CHICAGO JOURNALS

Deconstructing Life Cycle Expenditure

Mark Aguiar

Princeton University

Erik Hurst

University of Chicago

We revisit two well-known facts regarding life cycle expenditures: the "hump"-shaped profile of nondurable expenditures and the increase in cross-household consumption inequality. We document that the behavior of total nondurables masks surprising heterogeneity in the life cycle profile of individual consumption subcomponents. We provide evidence that the categories driving life cycle consumption either are inputs into market work or are amenable to home production. Using a quantitative model, we document that the disaggregated life cycle consumption profiles imply a level of uninsurable permanent income risk that is substantially lower than that implied by a model using a composite consumption good.

I.   Introduction

This paper reconsiders two prominent features of life cycle consumption expenditures. The first is the fact that expenditures are "hump" shaped

over the life cycle, peaking in middle age and then declining thereafter.[1] The second fact is that cross-sectional consumption inequality increases as individuals age (see Deaton and Paxson 1994; Attanasio and Jappelli 2000; Storesletten, Telmer, and Yaron 2004; Heathcote, Storesletten, and Violante 2005; Guvenen 2007). These patterns are depicted in figure 1, the details of which are discussed in Section IV. Both facts have had tremendous influence on economists' inferences about household preferences, the income process that households face, and the extent to which public and private insurance markets limit household exposure to risk.

   In this paper we revisit these two familiar facts by disaggregating nondurable expenditures into more detailed consumption categories. We show that there is substantial heterogeneity across consumption goods with respect to both the life cycle profile of mean expenditures and the evolution of the cross-household variance in expenditures. Specifically, we first replicate the standard finding that, after controlling for family composition, composite nondurable expenditures (excluding housing services) peak in middle age at a level roughly 25 percent higher than expenditures at 25 or 65. Similarly, we document that the cross-sectional variance in log nondurable expenditure doubles between ages 25 and 75. We then show that there is substantial heterogeneity in these patterns across different consumption categories. In particular, we document that the decline in nondurable expenditure after middle age is essentially driven by three categories: food, nondurable transportation, and clothing/personal care.[2] Moreover, these three categories account for a substantial portion of the increase in the cross-sectional variance of expenditures over the life cycle. All the other components of our composite nondurable measure (housing services, utilities, entertainment, domestic services, charitable giving, etc.) show no decline in expenditures after the age of 45 and exhibit

   [1] This literature documenting the hump-shaped profile of expenditures is large and extends back nearly 40 years. See, e.g., Thurow (1969), Heckman (1974), Carroll and Summers (1991), Attanasio and Weber (1995), Attanasio et al. (1999), Angeletos et al. (2001), Gourinchas and Parker (2002), and Fernandez-Villaverde and Krueger (2006). The hump shape holds for nondurable expenditures as well as total expenditures.
   [2] These three categories represent roughly 60 percent of nondurable expenditures excluding housing services and roughly 40 percent of nondurable expenditures including housing services.
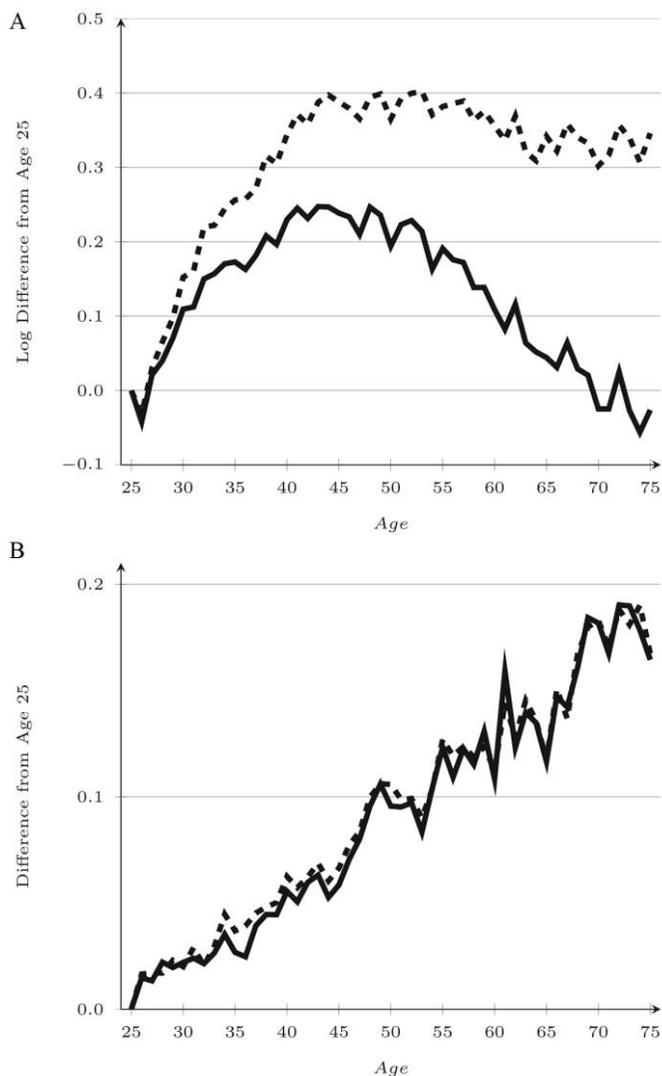
Fig. 1.—Life cycle profiles of nondurable expenditures. Panel *A* plots mean log expenditure by age conditional on cohort, normalized year, and family status controls. Each point represents the coefficient on the corresponding age dummy from the estimation of equation (4), with age 25 being the omitted group. Panel *B* plots the life cycle profile of the cross-sectional variance of log expenditure, conditional on cohort, year, and family composition controls. Specifically, we compute the cross-sectional variance of the residuals from the first-stage regression (eq. [4]) for each age-cohort pair and then remove cohort fixed effects to isolate the life cycle profile of cross-sectional variance (eq. [5]). Again, all deviations are from age 25. The solid (dashed) line represents total nondurable expenditures without (with) housing services. The sample size is 53,412 households covering the 1980–2003 waves of the CEX. See Appendix A for details on sample construction. All data are weighted to be nationally representative using the CEX core weights. See the text for definitions of nondurable and housing service expenditures.

little, if any, increase in cross-sectional variance over the life cycle between the ages of 45 and 65.

Canonical models of consumption emphasize movements in uninsurable permanent income as key to both the hump shape and the increase in cross-sectional dispersion. Models based solely on fluctuations in financial resources to explain the profiles predict that categories with larger income elasticities should display greater increases in cross-sectional dispersion and more pronounced hump shapes. However, the disaggregated data show no such pattern. For example, households increase spending on relative luxuries such as entertainment and charitable giving after middle age while they simultaneously reduce spending on food, clothing, and transportation. Similarly, the cross-sectional dispersion in the former categories all show declines over the life cycle. As a result, standard explanations for the life cycle expenditure profiles based on insurable income risk are not easily reconcilable with the disaggregated expenditure data.

The data do, however, support a prominent role for expenses that are closely linked to a household's opportunity cost of time. These categories consist of clothing and transportation, which can be categorized as inputs into market labor supply, as well as food away from home, which is amenable to home production. As the opportunity cost of time falls over the life cycle and households reduce their attachment to the labor force, expenditures on such "work-related" categories should fall even if there is no change in lifetime resources or preferences. As we show, such work-related expenses account for the entire decline in nondurable expenditures after middle age, coincident with the peak in market labor supply for the average household. Moreover, while inequality in composite nondurables increases throughout the life cycle by roughly 18 percentage points between ages 25 and 75, inequality in nondurable expenditure excluding food and work-related expenses increases by only 8 percentage points, with nearly all of the increase occurring prior to the age of 46 or after the age of 65.

To gain more insight into the importance of clothing, nondurable transportation, and food away from home as being work related, we perform a number of additional exercises. First, we document that the decline in expenditure on food away from home after middle age is associated with a decline in the frequency with which individuals patronize fast-food establishments or cafeterias, with no indication that individuals reduce their visits to restaurants with table service. This fact is consistent with the hypothesis that life cycle variation in expenditures on food away from home is driven by work-related meals. Second, we analyze time diaries and show that there is a large decline in time spent commuting to work after the age of 50. However, time spent on non-work-related traveling increases slightly over the second half of the life cycle. To the ex-

tent that transportation expenditures are proportional to transportation time, these results imply that the decline in transportation expenses is due entirely to a decline in work-related transportation. Finally, we estimate demand systems and document that controlling for labor supply eliminates nearly all the post-middle-age relative decline in spending on clothing and food away from home and much of the decline in transportation.

The patterns documented in this paper argue for a reassessment of the mapping of consumption to uninsurable permanent income. In particular, the differential patterns of "core" nondurable expenditures (which we define as nondurable expenditures excluding work-related expenses and food) and "home production" expenditures (work-related expenses and food) suggest that cross-household consumption inequality increases much less than suggested by total nondurables. In the final part of the paper, we quantify this claim by extending a standard incomplete markets life cycle model to include two consumption goods, one of which enters nonseparably with time. Using consumption data, we calibrate this two-good model to match the life cycle profiles of the first and second moments of total nondurable expenditure as well as for disaggregated subcomponents. For contrast, we also calibrate a canonical one-good, separable model using only total nondurable expenditures. We find that the uninsurable risk at the 20-year horizon is overstated by 25 percent when we ignore heterogeneity across consumption categories. This suggests that households face less uninsurable income risk—particularly during middle age—than suggested by the use of total consumption expenditures to discipline the model. Moreover, the implied long-run income risk from the two-good model is marginally below that estimated directly from wage data, while that of the one-good model exaggerates the role of persistent income shocks. In this sense, this paper complements recent studies that conclude that the canonical consumption models have overestimated the extent of uninsurable income risk later in the life cycle.[3]

This paper is organized as follows. Section II lays out a simple Beckerian framework that emphasizes the importance of consumption goods that are produced using both market expenditures and individual time to motivate our empirical work. Section III discusses the data set and empirical methodology we use. Section IV shows the descriptive results for the life cycle profiles of our disaggregated consumption categories. Section V shows further results highlighting that the life cycle patterns are, in fact, driven by individual home production or changes in work-related expenditures. Section VI introduces and calibrates a fully specified ver-

---

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

sion of the Beckerian model and discusses key implications for inference regarding uninsurable income risk. Section VII presents conclusions. Appendices A and B contain additional empirical results and details on the solution and estimation of the quantitative model.

II. Conceptual Framework

The predominant approach to studying life cycle consumption is to aggregate expenditure on different goods to construct a single index of consumption, with perhaps some distinction between durable and nondurable goods.[4] Given this, there are many papers that have attempted to explain the life cycle profile of mean total nondurable expenditure with rule-of-thumb behavior (Carroll and Summers 1991), imperfect household planning (Bernheim, Skinner, and Weinberg 2001), time-inconsistent preferences (Angeletos et al. 2001), precautionary savings coupled with impatience (Gourinchas and Parker 2002), and nonseparable preferences in utility between consumption and leisure (Heckman 1974). However, the use of a composite expenditure measure (such as total nondurable expenditures) makes it difficult to differentiate among the various stories that explain the profile of expenditure over the life cycle. In this section we discuss how using disaggregated expenditure data facilitates testing across such consumption theories.

As famously studied by Hicks (1939), the validity of using a "composite" consumption good relies on the assumption that relative prices across disaggregated consumption goods are stable (or an equivalent set of assumptions, as discussed in Deaton and Muellbauer [1980]). In the standard life cycle context, this implies that individuals at the same point in time—but at different points in their life cycle—face the same prices for each of the disaggregated consumption goods. One of the motivations for taking a close look at disaggregated data is that in a Beckerian model of consumption (Becker 1965), the relative prices across different consumption goods will not be stable over the life cycle, even if we control for market prices of purchased commodities. This follows from the fact that in the Beckerian model the true cost of consumption includes the value of time used to produce the good, which varies (idiosyncratically) over the life cycle. To set ideas, we now introduce a simple Beckerian framework so as to (i) illustrate that the total cost of different consumption goods should evolve differentially over the life cycle on the basis of the elasticity between time and expenditures in the production of that consumption good and (ii) compare the Beckerian model to standard models of life

[4] There are many demand system analyses that exploit disaggregated expenditure data. For example, such studies have used micro data to estimate key preference parameters or test implications of consumer optimization. To the best of our knowledge, ours is the first study to directly focus on the disaggregated expenditure behavior behind figs. 1*A* and 1*B*.

cycle expenditures, which assume that nondurable consumption goods differ only by their income elasticities.[5]

Assume that agents have time-separable, strictly concave utility over N consumption commodities, $c^1, c^2, \ldots, c^N$, defined as $u(c^1, c^2, \ldots, c^N)$. Each commodity in turn represents the combination of market expenditures, $x^1, x^2, \ldots, x^N$, and time inputs, $h^1, h^2, \ldots, h^N$, using technologies $c^n = f^n(x^n, h^n)$. For simplicity, we assume that the commodity production functions are constant returns to scale. 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. The price to the consumer of a unit of c^n is a function of the market price of x^n as well as the agent's opportunity cost of time. Agents maximize the present value of expected utility subject to a lifetime budget constraint.

At this point, there is no need to take a strong stand on the nature of the income process or asset markets that agents face, but we will do so in the fully specified model of Section VI. As motivation, we can focus on the static optimization in any one period conditional on the agent's total within-period expenditure X and available nonmarket time H:

$$\max_{x^n, h^n, c^n} u(c^1, \ldots, c^N)$$

subject to

$$\sum_n p^n x^n \leq X,$$

$$\sum_n h^n \leq H,$$

where p^n is the market price of input x^n. Let λ be the multiplier on the agent's budget constraint and let $w\lambda$ be the multiplier on the agent's within-period time constraint, using the fact that $\lambda > 0$ under standard assumptions. (While we hold labor fixed in discussing this part of the budgeting problem, if labor supply for the agent is interior, w will be pinned down by the agent's wage.) The first-order conditions for optimization imply

$$u_n f_1^n = \lambda p^n, \tag{1}$$

$$u_n f_2^n = \lambda w, \tag{2}$$

[5] The difference in income elasticities across goods is related to differences in the intertemporal elasticity of substitution across the goods in a world where the goods are separable in utility (see, e.g., Browning and Crossley 2000).

where $f_1^n = \partial f^n / \partial x^n$, $f_2^n = \partial f^n / \partial h^n$, and $u_n = \partial u / \partial c^n$.[6] These conditions imply that the consumer equates the technical rate of substitution in production of the consumption commodity to the real opportunity cost of time:

$$\frac{f_2^n}{f_1^n} = \frac{w}{p^n}. \tag{3}$$

The total response of x^n to a change in the agent's opportunity cost of time (w) can be decomposed into three separate effects. The first is a traditional income effect. To illustrate this effect, consider an increase in lifetime resources (a decrease in λ) holding w unchanged. For a fixed w, equation (3) and constant returns to scale imply that any change in c^n will be implemented by increasing x^n and h^n by the same proportion as consumption. That is,

$$\left. \frac{d \ln x^n}{d \ln \lambda} \right|_{dw=0} = \left. \frac{d \ln c^n}{d \ln \lambda} \right|_{dw=0}.$$

The amount by which c^n (and hence x^n) increases depends on the expenditure elasticity of that good. Under additive separability (or homotheticity), we have

$$\left. \frac{d \ln c^n}{d \ln \lambda} \right|_{dw=0} = \frac{u_n}{c^n u_{nn}}.$$

More generally, expenditures on luxury goods will respond more than expenditures on necessities.

In the Beckerian model, the response of x^n to a change in w involves two substitution effects: one between time and market inputs in the production of a fixed c^n and the other concerning the change in c^n across time. To be more concrete, and again assuming separability for transparent expressions, we have

[6] Consumer optimization implies an indirect (flow) utility function, $v(X, w, \{p^n\})$, that takes as arguments total expenditure $X = \sum_n x^n$, the price of time w, and market prices for x^n. Holding market prices constant, we can view this as a nonseparable utility function that takes expenditures and some measure of the price of time (usually market labor) as arguments. Such an approach has been successfully used to explain business cycles (Greenwood, Rogerson, and Wright 1995), female labor force participation (Mincer 1962), and retirement behavior (Aguiar and Hurst 2005), among many other questions. Heckman (1974) has proposed nonseparability between consumption and leisure to explain the hump-shaped consumption profile depicted in fig. 1A. While a reduced-form nonseparability is tractable and appealing, without strong functional form assumptions (or additional data, like disaggregated expenditure) it is difficult to distinguish the nonseparability hypothesis from other explanations of a given empirical pattern like figs. 1A and 1B. By using the Beckerian model—instead of a simple, reduced-form nonseparability across goods in the utility function—we can use the disaggregated data to help distinguish different stories that explain the life cycle profiles of expenditure.

