Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data*

Jesús Fernández-Villaverde
University of Pennsylvania

Dirk Krueger
University of Frankfurt, CEPR, and NBER

September 14, 2004

Abstract

This paper uses Consumer Expenditure Survey data and a seminonparametric statistical model to estimate life-cycle profiles of consumption, controlling for demographics, cohort, and time effects. We construct age profiles for total and nondurable consumption as well as expenditure patterns for consumer durables. Special emphasis is placed on the comparison of different approaches to control for changes in demographics over the life cycle. We find significant humps over the life cycle for total, nondurable, and durable expenditures. Changes in household size account for roughly half of these humps. Bootstrap simulations suggest that our empirical estimates are tight in that standard errors are small.

Keywords: Consumption, Life-Cycle Profiles, Durables, Nonparametric Estimation, Bootstrap, Partially Linear Models

JEL Classification: D12, D91, E21, C14, J10.

*All comments are welcomed to jesusfv@econ.upenn.edu or dirk.krueger@wiwi.uni-frankfurt.de. We would like to thank Andrew Atkeson, Michele Boldrin, Hal Cole, MariaCristina De Nardi, Thesia Garner, Narayana Kocherlakota, Lee Ohanian, Luigi Pistaferri, Edward Prescott, participants at seminars in numerous institutions, members of the Bureau of Labor Statistics, and especially Daron Acemoglu and two referees for many helpful comments, and the NSF under grants SES-0004376 and SES-0234267 for financial support. All remaining errors are our own.

1. Introduction

This paper uses Consumer Expenditure Survey (CEX) data to estimate life cycle profiles of consumption, controlling for demographics, cohort, and time effects. In addition to documenting profiles for total and nondurable consumption, we provide an age expenditure pattern for consumer durables.

Two reasons motivate our empirical study. First, we want to provide empirical life-cycle consumption profiles that can be used to assess the ability of quantitative life-cycle simulation models, pioneered by Auerbach and Kotlikoff (1987), to match the data. These models typically abstract from business cycle fluctuations, cohort effects, and differences in household size. Comparing model-generated life-cycle consumption patterns with their empirical counterparts therefore requires removing these effects. In this paper, special emphasis is placed on the comparison of different approaches to control for changes in demographics over the life cycle.

Second, we report life-cycle expenditure patterns for consumer durables, the most important asset in the median U.S. household's wealth portfolio. To the best of our knowledge, we are the first to exploit CEX data to construct these profiles, which can also be used to evaluate quantitative theoretical models that incorporate consumer durables. We undertake the analysis of nondurables and durables jointly since households' decisions to purchase nondurables and durables or to save in financial assets are intertwined by the period budget constraint. Furthermore, a household's ability to borrow may depend on its stock of consumer durables, pointing to further interdependence between the life-cycle profiles for nondurables and durables. By quantifying the size, timing, and correlation between the humps in nondurables and durables, our empirical results may shed light on which elements a successful model must possess to accurately account for the data.

Our main result is that total consumption expenditures as well as expenditures for nondurables and durables display a significant hump over the life cycle, even after accounting for changes in family size. If we measure the hump as the ratio of peak consumption to consumption at age 22, the size of the hump before demographic adjustment is roughly 1.6 and 1.3 thereafter. This finding is difficult to reconcile with the basic version of the life cycle model, augmented with consumer durables (i.e., separable utility between nondurables, durables, and leisure, no adjustment costs, and complete markets). The empirical evidence suggests the need for enriching this model with further elements, such as nonseparabilities in the utility function, different forms of adjustment costs or indivisibilities for consumer durables, or prudence in the light of idiosyncratic uncertainty.

Our paper builds on the sizable literature documenting empirical life-cycle consumption expenditure profiles, examples of which include, among many others, Carroll and Summers (1991), Carroll (1992), Deaton (1992), Kotlikoff (2001) and Gourinchas and Parker (2002). However, it offers the following new contributions. First, to the best of our knowledge, we are the first to employ the information on consumer durables from the Consumer Expenditure Survey to build life-cycle expenditure profiles of these items.

Second, we revisit the issue of controlling for family size and propose the use of household equivalence scales for this purpose. The recent contributions of Blundell *et al.* (1994), Attanasio and Browning (1995) and Attanasio *et al.* (1999) emphasize the importance of changes in household size to rationalize consumption expenditure profiles over the life cycle. These papers argue that demographics can explain, at least to a substantial degree, why consumption tracks income over the life cycle. Using household equivalence scales we find that demographics indeed play a large role, accounting for roughly half of the size of the hump in both expenditures on nondurables and durables. When employing household size adjustments implicitly estimated by other papers in the consumption life-cycle literature, we document that even with these alternative adjustment procedures, a sizable hump emerges.

Third, we control for cohort, time, and age effects in a flexible way by employing a seminonparametric partially linear model that imposes minimum conditions on the data. This procedure provides efficiency advantages in estimating age profiles compared to the use of dummy variables, yet is tractable and relatively straightforward to implement.

Finally, we perform bootstrap simulations to assess the precision of our estimates, an issue that has received little attention in the literature. We find that confidence intervals and bands are tight around our point estimates. This suggests that the hump cannot be explained purely by sampling uncertainty. In addition, an extensive sensitivity analysis shows that our main findings survive across a wide set of econometric specifications.

The rest of the paper is organized as follows. Section 2 describes the CEX data. Section 3 presents the specification of the estimated model of life-cycle consumption. It also explains in detail how we control for age, time and cohort effects, and for demographic changes. Section 4 discusses our empirical findings, with section 5 evaluating the precision of the estimates using the bootstrap. In section 6 we compare our results with those obtained employing alternative demographic adjustment procedures. Section 7 offers concluding remarks. Further details about the data, variable definitions, estimation, results, and robustness analysis are contained in a technical appendix, available at `www.econ.upenn.edu/~jesusfv/appen_consum.pdf`.

2. The CEX Data

We exploit the 1980-2001 Consumer Expenditure Survey, a widely used source of data on consumption expenditures (see Attanasio, 1998). We excluded the years 1982 and 1983 because of methodological differences in the survey. The CEX is a rotating panel. Each household is interviewed every three months over five calendar quarters, and every quarter 20 percent of the sample is replaced by new households. The CEX is designed to constitute a representative sample of the U.S. population, with a sample size of about 5000 households.

For the purpose of this paper, two issues with the way the CEX data are collected make it difficult to use them directly. First, the CEX records only data on consumption expenditure, and not on consumption services, our final object of interest. Second, the CEX lacks a significant panel dimension since it follows households for at most five quarters. In the remaining part of this section we discuss how we address both issues.

2.1. Expenditures versus Consumption

As mentioned before, the CEX does not report a measure of consumption services, arguably the object of strongest interest from the point of economic theory; it reports only *expenditures* on consumption goods. This distinction is not very relevant for the case of nondurable goods, but it is crucial when dealing with durables. For example, if the household buys a car today, it will receive a flow of transportation services over a long number of periods, despite the fact that expenditures are incurred (and show up in the data) only in the current period.

However, since data problems prevent us from reliably imputing services flows from information on the stock of consumer durables, we focus our analysis on expenditure data.[1] Quantitative life-cycle models that incorporate durables have predictions not only for service flows from durables but also for the timing of expenditures on these durables over the life cycle. Our results may serve as an empirical benchmark against which the predictions of these models can be evaluated.[2]

[1] The CEX provides only partial information for the value of the stock of durables. While the survey asks for an estimate of the current value of the owned residence and the original cost of vehicles, it only takes a physical, but not a value inventory, of major appliances owned by the household. The omission of these items may significantly underestimate the stock of durables for low-wealth households. Thus, since younger households tend to be wealth-poor, the omission may distort estimates of life-cycle consumption profiles for durables. Also, since we do not observe the initial stock of durables of the cohort and the sample length is small, we cannot use the perpetual inventory method to build estimates of the stock of consumer durables.

[2] In the technical appendix we exploit the information in the CEX on current values of owned residences, thus indirectly providing life cycle profiles of services from owned homes. If owner-occupied housing and other durables are complements, life-cycle profiles of housing services can serve as first approximations of profiles for other durables.

2.2. A Pseudopanel Approach

The second problem mentioned above is that the short panel dimension of the CEX makes the use of direct panel techniques problematic. An alternative is to exploit the repeated nature of the survey and build a pseudopanel (see Deaton, 1985). New households that enter the survey are a randomly chosen large sample of the U.S. population, and consequently, they contain information about the consumption means of the groups they belong to. This information can be used by interpreting the observed group means as a panel for estimation purposes. Also, a pseudopanel reduces the attrition problem, approximately averages out expectational errors and eliminates the need to control for individual effects since it aggregates across agents.

We use the age of the reference person to associate a household with a cohort. We define 10 cohorts with a length of five years, evaluate their means using CEX-provided population weights, and follow them through the sample to generate a balanced panel. Our choice of 10 cohorts is a compromise between the need for a large time series dimension to enrich the longitudinal aspect of the pseudopanel and the desire for a large cohort size to confidently assume that the sample means are good approximations for their population counterparts. Most of our cells have between 200 and 500 observations, on average, around 350.

