

Centro de Altísimos Estudios Ríos Pérez

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

Jesús Fernández-Villaverde

Dirk Krueger

CAERP

Documento de Trabajo #7

Working Paper #7

Consumption over the Life Cycle:

Some Facts from Consumer Expenditure Survey Data*

Jesús Fernández-Villaverde

Dirk Krueger

University of Pennsylvania

Stanford University and NBER

November 20, 2002

Abstract

This paper uses a seminonparametric model and Consumer Expenditure Survey data to estimate life cycle profiles of consumption, controlling for demographics, cohort and time effects. In addition to documenting profiles for total and non-durable consumption, we devote special attention to the age expenditure pattern for consumer durables. We find hump-shaped paths over the life cycle for total, for nondurable and for durable expenditures. Changes in household size account for roughly half of these humps. The other half remains unaccounted for by the standard complete markets life cycle model. Our results imply that households do not smooth consumption over their lifetimes. This is especially true for services from consumer durables. Bootstrap simulations suggest that our empirical estimates are tight and sensitivity analysis indicates that the computed profiles are robust to a large number of different specifications.

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 dkrueger@leland.stanford.edu. We would like to thank Andrew Atkeson, Michele Boldrin, Hal Cole, MariaCristina De Nardi, Thesia Garner, Narayana Kocherlakota, Lee Ohanian, Luigi Pistaferri, Edward Prescott and participants at seminars in numerous institutions and especially at the Bureau of Labor Statistics for many helpful comments. All remaining errors are our own.

1. Introduction

This paper uses a semiparametric model and 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 devote special attention to the age expenditure pattern for consumer durables.

Our main finding is that consumption expenditures for both nondurables as well as durables display a significant hump over the life cycle, even after accounting for changes in family size. This finding is clearly at odds with the basic predictions of the standard life cycle model¹. For instance, according to this model consumption expenditure profiles for nondurables should be smooth. If utility is equally desirable over time up to some discount factor, households want to equate marginal utilities across time and states of the world, possibly with some growth rate, depending on the interest rate and the discount factor. With isoelastic period utility, a natural benchmark to match long-run growth trends, consumption growth itself should be constant across time. This is not at all what we observe in the data.

Our empirical findings are even more strikingly at odds with standard theory in the case of expenditures on consumer durables. If consumption services from durables are proportional to their stock, separable from nondurables in the utility function, and interest rates are constant and equal to the time discount factor, a household would immediately build up its desired stock of durables. Further expenditures occur only to replace depreciation. As with nondurables, under this benchmark model one does not expect to observe the hump in expenditures we document empirically.

In order to assure that the estimated consumption life cycle profiles are not just artifacts of sampling uncertainty or due to a particular choice of our econometric model, we perform bootstrap simulations to assess the precision of our estimates and a number of robustness checks. 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, our sensitivity analysis shows that our main findings survive across a wide set of econometric specifications.

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

¹By the standard model we mean a model in which households have utility that is separable in nondurable consumption, service flows from durables and leisure, and face no idiosyncratic shocks (or have available a complete set of insurance claim to isolate consumption from these shocks) or liquidity constraints.

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 the removal of the time, cohort and demographic effects discussed above.

Second, we not only attempt to provide life cycle expenditure patterns for nondurable consumption, but also for expenditures on consumer durables. Any quantitative study of portfolio choice over the life cycle has to account for the fact that for most households the largest share of their portfolio consists of consumer durables, in particular houses. By documenting empirical life cycle expenditure patterns for durables we again aim at establishing a benchmark for judging the performance of quantitative models.

We undertake the analysis of nondurables and durables jointly since, at each point in time, a household faces a trade-off between purchasing nondurables, durables and saving in financial assets. Furthermore, the households' ability to borrow may depend on its stock of consumer durables. Therefore all these decisions are interdependent and the resulting life cycle profiles should be analyzed simultaneously. Also, by quantifying the size, timing and correlation between the humps in nondurables and durables our empirical results may shed light on which deviations from the standard benchmark model can most accurately account for the data. For example, the findings in this work provide the motivation for Fernández-Villaverde and Krueger (2002) who show that a life cycle model with consumer durables that provide both utility and collateral services is capable of generating life cycle consumption profiles quantitatively consistent with the data.

There is a long literature documenting empirical life cycle consumption expenditures on nondurables. Examples include, among many others, Carroll and Summers (1991), Carroll (1992), Deaton (1992), Kotlikoff (2001) and Gourinchas and Parker (2002). Our paper builds on this tradition, but offers the following new contributions.