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

where s_h^n denotes the cost share of time in the production of consumption good n, $h^n f_2^n / c^n$. The larger this share, the more relevant are time inputs in producing a unit of c^n. The first substitution effect is driven by the *intratemporal* elasticity of substitution, σ^n. Recall that σ^n measures the extent to which expenditures and time are substitutes or complements in the commodity production function. As σ^n increases, the consumer is more willing to substitute market inputs for time when the opportunity cost of time increases. The second substitution effect is the response of c^n to an increase in the composite price (including the price of time), holding λ constant. An increase in the price of time makes commodities for which time is an important input relatively expensive to consume. In response to this, agents have an incentive to shift consumption to other goods or periods for which the cost is lower. In a life cycle setting, the extent of this substitution is governed by the *intertemporal* elasticity of substitution, $-(u_n / c^n u_{nn})$. Whether market expenditures x^n ultimately increase or decrease with w (holding λ constant) depends on whether the intra- or intertemporal elasticity effect is greater.

With this framework in hand, we return to the life cycle profile of mean composite expenditure. If the composite measure of expenditure declines during the second half of the life cycle, it could be due to (i) agents having a high discount rate, (ii) agents experiencing an uninsured/unanticipated decline in lifetime resources (an increase in λ), (iii) agents being myopic or having time-inconsistent preferences, or (iv) agents experiencing a decline in their opportunity cost of time holding lifetime resources fixed. As noted above, this latter effect would occur only if the intratemporal elasticity of substitution for the composite good is large relative to the intertemporal elasticity of substitution for the composite good. Notice that the use of the composite consumption good obscures the distinction between these stories. However, using disaggregated data can help with such identification.

To see how disaggregated expenditure data can help distinguish among the above different stories, consider two consumption commodities that have different degrees of substitutability between time and market inputs in their production. In particular, let good m depend only on market expenditures $f^m = x^m$, while good n is a home-produced good that is produced with both time and market expenditures. For simplicity, assume that the two commodities enter utility separably, and assume that the intratemporal elasticity of substitution in f^n is greater than the intertemporal elasticity, making time and expenditures easily substitutable for the home-produced good. The fact that time plays a differential role in

the two consumption commodities makes the change in the relative expenditure on the two goods particularly informative about the nature of a shock to wages. Specifically, the income effect of an unanticipated/ uninsured permanent increase in the wage will generate increases in expenditure on both goods, with the magnitude depending on the relative income elasticity. Similar patterns of correlated expenditure changes would result if households were myopic or had time-inconsistent preferences. However, the substitution effect of an insurable change in the wage generates a change in expenditure on x^n and no change in expenditure on x^m. This lowers the correlation of the change in expenditure of the two goods. Therefore, the differences in first and second moments across goods of differing nonseparability with market labor are informative about whether innovations to wages have a strong, uninsurable permanent-income component or are easily smoothed using available asset markets and manifest primarily as changes in the price of time inputs into home production.

  In Section VI, we will formalize these simple insights so that we can revisit estimates of how much uninsurable risk households face. The disaggregated data that we document in the following sections are going to form the basis of our identification strategy. If part of the reason that life cycle expenditure is falling and the cross-sectional variance of expenditure is increasing after middle age is uninsurable permanent income shocks, this should show up for all consumption categories with positive income elasticities. Yet, as we show empirically in the following sections, disaggregated goods behave very differently with respect to their life cycle profiles of mean expenditure and the cross-sectional variance of expenditure. Much of the differences across goods can be explained by the extent to which time and expenditures are substitutable in the production of the ultimate consumption commodity. Using the data on the disaggregated goods allows us to isolate the movements in expenditure that are driven by uninsurable changes in wages (i.e., changes in λ) from the movements in expenditure that are driven by the nonseparabilities introduced through the commodity production functions.

III. Data and Empirical Methodology

To examine the life cycle profile of expenditure and the life cycle evolution of the cross-sectional dispersion, we use data from the Consumer Expenditure Survey (CEX). Specifically, we use the National Bureau of Economic Research CEX extracts, which include all waves from 1980 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 nonzero 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 10 percent of the households. When looking at smaller consumption aggregates in isolation (food away from home, domestic services, alcohol and tobacco, and the residual other nondurables), we bottom-code the expenditure data at one dollar and then take logs. Online Appendix C explores how this assumption affects the results. Finally, we focus our analysis on households in which the head is between the ages of 25 and 75 (inclusive). After we impose these restrictions, our analysis sample contains 53,412 households. When examining the life cycle profile of mean expenditures and cross-sectional dispersion, we limit our analysis to nondurables excluding health and education expenditures. Our measure of nondurables consists of expenditure on food (both home and away), alcohol, tobacco, clothing and personal care, utilities, domestic services, nondurable transportation, airfare, nondurable entertainment, net gambling receipts, business services, and charitable giving.[7] We also examine a broader measure of nondurables that includes housing services, where housing services are calculated as either rent paid (for renters) or the self-reported rental equivalent of the respondent's house (for homeowners). We exclude expenditures on education and health care from the analysis as the utility (or returns) from consuming these goods varies significantly over the life cycle. Likewise, we exclude all durables aside from housing given the difficulty in creating annual service flow measures for these expenditures. Our measure of nondurable expenditure plus housing services constitutes roughly 75 percent of household annual monetary outlays. The remaining portion of annual outlays can be attributed to expenditures on durables such as automobiles, home furnishing, and large entertainment durables (14 percent); health expenditures (5 percent); education expenditures (1 percent); and other expenditures that are difficult to classify (5 percent).[8]

---

[7] Appendix A contains additional details about the construction of the data set and sample selection. Additionally, the appendix provides examples of the types of expenditures that are included in each of the categories.

[8] These other categories include, among others, life insurance premiums, college dormitory fees, money allocated to burial plots, union dues, books, lodging expenses away from home, legal services, etc. Some of these categories were excluded because of the classification system introduced by Ed Harris and John Sabelhaus when creating the NBER CEX files. For example, the category of "books" includes money spent on books for leisure reading and books purchased for course work. Likewise, the category of "other lodging expenditures" includes both college dormitory expenses and vacation rentals. For consistency, we excluded from our analysis any category that included some health or education component. However, in the NBER working paper version of this paper (Aguiar and Hurst 2008), we examined these categories in greater detail. None of our results are changed if we include these measures in our nondurable expenditure measure. This is not surprising given that they constitute only a small fraction of total household expenditures.

A.   Estimating the Life Cycle Profile of Expenditure

When examining life cycle profiles of mean expenditure and cross-sectional dispersion, we adjust all expenditures for cohort and family composition effects. The CEX is a cross-sectional survey, and therefore, age variation within a single wave represents a mixture of life cycle and cohort effects. Moreover, expenditures are measured at the household level and not the individual level. Household size has a hump shape over the life cycle, primarily resulting from children entering and then leaving the household and from changing marriage and death probabilities over the life cycle. We identify life cycle from cohort variation by using the multiple cross sections in our sample and use cross-sectional differences in family composition to identify family composition effects.

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

$$\ln C_{it}^k = \beta_0^k + \beta_{\text{age}}^k \text{Age}_{it} + \beta_c^k \text{Cohort}_{it} + \beta_t^k D_t \\ + \beta_{fl}^k \text{Family}_{it} + \varepsilon_{it}^k, \tag{4}$$

where C_{it}^k is expenditure of household i during year t on consumption category k, Age_{it} is a vector of 50 1-year age dummies (for ages 26–75) referring to the age of the household head, Cohort_{it} is a vector of 1-year birth cohort dummies (1915–68), D_t is a vector of normalized year dummies to be described below, and Family_{it} is a vector of family structure dummies that include a marital status dummy, 10 household size dummies, and controls for both the number and age of household children aged 21 or under.[9] Specifically, we control for the number and age of household children by including dummy variables for the number of children in the following age categories: 0–2, 3–5, 6–13, 14–17, and 18–21. Moreover, for the latter two categories, we create separate indicators for male and female children. Our detailed family composition controls allow us to control flexibly for the potential that children of different ages and sex have different consumption needs or preferences.

As is well known, collinearity prevents the inclusion of a full vector of time dummies in our estimation of (4). In particular, as discussed in Hall (1968), age, year, and cohort effects are identified in repeated cross sections up to a log-linear trend that can be arbitrarily allocated across the three effects. To isolate age profiles, additional assumptions are required. We follow standard practice in the consumption literature (see Deaton 1997) by attributing consumption growth to age and cohort effects and use year dummies to capture cyclical fluctuations. Specifically, we restrict the year effects to (1) average zero over the sample period and (2) be or-

---

[9] For married households, we use the husband's age. See App. A for additional details of how we identify the household head in multiadult households.

thogonal to a time trend. Henceforth, we refer to the year dummies with these restrictions on their coefficients as normalized year dummies.[10]

We also account for changes in the relative price of each consumption category by deflating all categories into constant dollars using the relevant consumer price index (CPI) product-level deflator, if available. Otherwise, we use the relevant personal consumption expenditure deflator from the National Income and Product Accounts (NIPA). All data in the paper are expressed in 2000 dollars. We have also done the analysis using the aggregate CPI-U to deflate all categories and found that our results were robust to this alternative.

The coefficients on the age dummies, β^k_{age}, represent the impact of the life cycle conditional on cohort, normalized year, and family size fixed effects, all of which we allow to vary across expenditure categories. Each of these age coefficients should be interpreted as log deviations from the spending of 25-year-olds. These coefficients are the focus of our analysis as they represent the conditional mean expenditure at each point in the life cycle.

Two additional things should be noted about our estimation procedure. The first pertains to our choice of how to adjust the life cycle profile of expenditures for life cycle changes in family size. There is little consensus within the literature about the appropriate way to adjust for changes in family size. Moreover, the size of the hump in life cycle expenditures is sensitive to the family size controls.[11] One common alternative approach is to adjust for changes in family size over the life cycle by deflating expenditure in year t by a measure of adult equivalence scales in year t, where the equivalence scales are based on the household's family composition in that year.[12] The equivalence scale usually assigns a value of 1 to the first adult household member, a value of either 0.5 or 0.7 to each additional adult member, and a value of 0.3 or 0.5 to each child. Alternatively, the equivalence scale is some mathematical rule such as the square root of family size. We see three limitations to these methods. First, there is little consensus as to the exact value of the equivalence scales. It makes a difference for the life cycle profile of expenditure if each child is worth 0.3, 0.5, or 0.7 of an adult. Second, there is

[10] It should be noted that we estimated (1) with only cohort effects (and no time effects) and with 1-year time dummies (and no cohort effects). The conclusions of the paper are generally robust to either alternative specification. The one exception is housing services. Consumption of housing services has increased over our sample period, and the life cycle profile is sensitive to whether these increases represent cohort or time effects. This point is discussed in detail in App. C.

[11] See Fernandez-Villaverde and Krueger (2006) for a discussion of the various ways in which the literature has controlled for family size when estimating life cycle profiles of expenditures. Fernandez-Villaverde and Krueger also show how the hump in lifetime expenditures is quantitatively sensitive to the choice of family size controls.

[12] A common set of equivalence scales are provided by the Organization for Economic Cooperation and Development (OECD).

likely heterogeneity even within the categories. For example, a teenager almost certainly should be given a higher equivalence weight relative to a toddler. Given that the fraction of teenagers in the household varies over the life cycle, ignoring such heterogeneity will bias the true life cycle variation in expenditure. Finally, and most importantly for our purposes, the equivalence scales should almost certainly differ by good. The returns to scale in entertainment (television subscriptions, digital video discs, etc.) should be different from the returns to scale in clothing. Using a common equivalence scale for all categories would bias the differences in the underlying life cycle patterns across the consumption categories that we want to emphasize.

For this reason, in the main body of the paper we estimate the family size adjustments from the data. Our approach allows us to do this differentially across goods. The main drawback to our approach is that actual family size is not necessary exogenous to permanent income. For example, lower-income individuals are slightly more likely to have more children and are slightly less likely to be married. Differences in family size across households, therefore, will be partially proxying for differences in permanent income across households. Given this, our family size controls could be purging more than just family size from our regressions. We took this concern seriously. In Appendix C, we use the panel dimension of the Panel Study of Income Dynamics (PSID) to see how serious an issue this is for food expenditure. Food expenditure is the only measure of expenditure consistently measured within the PSID. Within the PSID, we can replace our current procedure of identifying the life cycle profile off of repeated cross sections controlling for both cohort and family size effects. We can then use a different procedure to recover the age profiles by exploiting the panel dimension and controlling for individual fixed effects as well as our family size controls. The results of the two procedures were nearly identical, suggesting that the bias introduced in our estimates of the life cycle expenditure profile resulting from the potential correlation between family size and permanent income is likely small.[13]

The second issue we wish to note pertains to the well-documented measurement error within the CEX. Over time, total spending measured by the CEX has fallen as a fraction of total spending measured by the NIPA. Moreover, Bee, Meyer, and Sullivan (2012) have shown that the

---

[13] Given the debate surrounding how to adjust for changing family size over the life cycle when estimating the life cycle profile of expenditure, we perform many additional robustness exercises. Primarily, we have redone all the main empirical analyses within our paper using the OECD equivalence scales to adjust for family size. These results are detailed in App. C. The change in equivalence scales does change the life cycle patterns of the composite consumption good. However, our main point of the paper is still preserved. Even with the OECD equivalence scales, the categories that we highlight—food, clothing, and nondurable transportation—behave very differently over the life cycle, in both mean and cross-sectional variance, than the other consumption categories.

deterioration has differed by consumption category. For example, there has been little deterioration in the ratio of CEX spending to NIPA spending between the mid-1980s and the late 2000s for the following categories: food at home, food away from home, rent and utilities, and cable and satellite television and radio services. However, the ratio of CEX spending to NIPA spending has fallen sharply for clothing, gas and energy expenditures, and child care services. Given that the trends in measurement error have evolved differentially for the different categories, we want to ensure that the patterns we are documenting are not driven by the differential trends in measurement error. We explore this potential issue in Appendix C. Specifically, we examine the robustness of our results so that for each category and in each year, average expenditure in the CEX matches its NIPA counterpart. We then redo all of our estimation on the rescaled data. As we show in Appendix C, the patterns we document in the subsequent sections are robust to such adjustments.

B. *Estimating the Life Cycle Profile of Cross-Sectional Expenditure Dispersion*

To estimate the life cycle profile of the cross-sectional expenditure dispersion, we start by computing $(\sigma^2)^k_{it}$, the variance of ε^k_{it} (the residuals from [4]) for each age and cohort. We then estimate the following equation:

$$(\sigma^2)^k_{it} = \alpha^k_0 + \alpha^k_{\text{age}}\text{Age}_{it} + \alpha^k_{\text{cohort}}\text{Cohort}_{it} + \eta^k_{it}. \tag{5}$$

The vector of age coefficients, α^k_{age}, for each consumption category, k, provides our estimates for the evolution of cross-sectional variance in expenditures over the life cycle. This method is essentially the same as the one used by Deaton and Paxson (1994).

IV. Empirical Patterns

Figures 1*A* and 1*B* plot the coefficients on Age$_{it}$ from equations (4) and (5), respectively. Within each figure, the solid line represents the results using nondurable expenditures without housing services. The dashed line represents the results using nondurable expenditures with housing services. Figure 1*A* replicates the well-documented profile of nondurable expenditures over the life cycle, with nondurable expenditures excluding housing services peaking in middle age at roughly 0.25 log points higher than the level of the 25-year-old expenditure and then declining by nearly 0.30 log points over the latter half of the life cycle.[14] Nondurable expenditures inclusive of housing services rise faster early in the life cycle but

---

[14] The patterns in fig. 1*A* are similar to what others have documented in the literature. As discussed above and in App. C, the extent to which the life cycle profiles differ across papers can be explained in large part by differences in how the papers control for family size.

then do not decline as significantly later in the life cycle. The gap between the two series represents the life cycle behavior in housing services. As discussed below with regard to finer disaggregation of expenditure, housing services behaves like utilities, entertainment, and several other nondurables by displaying no decline after middle age. The fact that housing services is a relatively large share of expenditure indicates that it has a clear influence on the overall trend.