3. Specification and Estimation of Life-Cycle Profiles

The most straightforward way to document consumption profiles over the life cycle is to use the pseudopanel to plot consumption against the age of the household's head. This simple procedure, however, faces two problems. First, households in these cohorts were born at different dates and may have experienced very different conditions during their lifetime. With positive long-run growth of real wages, for example, cohorts born at later dates have higher discounted lifetime earnings. Therefore, it is key to control for cohort effects. But even if we could observe one cohort over its entire life cycle, aggregate fluctuations would affect the cohort's consumption profiles. These effects should be attributed to time rather than aging. In subsection 3.1 we describe how to disentangle cohort and time effects from age effects, the primary object of interest of our analysis.

Second, the CEX reports consumption data for households and not for individuals. Since we want to provide empirical life-cycle consumption patterns to evaluate quantitative life-cycle models, which usually abstract from variations of household composition, it is crucial to separate changes in expenditures induced by changes in family size and changes induced by other factors. We describe in subsection 3.2 how we adjust the raw data for demographics.

3.1. Controlling for Age, Cohort, and Quarter Effects

We propose to relate age and consumption expenditures by a simple and flexible seminonparametric regression. In particular, we specify the partially linear model:

$$c_{it} = \pi_i cohort_i + \pi_t \gamma_t + m\left(age_{it}\right) + \varepsilon_{it} \tag{1}$$

where c_{it} is the cohort i average of log-consumption at time t, $cohort_i$ is a dummy for each cohort (except the youngest one), γ_t a dummy for each quarter, age_{it} is the age of cohort i at time t, measured in years, $m\left(age_{it}\right) = E\left(c_{it} | age_{it}\right)$ is a smooth function of age_{it}, and ε_{it} is an independent, zero mean, random error. The random term captures multiplicative measurement error in consumption expenditures (since the dependent variable is log-consumption) as well as unobserved cross-sectional heterogeneity.

This specification consists of two different components, a parametric part that includes cohort and quarter dummies, $\pi_i cohort_i + \pi_t \gamma_t$, and a nonparametric function of age, $m\left(age_{it}\right)$. This combination of a parametric and nonparametric approach achieves a satisfactory balance between flexibility and efficiency. A fully nonparametric approach is hopelessly inefficient in a small sample such as the pseudopanel from the size of the CEX. A pure parametric approach with age dummies, on the other hand, delivers a nonsmooth consumption profile that is difficult to use as an empirical benchmark. Furthermore, it is not robust to model misspecification problems.

We estimate the partially linear model using the two-step estimator proposed by Speckman (1988). This estimator combines ordinary least squares to estimate the parametric component with a standard kernel smoothing estimator to estimate the nonparametric component.[3]

Note that because time, age, and cohort effects are linearly dependent, it is not possible to separately identify them without further assumptions.[4] Following Deaton (1997), our identification scheme assumes that time effects are orthogonal to a time trend and that their sum is normalized to zero.

[3] Our estimator is described in detail in the technical appendix. The nonparametric component is estimated using a Nadaraya-Watson estimator with an Epanechnikov kernel. For our benchmark estimates we choose, based on cross-validation, a bandwidth parameter for the kernel of five years. We checked that the results are robust to this choice. Note that setting the bandwidth to one year is equivalent to estimating a model with age dummies. Thus, our model nests and generalizes this more traditional approach. The technical appendix also reports the estimation under this age-dummy specification and documents that the main results of the paper remain unchanged.

[4] Since we apply a nonlinear transformation to the age variable, time, cohort, and age are not perfectly collinear. However, these variables are so highly collinear that without further identifying restrictions, we would obtain extremely imprecise estimates.

3.2. Controlling for Family Size: Household Equivalence Scales

Households of different size plausibly face different marginal utilities from the same consumption expenditures. Economic theory predicts only that marginal utilities should be equated across time (up to some constant that depends of the discount factor and the interest rate) and not expenditures per se. Since household size displays a hump over the life cycle, the hump in consumption may largely be explained by changes in household composition, as argued in two influential papers by Attanasio and Weber (1995) and Attanasio *et al.* (1999).

Part of the objective of this paper is to quantify how much of the change in consumption over the life cycle is explained by demographics. We can attribute changes in expenditures on particular consumption items as a household ages either to changes in household size or to changes in consumption profiles. By using information on expenditure shares of households, one can construct household equivalence scales, which measure the change in consumption expenditures needed to keep the welfare of a family constant when its size varies.[5]

The simplest scale divides total expenditures by the number of household members to obtain *per capita* consumption. This scale therefore assumes that a household's technology to transform expenditures into consumption service flows exhibits constant returns to scale. Theory and evidence suggest otherwise. Lazear and Michael (1980) point to three mechanisms through which household size affects the rate of transformation between expenditures and services (family goods, economies of scale, and complementarities) and present data implying that their quantitative effects are important. Their findings suggest that more elaborate equivalence scales are needed to deflate household consumption expenditure by household size.

In this paper we borrow from rich previous work that exploits detailed information on expenditure shares to construct household equivalence scales. This literature documents, first, that economies of scale in household consumption exist. Second, opinions differ with respect to their size.

To summarize these differences we present a representative sample of household equivalence scales in table 3.2.1 (in which, for convenience, we assume the first two members of the household to be adults and the rest children). Columns 2 to 5 are based on expert evaluations, and columns 6 and 7 are econometric estimates based on observed choices.[6] Interestingly, the

[5]Early papers that deflate household consumption expenditure by a function of family size include Zeldes (1989), who adds adjusted food requirements as a regressor in some of his Euler equation estimates, and Blundell *et al.* (1994), who plot the life-cycle path of consumption, deflated by the number of adults plus 0.4 times the number of children in the household, for U.K. data.

[6]These are constructed, respectively, by the OECD (OECD, 1982), the Panel on Poverty and Family

two explicit econometric procedures deliver the biggest economies of scale.

Table 3.2.1: Different Household Equivalence Scales

Family Size	OECD	NAS	HHS	DOC	LM	Nelson	Mean
1	1	1	1	1	1	1	1
2	1.70	1.62	1.34	1.28	1.06	1.06	1.34
3	2.20	2.00	1.68	1.57	1.28	1.17	1.65
4	2.70	2.36	2.02	2.01	1.47	1.24	1.97
5	3.20	2.69	2.37	2.37	1.69	1.30	2.27

Since all reported estimates have advantages and drawbacks we choose their mean as our benchmark scale; it combines simplicity and a relatively conservative stand on the effect of household size. In section 6 we document that our main findings are quite robust to changes in the household equivalence scale.[7]

After choosing this equivalence scale, we take consumption expenditure measures C_{jt} from the CEX for household j at quarter t, use the demographic information of the household to obtain the equivalence scale es_{jt}, and adjust consumption to obtain $\hat{c}_{jt} = \log \frac{C_{jt}}{es_{jt}}$. Let $\tilde{c}_{it}$ denote the synthetic cohort i average of $\hat{c}_{jt}$, on which we then estimate the partially linear model

$$\widetilde{c}_{it} = \pi_i cohort_i + \pi_t \gamma_t + m\left(age_{it}\right) + \varepsilon_{it}. \tag{2}$$

4. Results

In this section we present the results of our estimation. First, we plot life-cycle profiles of total expenditure (figure 4.1), expenditures on nondurables (figure 4.2), and expenditure on

Assistance of the National Academy of Sciences (Citro and Michael, 1995), the Department of Health and Human Services (Federal Register, 1991), the Department of Commerce (U.S. Department of Commerce, 1991), Lazear and Michael (1980) and Nelson (1993). Since the latter estimates scales only for families of size 2 or higher, to complete the table we took 1.06 as the scale entry for households of size 2 from Lazear and Michael (1980). Beyond the results in the table, the literature presents a large number of alternative equivalence scales, such as Colosanto *et al.* (1984), Datzinger *et al.* (1984), Johnson and Garner (1995), Jorgenson and Slesnick (1987), Garner and de Vos (1995), and Phipps and Garner (1994). These scales stay within the bounds set by columns 2 and 8 of table 3.2.1.

[7] The use of equivalence scales to adjust for changes in household size is not free of problems. First, family size is endogenous. Second, a welfare analysis requires to infer *unconditional* preferences for a demographic structure and consumption, whereas usually only preferences for goods *conditional* on a particular demographic profile are studied. Ferreira *et al.* (1998) estimate a model that allows for endogenous choices in family size and obtain even larger economies of scale than Lazear and Michael (1980). Finally, for equivalence scales to be used successfully, they should not vary across household income levels. Pendakur (1999) finds that they satisfy this requirement.

durables (figure 4.3), controlling for cohort and time effects but *not* for family size. Total quarterly expenditures follow a clear hump; they increase from $3300 at the age of 22 to $5400 at the peak in the late forties, and decrease afterward. This pattern is well known and has been reported in the literature (see, e.g. the widely cited work by Carroll and Summers, 1991). More interestingly, similar humps appear if we separately plot nondurable consumption expenditure (figure 4.2) and expenditures on durables (figure 4.3) against the age of the household.[8] The hump in durables expenditures is, to the best of our knowledge, a fact that has not been documented before. Note that both humps, for durables and nondurables, are of similar magnitude (the increase from age 20 to the peak is around 80 percent) and that the peak occurs at the same stage in the life cycle, around the late forties.

