change Between 34 and 43</i>	<i>Percent Change Between 43 and 52</i>	<i>Percent Change Between 52 and 60</i>	<i>Percent Change Between 60 and 68</i>
Expenditures on Food	0.17	0.12	-0.01	-0.10	-0.08
Time Spent on Food Production	0.13	0.13	0.05	0.13	0.08
Predicted Opportunity Cost of Time	0.03	-0.01	-0.04	-0.15	-0.11
Time Spent on Entertainment	-0.02	0.00	0.04	0.11	0.20
Predicted Expenditures on Entertainment	0.00	0.00	0.01	-0.01	0.12
Actual Expenditures on Entertainment	0.42	0.32	0.01	-0.02	0.12

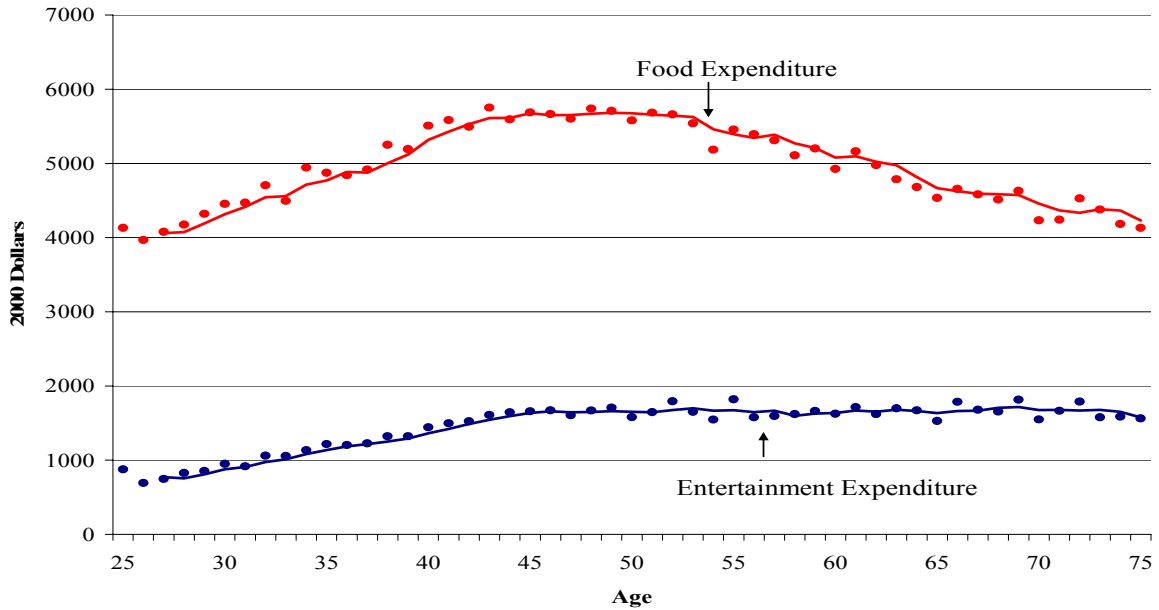
Note: See text for definition of each category and the appendix for data sources. For each category (food and entertainment), we regress log expenditures in 2000 dollars on four-year cohort dummies, marital status and family size dummies, as well as one year age dummies. The log changes are the difference in the coefficients on the age dummies at the respective ages in each column. For each age, we average over three years centered on the age indicated in the column head (e.g., the change between 43 and 52 is the average of 51-53 minus the average of 42-44). The exceptions are age 25, which represents the average of 25 through 27, and age 68, which is the average of 66 through 68. The first row reports the empirical lifecycle expenditure on food, while the last row reports the empirical expenditure on entertainment. For time allocation, we regress average time spent on each category on age dummies plus family size controls. We average the time allocated for each three-year end point, take logs, and then difference to compute the number reported in the table. The predicted opportunity cost of time is computed as the log change in food expenditure minus the log change in time allocated to food, this quantity divided by the assumed elasticity of substitution in food production of 1.5. The calibrated elasticity of production of entertainment is computed to match the predicted change in expenditures between age 60 and 68 to the observed change in expenditures between age 60 and 68. This elasticity is 0.9. The remaining rows of predicted expenditure on entertainment are computed by multiplying this elasticity times the predicted cost of time and then subtracting the change in time allocation.