Figure 1B shows the increase over the life cycle of the cross-sectional variance of log nondurable expenditures relative to the variance observed for 25-year-olds. The variance for nondurable expenditures with and without housing expenditures for 25-year-olds is 0.16 and 0.17, respectively. Between the ages of 25 and 75, the cross-sectional variance of nondurable expenditures increases by roughly 0.15 points, regardless of whether or not housing services are included in the measure of nondurable expenditures. These magnitudes are similar to the results reported by Guvenen (2007) and are consistent with the findings of others that the cross-sectional variance of expenditure increases by roughly 100 percent over the life cycle.[15] Additionally, most of the increase comes later in the life cycle (after the age of 40), leading some researchers to conclude that there is a prominent role for permanent income shocks during middle age.

The familiar patterns depicted in figures 1A and 1B mask substantial heterogeneity among less aggregated consumption categories. We begin with the following classification scheme involving three subaggregates: (i) clothing/personal care, food away from home, and nondurable transportation; (ii) food consumed at home; and (iii) all other nondurable expenditure categories including housing services.[16] We refer to the first group as "work-related" expenditures and the last measure as "core" nondurable expenditures. In the next section, we provide the evidence underlying the labeling of clothing, food away from home, and transportation as work-related expenses.

The mean and cross-sectional variances of these categories are depicted in figures 2A and 2B, respectively. Tables 1 and 2 summarize the life cycle profiles for the mean and variance of these three composite consumption goods. Additionally, table 1 shows the fraction of expenditures spent on each of the three categories (relative to total expenditures on the three categories combined) for the average household in our sample at age 25, age 45, and age 65. A few things are of note with respect to the results in figures 2A and 2B. In figure 2A, we see that the different expenditure categories display very different life cycle profiles for mean spending. Food

---

[15] The increase in inequality over the life cycle is somewhat larger than that documented in Heathcote, Perri, and Violante (2010). This again is due to differences in the adjustment for family size.

[16] As discussed below in n. 18, we exclude alcohol and tobacco from the latter measure.

A

B
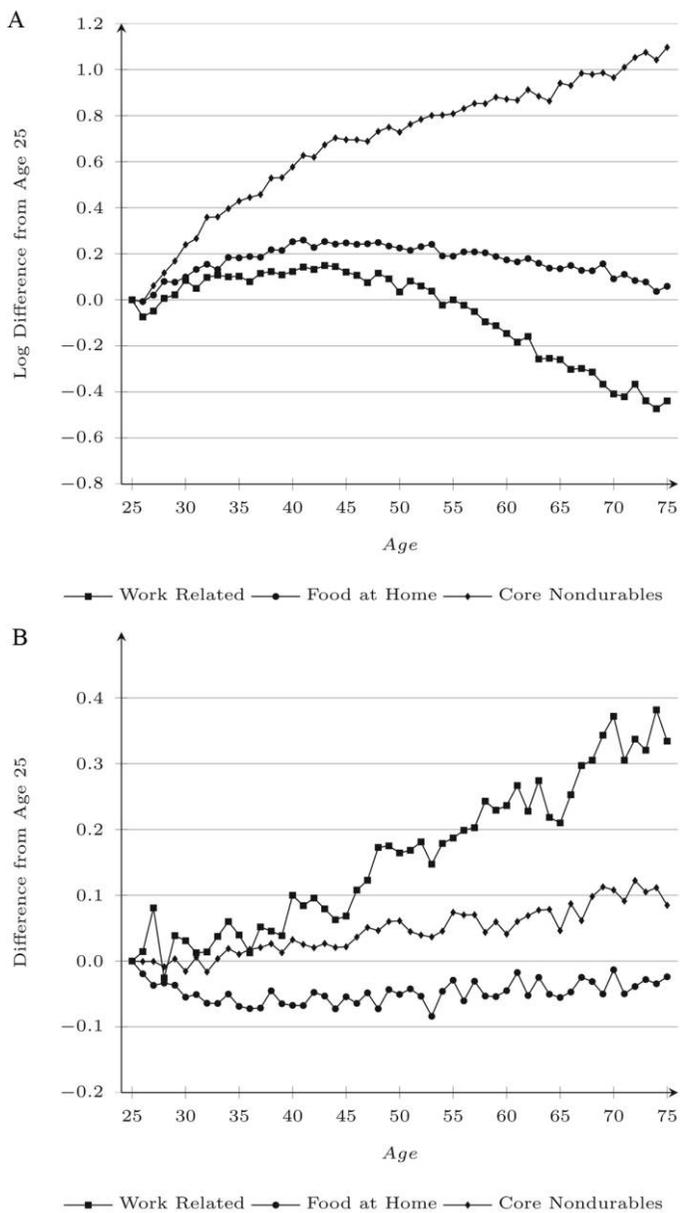
Fig. 2.—Life cycle profiles of three subaggregates. Panels *A* and *B* are identical to panels *A* and *B* of figure 1, respectively, except that we disaggregate nondurable consumption into three categories. The categories are food at home (circles); work-related expenses (squares), which include transportation, food away from home, and clothing/personal care; and core nondurables (diamonds), which include all other categories of total nondurable expenditure (including housing services but excluding alcohol and tobacco). See the caption of figure 1 for additional sample and estimation descriptions.

TABLE 1
SUMMARY OF MEAN CHANGE IN EXPENDITURE OVER THE LIFE CYCLE BY CONSUMPTION CATEGORY

Disaggregated Consumption Group	Share of Expenditures at Ages 25–27 (1)	Share of Expenditures at Ages 43–45 (2)	Share of Expenditures at Ages 64–66 (3)	Log Change in Expenditure between Ages 25 and 45 (4)	p-Value of Change (5)	Log Change in Expenditure between Ages 45 and 65 (6)	p-Value of Change (7)
Core nondurables	.52	.55	.59	.66	<.01	.21	<.01
Food at home	.17	.17	.17	.24	<.01	−.11	<.01
Work-related expenses	.31	.28	.24	.18	<.01	−.40	<.01
Total nondurables with housing (excluding alcohol and tobacco)	1.00	1.00	1.00	.44	<.01	−.01	.47
Total nondurables with housing (including alcohol and tobacco)				.39	<.01	−.07	<.01

NOTE.—This table summarizes the life cycle profiles of the mean expenditure for different broad consumption categories as shown in fig. 2A. In cols. 1–3, we report the share of expenditure at ages 25–27, 43–45, and 64–66, respectively. The remaining columns report the log change in mean expenditure between 25–27 and 43–45 and 64–67, respectively, along with the associated p-values of the test that the respective change is zero. The log change in expenditure between ages 25 and 45 is the coefficient on the 43–45 age dummy from the regression of log expenditure on 3-year age dummies and demographic controls, and the log change between 45 and 65 is the difference in coefficients across the respective age dummies. To smooth out some of the age-to-age variability, we used 3-year age dummies instead of 1-year age dummies.

454

TABLE 2
SUMMARY OF CHANGE IN CROSS-SECTIONAL VARIANCE OVER THE LIFE CYCLE BY CONSUMPTION CATEGORY

Disaggregated Consumption Group	Unconditiional Cross-Sectional Variance at Ages 25–27 (1)	Change in Residual Cross-Sectional Variance between Ages 25 and 45 (2)	p-Value of Change (3)	Change in Residual Cross-Sectional Variance between Ages 45 and 65 (4)	p-Value of Change (5)
Core nondurables	.28	.02	<.01	.05	<.01
Food at home	.34	−.04	<.01	.01	<.01
Work-related expenses	.41	.03	<.01	.16	<.01
Total nondurables with housing (excluding alcohol and tobacco)	.19	.05	<.01	.06	<.01
Total nondurables with housing (including alcohol and tobacco)	.18	.05	<.01	.07	<.01

NOTE.—This table summarizes the life cycle profiles of the cross-sectional variance of expenditure for different broad consumption categories as shown in fig. 2B. In col. 1, we report the pooled cross-sectional variance of expenditure at ages 25–27 for each of the three consumption categories. In cols. 2 and 4, we report the change in cross-sectional variance of expenditure between the ages of 25–27 and 43–45 and then again between 43–45 and 64–66, respectively. For the change in the cross-sectional variance over the life cycle, we performed an analysis similar to the one we used to generate fig. 2B. To smooth out some of the age-to-age variability, we used 3-year age dummies instead of 1-year age dummies. Columns 3 and 5 are the p-values for the changes shown in cols. 2 and 4.

455

at home most resembles the profile of the composite nondurable consumption measure excluding housing services. Food at home rises by roughly 25 log points between the ages of 25 and 45 before declining by roughly 20 log points by age 70. The life cycle patterns for core nondurables and work-related expenses are dramatically different from both the composite measure and each other. Core nondurables increase sharply up through middle age and then continue to increase steadily thereafter. Work-related expenditures, however, fall sharply (by roughly 60 log points) after middle age.

Additionally, figure 2*B* provides a striking reflection of the results pertaining to the life cycle profile of consumption inequality. The cross-sectional variance of core nondurable expenditures displays a life cycle pattern dramatically different from that of the cross-sectional variance of total nondurable expenditures as analyzed by Deaton and Paxson and others and replicated in figure 1*A* above. In particular, up through the age of 65, the cross-sectional dispersion in core nondurables increases by approximately 8 points, with nearly all of the increase coming prior to the age of 45 or after the age of 65. Given that the variance of core nondurables for 25-year-olds is 0.28, the cross-sectional dispersion of core nondurables increases by less than 30 percent over the life cycle. This is less than a third of the proportional increase in cross-sectional variance for total nondurables. The implication is that much of the increase in cross-sectional variance over the life cycle stems from work-related expenses and the associated covariances.[17] The sharp increase in inequality in expenditure on work-related expenses is clear in figure 2*B*. Note in particular that the variance of work-related expenses increases significantly after middle age, while core nondurables show no comparable increase. The cross-sectional variance of total nondurables increases by nearly 10 percentage points between the ages of 45 and 68 (fig. 1*B*), which represents nearly half of the increase in life cycle dispersion of total nondurables. All of the increase in variance between the ages of 50 and 68 in total nondurables is due to an increase in the variance of work-related expenditures (as well as the changing shares of goods over the life cycle and the associated covariances).

In summary, core nondurable expenditure displays a life cycle profile for both the mean and the cross-sectional variance dramatically different from that of the standard composite measure of nondurable expenditure. The results indicate that the prominent features of life cycle

---

[17] As discussed in App. C, the variance of the total can be decomposed into the variance of individual goods, the relative shares in expenditure, and the covariances. The changes in disaggregated variances are reported in fig. 2*B*, and the shares can be inferred from the average shares and the differential trends in mean expenditure. The covariances between the goods are also changing over the life cycle. We discuss the three separate covariances in App. C.

consumption, particularly after middle age, primarily reflect changes in work-related expenditures that move independently of core consumption categories.

Delving a little deeper, we now document that our three-good categorization is a reliable guide to the life cycle behavior of more disaggregated consumption categories. In figures 3 and 4, we plot mean expenditure and the cross-sectional dispersion separately for housing services, utilities, nondurable entertainment, nondurable transportation, food consumed at home, food consumed away from home, domestic services, clothing and personal care, and a residual "other" category. The other-nondurable category includes airfare spending, charitable giving, and net gambling receipts. Figure 4A depicts the goods that do not follow a hump shape but in fact increase steadily over the life cycle, and figure 4B collects those categories that exhibit declines after middle age.[18] Appendix tables A1 and A2 summarize the patterns shown in figures 3 and 4. It should be noted that expenditures in all subcategories displayed in figure 3 increase over the front half of the life cycle. The difference between the two groups of categories occurs after the mid-40s.

Figure 4 and table A2 reveal which categories drive the increasing cross-sectional variance of log expenditures over the life cycle. These categories include transportation, clothing and personal care, food away from home, and domestic services. From figure 4 and the top panel of table A2, we see that at the lower end, the cross-sectional variance of transportation expenditures is essentially flat through age 65 before increasing in the 70s. At the upper end, the variance of domestic service expenditures increases 2.6 log points between 25 and 65. In between, we have the variance of food away from home increasing 1.5 log points and clothing increasing 0.8 log points.

As seen from the disaggregated data, there is substantial heterogeneity across consumption categories with respect to both the life cycle profile of mean expenditures and the life cycle profiles of the cross-sectional variance. Spending on food away from home, clothing and personal care, and transportation drive both the decline in nondurable spending after middle age and the increase in the cross-sectional variance of log nondurable spending over the life cycle. One potential reason why these categories may behave differently over the life cycle is that food is amenable to home production, and clothing and transportation spending

---

[18] One declining category that is not included in either figure is alcohol and tobacco. This category behaves in a manner distinct from the other categories depicted in figs. 2A and 2B. Alcohol and tobacco expenditure falls continuously over the entire life cycle. Moreover, the decline in expenditure is very large: Spending on alcohol and tobacco falls by 1.35 log points between 25 and 45, another 1.69 log points between 45 and 60, and another 1.22 log points between 60 and 68. Even though alcohol and tobacco expenditure constitutes only 5 percent of composite nondurables, its large decline also contributes significantly to the overall decline in nondurable spending after middle age.
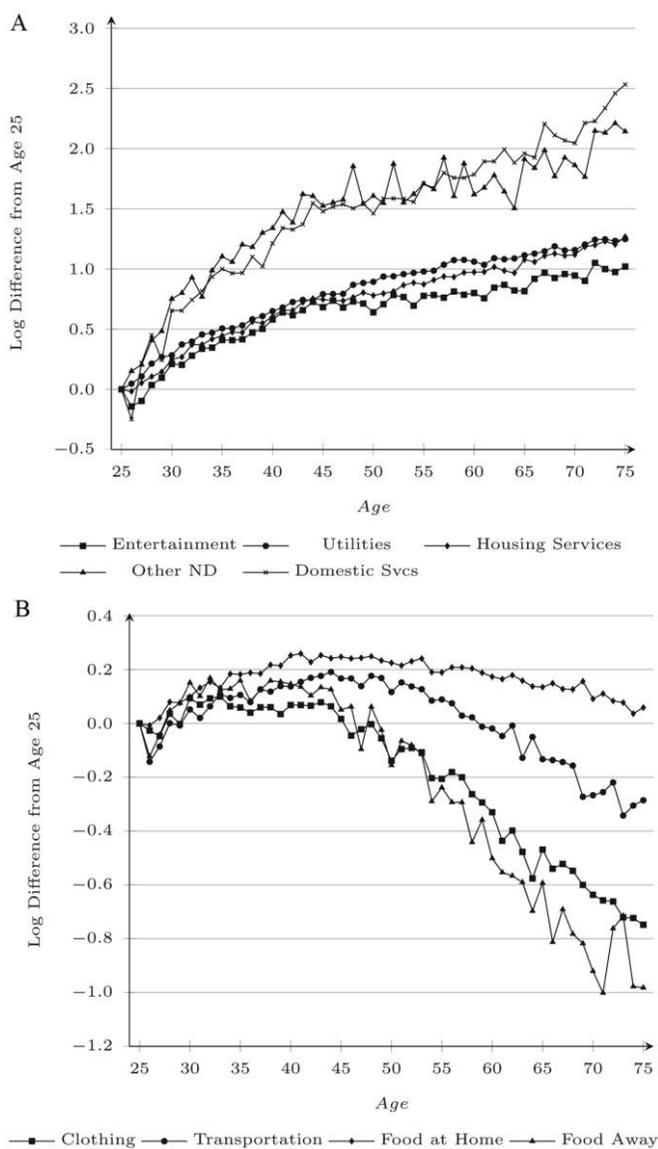
FIG. 3.—Life cycle profiles of disaggregated expenditure: means. This figure plots mean expenditure for disaggregated consumption categories by age conditional on cohort, normalized year, and family status controls. Each point represents the coefficient on the corresponding age dummy from the estimation of equation (4), with age 25 being the omitted group The consumption categories depicted in panel *A* are entertainment (squares), utilities (circles), housing services (diamonds), other nondurables (triangles), and domestic services (x's). The consumption categories depicted in panel *B* are clothing and personal care (squares), transportation (circles), food at home (diamonds), and food away from home (triangles). The sample is the same as for figure 1. See the text and Appendix A for a discussion of the consumption categories.