To control for changes in demographics we now use the equivalence scale discussed in section 3 and repeat the estimation of life-cycle profiles. Figure 4.4 plots total expenditure against household age, with controls for cohort and quarter effects. Two main findings deserve comment. First, comparing figure 4.4 to figure 4.1, we see that the size of the hump, measured as the ratio between consumption at the peak and at the beginning of the life cycle, is reduced by about 50 percent. Nevertheless, a sizable hump remains: adjusted quarterly consumption increases from around $2550 to nearly $3300 and then decreases to about $1800. Also the peak in consumption is postponed, close to age 50.

The quarter effects are small, with the exception of significantly negative values in 1992 and significantly positive values for the quarters in 1984 and in 1997 and 1998. This pattern is plausible as it agrees with standard business cycle dating. The cohort effects are fairly small as well. Different reasons may explain this finding. The strong performance of the stock market during the last two decades, which especially profited older cohorts, may have compensated the long-run real wage growth advantage (which was less than stellar in the 1980's and early 1990's) of younger cohorts with smaller holdings of financial assets. Also the recent increases in the skill premium may have helped older (and more skilled cohorts) in comparison with (at the current point of their life cycle) less skilled, younger cohorts.

Figure 4.5 plots demographics-adjusted expenditures on nondurables. Consumption grows until the fifties and then declines, especially around retirement age, suggesting that some of these consumption expenditures are related to work. Comparing this figure to figure 4.2 we also observe a reduction of the hump of about 50 percent. Figure 4.6 plots expenditures on consumer durables: yet again, we observe a clear hump, although expenditures are already relatively high at the beginning of the adult life cycle, owing to first purchases of durable

[8] See the technical appendix for our definition of total expenditures, durables and nondurables expenditures.

goods. Interestingly enough, the reduction of the hump is quite similar to the case of nondurables. These figures show that both expenditures on durables and nondurables have very similar patterns and peak at the same time.

The results in figure 4.5 show that, even if changes in demographic composition of households can account for around half of the hump in nondurable consumption and hence are crucial in understanding life-cycle profiles, the other half remains to be explained by factors not present in the standard complete markets life-cycle model of consumption. The profile in figure 4.6 is even more difficult to reconcile with this textbook model when we augment it by including consumer durables. In addition to complete financial markets, suppose that utility is separable in nondurable consumption and services from durables, and that the real interest rate is equal to the time discount factor (which, as the depreciation rate, is constant over time). Then the optimal life cycle profile of consumer durables is to immediately build up the desired stock and to simply replace depreciation from there on. We do not see anything like this in the data; rather, the process of durables accumulation appears to be incremental over the life cycle.

Our profile for expenditure on durables is, however, consistent with papers that have documented liquidity constraints in the purchases of consumption durables (Alessie *et al.*, 1997, Attanasio *et al.*, 2000, Barrow and McGranahan, 2000, and Eberly, 1994) and with work arguing for the importance of nonseparabilities in the utility function (Attanasio and Weber, 1995).[9]

To further investigate whether liquidity constraints may play a role in generating the humps in figures 4.4 to 4.6, we construct consumption profiles separately for different educational groups. We report the profiles for total expenditure in figure 4.7, where to enhance comparability we have normalized the profiles at 1 at age 22. We observe that for low-education households (high school degree or less), there still emerges a hump, although its size is smaller than in our full sample benchmark. For high-education (at least some college) households, the profile shows the opposite features: now the hump is bigger in size than for the full sample. Due to space constraints we report the graphs for nondurables and durables in the technical appendix. It is interesting to note that the hump for expenditures on durables disappears after demographic adjustment for the low-education group.

These results suggest that liquidity constraints may play an important role in shaping life-cycle consumption profiles. Since high-education households have steeper income profiles

[9]The evidence is also qualitatively consistent with the importance of in-kind intergenerational transfers of durables. However, data from the Health and Retirement Study suggest that these transfers are fairly small (see Cardia and Ng, 2000).

as documented in Attanasio *et al.* (1999), in the presence of liquidity constraints their consumption profile is expected to track income and be steeper as well. Our results provide suggestive, albeit indirect (and subject to several qualifications), evidence for the presence and importance of liquidity constraints.

To investigate whether nonseparabilities between consumption and leisure may explain part of the hump, in the technical appendix we compute life-cycle profiles of hours worked in the market sector and discuss their correlation with consumption. We find that nonseparabilities may explain part of the hump, but that this explanation faces several problems to *quantitatively* account for the size and the timing of the hump.

Finally, note that the presence of a hump is robust to further breakdowns of expenditures. For instance, even when plotting adult equivalent food consumption (a necessary good for which a higher degree of smoothing would be expected), we see a hump.[10]

5. Using the Bootstrap to Evaluate Sampling Uncertainty

Since we want to provide empirical life-cycle consumption profiles that can serve as a useful benchmark for quantitative work, our profiles should be precisely estimated. To assess this precision we use the bootstrap. Even though under relatively mild technical conditions the Speckman estimator is consistent and asymptotically normal, its small sample properties tend to be better reflected by the bootstrap than by asymptotic approximations.[11] This is especially true at both ends of the age profiles because of the low number of observations. We implement the bootstrap as follows. We build 500 new pseudopanels sampling with replacement from the CEX and using the weights provided by the survey. Then we apply the Speckman estimator to each of these new data sets.[12]

Figure 5.1a plots the 95 percent confidence interval for the age profile of adult-equivalent total expenditures, controlling for cohort and quarter effects. The size of the interval indicates

[10] Studying food consumption is interesting because it allows comparison with data from the Panel Study of Income Dynamics (PSID), a survey with a long panel dimension. This comparison is performed in Fisher and Johnson (2002), who show that data on food consumption from the PSID and the CEX agree on the presence and quantitative size of a hump over the life cycle. The technical appendix offers further information, including a discussion of the role of housing.

[11] The kernel estimates converge more slowly than $n^{-\frac{1}{2}}$ and the asymptotic distributions have unconventional expansions that are not powers of $n^{-\frac{1}{2}}$, making their use in finite samples difficult (Hall, 1992a).

[12] Since the small sample bias of the kernel estimator distorts the fitted values of the new estimates and therefore transmits the bias to the computed standard error, Hall (1992b) suggests choosing a new smoothing parameter h that implies undersmoothing relative to the point estimate. We carry out this bias-removal strategy with an undersmoothing factor of 0.8. Note that the resulting confidence intervals and bands will not be centered on the point estimates because of this undersmoothing.

that the age profile is precisely estimated. Figure 5.1b plots the widest confidence interval computed from all the bootstrap replications, i.e., the worst possible case generated in the 500 simulations. The most interesting figure, however, is 5.1c, which shows a 95 percent confidence band that covers the *whole* true curve (instead of each point separately, as in a confidence interval). Since any curve that can be plotted entirely inside this small band implies a significant hump, figure 5.1c strongly reinforces our confidence in the point estimates: the data indicate a hump in consumption of nondurables, with size between 20 percent and 65 percent and a peak between ages 45 and 50. Finally, figure 5.1d plots all 500 simulated profiles: without exception, all simulations generate a quantitatively significant hump in consumption life-cycle profiles. Similar results are reported in the technical appendix for expenditures on nondurables and durables. In all cases, the bootstrap strongly suggests that our findings are not merely a result of sampling uncertainty.

6. Comparison with Alternative Procedures

Controlling for changes in household size reduces the consumption hump by 50 percent. We now ask how robust this decomposition is and if alternative procedures proposed in the literature result in a complete removal of the hump via changes in household composition.

First we argue that the equivalence scale we use does matter for the size of the hump in consumption over the life cycle, but not for its existence. For this we estimated life-cycle profiles using all scales reported in table 3.2.1 and we always found a hump. A sample of the results is plotted in figure 6.1. It shows the life-cycle profile of consumption for three equivalence scales: the two extremes in our table 3.2.1 (OECD and Nelson) and the mean equivalence scale we use for our benchmark results. To facilitate the comparison of results, we have normalized all three consumption profiles to one at the beginning of the life cycle. We observe that all scales deliver a hump, with its size depending on the extent of economies of scale in household consumption. In fact, even *per capita* consumption has a hump. Demographics completely eliminate the hump *only* with decreasing returns to scale in household production, an assumption that contradicts all empirical evidence we are aware of and our intuition.[13]

An entirely different approach to controlling for demographics taken in some of the literature is to estimate an Euler equation for consumption, augmented by demographic regressors.

[13] The results are also robust to the use of different equivalence scales for durables and nondurables. For example, for durables expenditures the use of the scales on owned housing presented by Garner and Short (2001) and Nelson (1988) delivers even more pronounced humps.

The point estimates of the regressors define an implicit equivalence scale. Suppose we apply these implicit equivalence scales on our data set. How do the results compare to our benchmark profiles? Because of space constraints, we only implement the scales implicit in two important papers in this tradition: Attanasio *et al.* (1999) and Browning and Ejrnæs (2002).