Figure 1: Core and Total Nondurable Expenditures over the Lifecycle
Log Deviation from 25 Year Olds



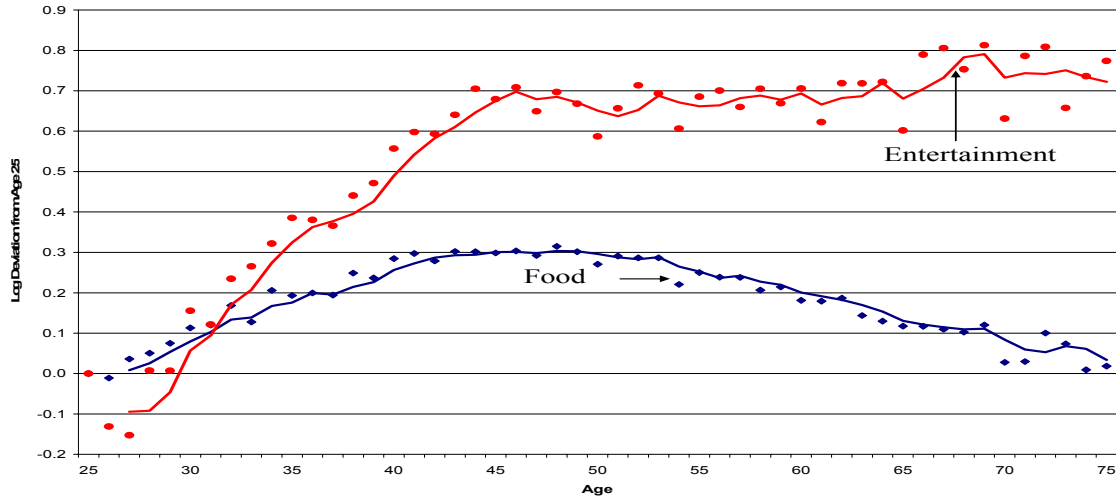
Notes: This figure plots mean expenditure by age conditional on cohort and family status. More specifically, each point represents the coefficient on the corresponding age dummy from the estimation of equation 1, with age 25 the omitted group. The solid lines depict three year moving averages. The bottom series (squares) is core nondurable expenditures and the top series (diamonds) is core nondurables plus housing services. See text for definitions of core nondurables and housing services.

Figure 2a: Food and Entertainment Spending over the Lifecycle



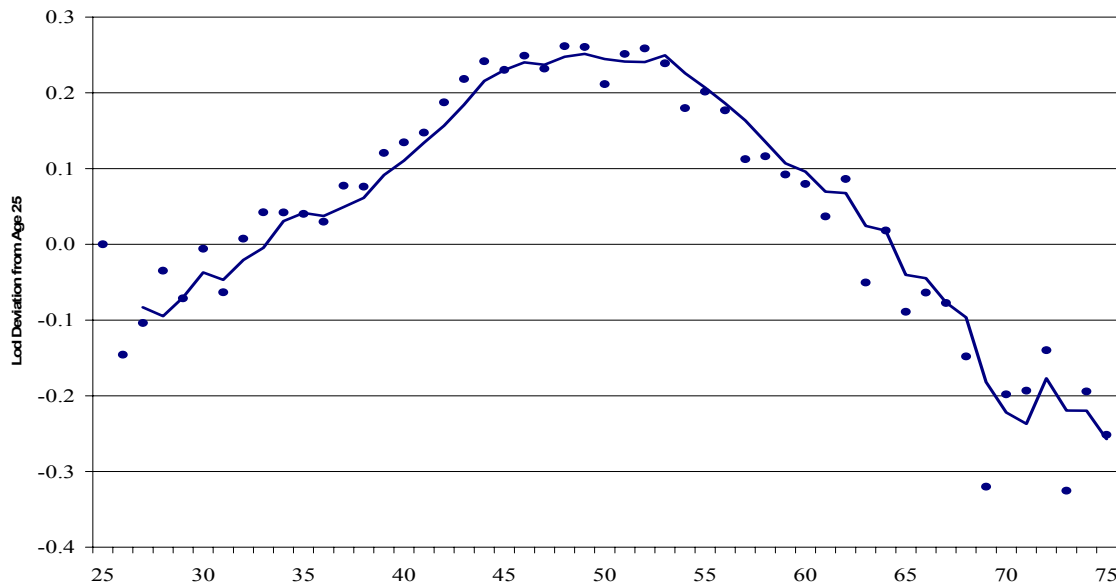
Notes: This figure plots mean expenditure on food and nondurable entertainment by age conditional on cohort and family status. More specifically, each point represents the coefficient on the corresponding age dummy from the estimation (in levels deflated to year 2000 dollars) of equation 1, with age 25 the omitted group. Each series is anchored at age 25 using mean expenditures of 25 year olds in the sample.