458

A

B
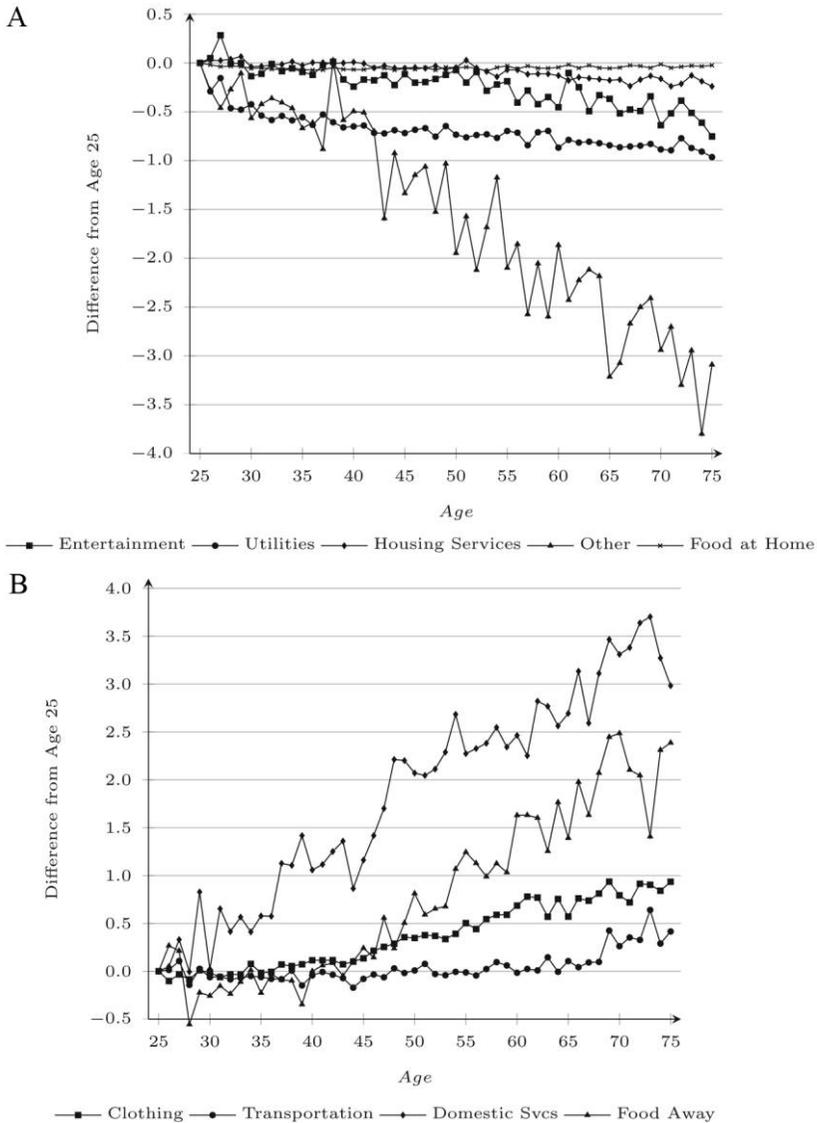
Fɪɢ. 4.—Life cycle profiles of disaggregated expenditure: variances. This figure depicts the life cycle profile of the cross-sectional variance of disaggregated log expenditure, conditional on cohort, year, and family composition controls. Specifically, we compute the cross-sectional variance of the residuals from the first-stage regression (eq. [4]) for each age-cohort pair and then remove cohort fixed effects to isolate the life cycle profile of cross-sectional variance (eq. [5]). Again, all deviations are from age 25. The consumption categories depicted in panel A are entertainment (squares), utilities (circles), housing services (diamonds), other nondurables (triangles), and food at home (x's). The consumption categories depicted in panel B are clothing and personal care (squares), transportation (circles), domestic services (diamonds), and food away from home (triangles). The sample is the same as for figure 1. See the text and Appendix A for a discussion of the consumption categories.

459

are complements to market work. In the next section, we discuss such evidence.

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

As discussed in Section II, to the extent that the opportunity cost of time evolves over the life cycle, one would predict changes in spending to occur within categories for which nonmarket work time and expenditures are substitutes. In this section, we document that much of the life cycle variation in spending on food, nondurable transportation, and clothing is accounted for by changes in labor supply. We do this in two ways. First, we use alternative data sets to shed light on the nature of expenditure in these categories, with a focus on changes over the life cycle. Second, we estimate a demand system to quantify the impact of labor supply on disaggregated expenditure categories. For reference, Appendix figure A1 shows the mean and the variance of the life cycle profiles of the labor supply of household heads from the CEX. Our analysis sample for this exercise is identical to the sample used above to document the life cycle consumption profiles. We show two measures of labor supply: the fraction of heads working (solid line) and the normal hours per week worked by the head (dashed line). This latter measure is not conditioned on working. Given that the decline in work hours starts for individuals around the age of 50, it is not surprising to find that work-related expenditures (and total nondurable expenditures) should start to decline around the age of 50. Likewise, given that the increase in the variance of labor force participation starts around the age of 50, it is not surprising to see the variance of work-related expenditures start to increase around the age of 50.

A.   Food Expenditures over the Life Cycle

In Aguiar and Hurst (2005, 2007a), we explored the differences between food expenditures and food intake. Using data from the Continuing Survey of Food Intake of Individuals (CSFII), which measures food intake at the individual level using detailed food diaries (including the quality of food consumed), the 2005 paper shows that food intake does not decline over the life cycle despite the decline in expenditures after middle age. On the contrary, using the detailed data on the quantity and quality of food consumed, we find that food intake actually increases after middle age. In the 2007a paper, we estimate a model of home production and food shopping to explain the differences between food expenditures and food intake. Using a variety of different data sources, that paper documents that, after middle age, individuals allocate more time to preparing meals

and shopping for food and, as a result, pay lower prices for a constant quality food basket.

Figure 5 sheds additional light on the margins of substitution that take place with respect to food spending over the life cycle. Using data from the CSFII, we measure an individual's propensity to eat away from home at various types of eating establishments. The primary design of the CSFII is to measure food intake via food diaries. The respondents were asked to provide very detailed comments about what they consumed, when they consumed it, and where they purchased it. We construct a variable called "eating away from home," which takes the value of one if the respondent reported purchasing food at a restaurant with table service, a restaurant without table service (i.e., establishments such as fast-food chains), a cafeteria, or a bar/tavern. On average, respondents in the CSFII spend roughly 2.5 days in the sample (some 2 days, others 3 days). For the entire sample, 64 percent of individuals reported eating away from home at least once during their time in the sample:[19] 38 percent eat at fast-food establishments, 33 percent eat at restaurants with table service, 10 percent eat at cafeterias, and 6 percent eat at bars. The percentages sum to more than 64 percent given that some individuals eat at multiple establishments during their time in the sample.

Figure 5 depicts the life cycle profile of the propensity to eat at the various types of restaurants. As with the expenditure data, we adjust the propensity to eat away from home for changing family composition, and all comparisons are made relative to households in their late 20s (25–29). Family controls consist of dummies for household size and four region dummies. The two waves of the CSFII include diaries from 1989–91 and 1994–96, which we pool as a single cross section and include year dummies. The overall pattern is similar to that of expenditures on food away from home, especially as it relates to the declines after middle age. In particular, the propensity to eat away from home falls by nearly 23 percentage points for individuals in their late 60s relative to individuals in their late 40s. However, the entire decline is due to a declining propensity to eat at fast-food restaurants and cafeterias. There is no decline in the propensity for individuals to eat at restaurants with table service as they age. This finding is consistent with the premise that the decline in food expenditures reflects households switching toward home production as

---

[19] The CSFII is a large nationally representative survey of individuals (as opposed to households). As in Aguiar and Hurst (2005), we use the survey waves conducted in 1989–91 and 1994–96 for our analysis. It should be stressed that the CSFII is essentially a cross section during one time period. As a result, we are not able to control for both cohort and age simultaneously when analyzing the data. For our analysis, we restricted the sample to 25–75-year-olds. Our total sample size used for the results in fig. 4 was 6,615 individuals. See App. A for a more detailed description of the CSFII data.
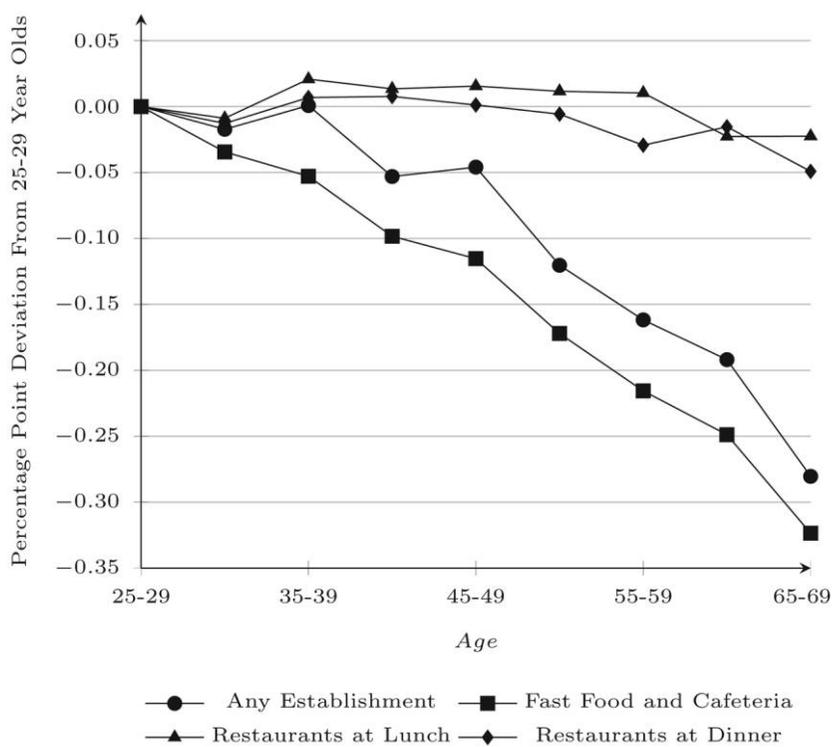
Fig. 5.—Propensity to eat away from home by establishment. Data come from the Continuing Survey of Food Intake of Individuals (CSFII) for the years 1989–91 and 1994–96. The figure plots the life cycle profile of the propensity to eat at different types of establishments, where propensity is measured by a dummy variable that takes a value of one if the respondent patronized the type of establishment during the course of the food diary and zero otherwise. "Any establishment" refers to the propensity for individuals to eat at a restaurant, a fast-food establishment, a cafeteria, or a bar/tavern during their time in the sample. "Fast food and cafeteria" is differentiated from "restaurants" by whether the establishment has table service. For restaurants "at dinner" and "at lunch," we differentiate whether the establishment was frequented for dinner or lunch, respectively. We regress each dummy variable on 5-year age dummies as well as family composition controls and plot the coefficients on the age dummies. All life cycle coefficients should be interpreted as linear probability deviations from 25–29-year-olds. See the text for a discussion of the family size controls. All data are weighted to be nationally representative using the CSFII survey weights. See the text and Appendix A for additional details of the CSFII sample.

their opportunity cost declines past middle age. The shift toward home production results in households purchasing fewer meals from fast-food establishments and cafeterias, which are close substitutes for home-produced food. The propensity to eat at restaurants with table service, which may provide additional utility beyond the food consumed, remains constant during the latter half of the life cycle.

B.   Transportation and Clothing Expenditures over the Life Cycle

Spending on clothing and transportation has long been viewed as a complement to market work (see, e.g., Nelson 1989; DeWeese and Norton 1991; Banks, Blundell, and Tanner 1998; Cogan 2001; Battistin et al. 2009). In order to work, households have to purchase additional clothing and must pay additional transportation costs associated with commuting. Lazear and Michael (1980), among others, have argued that certain costs of employment, such as costs of transportation to work and requisite clothing expenditures, be netted out of income when computing welfare calculations across people.

Spending on broad categories such as transportation and clothing likely includes components of spending that are associated with work, but this spending is also bundled with nonwork spending. For example, transportation expenditures reflect the need to commute to work as well as travel for other (leisure) purposes. While the expenditure data set does not distinguish costs due to work travel from costs due to nonwork travel, we can use time diaries from the pooled 2003–5 American Time Use Survey (ATUS) to gauge the relative importance of each.[20] The detailed categories of the ATUS allow us to identify time spent traveling to and from work separately from time spent traveling for other reasons (including going to the grocery store, going to visit friends, going to the movies, etc.). The average individual between the ages of 25 and 75 spends 9.0 hours per week traveling, with 2.3 hours per week associated with commuting to and from work. For those who work, work-related travel represents roughly one-third of all time spent traveling.

Figure 6 shows the life cycle profile of travel time after adjusting for changing family composition. The family composition controls include a marital status dummy, dummies for household size, and a dummy for whether the household has a child under the age of 5. The life cycle profile is expressed as an hours per week deviation from households aged 25–29. Consistent with the decline in transportation expenditures over the life cycle starting for households in their early 50s documented in figure 2B, the decline in transportation travel time also starts for individuals in their early 50s. However, as seen from figure 5, the entire decline in travel time occurs as a result of a decline in traveling to and from work. Nonwork travel time actually increases over the second half of the life cycle. If transportation expenditures are roughly proportional to trans-

---

[20] The ATUS is a nationally representative survey that uses time diaries to measure how individuals allocate their day. For a detailed account of the ATUS, see Aguiar and Hurst (2007a). For this analysis, we restrict the sample to only households between the ages of 25 and 75. Our total sample size was 38,876 individuals. See App. A for additional details about the ATUS, our sample selection, and our definition of variables.
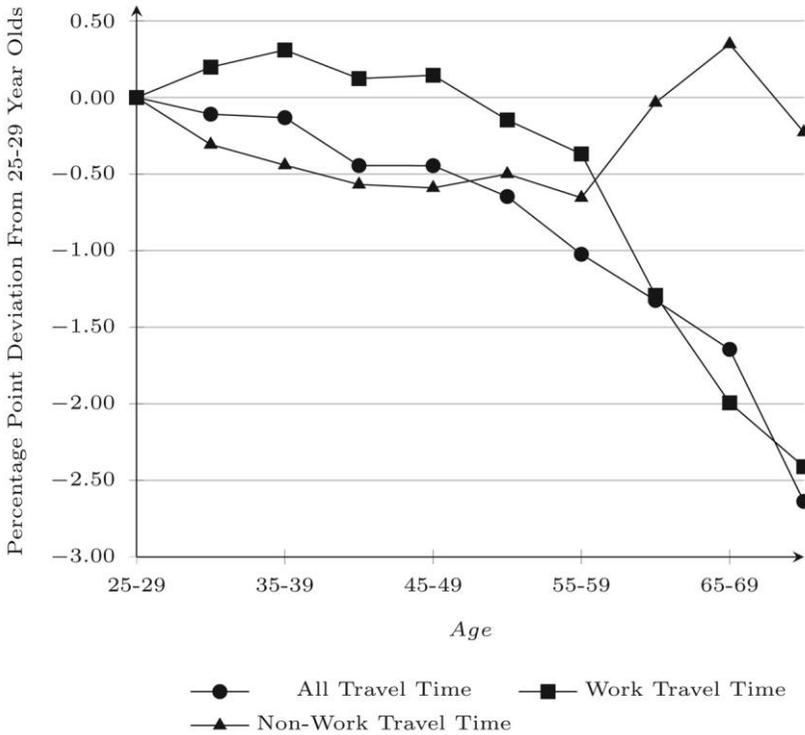
F IG. 6.—Travel times over the life cycle. Data come from the 2003–5 American Time Use Sample (ATUS). The figure plots the life cycle profile of the average time spent "traveling" (in hours per week) adjusted for family composition. "All travel time" refers to the amount of time individuals spend traveling to/from work (i.e., commuting time) and all other travel time. We regress hours per week on 5-year age dummies as well as family composition controls. The figure depicts the coefficients on the age dummies. See the text for a discussion of the family size controls. All age coefficients should be interpreted as hours per week deviations from 25–29-year-olds. All data are weighted to be nationally representative using the ATUS survey weights. See the text and Appendix A for additional details of the ATUS sample.

portation time, the data from the time use surveys suggest that the decline in transportation spending over the life cycle stems from the decline in time spent commuting to work. Again, this is consistent with the fact that transportation expenditures, and particularly their fluctuations over the life cycle, have a substantial work-related component.