First, to the best of our knowledge, we are the first to exploit the information on consumer durables from the Consumer Expenditure Survey to build life cycle expenditure profiles on 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 household size changes to rationalize consumption expenditure profiles over the life cycle. These papers argue that demographics can explain, at least in a substantial part, why consumption tracks income over the life cycle. Using several 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. This leaves the other half of the hump unaccounted for by the standard complete markets life cycle model. Liquidity constraints, prudence in the light of idiosyncratic uncertainty and nonseparabilities in the utility function are potential candidate explanations for these findings.

Third, in addition to accounting for demographics we also control for cohort, time and age effects in a flexible way that imposes minimum conditions on the data. In particular, we specify a seminonparametric partial linear model which we estimate with a Speckman estimator. This procedure provides efficiency advantages, compared to the use of dummy variables, to estimate age profiles, yet its application for our problem is relatively straightforward.

Finally, as mentioned above, we use the bootstrap to evaluate the tightness of our estimates, an issue that has received little attention in the literature. Our simulations are useful in establishing that life cycle consumption profiles are estimated very precisely.

The rest of the paper is organized as follows. Section 2 describes the CEX data that we use as well as our procedure to build a pseudopanel of households from the raw data. Section 3 presents the specification of the estimated model of life cycle consumption and it explains in detail how we control for age, time and cohort effects and for demographic changes. Section 4 discusses our main empirical findings. In section 5 we document how robust our main findings are to sampling uncertainty and different econometric specifications. Section 6 concludes by suggesting some implications of our results and pointing to future research. The appendix contains further details about the data and variable definitions used in the paper.

2. The CEX Data

To document our facts about consumption over the life cycle we use the 1980-1998 data from the Consumer Expenditure Survey². During the last few years, following the work of Cutler and Katz (1991), Mace (1991) and Attanasio and Weber (1995), among others, the CEX has become one of the main sources for empirical work on consumption. The survey is a rotating panel. Each household is interviewed every three months over five calendar quarters and every quarter 20% of the sample is replaced by new households³. In the initial interview

²Our sample is only limited by data availability. Prior to 1980, the CEX was conducted about every 10 years and not on a regular basis. Data for years after 1998 are not yet released. We excluded the years 1982 and 1983 because of methodological differences in the survey. See Attanasio (1998) for details.

³The CEX definition of a household is a consumer unit that consists of any of the following: (1) all members of a particular household who are related by blood, marriage, adoption, or other legal arrangements; (2) a person living alone or sharing a household with others or living as a roomer in a private home or lodging house

information on demographic characteristics and on the inventory of major durable goods of the consumer unit is collected. Further expenditure information is gathered in the second through the fifth interviews using uniform questionnaires. In total, around 5000 households participate in the survey each quarter. The CEX is designed to constitute a representative sample of the U.S. population.

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 object of interest. Second the CEX lacks a significant panel dimension since it only follows households for at most 5 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 stongest interest from the point of economic theory; it only reports *expenditures* on consumption goods. This distinction is not very relevant for the case of nondurable goods for which it is plausible to equate the flow of consumption services to expenditures in the current period, but is crucial when dealing with consumer 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 only incurred (and show up in the data) once, namely in the current period. As an alternative to the analysis of expenditures one could impute service flows from the stocks of durables. However, the CEX only provides partial information for the value of that stock. While the survey asks for an estimate of the present value of the owned residence and the original cost of vehicles, it only takes a physical, but not a value inventory, of major household appliances owned by the household. The omission of these items may significantly underestimate the stock of durables for low-wealth households and thus, since younger households tend to be wealth-poor, distort estimates of life cycle consumption profiles for durables⁴.

These difficulties lead us to focus our analysis on expenditure data. Quantitative life cycle

or in permanent living quarters in a hotel or motel, but who is financially independent; or (3) two or more persons living together who use their incomes to make joint expenditure decisions. Financial independence is determined by the three major expense categories: housing, food, and other living expenses. To be considered financially independent, at least two of the three major expenditure categories have to be provided entirely or in part by the respondent.

⁴Also note that, since we do not observe the initial stock of durables of the cohort and the sample length is fairly small, we cannot use the *perpetual inventory method* to build estimates of the stocks of consumer durables.

models which 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. Nevertheless, even though our main focus is on expenditures, we also exploit the information in the CEX on present values of owned residences, thus indirectly providing life cycle profiles of services from owned homes⁵.

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. New households that enter the survey are a randomly chosen large sample of the US population and, consequently, they contain information about the consumption means of the groups they belong to. This information can be exploited by interpreting the observed group means as a panel for estimation purposes. This method, pioneered by Browning, Deaton and Irish (1985) and developed by Deaton (1985) and Heckman and Roob (1985), is known as the pseudopanel or synthetic cohort technique⁶.