Figure 2b: Entertainment vs Food over the Lifecycle
Log Deviation from 25 Year Olds



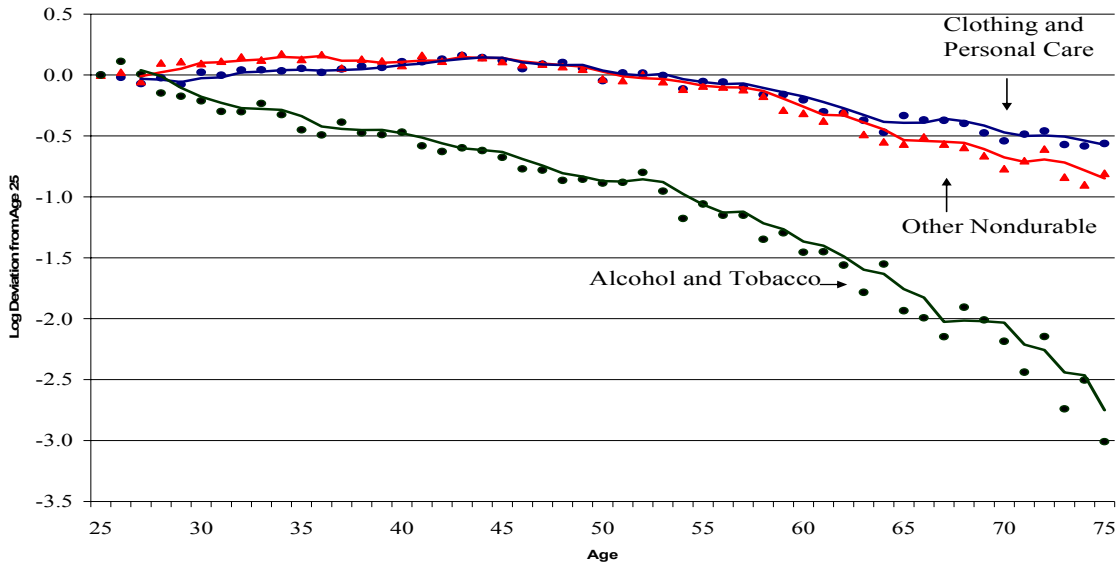
Notes: This figure plots mean expenditure on food and nondurable entertainment by age conditional on cohort and family status. More specifically, each point represents the coefficient on the corresponding age dummy from the estimation of equation 1, with age 25 the omitted group. The solid lines depict three year moving averages. The diamonds (blue) depict food expenditures, and the squares (red) depict nondurable entertainment expenditures.

Figure 3a: Nondurable Transportation Expenditures over the Lifecycle
Log Deviation from 25 Year Olds



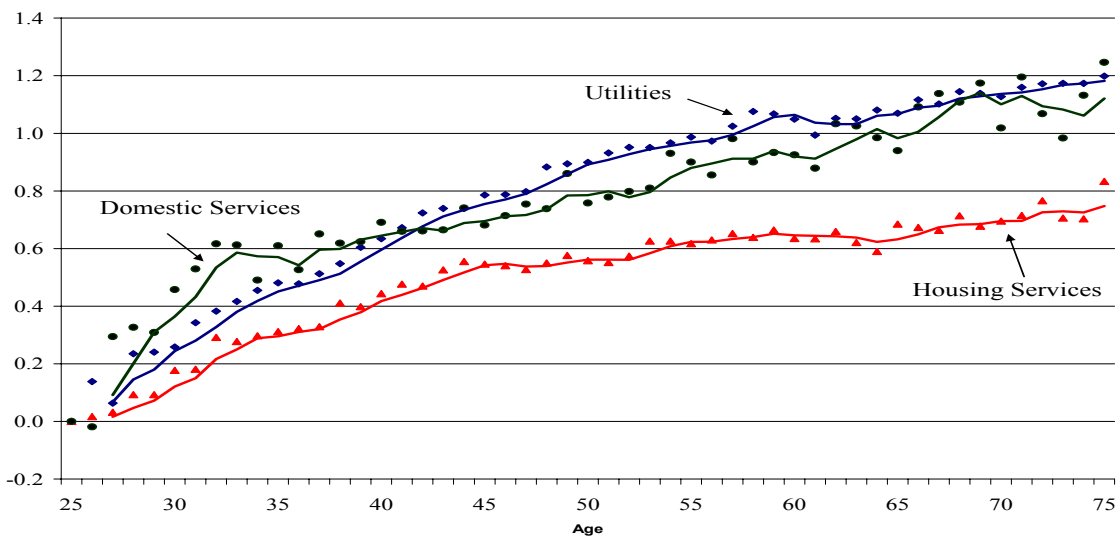
Notes: This figure plots mean expenditure on nondurable transportation by age conditional on cohort and family status. More specifically, each point represents the coefficient on the corresponding age dummy from the estimation of equation 1, with age 25 the omitted group. The solid lines depict three year moving averages.