Using Attanasio *et al.*'s (1999) point estimates for demographics, their implicit equivalence scale, for the case where the first two members of the household are adults and the rest are children of age less than 16, is $\{1, 1.57, 1.80, 2.04, 2.28\}$. Of course, different family structures lead to alternative equivalence scales. This scale is quite similar to the one we employed, although ours implies more economies of scale for couples: 1.34 versus 1.57 (remember the interpretation of household equivalence scales: two persons need \$1.34 to obtain the same level of utility as one person living alone with \$1). For bigger families both equivalence scales indicate roughly the same magnitude of economies of scales.

This similarity is reflected in figure 6.2, where we plot consumption life-cycle profiles for our benchmark scale and for Attanasio *et al.*'s scale, using our econometric methodology. To enhance comparability we have normalized both profiles so that consumption at age 22 equals 1. For both scales we observe a significant hump in life-cycle consumption, although the size of the hump is somewhat smaller (about 40 percent smaller) with the scale implied by Attanasio *et al.*'s results. The ratio of peak consumption to that of age 22 is 1.3 with our scale and 1.18 with Attanasio *et al.*'s scale. It is remarkable that during the first 15 years of the life cycle, both profiles are nearly identical and only diverge after age 36, when singles (for which both scales are equivalent) become rarer.

Consequently, employing Attanasio *et al.*'s empirical strategy results in a sizable, albeit somewhat smaller hump. This result is not at all at odds with their paper, since the authors are explicit about the fact that they require income uncertainty and precautionary savings motives, *in addition* to demographics, to match the data.

The second paper we discuss is Browning and Ejrnæs' (2002), who argue, using U.K. data, that controlling for the age of the children in the equivalence scale is crucial to account for the hump on consumption. Repeating Browning and Ejrnæs' procedure for U.S. data or even simply applying their estimated scales is impossible because the CEX data files report ages of children in two different bins only (children of age less than 2, children of ages between ages 2 and 15). All others members of the household are considered adults. Given these data limitations we approximate Browning and Ejrnæs' procedure in a crude but implementable way. We assume that children between ages 2 and 15 are equivalent to 0.2 adults, and from age 16 on they are equivalent to one adult. Infants of age less than 2 count as zero adults.

With our new equivalence scale, which differentiates children by their age, we re-estimate our partially linear model. Figure 6.2 presents the resulting consumption profile, in conjunction with the profile obtained with our benchmark scale. Both profiles are qualitatively similar in that they display a significant hump in life-cycle consumption, but the size of the hump is 50 percent smaller with the child-age adjusted scale. The two profiles start to diverge at around age 45, the point of the life cycle at which households with children of ages 16 and older start to dominate the sample. Our approximation of the equivalence scale of Browning and Ejrnæs features a drastic increase from 2.12 to 2.57 for a household with two adults and a child at exactly the time when the child turns 16. This suggests that the divergence of the two graphs is partially due to our imperfect implementation of their approach since for data reasons we have no choice but to treat all household members 16 years and older as adults.[14]

7. Conclusions

In this paper we document the life-cycle profiles of consumption, with special emphasis on the distinction between expenditures on durables and nondurables. We find that both expenditures on nondurables and durables have a sizeable hump, roughly 50 percent of which is accounted for by changes in household demographics. The other half remains to be explained by factors not present in the standard complete markets life-cycle model of consumption, one of the main workhorses of modern macroeconomics. The failure of this textbook model is especially serious for expenditures on durables. Instead of immediately building up their stock of durables and then just compensating for depreciation, households in our data continue to increase expenditures until quite late in their life cycles.

A number of possible deviations from the basic life-cycle model may qualitatively account for the humps documented in this paper. First, one may relax the assumption of separability between leisure and consumption. A second departure is the introduction of uninsurable idiosyncratic uncertainty (e.g., with respect to labor income or lifetime horizon) into a model where households are prudent. Finally, one may argue for the importance of liquidity constraints that prevent young households from borrowing against future (higher) labor income to finance higher current consumption. These features, in conjunction with nonconvex adjustment costs and indivisibilities for consumer durables, may help to rationalize the empirical consumption profiles we have documented in this paper.

[14] The technical appendix offers further details on the advantages and drawbacks of the regression approach as well as the results of using alternative demographic adjustment methods.

Given the similar timing and size in the humps for expenditures on nondurables and durables a successful model will likely incorporate consumer durables into standard consumption models for nondurables. Examples of attempts to construct those models and to derive their quantitative implications include Díaz and Luengo-Prado (2002), Fernández-Villaverde and Krueger (2002), and Laibson *et al.* (2001).

References

[1] Alessie, R., M. Devereux and G. Weber (1997), "Intertemporal Consumption, Durables and Liquidity Constraints: A Cohort Analysis". *European Economic Review* 41, 37-59.

[2] Attanasio, O. (1998), "Cohort Analysis of Saving Behavior by U.S. Households". *Journal of Human Resources* 33, 575-609.

[3] Attanasio, O.. and M. Browning (1995), "Consumption over the Life Cycle and over the Business Cycle". *American Economic Review* 85, 1118-1137.

[4] Attanasio, O. and G. Weber (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence for the Consumer Expenditure Survey". *Journal of Political Economy* 103, 1121-1157.

[5] Attanasio, O., J. Banks, C. Meghir and G. Weber (1999), "Humps and Bumps in Lifetime Consumption". *Journal of Business and Economics* 17, 22-35.

[6] Attanasio, O., P. Goldberg and E. Kyriazidou (2000), "Credit Constraints in the Market for Consumer Durables: Evidence from Microdata on Car Loans". *NBER Working Paper* 7694.

[7] Auerbach, A. and L. Kotlikoff (1987), *Dynamic Fiscal Policy,* Cambridge University Press.

[8] Barrow, L. and L. McGranahan (2000), "The Effects of the Earned Income Credit on the Seasonality of Household Expenditures". *National Tax Journal,* forthcoming.

[9] Blundell, R., M. Browning and C. Meghir (1994), "Consumer Demand and the Life-Cycle Allocation of Household Expenditures". *Review of Economic Studies* 61, 57-80.

[10] Browning, M. and M. Ejrnæs (2002), "Consumption and Children". *Mimeo*, Institute of Economics, University of Copenhagen.

[11] Cardia, E. and S. Ng (2000), "How Important are Intergenerational Transfers of Time? A Macroeconomic Analysis". *Mimeo*, John Hopkins University.

[12] Carroll, C. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence". *Brookings Papers on Economic Activity*, 23, 61-156.

[13] Carroll, C. and L. Summers (1991), "Consumption Growth Parallels Income Growth: Some New Evidence" in D. Bernheim and J. Shoven (eds), *National Saving and Economic Performance*, University of Chicago Press for the NBER.

[14] Citro, C. and R. Michael (1995), *Measuring Poverty: A New Approach.* National Academy Press.

[15] Colosanto, D., A. Kapteyn, J. van der Gaag (1984), "Two Subjective Measures of Poverty. Results from the Wisconsin Basic Needs Survey". *Journal of Human Resources* 19, 127-137.

[16] Datzinger S, van der Gaag, M. Taussig and E. Smolensky (1994), "The Direct Measurement of Welfare Levels: How Much Does it Cost to Make Ends Meet?". *Review of Economics and Statistics*, 66, 500-505.

[17] Deaton, A. (1985), "Panel Data from Time Series of Cross-Sections". *Journal of Econometrics* 30, 109-126.

[18] Deaton, A. (1992), *Understanding Consumption*. Oxford University Press.

[19] Deaton. A. (1997), *The Analysis of Households Surveys*. John Hopkins University.

[20] Díaz, A. and M. Luengo-Prado (2002), "Durable Goods and the Wealth Distribution". *Mimeo*, Universidad Carlos III.

[21] Eberly, J. (1994), "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases". *Journal of Political Economy* 102, 403-436.

[22] Federal Register (1991), 56, n 34.

[23] Fernández-Villaverde, J. and D. Krueger (2002), "Consumption and Saving over the Life Cycle: How Important are Consumer Durables". Mimeo, University of Pennsylvania. Available at `www.econ.upenn.edu/~jesusfv`

[24] Ferreira, M., R. Buse and J. Chavas (1998), "Is There Bias in Computing Household Equivalence Measures". *Review of Income and Wealth*, 44, 183-198.

[25] Fisher, J. and D. Johnson (2002), "Consumption Mobility in the United States: Evidence from Two Panel Data Sets". *Mimeo*, Bureau of Labor Statistics.

[26] Garner, T. and K. de Vos (1995), "Income Sufficiency vs. Poverty: Results from the United States and The Netherlands". *Journal of Population Economics* 8, 117-134.

[27] Garner, T. and K. Short (2001), "Owner-Occupied Shelter in Experimental Poverty Measurement with a 'Look' at Inequality and Poverty Rates". *Mimeo*, Bureau of Labor Statistics.

[28] Gourinchas, P. and J. Parker (2002), "Consumption over the Life Cycle". *Econometrica* 70, 47-89.