In order to associate a household with a particular cohort we use the age of the reference person⁷. We define ten cohorts with a length of 5 years and we follow them through the whole sample, generating a balanced panel. With our choice of the number of cohorts we trade off the need for a large time series dimension to exploit the longitudinal aspect of the pseudopanel for the desire of a large cohort size to confidently assume that the sample means are good approximations to their population counterparts. Most of our cells have between 200 and 500 observations, with a mean of around 350.

We use our pseudopanel to estimate life cycle profiles of total consumption expenditures,

⁵If owner-occupied housing and other durables are complements, life cycle profiles of housing services can serve as first approximations of profiles for other durables, with the approximation the better, the stronger the complementarity between housing and services from other durables.

⁶Beyond increasing the time dimension of the data, a pseudopanel presents other advantages over pure panels. First, it eliminates the attrition problem. Most longitudinal surveys have low responses rates and these rates deteriorate when the same households are interviewed repeatedly. More importantly, attrition introduces a sample selection problem because of its possible correlation with explanatory variables. Second, the long temporal dimension of the pseudopanel approximately averages out expectational errors. Third, with pseudopanel we are not required to control for individual effects, as we would when using pure micro data, a difficult task when dealing with such a short panel as the CEX. Pseudopanel smooth out within-cohort heterogeneity by aggregating across agents. Note that this aggregation does not remove common cohort effects, but these are more easily controlled in the longer sample created by the pseudopanel.

⁷The reference person of the consumer unit is the first member mentioned by the respondent when asked to “Start with the name of the person or one of the persons who owns or rents the home”. It is with respect to this person that the relationship of the other consumer unit members is determined.

nondurable consumption expenditures and expenditures on durable goods. To obtain a preliminary feeling for the basic characteristics of the data we plot, in figure 2.2.1, total quarterly consumption expenditures, in thousands of 1982-84 dollars for our pseudopanel. Each line represents observations for one of the cohorts. In the same way we plot, in figure 2.2.2, expenditures on nondurables and in figure 2.2.3 expenditures on consumer durables.

2.3. Pseudopanel versus Pooled CEX Data

An alternative to the pseudopanel approach is to rely exclusively on the cross-sectional nature of the CEX survey and pool all observations. This alternative is closely related to a pseudopanel construction, except in two aspects. First, most estimations undertaken with pseudopanels do not weight each cohort by the size of observations in each cell. A pool, in practice, does exactly this since each household counts as one observation (unless also weighted). For relatively homogeneous cohort sizes the practical impact of this difference is minor. Also, from an estimation efficiency perspective it is ambiguous whether the weighted estimates are to be preferred. The second difference is that in a balanced pseudopanel (as in our case) we disregard observations from those cohorts that are not covered in the data during the entire sample period (i.e. those that are too young in 1980 to be already in the CEX) and from those cohorts that were already of advanced age at the beginning of the sample period and thus have small cohort sizes because of mortality. Since both types of households are relatively rare, their quantitative impact is small, and it is not clear to us that the information they provide compensates for the problems associated with their inclusion⁸. Nevertheless, for completeness below we also present selected results from pooled data and compare them to those obtained with the pseudopanel.

3. Specification and Estimation of Life Cycle Profiles

The most straightforward way to document consumption profiles over the life cycle may be to plot consumption against the age of the household head. This simple procedure, however, faces at least two problems. First, our data do not allow us to observe the profile of the same cohort for more than for 19 years. As explained above we address this problem by pooling information from different cohorts using a pseudopanel. Second, households in these cohorts

⁸It is also worth noting that the size of a pseudopanel (680 observations) is much smaller than the size of the pooled data (345436 observations). This smaller size has considerable computational advantages, in particular when carrying out the bootstrap to derive standard errors of estimated life cycle profiles.

were born at different dates and may have experienced very different conditions during their lifetime (think, e.g., about the cohort that went through the Great Depression during their prime working age). Also, with positive long-run growth of real wages, cohorts born at later dates have higher discounted lifetime earnings. Therefore, when using synthetic cohorts to measure pure life-cycle effects on consumption expenditures, it is crucial to control for cohort effects. But even if we observe one cohort over its entire life cycle, aggregate fluctuations will affect consumption profiles of the cohort; and these effects should not be attributed to aging, but rather to time effects. In subsection 3.1 we describe how to disentangle cohort and time effects from age effects, the object of interest of our analysis.