C.   The Relationship between Spending and Work Status

Given the potential importance of work-related expenses to drive changes in expenditure over the life cycle, a natural approach would be to directly control for work status when estimating the life cycle profile of mean

expenditures or dispersion. A difficulty with simply adding controls for employment status to regression (4) is the fact that labor supply is closely associated with permanent income. For example, lower-wage workers in the time frame of our sample tend to work fewer hours than high-wage workers (see Aguiar and Hurst 2007b). Without a panel, controls for labor supply will also proxy for permanent income. However, using the standard tools of demand system analysis, we can explore the effect of labor supply on how expenditure is allocated across different goods, conditional on a given level of total expenditure. That is, by including total expenditure, we can isolate the effect of labor supply from variation in permanent income across households.

Specifically, we estimate the following:

$$
\begin{aligned}
s_{it}^k = \omega_0 &+ \omega_{\text{age}}\text{Age}_{it} + \omega_c\text{Cohort}_{it} + \omega_f\text{Family}_{it} \\
&+ \sum_k \omega_p^k \ln P_t^k + \omega_p \ln P_t + \omega_x \ln X_{it} + \omega_l L_{it} + \varepsilon_{it}^k,
\end{aligned}
\tag{6}
$$

where X_{it} is our measure of total nondurable spending and is defined as the sum of spending across core nondurables, work-related expenses, and food at home for household i in period t; s_{it}^k is the share of spending on consumption category k out of X_{it} for household i in period t.[21] By definition, the shares across the different consumption categories sum to one for each household. The age, cohort, year, and family status controls are the same as in equation (4). We include as additional controls the log price index of each of our subaggregates (P^k) as well as the log of the overall price index (P). These variables, together with the normalized year dummies, control for changes in relative prices across the consumption categories. Finally, we include a vector of controls describing household labor supply (L).

The fact that total expenditure appears on the right as a control and in the denominator on the left (as well as the sum of the individual goods) makes this specification vulnerable to measurement error. We follow the standard practice of instrumenting X with income and education.[22] There is also the issue of potential endogeneity of labor supply and that labor supply shocks are correlated with the residual shocks to relative expenditure shares. Unfortunately, we lack readily available instruments for labor

[21] Note that eq. (6) is a close parallel to the almost-ideal demand system of Deaton and Meullbauer (1980), conditioned on work status, family size, cohort, and age. We impose the restriction that the overall price index is given by the CPI-U but do not impose restrictions related to consumer optimization such as symmetry and homogeneity. The inclusion of work status controls to form a conditional demand system follows the important work of Browning and Meghir (1991) and Blundell, Browning, and Meghir (1994).

[22] Specifically, we instrument using log total household family income (summing together both labor and transfer income and bottom-coding income at one dollar), an indicator for whether income has been bottom-coded, income squared, income cubed, and education dummies.

supply, as the income and education controls used for measurement error would be prone to endogeneity issues similar to those of labor supply. While we view measurement error as the primary concern, we recognize the inability to formally test and control for the orthogonality of labor supply.

Using equation (6), we answer two different questions. First, among younger households (those under the age of 50), we assess whether work status is associated with spending on different consumption goods. If there are work-related consumption needs, we would predict that, all else equal, an increase in household labor supply would be positively associated with spending on those categories. Second, we use (6) to assess how much of the decline in spending after middle age on work-related consumption categories can be attributed to changes in household labor supply. In particular, we estimate (6) both with and without controls for labor supply and see how the age coefficients change.

To save space, we only highlight the results of our estimation in the text. However, a full discussion of our results, including all relevant tables and figures, can be found in Appendix C. We begin estimating (6) on a sample of married households in which the head is 50 years old or younger. For this specification, L includes two dummy variables: one indicating whether the husband is currently employed and another indicating whether the wife is currently employed. We estimate this specification separately for each of the disaggregated consumption categories shown in table A1. When we do this, we find that there are only three consumption categories for which the share of spending is positively associated with household labor supply. The three categories are nondurable transportation, food away from home, and clothing. Given the adding-up constraint, the share of spending on all other categories was negatively related to employment status. These simple demand system estimates confirm what we discussed above: spending on clothing, nondurable transportation, and food away from home is positively associated with household labor supply.

We also estimated (6) on a sample of all married households between the ages of 25 and 75. We then asked how much of the declining share of spending on food away from home, nondurable transportation, and clothing after middle age can be explained by changing work status. In this analysis, our vector of work status controls includes both whether the head and spouse were working and detailed controls for the hours worked conditional on working. Our estimates suggest that essentially all the decline in clothing and food away from home after middle age and 40 percent of the decline in nondurable transportation after middle age are due to changes in work status after middle age.

Collectively, the results in this section show that most of the declines in clothing, food, and nondurable transportation during the latter half of the life cycle are due to reductions in work status, which results in in-

creased nonmarket time inputs or a reduction in work-related expenses. Given this, it is not surprising that these categories display different life cycle profiles of mean spending and cross-sectional dispersion relative to all other categories after middle age.

VI. Quantitative Implications of Disaggregated Expenditure

In the previous sections, we have documented heterogeneity in consumption profiles across disaggregated commodity classes. In this section, we turn to drawing some broader lessons from these patterns. That is, what does looking at disaggregated commodities teach us that is not apparent from total nondurable expenditures? In Section II, we discussed the fact that differing consumption theories can match the same aggregate expenditure facts but could potentially be differentiated using disaggregated consumption data. For example, theories that stress poor planning in explaining the decline of expenditure at retirement (or with income in general) implicitly suggest that all expenditures should fall with income, with the magnitude of decline governed by the good's income elasticity. However, as we have documented above, many consumption categories continue to increase throughout the life cycle. The broader point that we want to make is that disaggregation can assist in identification when a consumption theory is suitably extended to include subcategories of consumption goods.

We take a first step at highlighting the power of using disaggregated data by revisiting the canonical incomplete markets model that has been the primary prism for viewing consumption data at least since Deaton (1991). To this end, we present an augmented model of consumption in which agents must insure idiosyncratic labor risk using a single risk-free bond, subject to a borrowing constraint. We then ask, through the lens of the model, whether using disaggregated consumption data delivers different estimates pertaining to the nature of uninsurable income risk faced by households.

A. Environment

We consider two versions of the model, a standard "one-good" formulation and an extended model with two consumption commodities ("two-good"), one of which is produced using nonmarket time. We collapse our model from the three goods depicted in figure 2 to two goods for tractability and ease of exposition. The home-produced good will comprise both food at home and work-related expenses, both of which show significant declines when individuals leave the labor force. This will also reflect the fact that work-related expenses are complements to market work (or substitutes for nonmarket time) but do not fall entirely to zero at retirement. We will refer to this composite good as "home production/

work-related," or just home production for short. We describe both environments together as the two-good framework nests the one-good model. Much of the model and its solution are standard, so we defer many details to Appendix B and focus in the text on the key deviations from the benchmark.

Agents have preferences over two consumption commodities. Specifically, agents have flow utility over core consumption c_1 and home-produced/work-related consumption c_2 according to the function $u(c_1, c_2)$. The home-produced good combines market inputs x_h and time input h according to the home production function: $c_2 = f(x_h, h)$. We assume that f is strictly concave and homothetic. Note that h captures all nonmarket activities, including leisure. In our numerical implementation, we use the following functional forms:

$$u(c_1, c_2) = \theta \frac{c_1^{1-\gamma}}{1-\gamma} + (1-\theta) \frac{c_2^{1-\sigma}}{1-\sigma},$$
$$c_2 = f(x_h, h) = x_h^{\psi} h^{1-\psi}.$$

The utility function is additively separable between the two goods, which implies that core consumption is separable from time allocation. This highlights the distinction between core and home-produced consumption discussed above.[23] The home production function is Cobb-Douglas. This is a common choice in the home production literature. Estimates of the elasticity of substitution between time and goods in home production tend to be around one or slightly above (see the discussion in Aguiar and Hurst [2007a]). The constant returns to scale assumption in home production is not restrictive given the power utility specification. The standard one-good model is obtained by setting $\theta = 1$.[24]

The rest of the model is largely standard. A unit-continuum of agents live for $T + 1$ periods, indexed by $t = 0, \ldots, T$, and discount flow utility at the rate β. We assume that there is no mortality risk and that agents invest in only a risk-free asset, which carries a risk-free rate r, subject to a borrowing constraint. We consider a stationary environment in which aggregate variables such as r are constant over the life cycle. The only uncertainty concerns an agent's idiosyncratic return to labor.

There are two sources of idiosyncratic labor income risk. The first is a labor productivity shock z, which we assume follows a Markov process

---

[23] Note that we do not impose homotheticity in the utility function, so in the presence of growth in market productivity we would need to allow θ (or, equivalently, relative home production productivity) to adjust accordingly.

[24] Another case that can be interpreted as the standard model is the one in which $\gamma = \sigma = 1$. This log-log specification allows for two goods that are both separable from leisure. The $\theta = 1$ specification ignores the intensive labor supply margin, which is the common—but not exclusive—assumption in the precautionary savings literature.

plus a common deterministic, age-related component. In particular, for agent i at age t we have

$$z_t^i = b_1 t + b_2 t^2 + \bar{z}^i + \alpha_t^i + \varepsilon_t^i,$$
$$\alpha_t^i = \rho \alpha_{t-1}^i + v_t^i,$$

where b_1 and b_2 define the (common) age-specific deterministic component of income, $\bar{z}^i$ is an individual-specific fixed effect, α_t is a persistent component of productivity that follows an AR(1), and ε is a transitory (independent and identically distributed [iid]) component. The fixed effect $\bar{z}^i$ is iid across individuals and the shocks (ε_t^i, v_t^i) are independent of each other (and $\bar{z}^i$) and iid across i and t. Each is drawn from a Normal distribution with respective variance $\sigma_i^2, i = \bar{z}, \varepsilon, v$. Henceforth, whenever possible we drop the i notation. Let $e^z n$ denote the efficiency units generated from n units of labor input, and let w denote the (aggregate) market wage per efficiency unit of labor.

The second source of labor risk concerns retirement (or disability or both). In particular, let R_t be a random variable that takes on the values of zero or one. Every agent is born with $R_0 = 0$. Conditional on $R_t = 0$, there is an age-dependent hazard that next period $R_{t+1} = 1$, which is an absorbing state. Early in the life cycle we can interpret this shock as a health or disability shock, while later in the life cycle this captures retirement. For simplicity, we model the exact timing of retirement as an exogenous shock rather than as a choice variable, although we calibrate to actual hazard rates so agents recognize that there is a high probability they will retire at certain points in the life cycle (like age 65). This captures the fact that retirement is anticipated in general, but the exact timing of retirement may be induced by a health or labor shock. Once retired, agent i lives off of financial wealth, a_i, and social security benefits, S_i. As a parsimonious proxy for disability insurance, we allow agents who have an early retirement shock to receive government transfers under the same social security system as retirees.

Let s denote the relevant state variables for an agent: $s = (a, \alpha, \epsilon, R, t, \bar{z})$. We approximate the social security payroll tax as a linear tax on income, τ, and assume that this tax is paid by the firm so that w is the after-tax wage rate received by agents.[25] The agent's problem in recursive form is therefore

$$V(s) = \max_{\{c_1, c_2, x_h, h, n, a'\}} u(c_1, c_2) + \beta \mathbb{E}\{V(s')|s\}$$

---

[25] In our calibration, we assume that social security benefits are indexed by an agent's fixed effect $\bar{z}_i$ and thus are not a separate state variable in the agent's problem.

subject to

$$a' = \begin{cases} (1+r)a + we^z n - c_1 - x_h & \text{if } R = 0 \\ (1+r)a + S - c_1 - x_h & \text{if } R = 1, \end{cases}$$

$$c_2 = f(x_h, h),$$

$$1 = n + h,$$

$$a' \geq \underline{a},$$

$$n \geq \underline{n} \qquad \text{if } R = 0; \ h \geq 0.$$

The first constraint is the budget constraint for workers and retirees, where we normalize units so that all expenditure categories carry a price of one; the next constraint is the home production technology; the third constraint is a time constraint, with the total time endowment normalized to one; the fourth constraint is the borrowing constraint; and the final two constraints are a minimum-work constraint that potentially rules out working part-time and will be useful in the quantitative implementation and a nonnegativity constraint on home production time. In the standard one-good formulation with $\theta = 1$, there is no operable intensive margin for labor supply, and we set n equal to one-third while employed and zero while retired.

We close the model by assuming an interest rate of 4 percent and discipline the equilibrium by targeting an aggregate wealth to aggregate income ratio of 3.1. The wealth-to-income target corresponds to that of the bottom 99 percent of the wealth distribution in the US economy, based on the 1992 Survey of Consumer Finances (see Diaz-Gimenez, Quadrini, and Rios-Rull 1997), and is the same used by Storesletten et al. (2004). Excluding the top percentile is necessitated by the fact that the CEX does not contain a representative sample of the extremely wealthy, while the top percentile of the wealth distribution holds roughly 30 percent of the economy's wealth. The target of 3.1 ensures that agents in our simulated economy do not accumulate counterfactually large asset positions.

B. Calibration

The goal of our quantitative exercise is to understand what the disaggregated consumption profiles tell us about the key preference and income process parameters and how this differs from lessons drawn from the one-good model. To implement this, we estimate the age-dependent deterministic component of income and the retirement/disability hazard rate directly from a sample of the PSID, which is described in Section D of Appendix A. We also use the PSID sample to pin down the age 25 cross-sectional variance of wages of 0.3, which restricts the sum of

variances $\sigma_{\tilde{z}}^2 + \sigma_{\varepsilon}^2 + \sigma_{v}^2 = 0.3$. Appendix B contains the details of our estimation and numerical solution of the model. The remaining parameters are selected so that the moments of the stationary distribution of the simulated model match target moments in the data. In matching life cycle profiles, we consider $t = 0$ as age 25 and set $T = 75$.

We match moments to calibrate the income process parameters (ρ, σ_{ε}, σ_{v}, $\sigma_{\tilde{z}}$), the preference parameters (β, γ, σ, θ), and the technology parameters (ψ, $\underline{n}$). The target moments consist of a real interest rate of 4 percent; an aggregate wealth-to-income ratio of 3.1; an average labor supply for prime-age workers (model ages 0–30) conditional on employment of one-third; the life cycle profile of mean log expenditure on core, home production/work-related goods, and the total; and the life cycle profile for the variance of (log) core, home-produced/work-related, and total expenditure, as well as the covariance of core and home-produced/work-related expenditure.[26] For the one-good model, we drop the labor target and the life cycle profiles of disaggregated core and home production expenditure and retain the profiles for total expenditure. We minimize the sum of the squared deviations between the model and the data across all targets.[27]

Note that we use empirical income data to estimate only the common deterministic component of wages and the initial cross-sectional variance. All other income processes are estimated through the model using the consumption data. To the extent that the implied income processes differ from the observed, the gap should be interpreted as a combination of measurement error in income (most relevant for the transitory component) and model misspecification (with the particular vulnerability due to the model's relatively parsimonious characterization of insurance contracts and income shocks).

C. Results: Implications for Income Risk

Table 3 reports the calibrated parameter values, and figures 7 and 8 report the simulated profiles and their empirical counterparts. The left-hand panels of figures 7 and 8 contain the mean profiles and the right-hand panels contain the variances. We begin our discussion with the one-good model and focus on the discount factor and the income risk parameters, which are present in both models. We defer discussion of the other parameters and additional results from the two-good model to Appendix B.

[26] Core includes housing services in the benchmark. We have recalibrated the model dropping housing services. The main quantitative point regarding permanent income risk described below remains essentially unchanged.