[29] Hall, P. (1992a), *The Bootstrap and Edgeworth Expansion*. Springer-Verlag.

[30] Hall, P. (1992b), "Effect of Bias Estimation on Coverage Accuracy of Bootstrap Confidence Intervals for a Probability Density". *Annals of Statistics* 20, 675-694.

[31] Johnson, D. and T. Garner (1995), "Unique Equivalence Scales: Estimation and Implications for Distributional Analysis". *Journal of Income Distribution* 4, 215-234.

[32] Jorgenson, D. and D. Slesnick (1987), "Aggregate Consumer Behavior and Household Equivalence Scales". *Journal of Business Economics and Statistics* 5, 219-232.

[33] Kotlikoff, L. (2001), *Essays on Saving, Bequest, Altruism and Life-Cycle Planning*. MIT Press.

[34] Laibson, D., A. Repetto and J. Tobacman (2001), "A Debt Puzzle". *Mimeo*, Harvard University.

[35] Lazear, E. and Michael, R. (1980), "Family Size and the Distribution of Real Per Capita Income". *American Economic Review* 70, 91-107.

[36] Nelson, J. (1988), "Household Economies of Scale in Consumption: Theory and Evidence," *Econometrica* 56, 1301-1314.

[37] Nelson, J. (1993), "Independent of a Base Equivalence Scales Estimation Using United States Micro-Data". *Annales d'Économie et de Statistique* 29, 43-62.

[38] Organization for Economic Cooperation and Development (1982). *The OECD List of Social Indicators.*

[39] Pendakur, K. (1999), "Semiparametric Estimates and Test of Base-Independent Equivalence Scales". *Journal of Econometrics* 88, 1-40.

[40] Phipps, S. and T. Garner (1994), "Are Equivalent Scales the Same for the United States and Canada?". *Review of Income and Wealth* 40, 1-17.

[41] Speckman, P. (1988), "Kernel Smoothing in Partial Linear Models". *Journal of the Royal Statistical Society B* 50, 413-436.

[42] U.S. Department of Commerce (1991). *Trends in Relative Income: 1964 to 1989.* Series P60-177.