A second problem is that the CEX reports consumption data for households and not for individuals in that household. For our purpose it is crucial to separate changes in consumption expenditure induced by changes in family size and changes in other economic circumstances, since most quantitative life cycle models of consumption abstract from variations of household size and composition as the household ages. To provide useful empirical benchmark life cycle consumption patterns we adjust the raw data for demographics; section 3.2 describes our procedure to that effect.

3.1. Controlling for Age, Cohort and Quarter Effects

A simple and flexible way to relate age and consumption expenditure is to use a nonparametric regression of consumption on age. In particular we could estimate the model:

$$c_{it} = m(\text{age}_{it}) + \varepsilon_{it} \tag{1}$$

where c_{it} is the consumption expenditure of the household i at time t , age_{it} is the age of the household i at time t , measured in years, $m(\text{age}_{it}) = E(c_{it}|\text{age}_{it})$ is a smooth function of age_{it} and ε_{it} is an independent, zero mean, random error. This estimation model, however, does not include cohort or time effects. Since, as previously discussed this omission may be of importance, we modify our estimation equation (1) to account for these effects.

Consequently we estimate the partial linear model

$$c_{it} = \pi_i \text{cohort}_i + \pi_t \gamma_t + m(\text{age}_{it}) + \varepsilon_{it} \tag{2}$$

where cohort_i is a dummy for each cohort (except the youngest one) and γ_t a dummy for each quarter. Using dummies for cohort and quarter effects has the advantage of allowing the data to pick the preferred profiles for these effects without imposing functional forms that

tend to be rejected by the data (see, e.g., Heckman and Vytlacil (2001) for evidence against parametric assumptions).

Because time, age and cohort effects are linearly dependent it is not possible to separately identify all of them without further assumptions⁹. Our identification scheme is an adaptation of the one outlined by Deaton (1997). The main idea is to impose the assumptions that time effects are orthogonal to a time trend and that their sum is normalized to zero¹⁰. This identification amounts to attributing linear trends in the data to a combination of age and cohort effects. Attanasio (1998) discusses the advantages of this procedure in detail.

The partial linear model (2) can be easily estimated using the two step estimator proposed by Speckman (1988)¹¹. In this procedure the nonlinear part $m(age_{it})$ (the cross-sectional conditional expectation) is estimated with a Nadaraya-Watson estimator of the form

$$\widehat{m}_h(age) = \frac{\sum_{i=1}^n \sum_{t=1}^T K_h(age - age_{it}) * c_{it}}{\sum_{i=1}^n \sum_{t=1}^T K_h(age - age_{it})} \quad (3)$$

where $K_h(u) = \frac{0.75}{h} \left(1 - \left(\frac{u}{h}\right)^2\right) I\left(\left|\frac{u}{h}\right| \leq 1\right)$ is an Epanechnikov kernel¹² and h is the bandwidth parameter¹³.

Two sources of errors in variables may affect our results. First, because of sampling variance, the observed consumption means may differ from the cohort means. Since this error only affects the left-hand side variable (it is plausible that the average age is measured with high accuracy; in all cells, age samples averages are very close to the age interval midpoints), it only increases the variance of the residuals, provided that the error has a zero mean. Second, consumption data may suffer from large measurement errors. If these errors are linear and have zero cohort mean, the pseudopanel helps us because aggregation over the cohort sample will average them out.

⁹Since 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.

¹⁰Thus there are no dummies for the first two quarters. These effects are recovered using the orthogonalization and normalization conditions discussed in the main text.

¹¹The idea of the procedure is to first regress c_{it} on $m(age_{it})$ to obtain residuals ν_{it} and a smoothing matrix S . We apply this matrix to cohort and time effects, and project ν_{it} on the so corrected effects. The constructed new adjusted values for consumption are then nonparametrically regressed on age.

¹²Härdle (1990) discusses in detail the advantages of an Epanechnikov kernel for applications like ours. Beyond Härdle's arguments, the approximate lack of bias of this kernel in small samples will prove useful when applying bootstrap methods below.

¹³For our benchmark estimates we choose a bandwidth parameter of $h = 5$ years; in section 5.5 we defend this choice and perform sensitivity analysis with respect to this parameter.

3.2. Controlling for Family Size: Household Equivalence Scales

Interpreting the results from estimating (2) directly on the data is difficult. Not only consumption expenditure, but also, as shown in Figure 3.2.1, household size displays a hump over the life cycle. Households of different size plausibly face different marginal utilities from the same consumption expenditure, and economic theory only predicts that marginal utilities should be equated across time (up to some constant depending on the discount factor and the interest rate) and not expenditures per se. Consequently, changes in household size may go a long way towards explaining the hump in consumption, as argued in two influential papers by Attanasio and Weber (1995) and Attanasio *et al.* (1999).