[27] Given that the life cycle changes in mean expenditure are nearly an order of magnitude larger than the variances, we down-weight the squared differences for the mean profiles by a factor of 10. The trade-off between matching mean and variances is relevant only for home production expenditure, as discussed in App. B.

TABLE 3
CALIBRATED PARAMETERS

Parameter	Description	One-Good Model	Two-Good Model
	Preferences and Home Production		
β	Discount factor	.961	.964
ψ	Market share in home production	NA	.156
σ	CRRA for c_2 (home-produced)	NA	3.628
γ	CRRA c_1 (core)	1.486	.531
θ	Utility share on core	1.000	.658
	Labor Productivity		
$\underline{n}$	Minimum labor requirement	NA	.333
σ_ε^2	Transitory variance	.119	.127
σ_v^2	Persistant variance	.018	.016
ρ	Persistence parameter	.977	.960
σ_z^2	Fixed-effect variance	.166	.159
	Relative log wage variance at 20-year horizon*	2.67	2.08

NOTE.—CRRA = constant relative risk aversion.
* The relative log wage variance at the 20-year horizon equals

$$\frac{\sigma_\varepsilon^2 + [(1 - \rho^{42})/(1 - \rho^2)]\sigma_v^2}{\sigma_\varepsilon^2 + \sigma_v^2}.$$

Figure 7 shows that the one-good model can match the life cycle profiles closely. Table 3 indicates that this is done by setting the agent's discount factor to 0.96 versus an interest rate of 4 percent. This equivalence of discount rate and interest rate reflects that nondurables including housing services and excluding alcohol and tobacco do not decline significantly after middle age. The slight curvature in the profile reflects the presence of borrowing constraints and the induced desire to build up precautionary savings. The cross-sectional variance of expenditure can also be matched quite well. Table 3 shows that this is done with a transitory variance of 0.119 and a persistent innovation variance of 0.018. The requirement that the cross-sectional wage variance is 0.3 at $t = 0$ implies that the fixed-effect variance is 0.166. The calibrated persistence parameter for income is 0.977.

For the two-good model, figure 7 indicates that this model is also able to replicate the life cycle profiles of aggregate expenditure.[28] In regard to the disaggregated consumption categories, figure 8 shows that the two-good model matches the steady rise in both the mean and variance of core expenditure, although the model predicts a slight flattening of mean expenditure at the end of the life cycle relative to the data. In regard to home production/work-related expenditure, it matches the

[28] While the two-good model has three additional parameters, it also is calibrated to match the disaggregated life cycle profiles (each with 50 age-specific moments) as well as average labor supply. In this sense, the two-good model is not matching the aggregate profiles using extra degrees of freedom.

FIG. 7.—Simulated life cycle profiles of aggregated expenditures: means (left) and variances (right). This figure depicts the empirical and simulated life cycle profiles of aggregated consumption (core plus home production expenditure). The dashed line is the data, the dotted line is the one-good model, and the solid line is the two-good model. The left panel depicts means and the right panel depicts cross-sectional variances. The means are log deviations from age 25 and the variances are level differences from age 25. The empirical series is conditional on the same controls as the series in figure 1.

Core Expenditure

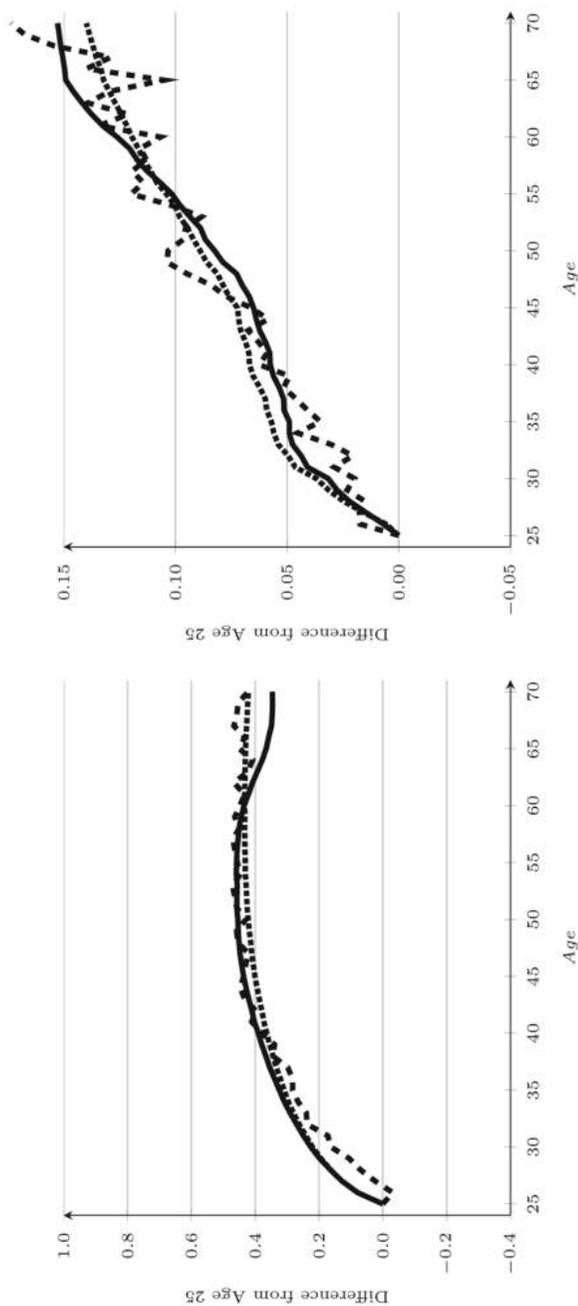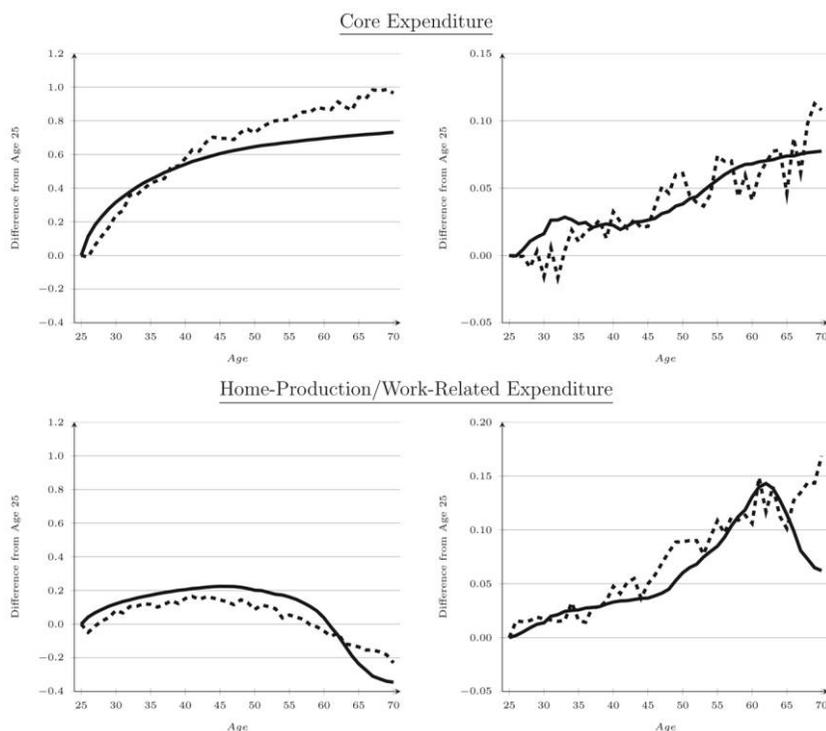Home-Production/Work-Related Expenditure

FIG. 8.—Simulated life cycle profiles of disaggregated expenditures: means (left) and variances (right). This figure depicts the empirical and simulated life cycle profiles of consumption. The dashed line is the data and the solid line is the two-good model. The left panels are means and the right panels are cross-sectional variances. The means are log deviations from age 25 and the variances are level differences from age 25. The empirical series is conditional on the same controls as the series in figure 1.

hump closely through age 60 and then slightly overstates the decline after the peak retirement years. Note in table 3 that the intertemporal elasticity of substitution for the home-produced good is $1/3.6 = 0.3$, while the elasticity of substitution of the home production function is set to one. From the discussion of Section II, this implies that home production expenditure will be positively correlated with the opportunity cost of time over the life cycle. In regard to the variance of home production/work-related expenditure, it matches the fact that home production expenditure is roughly flat early in the life cycle and then increases in middle age. However, the simulations indicate a late–life cycle mini "hump" in the home production expenditure variance, which is not clearly seen in the data. Given that labor income is the only source of risk in the model, late–life cycle increases in variance will be difficult to match when most agents have retired. This suggests room for another source of shocks late in the life cycle, a natural candidate for which is medical expenses (Palumbo 1999; DeNardi, French, and Jones 2010).

These discrepancies are discussed further in Appendix B. This appendix also contains a discussion of the covariance between core and home production expenditures.

In terms of comparison with the one-good model, the discount factor is the same, that is, is close to the interest rate. The fact that agents are relatively patient is clearly not at odds with the decline in home production expenditure.

A key finding of our comparison of the two models is that the two-good model matches the aggregate profiles with less persistent income risk. In particular, the implied innovation variance (0.016) and the persistence parameter (0.960) are lower for the two-good model. To see the difference in implied long-run risk clearly, table 3 reports the variance of income at a 20-year horizon relative to the initial income variance, conditional on the individual fixed effect.[29] This ratio summarizes the evolution of risk a young agent faces as he or she looks forward toward middle age. The one-good model indicates that the variance increases by a factor of 2.67, while the two-good model increases only by a factor of 2.08. That is, the one-good model overstates midlife income risk by 25 log points compared to the two-good model.

For comparison, the wage data from the PSID imply a ratio of 2.21. More specifically, estimates from the PSID for the wage process parameters are $\sigma_\varepsilon^2 = 0.156$, $\sigma_v^2 = 0.017$, $\rho = 0.976$, and $\sigma_{\bar{z}}^2 = 0.129$. These estimates were obtained by matching the life cycle profile of cross-sectional wage variance between age 25 and age 60 as well as the 1-year autocovariance of wages between ages 26 and 60, weighting each moment equally. See Section D of Appendix A for details. The implied wage process for the two-good model matches the observed process from the PSID quite well and is closer to the observed process than the one-good model. In this sense, there is less "missing insurance" when consumption is viewed through the multigood/nonseparable framework. In fact, the one-good model overpredicts income risk, which is hard to reconcile with the limited insurance opportunities in the model. That is, the increase in aggregated consumption volatility requires a counterfactually large increase in wage risk over the life cycle despite the limited insurance opportunities in the model. In contrast, the slight underprediction of risk implied by the disaggregated consumption data suggests that self-insurance is, to a first approximation, a useful description of insurance opportunities for the average consumer.

The difference in implied wage risk between the one-good and two-good models exists despite the fact that both models match the pro-

---

[29] Specifically, we calculate

$$\frac{\mathbb{E}(\varepsilon_{20}^i + \alpha_{20}^i)^2}{\mathbb{E}(\varepsilon_0^i + \alpha_0^i)^2} = \frac{\sigma_\varepsilon^2 + [(1 - \rho^{42})/(1 - \rho^2)]\sigma_v^2}{\sigma_\varepsilon^2 + \sigma_v^2}.$$

files of aggregate expenditure quite well. The difference reflects that the two-good model allocates some of the expenditure response to income risk to the substitution effect highlighted in Section II, while the one-good model uniformly attributes all consumption variation to permanent income shocks. The two-good model identifies the income risk by comparing the life cycle profiles for core expenditures to that of the home-produced good. The sensitivity of home production/work-related expenses to labor market status generates an additional source of consumption variation in the two-good model, augmenting permanent income shocks. In this manner, the profiles of the disaggregated expenditure categories provide additional information that is missing from the one-good exercise. As noted in Section II, there is a one-good reduced form for the two-good model in which there is a nonseparability between the one good and leisure. The disaggregated data discipline the extent of this nonseparability. The fact that agents exit the labor market entirely, as well as the smaller role played by the intensive margin, is reflected in the cross-sectional variance of home production expenditure as well as the covariance with core expenditure (we discuss the covariance in App. B). This generates the empirical profile for the variance of expenditure with a mixture of permanent income shocks and the substitution of time for expenditure, while the one-good model relies exclusively on the former.

The main takeaway from the quantitative exercise is that the heterogeneity in life cycle expenditure across disaggregated consumption goods can be useful in identifying the source and size of uninsurable income risk. Both the one-good and the two-good models can match the behavior of total expenditure quite closely. However, the information contained in the subaggregates (along with the structure of the model) provides a distinct perspective on the nature of the underlying income process. This exercise indicates that permanent risk is overstated by roughly 20 percent when one ignores the information contained in disaggregated expenditure.

VII.   Conclusion

In this paper, we highlighted the importance of using disaggregated consumption data to understand the behavior of the composite consumption good. In particular, we first documented that there is a tremendous amount of heterogeneity across goods in the life cycle profiles of the mean and cross-sectional variance of expenditure. In particular, the life cycle profiles of clothing, food, and nondurable transportation differed markedly from the profiles of other goods. For example, mean spending on these goods falls sharply after middle age, while spending on all other goods does not fall (or even increases) after middle age. Additionally, the cross-sectional variance of expenditure increases dramatically for cloth-

ing, food, and nondurable transportation between the ages of 45 and 65. No such increase is found in the other goods.

Second, we provide evidence showing that the differences in the profiles across goods can be explained by clothing, nondurable transportation, and food away from home being work-related expenditures and food at home being amenable to home production. As the propensity to work decreases starting in middle age, it is not surprising to see spending on work-related expenditures and home-produced goods fall. Likewise, as the variance of hours worked increases across households starting in the late 40s through the mid-60s, it is not surprising for the variance of cross-sectional expenditure on these goods to increase during this period of the life cycle.

The third innovation of the paper is to discuss how the disaggregated expenditure data can be used to test among and refine consumption theories. Many theories can match the given life cycle profiles of a composite nondurable good. However, many of those theories (implicitly or explicitly) have different implications for the life cycle profiles of disaggregated consumption goods. For example, theories that stress uninsurable income shocks or intertemporal substitution predict that the life cycle profiles of luxury goods should differ from the profiles of necessities.[30] Other theories predict that the life cycle profiles for goods that are amenable to home production or are complements to market work should differ from the profiles of other goods. While both theories may predict that mean spending on a composite nondurable good should fall after middle age and that the cross-sectional variance of expenditure on a composite nondurable good should increase with age, the implications for the disaggregated goods will differ.

In the final part of the paper, we show one such application that highlights the importance of using disaggregated expenditure data. The application we focus on is computing the amount of permanent income risk faced by households. Traditionally, this statistic is disciplined in life cycle consumption models by the change in both the mean and the cross-sectional variance of spending over the life cycle. But, as we discussed above, if there are work-related expenditures, home-produced goods, or nonseparabilities between consumption and leisure, the mean and cross-sectional variance of expenditure will also be determined by the cross-sectional variance in hours worked. The use of the disaggregated data allows us to isolate and quantify the relative importance of each mechanism. We find that within our multigood framework, the estimated increase in variance of income faced by households over a 20-year horizon is 25 percent lower relative to the estimate from an otherwise similar one-good model. Moreover, the multigood model's estimate of permanent

---

[30] In this regard, the time series of disaggregated data can similarly shed light on the evolution of income inequality over time, as in Aguiar and Bils (2011).

income risk closely matches actual estimates of permanent income risk calculated using household panel data on income.

Our work also highlights the potential importance of looking at the covariances across the disaggregated consumption categories. While our model does relatively well at matching the life cycle profiles of mean expenditures and the cross-sectional variances of both goods in our two-good model, it does less well at matching the life cycle profile of the covariance between the two goods. A fruitful line for future research is to shed more light on the covariance in expenditures across goods. Finally, our work stresses the importance of other risks that households face late in life. The cross-sectional variance in expenditure on both core expenditures and work-related expenditures increases after the age of 65, when the variance in hours worked is declining. A natural candidate to explain this risk is medical expense shocks (e.g., Palumbo 1999; DeNardi et al. 2010). Exploring the disaggregated expenditure response to health shocks may deepen our understanding of the relative importance of uninsurable risk at the end of the life cycle.