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

Figure 4.1: Total Expenditure

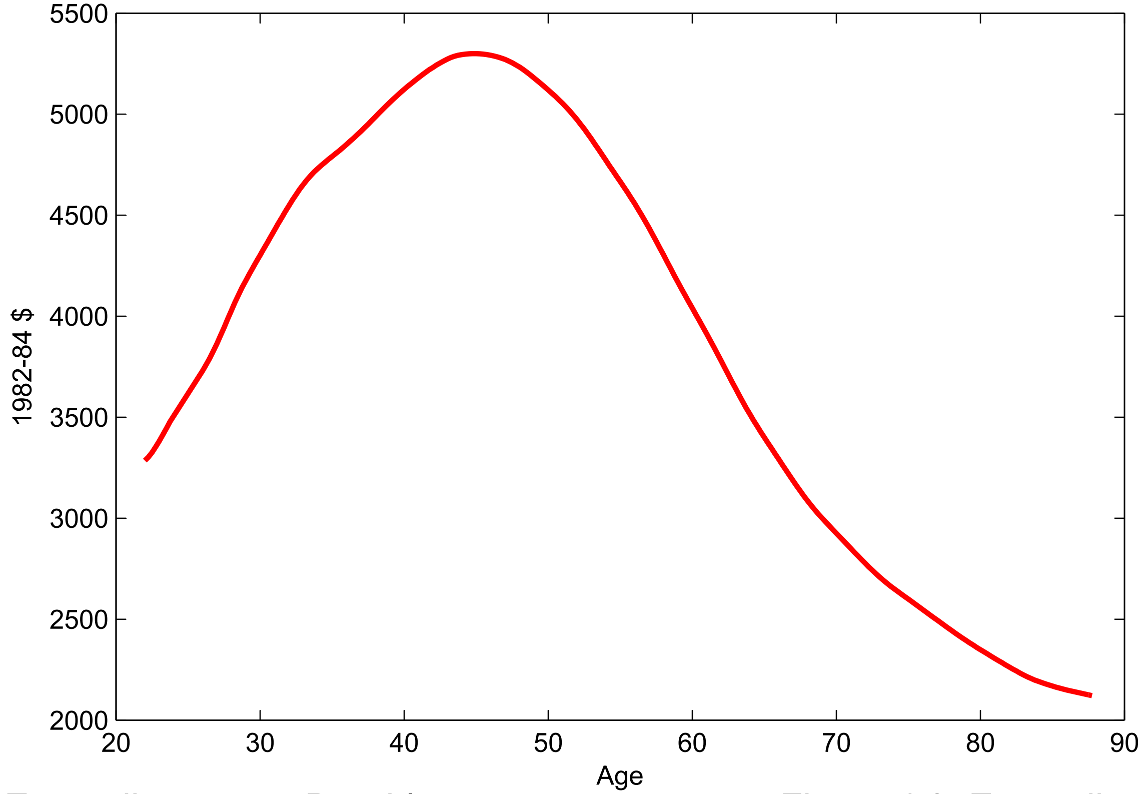

Figure 4.2: Expenditures non Durables

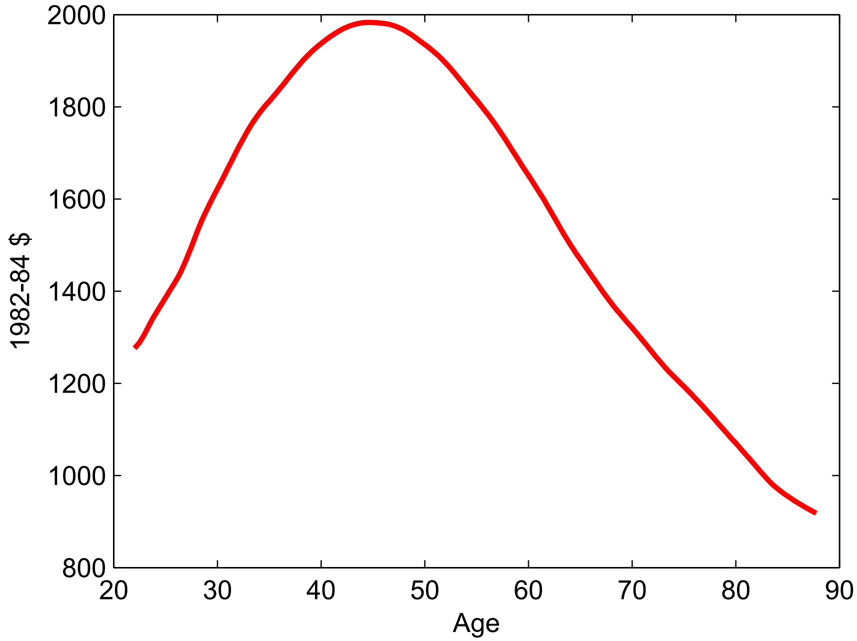

Figure 4.3: Expenditures Durables

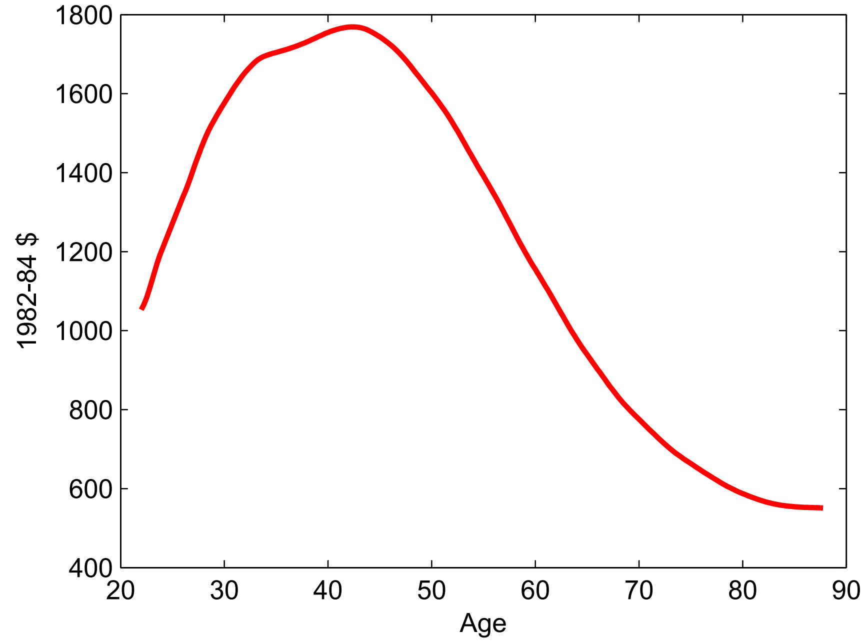

Figure 4.4: Total Expenditure, Adult Equivalent

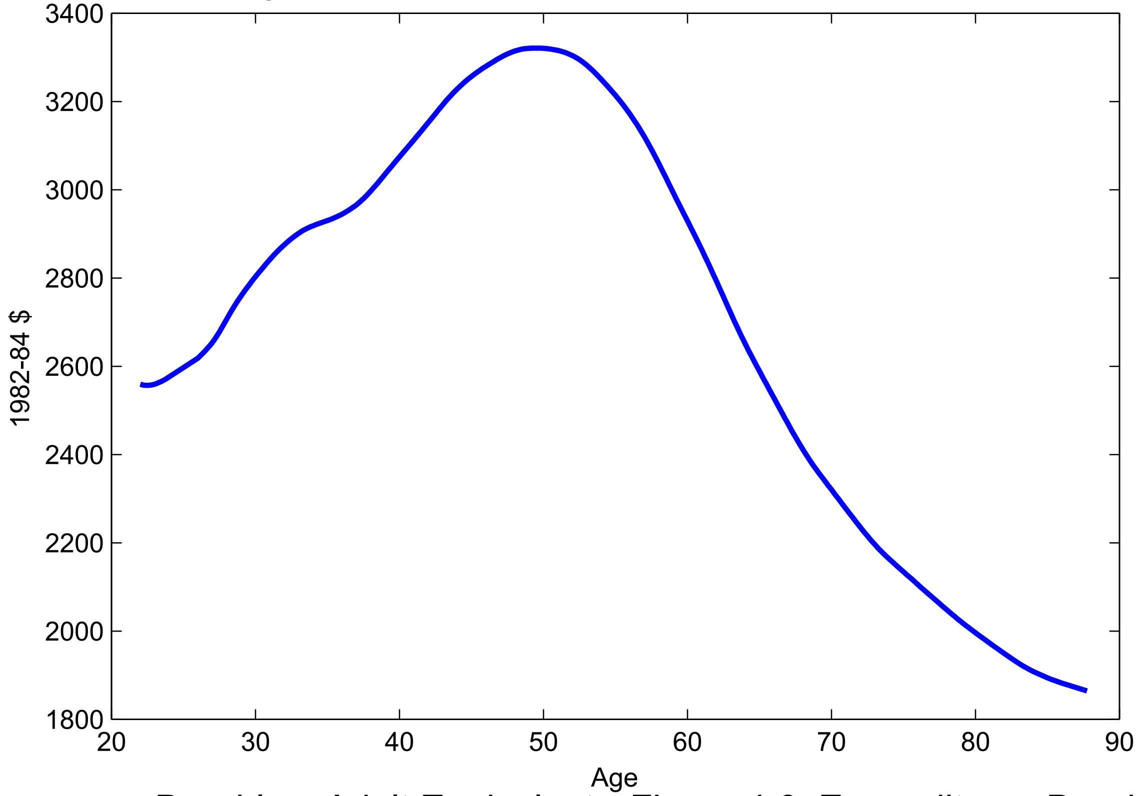