We therefore want to quantify how much of the change in consumption over the life cycle is explained by demographics. The information contained in expenditure shares helps in answering this question. 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. This accounting exercise is undertaken by household equivalence scales which quantify the change in consumption expenditure needed to keep welfare of a family constant when its size varies¹⁴.

The simplest scale divides total expenditures by household size to obtain *per capita* household consumption. This scale therefore assumes that a household's technology to transform expenditures into service flows exhibits constant returns to scale. Theory and evidence suggests otherwise. Lazear and Michael (1980) point to three different mechanisms through which household size affects the rate of transformation between expenditures and services (family goods, scale economies 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.

We do not undertake the construction of a new household equivalence scale ourselves. Rather we borrow from rich previous work. Our reading of this literature is as follows. First, there is agreement about the existence of economies of scale in household consumption. Second, there exist differences of opinion 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

¹⁴Early papers that deflate household consumption expenditure by a function of family size include Zeldes (1989), who includes adjusted food requirements as a regressor in some of his Euler equation estimates, and Blundell, Browning and Meghir (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.

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 household choices¹⁵. Interestingly, the two explicit econometric procedures deliver point estimates that tend to be lower than the expert scales¹⁶.

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 size¹⁷. In section 5.2 below we document how robust our main findings are to changes in the household equivalence scale. Note that the chosen scale is very close to the equivalence scale of the HHS, the estimates of Johnson and Garner (1995) and to the constant-elasticity equivalence scales used by Atkinson *et al.* (1995), Buhmann *et al.* (1988) and Johnson and Smeeding (1998), among others.

After choosing a particular equivalence scale we take consumption expenditure measures c_{it} from the CEX, use the demographic information of the household to obtain the equivalence scale es_{it} and adjust consumption to obtain $\tilde{c}_{it} = \frac{c_{it}}{es_{it}}$. With the adjusted consumption we then estimate the partial linear model

$$\tilde{c}_{it} = \pi_i cohort_i + \pi_t \gamma_t + m(age_{it}) + \varepsilon_{it} \quad (4)$$

¹⁵These are constructed, respectively, by the OECD (OECD (1982)), the Panel on Poverty and Family 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 (US Department of Commerce (1991)), Lazear and Michael (1980) and Nelson (1993). Since the latter estimates only scales 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).

¹⁶Beyond 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 the columns 2 and 8 of table 3.2.1.

¹⁷To evaluate our choice it is important to remember, as explained in the appendix, that our measures of consumption do not include expenditures on either health or education, two major causes of increases in expenditures for households with children.

and interpret the generated consumption life cycle profiles¹⁸.

4. Results

In this section we present the results of our estimation of (4). First we examine the data with minimal econometric treatment. In figure 4.1.1 we plot the profile of total quarterly consumption expenditure against the age of the household without any adjustment by family size, cohort or age effects. Consumption follows a clear hump; it increases from \$3800 at the age of 20 to \$6900 at the peak in the late forties, and decreases afterwards. This pattern is well known and has been repeatedly 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 4.1.2) and expenditures on durables (figure 4.1.3) against the age of the household¹⁹. The hump in durables expenditures, is, to the best of our knowledge, a fact that has not been documented before in the literature. Note that both humps, for durables and nondurables, are of similar magnitude (the increase from age 20 to the peak is around 80%) and that the peak occurs at the same stage in the life cycle, around the late forties.

To control for changes in demographics we use the equivalence scale discussed in section 3 and repeat the estimation of life cycle profiles. Figure 4.1.4 plots total expenditure against household age, with and without controls for cohort and quarter effects. Two main findings deserve comment. First, comparing figure 4.1.4 to figure 4.1.1 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 a bit less than 50% but the hump still clearly appears. Adjusted quarterly consumption increases from around \$2900 to nearly \$4300 and then decreases to about \$2300. Also note that the peak in consumption is postponed, close to the age of fifty.

Second, as seen from the comparison of both profiles (with and without quarter and cohort effects) these additional controls have limited explanatory power. The quarter effects are in

¹⁸We need to point out several potential problems with our method. First, family size is endogenous. A procedure that explains the increase in consumption by the increase in family size without explicitly controlling for this endogeneity may find biased estimates for consumption profiles. For instance, individuals that are more productive in raising children are more likely to have larger families. Second, as pointed out by Pollak and Wales (1979), 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 (i.e. should be base independent). Using a flexible semiparametric estimator of demand functions Pendakur (1999) does not reject base independence.