Figure 3b: Declining Categories over the Lifecycle
Log Deviation from 25 Year Olds



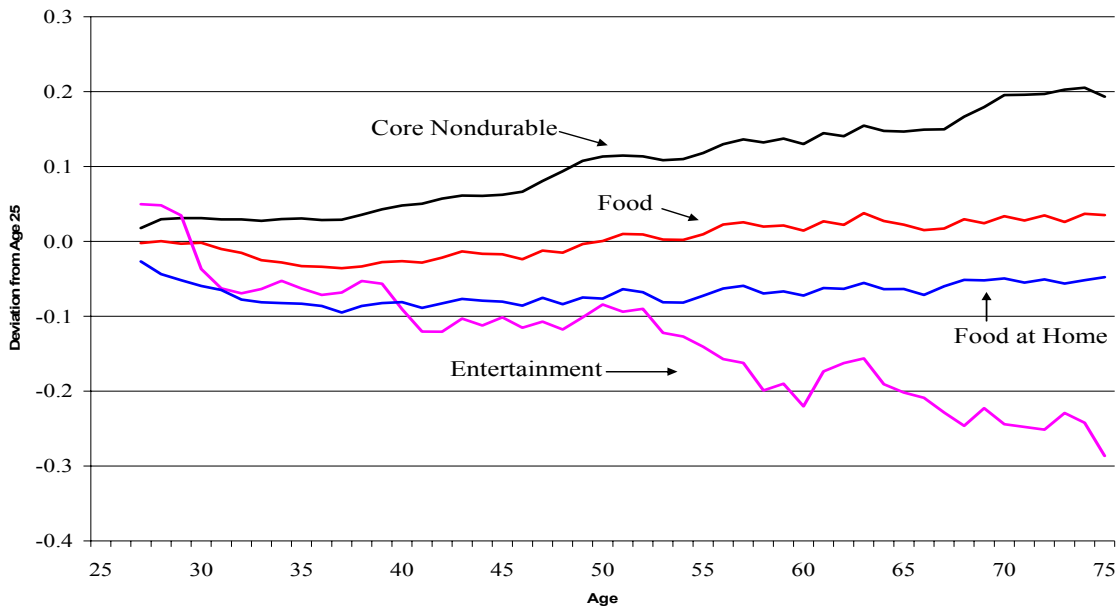
Notes: This figure plots mean expenditure on clothing and personal care, alcohol and tobacco, and other nondurable expenditures. More specifically, each point represents the coefficient on the corresponding age dummy from the estimation of equation 1, with age 25 the omitted group. The solid lines depict three year moving averages. The circles (blue) depict clothing and personal care, the triangles (red) depict other nondurables, and the diamonds (black) depict alcohol and tobacco. See text for definition of categories.

Figure 3c: Increasing Categories over the Lifecycle
Log Deviation from 25 Year Olds

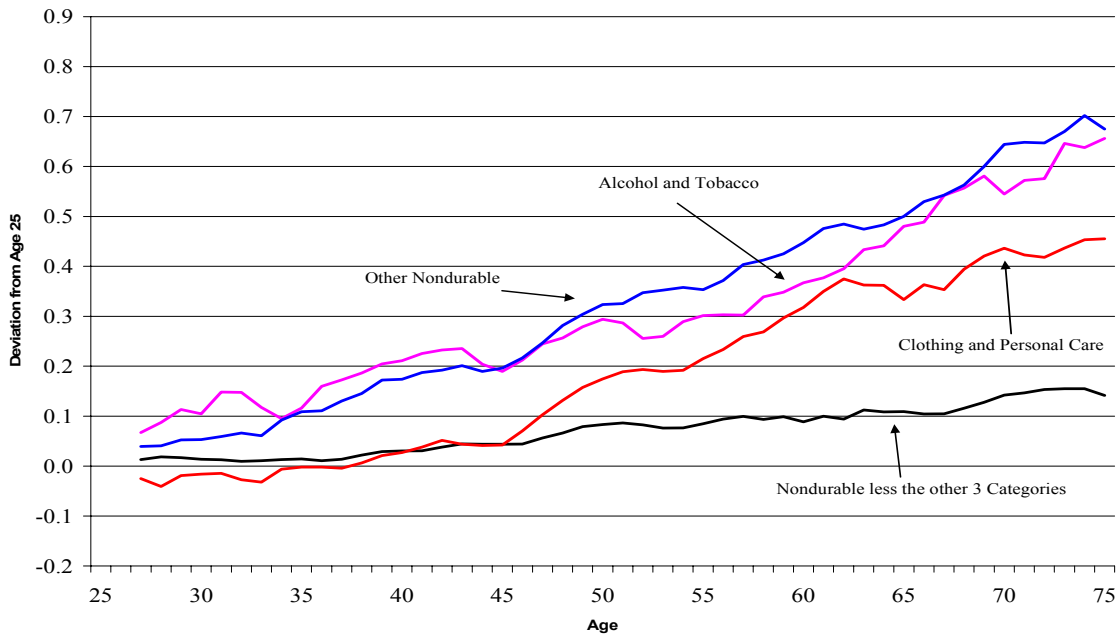


Notes: This figure plots mean expenditure on domestic services, utilities, and housing services. More specifically, each point represents the coefficient on the corresponding age dummy from the estimation of equation 1, with age 25 the omitted group. The solid lines depict three year moving averages. The squares (black) depict domestic services, the diamonds (blue) depict utilities, and the triangles (red) depict housing services (rent or rental equivalence). See text for definition of categories.