Appendix A

Data

In this appendix we discuss data sets and sample restrictions. All data and Stata programs are provided as supplementary material online. Additional robustness exercises can be found in online Appendix C.

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 3 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).[31]

---

[31] 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.

We use the extracts compiled by Ed Harris and John Sabelhaus and provided by the NBER (http://www.nber.org/data/ces_cbo.html). Harris and Sabelhaus aggregate expenditures into 47 categories, which are listed in the supplementary documentation. The Harris and Sabelhaus data set includes households whose first interview was conducted between the first quarter of 1980 and the second quarter of 2003. Owing to changes in the survey methodology, data from the last two quarters of 1985 and 1995 are omitted.[32] 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 ages 25 and 75. To obtain reliable estimates of cohort effects, we also restrict attention to cohorts with at least 10 years of data. In particular, we restrict the sample to households born between 1915 and 1968, that is, to households whose head is at most 65 in 1980 and at least 35 in 2003. 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, their restricted sample 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 53,412 households. This is our main sample for analysis.

B. Data from American Time Use Survey

We use the 2003, 2004, and 2005 waves of the American Time Use Survey (ATUS) conducted by the US 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 respondents' 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 25–75, 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 age controls, a dummy for marital status, and 10 family size dummy variables, and we report the coefficients on the age controls.

The ATUS reports time allocation using over 400 detailed activity codes. For our analysis we focus on three aggregates: total travel time (classification category

---

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

17 in 2003 and 2004 and classification category 18 in 2005), travel associated with work (subcategory 4 out of total travel time), and all other travel time.

C. Data from Continuing Survey of Food Intake of Individuals

For the analysis in figure 5, we use data from the Continuing Survey of Food Intake by Individuals (CSFII) collected by the US Department of Agriculture. The survey is cross-sectional in design and is administered at the household level. We pool the two most recent cross-sectional surveys: the first interviewed households between 1989 and 1991 (CSFII_89) and the second interviewed households between 1994 and 1996 (CSFII_94).

The CSFII_89 and CSFII_94 were designed to be nationally representative. On the basis of sample averages, the demographic coverage of the CSFII closely tracks that of the PSID. The 1989 data also include an additional data set that over-samples low-income households. We exclude the oversample from our analysis. When analyzing individual-level data, we restrict our analysis to household heads.

Each household member in the CSFII data also filled out detailed food diaries, recording his or her total food intake during a particular 24-hour period, with the CSFII_89 collecting 3 days and CSFII_94 2 days of diaries, respectively. As part of their entries, they had to record where their food was purchased. We focus on the food purchased at nongrocery establishments. In particular, we examine only food purchased at restaurants with table service (restaurants), restaurants with counter service (fast-food establishments), cafeterias, and bars. Collectively, we refer to these categories as food purchased away from home.

The data sets track standard economic and demographic characteristics of the survey respondents including age, educational attainment, race, gender, occupation, employment status, hours worked, retirement status, family composition, geographic census region, whether the household lives in an urban area, homeowner status, and household income. The survey also asks respondents detailed questions regarding health status, health knowledge, and preference for nutrition.[33]

D. Panel Study of Income Dynamics

When calibrating the model in Section VI and comparing the implied wage process to the observed wage process, we use additional data from the PSID. The PSID data set is that used in Kaplan (2012), and we thank Greg Kaplan for kindly providing the data set. Kaplan's data set contains a detailed appendix on the underlying data. The data cover survey years 1968–2007. Since 1997, the survey has been conducted every 2 years. The baseline sample includes household heads aged 25–75. This consists of 10,739 individuals for a total of 113,464 observations.

To calibrate the deterministic component of wages used in the model, we regress log real wages on age, age squared, a sequence of normalized year dummies that capture business cycle fluctuations (the same dummies from regression [4]

---

[33] See the data appendix of Aguiar and Hurst (2005) for a detailed discussion of the CSFII survey methodology and a comparison of the sample demographics in the CSFII to the sample demographics from other large household-based surveys.

used with the consumption data), and an individual fixed effect. For the wage process, we restrict attention to household heads between the ages of 25 and 60 with nominal earnings less than $1 million and annual hours between 520 and 5,200 hours. This is a sample of 9,261 individuals with a total of 85,277 observations. Hourly wages are computed by dividing annual earnings (income from wages, salaries, commissions, bonuses, overtime, and the labor part of self-employment income) by annual hours (sum of annual hours worked on the main job and on extra jobs plus annual hours of overtime). Hours are computed using information on usual hours worked per week and the number of weeks actually worked in the previous year. Kaplan (2012) fits a Pareto distribution to impute top-coded earnings data. Nominal earnings are deflated using the CPI-U. The estimated coefficient on the linear term (with age 25 normalized to zero) is 0.0317 (standard error 0.0007); the estimated coefficient on the quadratic term is -0.00073 (standard error 0.00002). In the quantitative model, we use 0.03 and -0.0007 for the respective coefficients on the deterministic trend.

In Section VI, we compare the model's implied wage process with that observed in our PSID sample. The latter parameters were estimated as follows. Using the PSID sample, we regressed log real wage on a full set of age dummies, normalized year dummies, and a full set of cohort dummies. This is essentially (without family size controls) the specification used for consumption (4). We then extract the residuals from this regression to obtained normalized wages. Note that we extract cohort means to control for trend (aggregate) growth in productivity and use year dummies to capture aggregate business cycles. This leaves the residual individual fixed effects in the normalized wages, consistent with the model's wage process. Similarly, the full set of age dummies ensures that the stochastic component of residual log wages has mean zero at each age, also consistent with the model's wage process. For each age between 25 and 60, we compute the cross-sectional variance and the first-order autocovariance of individual wages. Note that the biannual survey years will not have observations for the autocovariance. This yields 72 moments. The age 25 residual variance of 0.3 is used to calibrate the model of Section VI. Other than this and the quadratic discussed in the previous section, the calibration does not rely on the residual wage series. We use this series for comparison purposes only. We estimate the four wage process parameters using equally weighted generalized method of moments. As reported in the text, the estimated parameters are $\sigma_\varepsilon^2 = 0.156$, $\sigma_v^2 = 0.0167$, $\rho = 0.976$, and $\sigma_{\bar{z}}^2 = 0.144$.

To calculate retirement/disability hazard rates, we consider PSID male household heads between the ages of 25 and 75 who are either employed, unemployed/looking, retired, or disabled. This excludes students, homemakers, and those with a noncategorized employment status. This comprises 7,592 individuals and 89,422 observations. From this population, we compute the fraction who are working or are unemployed (i.e., not retired or disabled) at each age. To smooth this series, we take a 5-year centered moving average (truncating the 5-year window at the youngest and oldest ages). Using this smoothed series, we calculate the hazard rate of exiting the labor force at a particular age t as the percentage decline in the fraction working between age $t-1$ and age t. We assume that all agents are working at age 24 to initiate the series (the fraction working/unemployed at 25 is greater than 0.99). The hazard rate is depicted in figure A3.
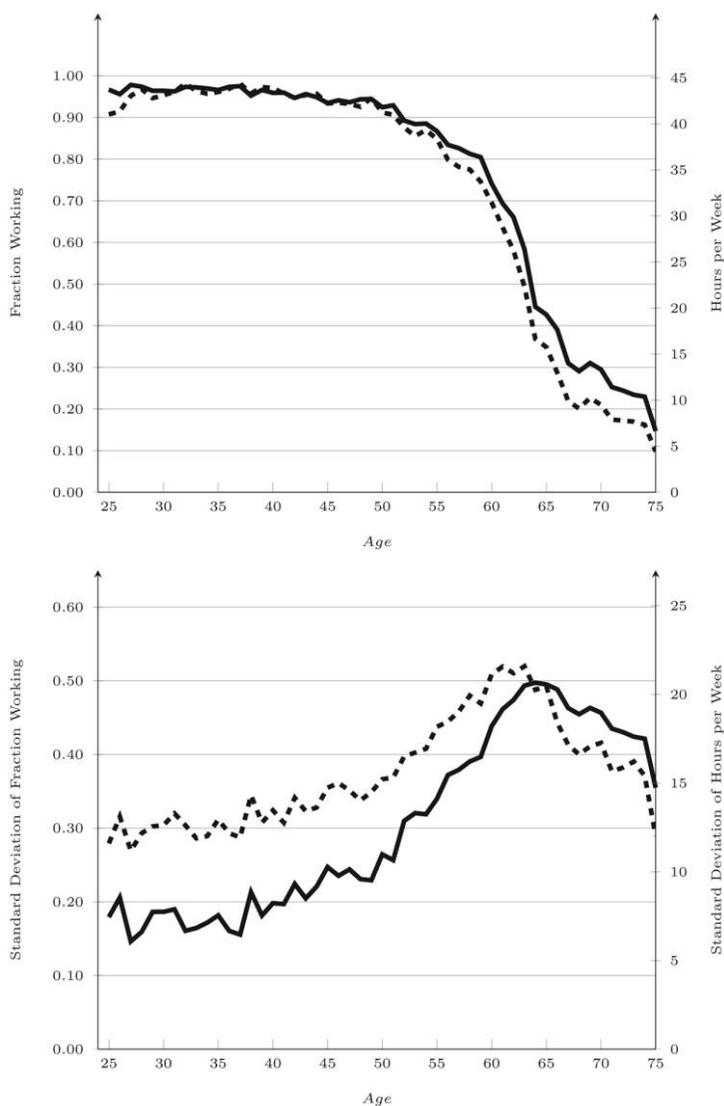
E.   Additional Figures and Tables

Fig. A1.—Employment and hours over the life cycle. The top panel shows the life cycle profile of the propensity to work (solid line, left axis) and average hours per week worked (dashed line, right axis) for household heads. The average hours per week series is not conditional on working. No other controls are used to adjust these series. The sample is identical to the sample described in the caption of figure 1. The bottom panel shows the corresponding life cycle profile of the standard deviation of the propensity to work (solid line, left axis) and average weekly work hours (dashed line, right axis) for household heads.
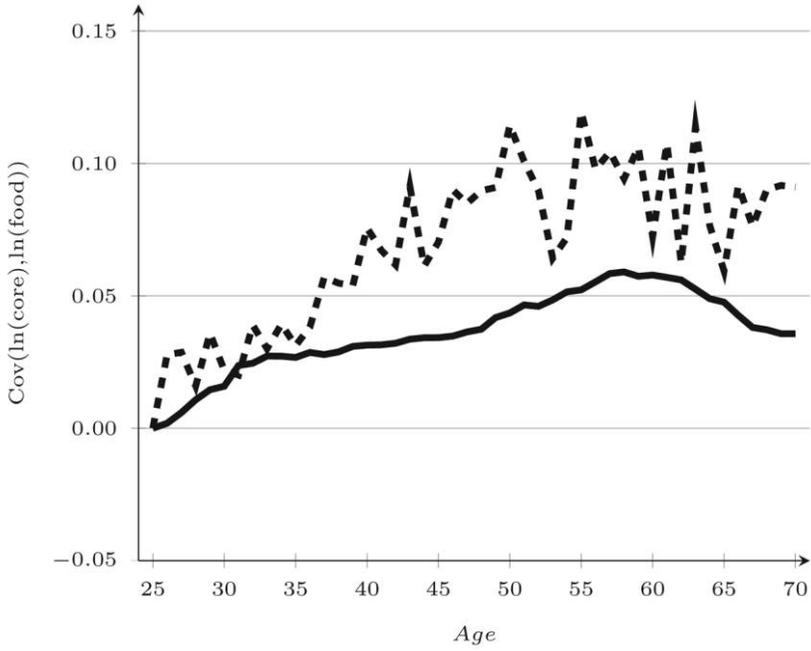
Fig. A2.—Covariance of core and home production/work-related expenditure. This figure depicts the simulated (solid) and empirical (dashed) profiles for the covariance between core expenditure and home production/work-related expenditure. Both series are log deviations from age 25.

Fig. A3.—Empirical retirement hazard rate. This figure depicts the empirical hazard used in the model solution and simulation. The data source and details of the calculations can be found in Section D of Appendix A.

484

TABLE A1
SUMMARY STATISTICS FOR DISAGGREGATED CONSUMPTION CATEGORIES
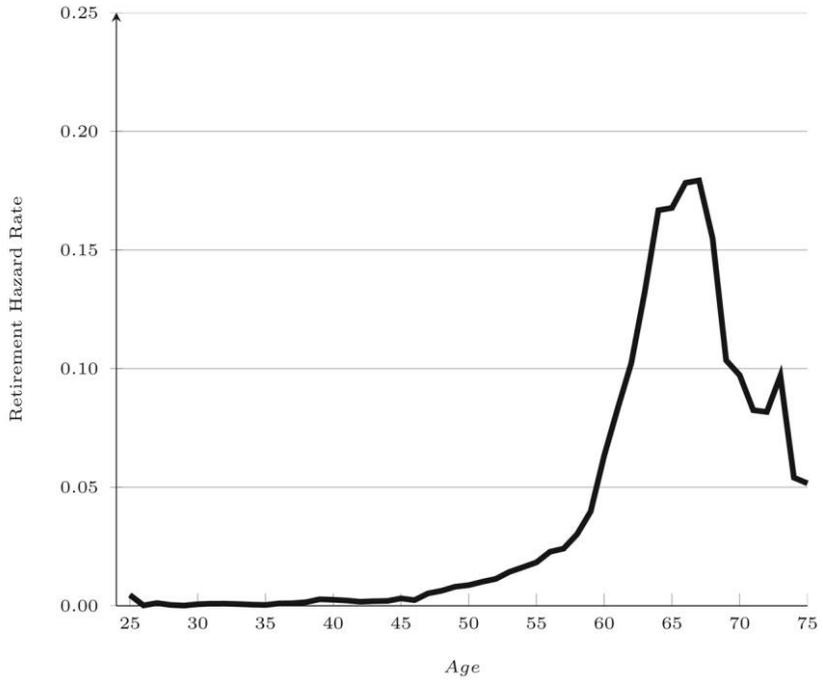
Disaggregated Consumption Group	Share of Expenditures at Ages 25–27 (1)	Share of Expenditures at Ages 43–45 (2)	Share of Expenditures at Ages 64–66 (3)	Log Change in Expenditure between Ages 25 and 45 (4)	p-Value of Change (5)	Log Change in Expenditure between Ages 45 and 65 (6)	p-Value of Change (7)
Decreasing categories:							
Food at home	.24	.25	.26	.24	<.01	−.11	<.01
Transportation	.20	.20	.18	.26	<.01	−.28	<.01
Clothing and personal care	.13	.12	.10	.08	<.01	−.40	<.01
Food away from home	.10	.09	.08	.16	<.01	−.79	<.01
Alcohol and tobacco	.06	.04	.04	−1.27	<.01	−2.47	<.01
Nondecreasing categories:							
Housing services	.49	.54	.63	.72	<.01	.28	<.01
Utilities	.15	.16	.19	.69	<.01	.34	<.01
Entertainment	.05	.07	.06	.77	<.01	.16	<.01
Other nondurable	.03	.05	.06	1.44	<.01	.14	.17
Domestic services	.03	.03	.04	1.45	<.01	.45	<.01

NOTE.—This table summarizes the life cycle shares and log changes in expenditure on different disaggregated consumption categories. The columns correspond to the ones in table 1.

TABLE A2

SUMMARY OF CHANGE IN CROSS-SECTIONAL VARIANCE OVER THE LIFE CYCLE BY CONSUMPTION CATEGORY

Disaggregated Consumption Group	Unconditional Cross-Sectional Variance at Ages 25–27 (1)	Change in Residual Cross-Sectional Variance between Ages 25 and 45 (2)	p-Value of Change (3)	Change in Residual Cross-Sectional Variance between Ages 45 and 65 (4)	p-Value of Change (5)
Decreasing categories:					
Food at home	.34	−.04	<.01	.01	<.01
Transportation	.79	−.15	<.01	.15	<.01
Clothing and personal care	.65	.15	<.01	.59	<.01
Food away from home	1.88	−.08	<.01	1.60	<.01
Alcohol and tobacco	6.20	1.61	<.01	3.21	<.01
Nondecreasing categories:					
Housing services	.42	−.06	<.01	−.13	<.01
Utilities	.89	−.55	<.01	−.12	<.01
Entertainment	1.47	−.28	<.01	−.24	<.01
Other nondurable	9.61	−1.03	<.01	−1.50	<.01
Domestic services	6.73	.98	<.01	1.62	<.01

NOTE.—This table summarizes the cross-sectional variance at ages 25–27 and the change in variance over the life cycle for different disaggregated consumption categories. The columns correspond to the ones in table 2.