¹⁹See the appendix for our definition of total expenditures, and durables and nondurables expenditures.

general small, with the exceptions 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 nearly flat. 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 80's and early 90's) of younger cohorts with smaller financial asset holdings. Also the recent changes in relative wages may have helped older, and more skilled cohorts in comparison with (at the current point of their life cycle) less skilled younger cohorts.

Figure 4.1.5 plots expenditures on nondurables. We see that consumption grows until the fifties and then declines, especially around retirement age, suggesting that some consumption expenditures are related to work. Comparing this to Figure 4.1.2 we observe a reduction of the hump of around a 50%. Figure 4.1.6 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, due to first purchases of durable goods. Interestingly enough, the reduction of the hump is quite similar to the case of nondurables (around 52%). These figures show, as for the raw data, that both expenditures on durables and nondurables have very similar patterns and peak at the same time. Thus, even if changes in demographic composition of households can account for around half of the hump in consumption and hence are crucial in understanding life cycle profiles, other departures from the standard life cycle model of consumption are needed to quantitatively account for the data²⁰.

Figure 4.1.6 possibly contains the most surprising result of this paper. Suppose that financial markets are complete, 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 thereon²¹. We do not see anything like this in the data; rather the process of durables

²⁰The presence of a hump is, in fact, robust to further breakdowns of expenditures. For instance, even if we focus on adult equivalent food consumption (a necessary good for which a higher degree of smoothing would be expected), we see a hump, with quarterly expenditures beginning at around \$475, increasing until \$630 and later falling to \$385. Studying food consumption is also 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 about the presence (and quantitative size) of a hump over the life cycle. In subsection 5.7 we offer further details about the breakdown of expenditures on consumer durables.

²¹>From a quantitative point of view, model simulations indicate that this result does not depend on the assumptions about interest rates (unless empirically unreasonable processes are assumed), but does depend

accumulation appears to be incremental over the life cycle²².

This finding is consistent with related evidence in the literature suggesting that households cannot perfectly smooth their consumption of services from durables. Alessie *et al.* (1997), Attanasio *et al.* (2000), and Eberly (1994) provide evidence of credit constraints for car purchases, Barrow and McGranahan (2000) document a spike in purchases of durables by low income households at the time Earned Income Credit checks are received and Browning and Crossley (1999) present evidence that expenditures on small consumer durables are cut back during unemployments spells. Finally Fisher and Johnson (2002) compute imputed services from a subset of durables using CEX data (and additional assumptions) and document a hump for these services, suggesting a lack of consumption smoothing over the life cycle.

In the next section we analyze how robust our findings are. First we evaluate the precision of our estimates by constructing confidence bands around the point estimates using the bootstrap. Then we assess how our results change if we use alternative methods to control for household size, with particular attention devoted to the effects the age of children of the household have on life cycle consumption profiles. After discussing two important technical details in our estimation we conclude the section with a more detailed analysis of life cycle expenditure pattern on consumer durables.

5. Robustness Analysis

5.1. 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. In order to assess this precision we present standard errors constructed with bootstrap methods. 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 approximations²³. This is especially true at both ends of the age profiles because of the low number of observations.

crucially on the separability and complete markets assumption.

²²This statement does depend on the assumed absence of in-kind intergenerational transfers of durables. Their presence at the beginning of the life cycle could rationalize low expenditures on durables, since these have been provided by the intergenerational transfer. However, data from the Health and Retirement Study (HRS) suggest that these transfers are fairly small (see Cardia and Ng (2000)).

²³In particular, 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 problematic. See Hall (1992a) for details.

To implement the bootstrap we adapt Härdle’s (1990) procedure to the Speckman estimator. We generate 500 samples of 680 observations each, drawing with replacement from the original pseudopanel. Then we apply the Speckman estimator to each of these new datasets. In this step care needs to be taken because 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 to choose a new smoothing parameter h that implies undersmoothing relative to the point estimate²⁴. We carry out this bias-removal strategy with an undersmoothing factor of 0.8.