Figure 4: Dispersion in Log Expenditure over the Lifecycle
 Panel A: Standard Deviation of Core Nondurable, Food, and Entertainment

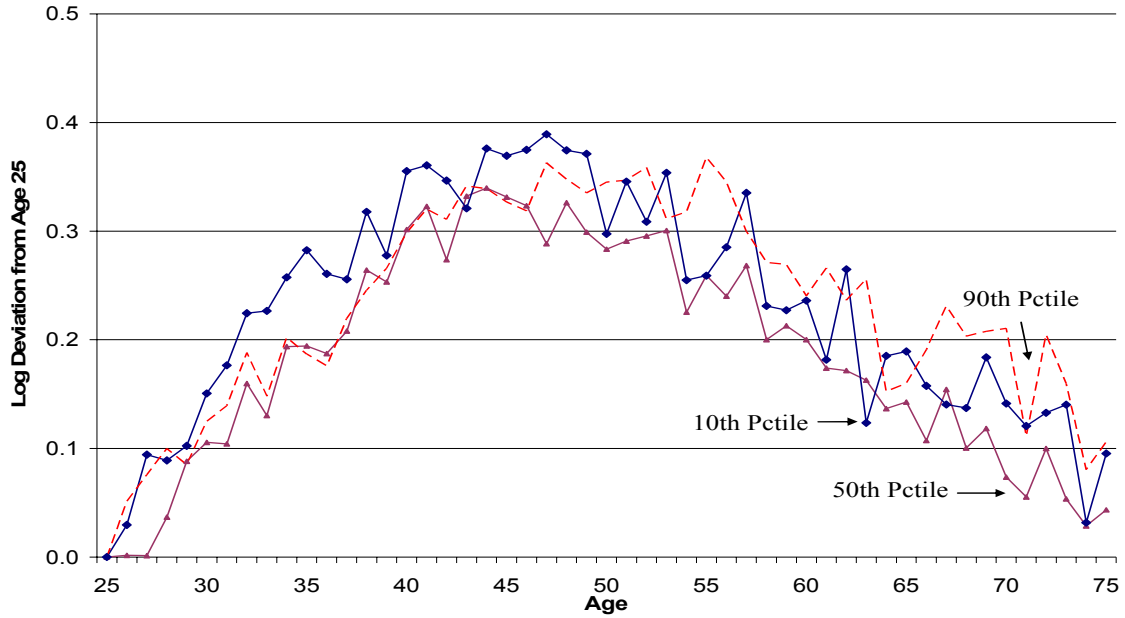


Panel B: Standard Deviation of Alcohol and Tobacco, Clothing and Personal Care, Other Nondurables, and Core Nondurables minus these Categories

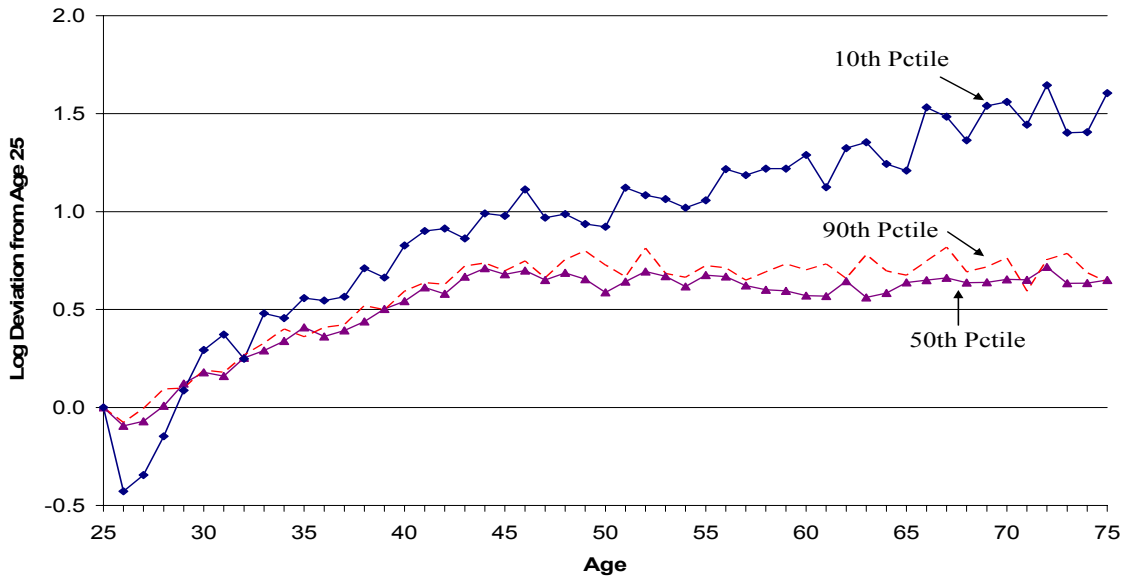


Notes: Panel of A of this figure depicts the standard deviation of log expenditure on core nondurables, food, and entertainment, conditional on cohort and family status. Specifically, we compute the standard deviation of the residuals at each age and cohort from the regression of log expenditures on age, cohort, and family status dummies (equation 1), and then remove cohort fixed effects from the age-specific standard deviations. The figure plots the 3-year moving average of the difference between the standard deviation at each respective age and the standard deviation at age 25. Panel B replicates panel A for alcohol and tobacco, clothing and personal care items, other nondurables, and core nondurables minus these latter three categories.