Appendix B

Model

In this appendix we provide further details on the quantitative model's solution and implications.

A. Model Solution Details

To begin, recall that the consumer's problem is

$$V(s) = \max u(c_1, c_2) + \beta \mathbb{E}\{V(s')|s\}$$

subject to

$$a' \leq \begin{cases} (1+r)a + we^z n - c_1 - x_h & \text{if } R = 0 \\ (1+r)a + S - c_1 - x_h & \text{if } R = 1, \end{cases}$$

$$c_2 \leq f(x_h, h),$$

$$1 \geq n + h,$$

$$a' \geq \underline{a},$$

$$n \geq \underline{n} \qquad \text{if } R = 0; \ h \geq 0.$$

Under our assumptions on utility and technology, this is a standard optimization problem with a concave objective and convex constraints. The first three constraints will be satisfied with equality at the optimum, and the Inada conditions on u and f ensure that $h \geq 0$ does not bind. We can substitute the constraint $c_2 = f(x_h, h)$ into the utility function and let $\lambda > 0$ be the budget constraint multiplier, $\phi\lambda$ the time constraint multiplier, μ the borrowing constraint multiplier, and $\zeta\mu$ the $n \geq \underline{n}$ multiplier. The first-order conditions are

$$u_1 = \lambda,$$

$$u_2 f_1 = \lambda,$$

$$u_2 f_2 = \phi\lambda,$$

$$we^z = \phi - \zeta,$$

$$\beta \mathbb{E} V_{a'} = \lambda - \mu.$$

The envelope condition is $V_a = (1+r)\lambda$. This implies the Euler equation

$$u_1(c_1, c_2) \geq \beta(1+r)\mathbb{E} u_1(c_1', c_2'),$$

with strict inequality implying $a' = \underline{a}$. We also have the static optimization condition:

$$\frac{f_2}{f_1} = \phi.$$

The term on the right is the shadow cost of time. When $n > \underline{n}$ (i.e., $\zeta = 0$), the first-order condition for n implies that ϕ equals the wage we^z.

To get a sense of how fluctuations in the price of time influence allocations, consider a change in the log wage (z) that leaves λ unaffected. Assuming interiority of labor, the static optimality condition plus our Cobb-Douglas assumption on f imply

$$\frac{d\ln x_h}{dz}\bigg|_\lambda = 1 + \frac{d\ln(1-n)}{dz}\bigg|_\lambda,$$

where we have used $h = 1 - n$. The last term on the right is the negative of the Frisch elasticity of "leisure," or nonmarket time. As this elasticity approaches zero, labor supply becomes inelastic and market inputs become the favored margin of adjustment. Therefore, a large response of market expenditures to the price of time goes hand in hand with a small Frisch elasticity. We can use our other functional form assumptions on utility and the first-order condition $u_2 f_2 = \phi\lambda$ to derive (when $n > \underline{n}$)

$$\frac{d\ln x_h}{dz}\bigg|_\lambda = \frac{(1-\psi)(\sigma-1)}{\sigma}. \tag{B1}$$

If $\sigma < 1$, then an increase in market productivity leads to a decline in expenditure on x_h. A low σ implies a willingness to substitute consumption of c_2 over time, which dominates the static choice between market and time inputs. Recall that this latter margin has an elasticity of one under Cobb-Douglas, which is why one is the relevant cutoff for σ. If $\sigma > 1$, then the consumer would rather not postpone consumption of c_2 until it is cheaper, relying instead on the static substitution between time and goods.

With the same logic, a high σ implies a low Frisch elasticity of labor supply. In particular, our functional forms imply

$$\frac{d\ln(1-n)}{dz}\bigg|_\lambda = -\frac{1-\psi(\sigma-1)}{\sigma}.$$

The negative of the right-hand side is decreasing in σ, implying that the Frisch elasticity of nonmarket time is decreasing in the curvature of utility over c_2. An increase in z leads to an increase in the ratio of market to nonmarket inputs $x_h/(1-n)$. The more an agent accommodates an increase in wage by increasing market inputs x_h (rather than postponing c_2), the less it reduces nonmarket time. We shall return to this trade-off when we discuss the multigoods model's implications for x_h below.

To solve the problem numerically, we consider a grid of assets and discretize the persistent shock α and the transitory shock ϵ. For the latter, we use Tauchen's approximation with a five-state discrete Markov chain for each process. For the former, we allow for 40 grids, with a nonuniform distribution to ensure denser coverage over the strongly concave region near the borrowing constraint $\underline{a} = 0$. The fixed effect $\bar{z}$ can take on two values: $\pm\sigma_{\bar{z}}$. Given two retirement states and 51 ages, our state variable $s = (a, \alpha, \epsilon, R, t, \bar{z})$ takes on 204,000 values. Recall that social security payments are indexed to the fixed effect $\bar{z}$, so it is not a separate

state variable. We solve the consumer's problem working backward from the last period of life, using the Euler equation and linearly interpolating across the asset grid. With the consumer's problem solved, we simulate 10,000 life cycle paths.

We calibrate the deterministic component of the income process and the retirement hazard using a sample from the PSID, as described in Section D of Appendix A. The retirement hazard is depicted in figure A3. We set the social security payments to 40 percent of expected lifetime income, with the expectation conditional on the fixed effect. The remaining parameters are calibrated through simulation. For each simulation, we compare the simulated moments to their empirical counterparts, as described in the text. We minimize the squared difference between the model and empirical moments using a simplex search, experimenting with a wide range of initial simplexes.

B. Additional Results for the Two-Good Model

The primary focus of our quantitative exercise is to contrast a single-good model with the multigood framework introduced in Section II, with the goal of understanding predictions for implied income risk. The multigood model has a number of additional predictions beyond those nested in the one-good benchmark that can be compared with the data. For completeness, in this section we discuss the parameters and prediction that are unique to the two-good model.

The model matches the core expenditure quite well, although the profile of the variance is slightly steeper at the start of the life cycle and flattens out later in the life cycle. This latter effect is somewhat unavoidable given that retirement implies no additional uncertainty. In particular, within the model there is no risk once an agent stops working. The large literature on late-life consumption has emphasized the importance of uninsurable medical expenses, which we have omitted given our primary focus on labor income risk. The model underpredicts the increase in mean core expenditure late in the life cycle, for similar reasons. With little remaining risk, consumption will have a relative flat slope given that $\beta \approx (1+r)^{-1}$.

In regard to home production/work-related expenses, the model matches the increase through late middle age in both the mean and the variance. It somewhat overpredicts the decline in mean expenditure late in the life cycle. Moreover, the fact that retirement is bunched around ages 50–65 leads to the hump in expenditure inequality later in the life cycle.

The relatively high value of σ is consistent with the intuition provided in Section II. Specifically, as noted in the text, the intertemporal elasticity of substitution (IES) for the home-produced good of 0.3 is larger than the static elasticity of substitution in home production of one. This implies that expenditure will be positively correlated with the price of time (holding lifetime resources constant). The relatively low value of γ implies that core utility is more easily substituted across time (and responds more to permanent income shocks). The two parameters together are also consistent with standard intuition as follows. If we consider the weighted sum of the coefficients of risk aversion (weighted by average lifetime expenditure shares in the model of 0.64 for core and 0.36 for home production), the parameters imply an overall risk aversion parameter of

1.6, which is in line with most macro estimates and close to the 1.5 of the one-good model. On the other hand, if we take a weighted average of intertemporal elasticities (the inverses of γ and σ), the overall IES is 1.3, which is a little higher than the range of standard estimates.

As discussed above, the same parameters that govern the response of home production expenditure to wage movements also govern the response of market hours, reflecting the fact that home production accounts for all nonmarket hours. In particular, recall that the Frisch elasticity of nonmarket time is $[1 - \psi(1 - \sigma)]/\sigma$. Evaluated at the parameters reported in table 3, this elasticity is 0.39; or, at $n = 1/3$, this corresponds to a Frisch elasticity of labor of 0.78, which is in line with many empirical estimates of this parameter. This is perhaps surprising given that we are not using hours data, making this an overidentification diagnostic of the model. However, as the Frisch elasticity of labor is fairly small, the substitution effect of wage changes on expenditure is quite significant.

This trade-off has the following implications for our calibration. Early in the life cycle, home production/work-related expenditure is not fanning out in the data. The model replicates this with a low Frisch elasticity of labor supply, preventing a sharp increase in hours dispersion and the associated variance in home production expenditure. The relatively high curvature parameter also mitigates the income effect. The increase in variance after age 50 is then matched by letting retirement have a large substitution effect on home production expenditure. As explained above, the combined low Frisch elasticity and a large retirement effect on the variance are mutually consistent. However, quantitatively, they lead to an overstatement of the mean decline at retirement.[34]

Finally, the multigood model somewhat underpredicts the increase in the covariance of core and home-produced consumption (fig. A2). This is related to the above discussion regarding the retirement effect on mean expenditure. To see this, consider a log-linear approximation for the deviation of household i's total expenditure from the cross-sectional mean:

$$\ln C_{it} - \ln \overline{C}_t \approx s_t^k (\ln C_{it}^k - \ln \overline{C}_t^k),$$

where s_t^k is the share of good $k = \{\text{core, home production}\}$ in total consumption at age t. This implies that

$$\text{Var}(\ln C_{it}) \approx \sum_k (s_t^k)^2 \text{Var}(\ln C_{it}^k) + 2 s_t^k s_t^{k'} \text{Cov}(\ln C_{it}^k, \ln C_{it}^{k'}).$$

That is, the trend in the cross-sectional variance of log total expenditure is due to changes in the variances of disaggregated log expenditure, shifts in shares over the life cycle, and changes in the covariance across the disaggregated goods. The

[34] We should remark as well that $\underline{n}$ is 0.33, which is only fractionally below the target labor supply of prime-age workers of $1/3$. The mean labor supply in the simulation is 0.35, which overshoots the target. The high lower bound on hours keeps nonretirees from sharply dropping their hours later in the life cycle as wages begin to fall.

model overstates the decline in the share of home production expenditure, and given the low level of variance for food, this pushes up the increase in total expenditure, all else equal. The overstatement of the decline in home production expenditure also makes it difficult to match the steady increase in covariance between core and home production expenditures.

References

Aguiar, Mark, and Mark Bils. 2011. "Has Consumption Inequality Mirrored Income Inequality?" Working Paper no. 16807, NBER, Cambridge, MA.

Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *J.P.E.* 113:919–48.

———. 2007a. "Life-Cycle Prices and Production." *A.E.R.* 97:1533–59.

———. 2007b. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *Q.J.E.* 122:969–1006.

———. 2008. "Deconstructing Lifecycle Expenditure." Working Paper no. 13893, NBER, Cambridge, MA.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evidence." *J. Econ. Perspectives* 15:47–68.

Attanasio, Orazio, James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Lifetime Consumption." *J. Bus. and Econ. Statis.* 17:22–35.

Attanasio, Orazio, and Tullio Jappelli. 2000. "The Lifecycle Hypothesis and Consumption Inequality." Working Paper no. W97/17, Inst. Fiscal Studies, London.

Attanasio, Orazio, and Guglielmo Weber. 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *J.P.E.* 103:1121–57.

Banks, James, Richard Blundell, and Sarah Tanner. 1998. "Is There a Retirement Savings Puzzle?" *A.E.R.* 88:769–88.

Battistin, Erich, Agar Brugiavini, Enrico Rettore, and Guglielmo Weber. 2009. "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach." *A.E.R.* 99 (5): 2209–26.

Becker, Gary. 1965. "A Theory of the Allocation of Time." *Econ. J.* 75:493–517.

Bee, Adam, Bruce Meyer, and James Sullivan. 2012. "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?" Working Paper no. 18308, NBER, Cambridge, MA.

Bernheim, Douglas, Jonathan Skinner, and Steven Weinberg. 2001. "What Accounts for the Variation in Retirement Wealth among US Households?" *A.E.R.* 91:832–57.

Blundell, Richard, Martin Browning, and Costas Meghir. 1994. "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." *Rev. Econ. Studies* 61:57–80.

Browning, Martin, and Thomas Crossley. 2000. "Luxuries Are Easier to Postpone." *J.P.E.* 108:1022–26.

Browning, Martin, and Costas Meghir. 1991. "The Effects of Male and Female Labor Supply on Commodity Demands." *Econometrica* 59:925–91.

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

Cogan, John. 2001. "Fixed Costs and Labor Supply." *Econometrica* 49:945–63.

Cunha, Flavio, James Heckman, and Salvador Navarro. 2005. "Separating Uncertainty from Heterogeneity in Life Cycle Earnings." Working Paper no. 11024, NBER, Cambridge, MA.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59: 1221–48.

———. 1997. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy.* Baltimore: Johns Hopkins Univ. Press (for World Bank).

Deaton, Angus, and John Meullbauer. 1980. "An Almost Ideal Demand System." *A.E.R.* 70:312–26.

Deaton, Angus, and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *J.P.E.* 102:437–67.

DeNardi, Mariacristina, Eric French, and John Baily Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *J.P.E.* 118:37–75.

DeWeese, Gail, and Marjorie Norton. 1991. "Impact of Married Women's Employment on Individual Household Member Expenditures for Clothing." *J. Consumer Affairs* 25:235–58.

Diaz-Gimenez, Javier, Vincenzo Quadrini, and Jose-Victor Rios-Rull. 1997. "Dimensions of Inequality: Facts on the U.S. Distribution of Earnings, Income, and Wealth." *Fed. Reserve Bank Minneapolis Q. Rev.* 21 (2): 3–21.

Fernandez-Villaverde, Jesus, and Dirk Krueger. 2006. "Consumption over the Life Cycle: Some Facts from CEX Data." *Rev. Econ. and Statis.* 89:552–65.

Gourinchas, Pierre-Olivier, and Jonathan Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70:47–89.

Greenwood, Jeremy, Richard Rogerson, and Randall Wright. 1995. "Household Production in Real Business Cycle Theory." In *Frontiers of Business Cycle Research*, edited by Thomas Cooley. Princeton, NJ: Princeton Univ. Press.

Guvenen, Fatih. 2007. "An Empirical Investigation of Labor Income Processes." Working Paper no. 13394, NBER, Cambridge, MA.

Hall, Robert. 1968. "Technical Change and Capital from the Point of View of the Dual." *Rev. Econ. Studies* 35:35–46.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni Violante. 2010. "Unequal We Stand." *Rev. Econ. Dynamics* 13:15–51.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni Violante. 2005. "Two Views of Inequality over the Life Cycle." *J. European Econ. Assoc.* 3:765–75.

Heckman, James. 1974. "Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption over the Life Cycle." *A.E.R.* 64:188–94.

Hicks, J. R. 1939. *Value and Capital.* Oxford: Oxford Univ. Press.

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *A.E.R.* 101 (7): 2923–54.

Kaplan, Greg. 2012. "Inequality and the Lifecycle." *Quantitative Econ.* 3:471–525.

Lazear, Edward, and Robert Michael. 1980. "Real Income Equivalence among One-Earner and Two-Earner Families." *A.E.R. Papers and Proc.* 70:203–8.

Mincer, Jacob. 1962. "Labor Force Participation of Married Women." In *Aspects of Labor Economics*, 63–106. Princeton, NJ: Princeton Univ. Press (for NBER).

Nelson, Julie. 1989. "Individual Consumption within the Household: A Study of Expenditures on Clothing." *J. Consumer Affairs* 23:21–45.

Palumbo, Michael. 1999. "Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle." *Rev. Econ. Studies* 66:395–421.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing over the Lifecycle." *J. Monetary Econ.* 51:609–33.

Thurow, Lester. 1969. "The Optimum Lifetime Distribution of Consumption Expenditures." *A.E.R.* 59:282–97.