Figure 5.1.1a) plots the 95% confidence interval for the age profile of adult-equivalent expenditures on nondurables, controlling for cohort and quarter effects. The size of the interval indicates that the age profile is estimated precisely. This result nothing but restates the fact that, as shown in the plot for the raw data, most observations group themselves quite naturally following a hump. Note that the interval (and the intervals and bands in the other figures) is not centered around the point estimates because of undersmoothing. Figure 5.5.1b) 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 5.5.1c) which shows a 95% confidence band, which covers the *whole* true curve (instead of each point separately, as in a confidence interval). Its interpretation is that, if we were able to get repeated samples from the original population of households, 95% of the associated age-consumption profiles would lie entirely within the band. Since any curve that can be plotted entirely inside this small band implies a significant hump, figure 5.5.1c) strongly reinforces our confidence in the point estimates: the data indicate a hump in consumption of nondurables, with size between 20% and 65% and a peak between age 45 and 50. Finally, figure 5.5.1d) plots all 500 simulated profiles: without exceptions all simulations generate a quantitatively significant hump in consumption life cycle profiles.

Figures 5.1.2 a)-d) repeat the same exercise for adult-equivalent expenditures on durables, again controlling for cohort and quarter effects. As before, our life cycle consumption profiles are quite precisely estimated. Similar figures (not included in the paper) are obtained for specifications without cohort and quarterly effects, with either only quarter or cohort effects,

²⁴See Horowitz (2001) for a theoretical explanation. Undersmoothing is achieved with a choice of a new smoothing parameter $h' = e \cdot h$, where $e < 1$ such that $nh'^{r+1} \rightarrow 0$ as $n \rightarrow \infty$ (here $r \geq 2$ is an even integer). As shown in Hall (1992a), using the Edgeworth expansion of a properly defined pivotal statistics, the bootstrap estimator of the confidence interval will be accurate up to $O((nh')^{-1})$. This asymptotic result does not provide clear advice for the appropriate choice of e in small samples. We tried several values of e without obtaining large differences in the results.

defining durables and nondurables in different ways, including and excluding housing, correcting for family size in different ways and other specifications we experimented with. We conclude that sampling uncertainty is unlikely to change our main findings²⁵.

5.2. Different Household Equivalence Scales

The choice of the equivalence scale we use does matter for the size of the hump in consumption over the life cycle, but not for its existence. We estimated life cycle profiles using all scales reported in table 3.2.1. and we always found the hump. To show a sample of the results we plot in figure 5.2.1. the life cycle profile of nondurables 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 particular, the use of the Nelson’s equivalence scale increases the hump, measured as the ratio from the peak of consumption to the consumption at the beginning of the life cycle, from 1.32 in our benchmark case to 1.50. On the other hand, the use of the OECD’s equivalence scale reduces the hump to 1.20. So even if this hump is quite smaller, it is still clearly present. In fact even *per capita* consumption has a hump. Demographics completely eliminate the hump only with decreasing returns to scale, an assumption that contradicts all empirical evidence we are aware of, and our intuition²⁶.

²⁵We check the robustness of our bootstrap results in two ways. First we implement a strategy with replacement subsampling to account for possible heteroskedasticity of errors. Subsampling generates results that are more robust since the simulations are less sensitive to violations of regularity conditions. The drawback of subsampling is larger variability and slower convergence of the estimates (see Horowitz (2001)). Confidence intervals and bands are slightly bigger than in the benchmark bootstrap, but otherwise the results are nearly identical. See, as an example, figure 5.1.3 for nondurables, controlling for cohort and quarter effects. Second, we implement a block sampling strategy to deal with possible data dependence since even with quarter and cohort effects, a basic resampling strategy may miss complicated dependences in the data. We experiment with several block sampling schemes; in all cases our findings are nearly indistinguishable from the main results. For instance, figure 5.1.4. reports outcomes from a cohort-based block sample scheme for nondurables with cohort and quarter effects.

²⁶We also investigate how our results change when we use separate equivalence scales for durables and nondurables. 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 since they estimate large economies of scale for housing.

5.3. Controlling for the Age of the Children

All equivalence scales used so far distinguish between adults and children in the household, but not between children of different ages. It is plausible, however, that a teenager has higher consumption needs than an infant, so that a larger increment in consumption expenditures is required for a teenager to keep the household at the same level of utility. If this is the case we may obtain biased results with our child-age-independent scale. In this subsection we explore this hypothesis.

A first alternative is to compute the variation in the cost of children by ages needed to keep a constant standard of living for the family. Lindert (1978 and 1980) estimates that these costs are fairly stable across the age of a child, particularly once education and health expenditures are excluded (as we do in our exercise, unless we use total consumption). With his equivalence scale for children we therefore would obtain results very similar to those in our benchmark estimation.