Figure 5: Lifecycle Evolution of Expenditure Distribution
 Panel A: Food Expenditure

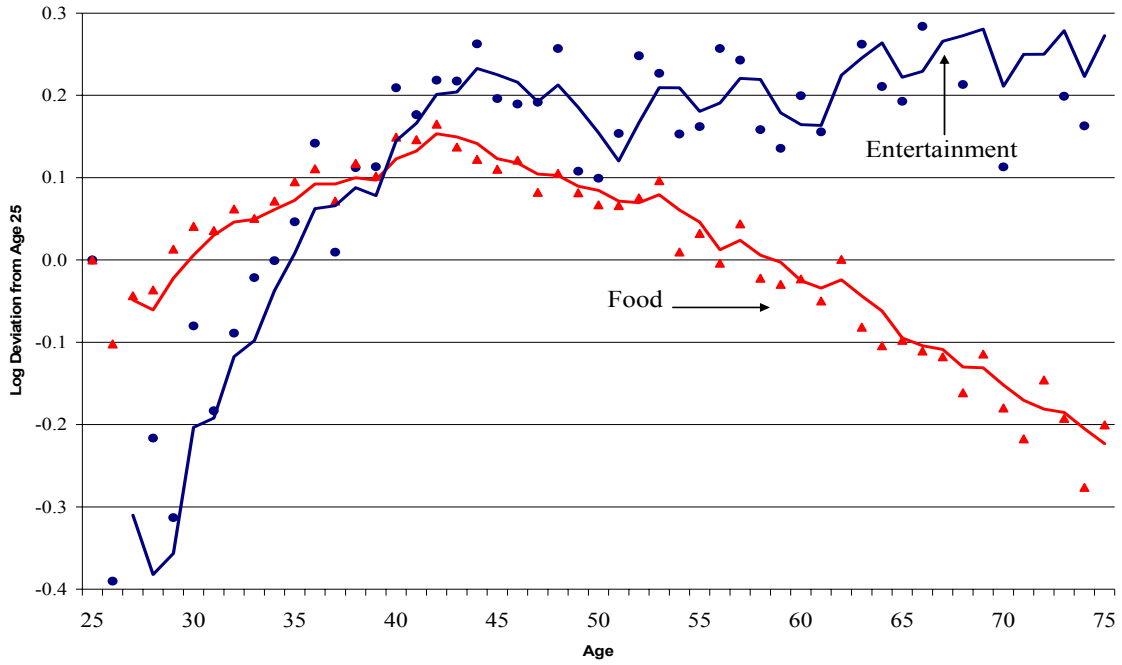


Panel B: Entertainment Expenditure

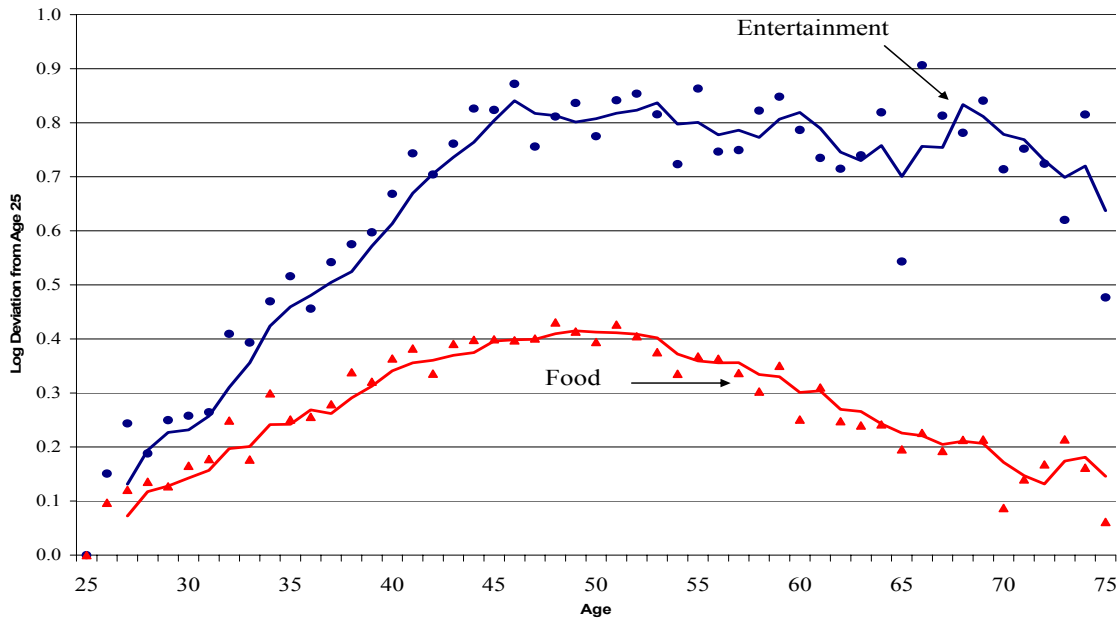


Notes: This figure depicts the coefficient on age dummies for quantile regressions of log expenditures on age, cohort, and family status dummies, with age 25 the omitted age group. The quantile regressions were performed for the 10th, 50th, and 90th percentiles for food (panel A) and entertainment (panel B) expenditures. For reference, the 10th, 50th, and 90th percentiles at age 25 (pooling all cohorts) are 7.5, 8.2, and 8.8 for log food expenditure and 4.9, 6.5, and 7.5 for log entertainment expenditure.