Figure 4.5: Expenditures non Durables, Adult Equivalent

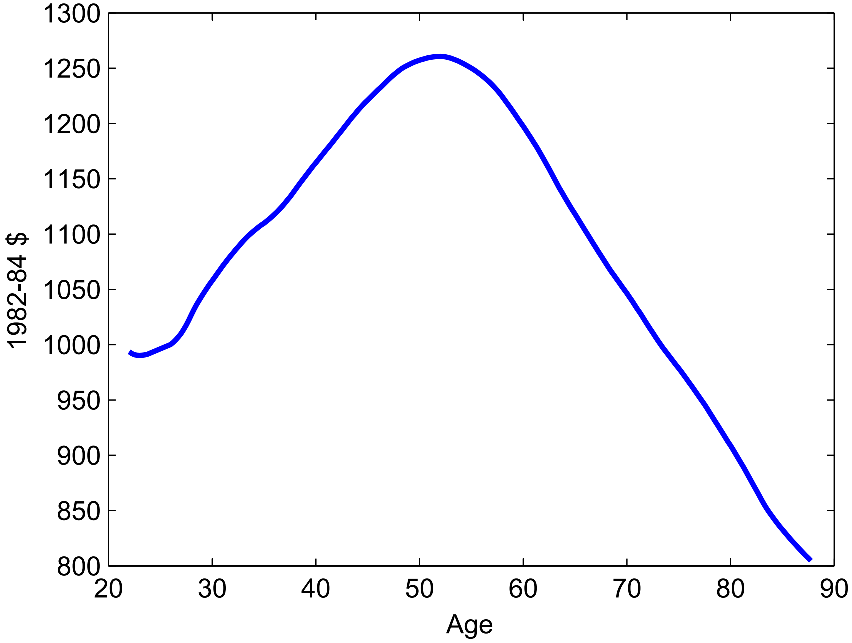

Figure 4.6: Expenditures Durables, Adult Equivalent

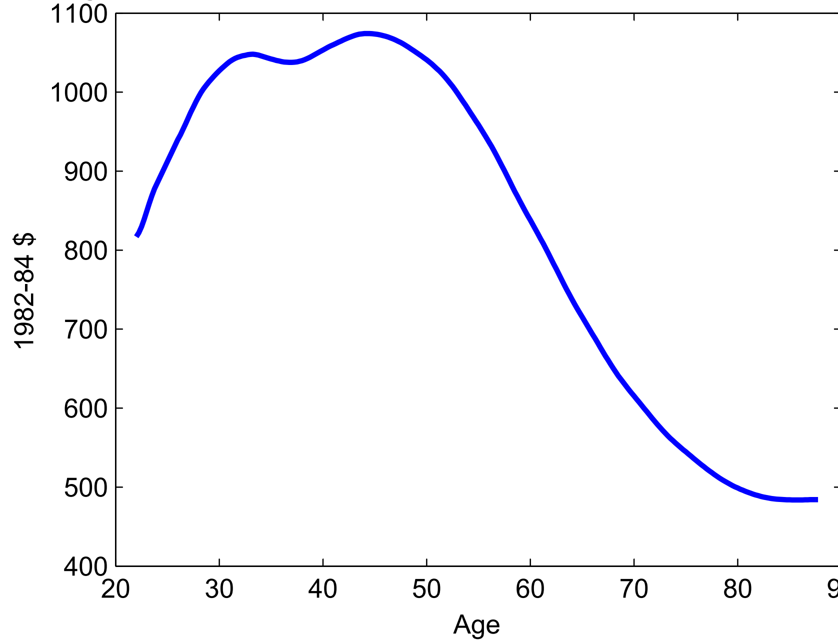

Figure 4.7: Total Expenditure, Adult Equivalent, by Education Groups

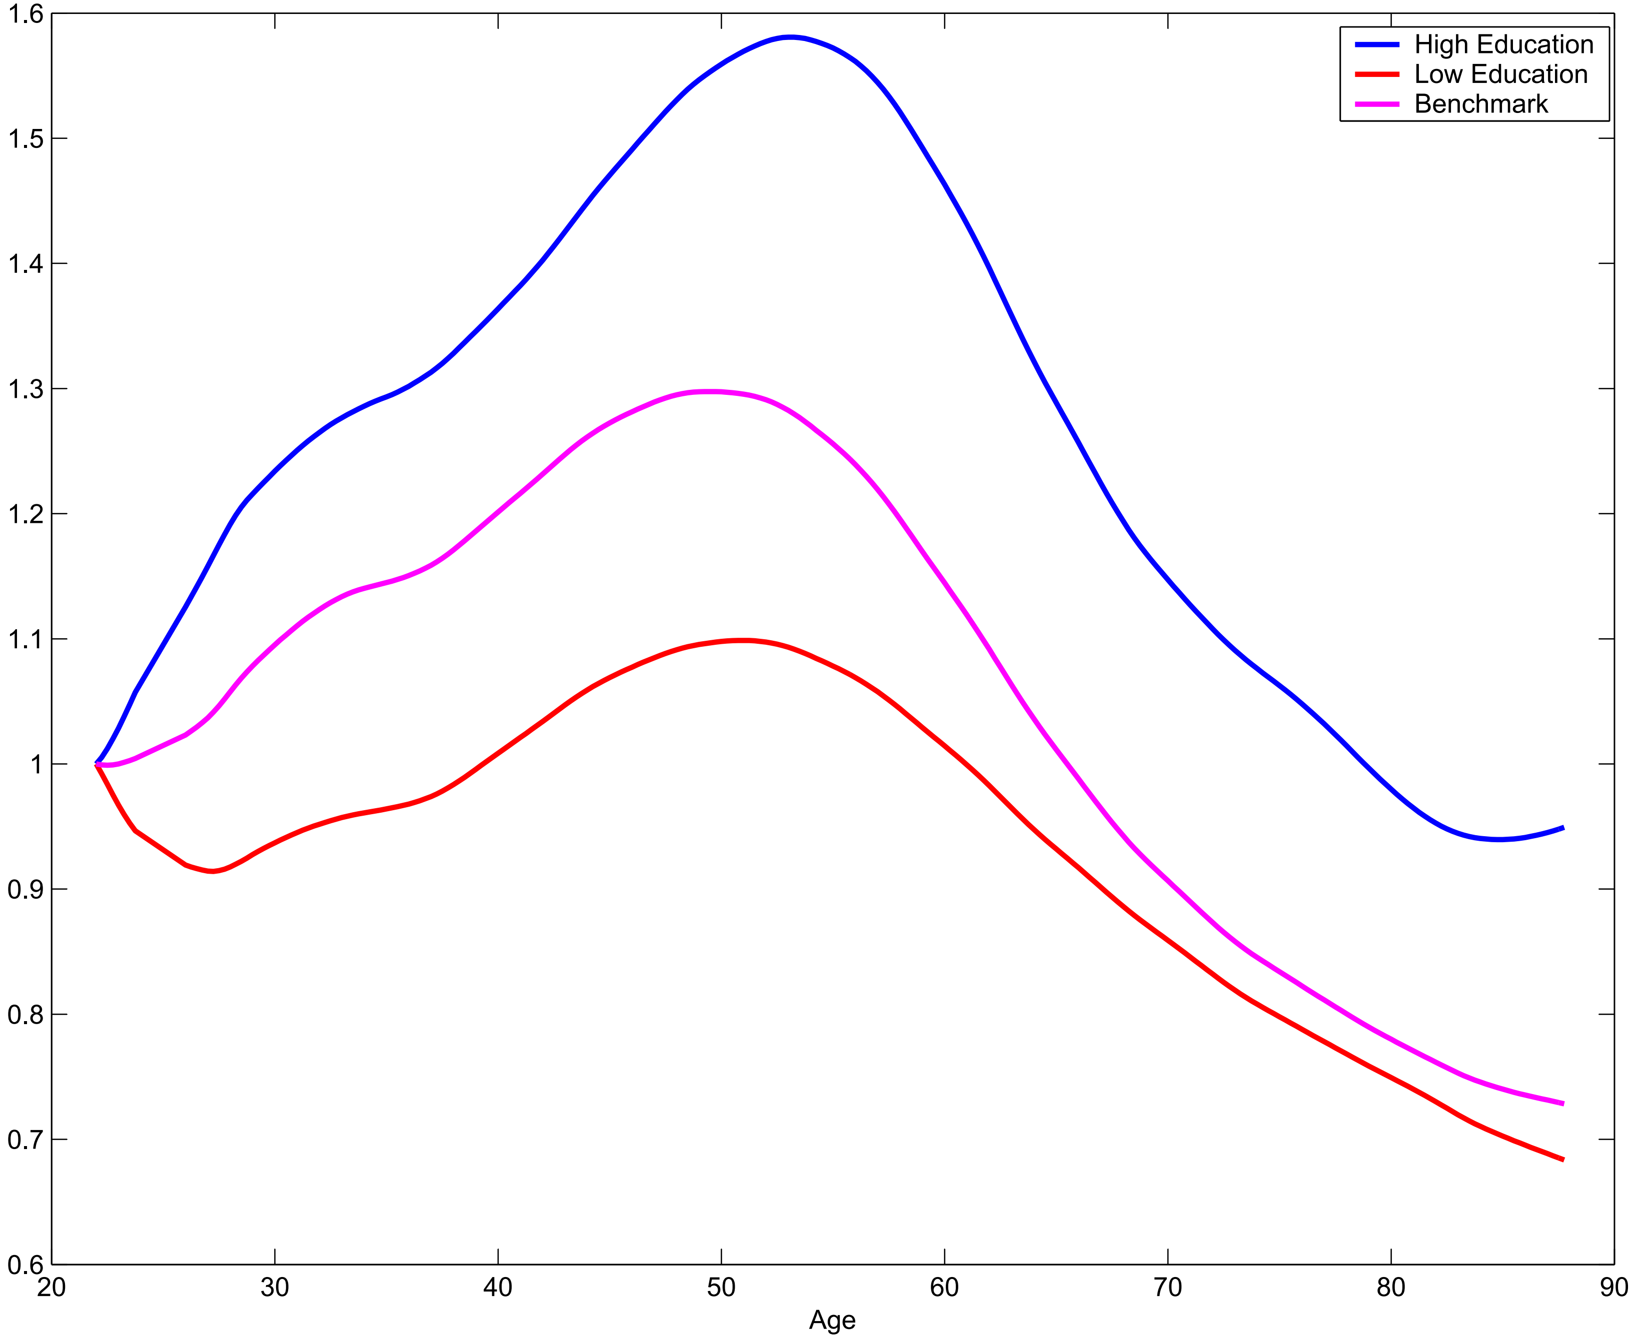

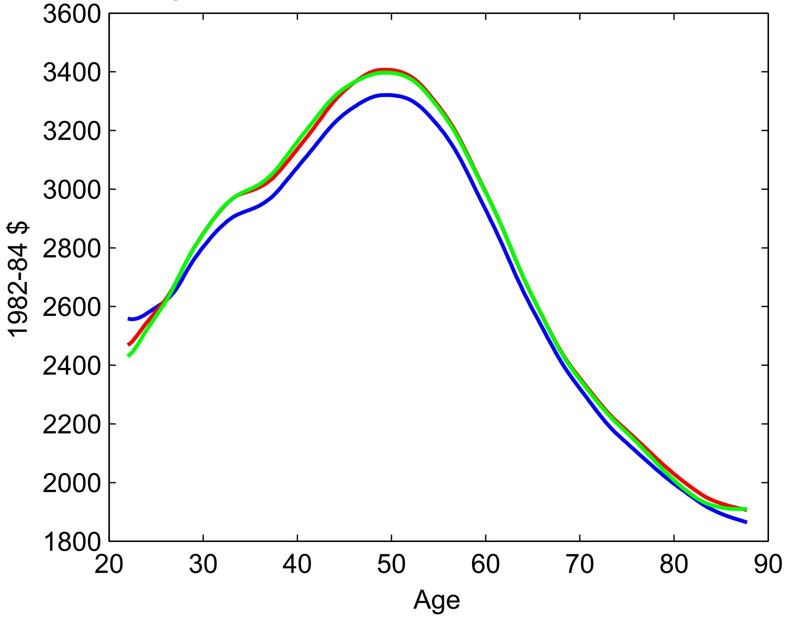

Figure 5.1a: 95% confidence interval