Figure 6: Entertainment vs Food by Educational Attainment
 Panel A: Less Educated Household Heads

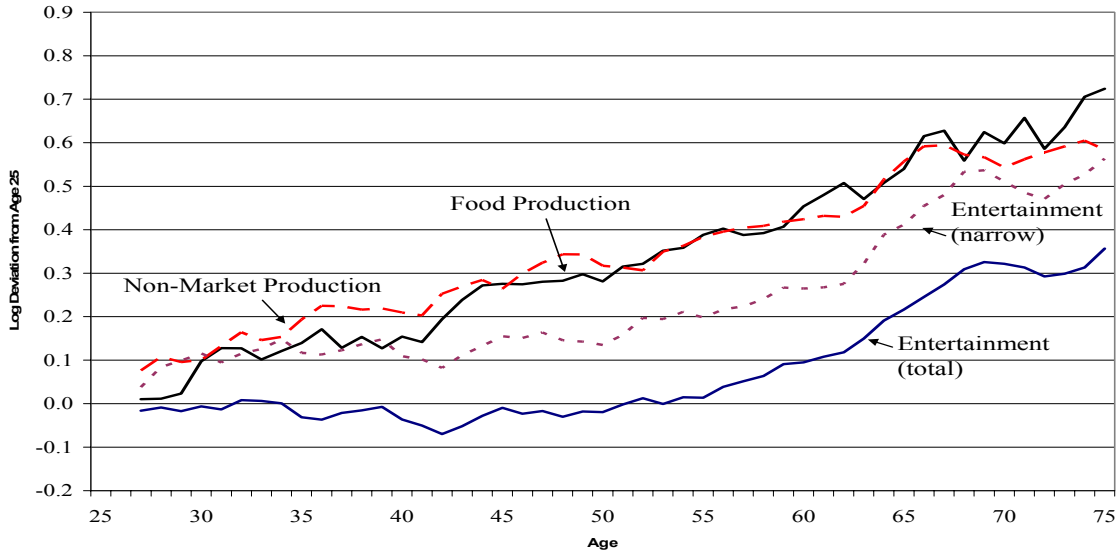


Panel B: More Educated Household Heads



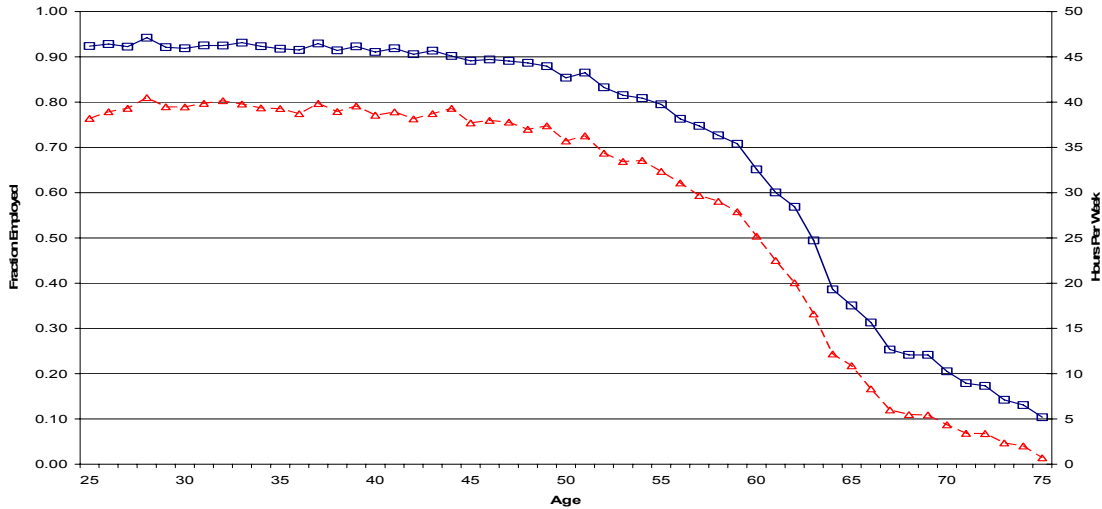
Notes: Panels A and B of Figure 6 replicate figure 2b for household heads with a high school education or less (panel A) and for household heads with more than a high school education (panel B). The triangles (red) depict food expenditures, and the squares (red) depict nondurable entertainment expenditures.

Figure 7: Time Spent on Non-Market Production, Food Production and Entertainment over the Lifecycle



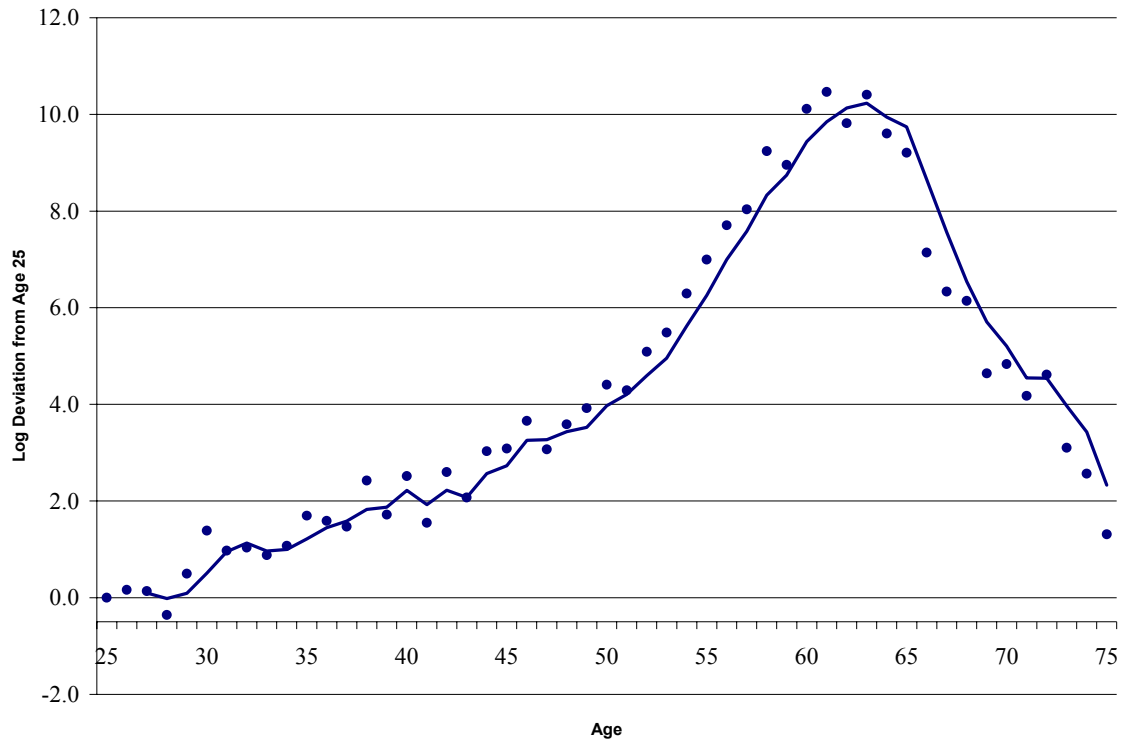
Notes: Source: American Time Use Survey 2003-2005. Food production time consists of time spent preparing and cleaning up after meals; non-market production time consists of all housework (excluding child care), including food production, cleaning house, and home maintenance; entertainment time consists of such activities as watching tv, going to movies, socializing, and sporting events. The figure depicts the percentage change between average time spent on each activity for respondents of the age corresponding to the horizontal axis and the average time spent by 25 year old respondents. In computing averages by age, we control for family size and marital status as in specification (1) excluding cohort dummies given the single cross section. Note that this figure depicts the log difference of the averages (which include respondents who reported zero time on the activity), and not the difference in average log time.

Figure A1: Fraction Employed and Hours Worked over the Lifecycle Deviation from Age 25



Note: This figure depicts the fraction of our CEX sample who are employed (squares and solid line, left axis) and the average number of market work hours per week performed by household heads (triangles and dashed line, right axis), both in deviations from the average for heads aged 25 years. Hours per week is computed unconditional on employment (i.e., includes zeros), but conditional on ...

Figure A2: Standard Deviation of Work Hours over the Lifecycle
Log Difference from Age 25



Notes: This figure depicts the cross-sectional standard deviation at each age of market work hours in the CEX sample, expressed as log deviation from the standard deviation of work hours for 25 year old household heads. Standard deviations are unconditional on employment status, family size, or cohorts.