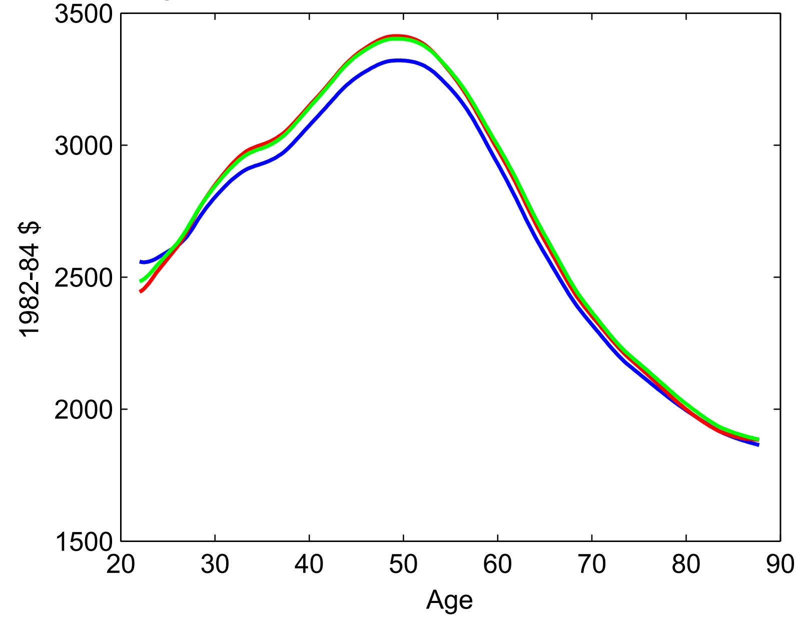

Figure 5.1b: Widest confidence interval

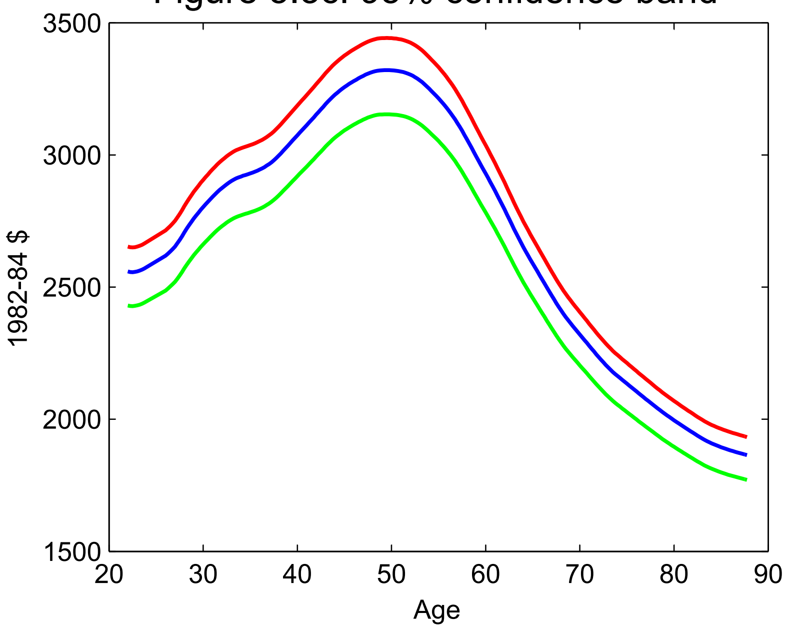

Figure 5.3c: 95% confidence band

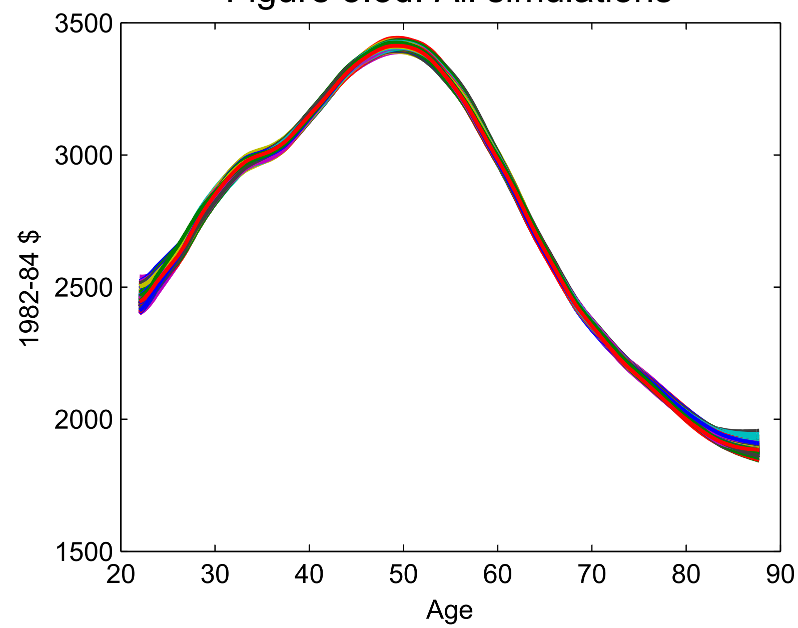

Figure 5.3d: All simulations

Figure 6.1: Comparison of Different Equivalence Scales I

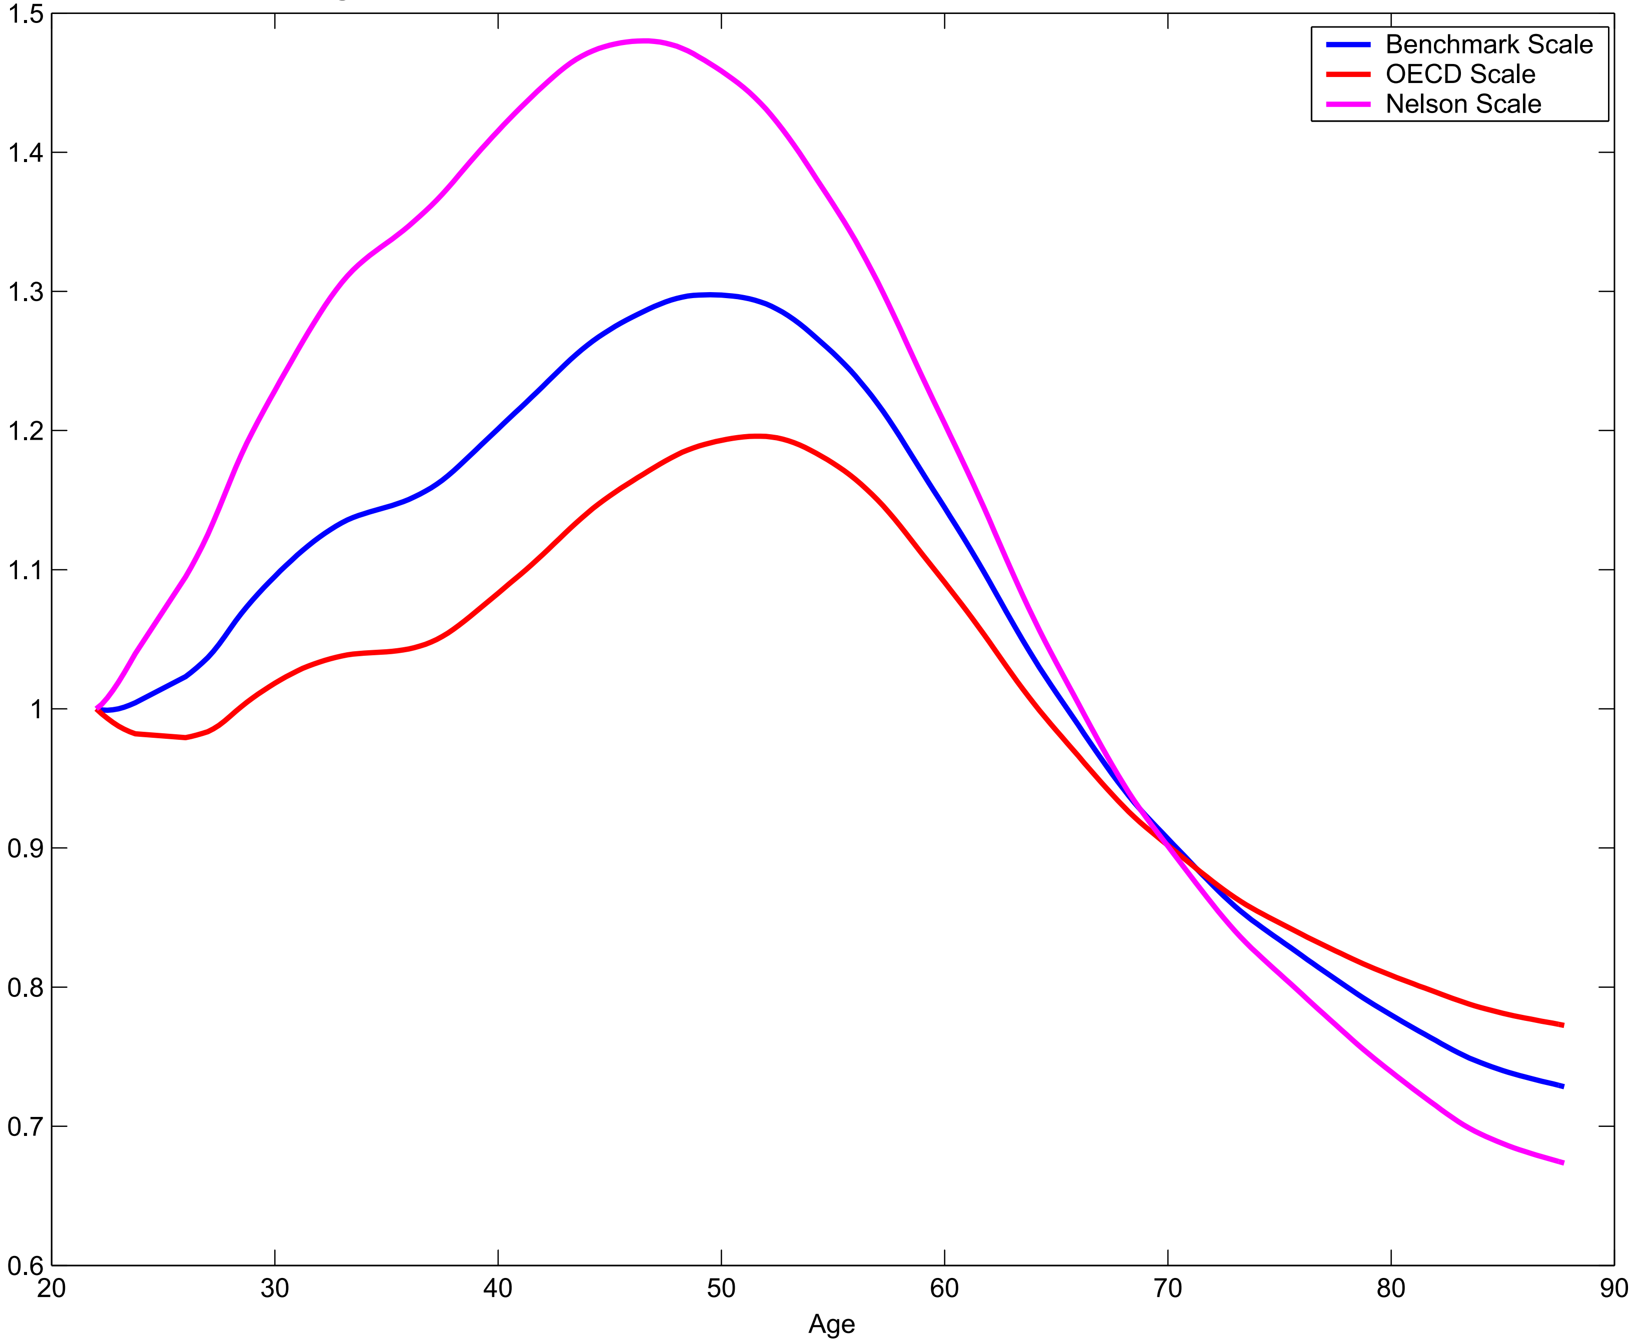

Figure 6.2: Comparison of Different Equivalence Scales II