Browning and Ejrnæs (2002) propose a different approach in a paper that forcefully argues for explicitly taking account of the age of children in the household when deflating consumption by household size. First they define the number of equivalent adults, n_{it} , in household i at time t as

$$n_{it} = n_{it}^a + \sum_{j=1}^{n_{it}^c} \left(\mu_0 + \mu_1 \left(\frac{age_j}{20} \right) + \mu_2 \left(\frac{age_j}{20} \right)^2 + \mu_3 \left(\frac{age_j}{20} \right)^3 \right) \quad (5)$$

where n_{it}^a is the number of adults, n_{it}^c the number of children and $(\mu_0, \mu_1, \mu_2, \mu_3)$ are parameters in the polynomial used to approximate the age effects of children. To ensure that a 20 year old “child” is equivalent to an adult, they impose the restriction $\mu_3 = 1 - \mu_0 - \mu_1 - \mu_2$. In a second step the authors allow for economies of scale by specifying $es_{it} = (n_{it})^\delta$, where es_{it} is the final equivalence scale for household i at date t .

The parameters $(\delta, \mu_0, \mu_1, \mu_2, \mu_3)$ can be estimated once $\ln(es_{it}) = \delta \ln(n_{it})$ has been substituted into an Euler equation for the household that relates first differences in household log-consumption $\Delta \ln(c_{it})$ to first differences in household scales $\Delta \ln(es_{it})$. Browning and Ejrnæs (2002) report point estimates for these parameters, using *British* data from the Family Expenditure Survey (FES). Their results imply that the hump in life cycle consumption is accounted for by demographics of the household, now including the age of children.

Repeating Browning and Ejrnæs’ procedure for U.S. data or even simply applying their estimated scales is impossible because the CEX family data files only report the age of the children in two different bins (less than two years old and between two and 15 years old)

and consider all others members of the household as adults. Given these data limitations we approximate Browning and Ejrnæs' procedure in a crude but implementable way. We assume that children between two and 15 years old are equivalent to 0.2 adults and from age 16 on they are equivalent to one adult. Infants of age less than two count as zero adults.

Browning and Ejrnæs divide their data into a group of more and a group of less educated households and estimate the parameters $(\mu_0, \mu_1, \mu_2, \mu_3)$ for these groups separately. Figure 5.3.1 compares our approximation to their original estimates of children age-adjusted equivalence scales. We see that our simple scheme matches their estimates quite well. Note that the higher the children equivalence scales, the more costly are children and the stronger is the reduction of the life cycle hump in consumption due to demographics. Therefore our approximation, by overstating the cost of infants and teenagers above the age of 15 and by roughly matching the mean for children between the ages of two and 15 favors the hypothesis that all of the hump in the data can be attributed to changes in family size and composition²⁷.

In order to implement their approach we also need to specify a value for the parameter δ that controls the degree of increasing returns to scale in consumption with respect to household size. We choose $\delta = 0.86$, the mean of Browning and Ejrnæs' estimates for the two education groups they consider²⁸.

With our new adult equivalence scales adjusted by the age of children, we re-estimate our partial linear model. A sample of our findings is presented in figure 5.3.2, where we show total expenditures with cohort and quarter effects. In order to facilitate interpretation we also plot our benchmark profiles and normalize both graphs so that consumption at the beginning of the life cycle equals to 1 in both cases. The original and the newly estimated profile are qualitatively similar, with both of them featuring a clear hump over the life cycle, but show quantitative differences. In particular, the divergence of the two profiles becomes significant around the age of 45, where households with children of ages 16 and older begin to be frequent in the data. The equivalence scale motivated by Browning and Ejrnæs, for example, features a drastic increase from 2.12 to 2.57 for a household with two adults and a child when it turns 16. This indicates that the divergence of the two graphs is partially

²⁷Note, however, that Browning and Ejrnæs obtain their results with British data. It is an open question whether we would obtain similar estimates if we could use U.S. data with more detailed information about the age of children.

²⁸Under the assumption that the first two members of a household are adults and the remaining are children between the ages of two and fifteen, the implied equivalence scale is $(1, 1.82, 1.97, 2.12, 2.27, \dots)$. Comparing this scale to the ones in table 3.2.1 we observe that it is similar to most scales, with the exception of the effect induced by increasing family size from one to two adults, which is much larger for the scale implied by Browning and Ejrnæs' estimates.

due to our imperfect implementation of their approach (for data reasons we have to treat all household members above the age of 16 as adults), and partially due to the fact that Browning and Ejrnæs estimate increasing returns to scale weaker than those implied by all scales in table 3.2.1.

We summarize our results as suggesting that, for U.S. data from the CEX, controlling for the age of children in a detailed way is unlikely to remove the hump in life cycle consumption in the raw data, but may reduce its size somewhat.

5.4. Different Alternatives to Control for Family Size

There are at least four alternatives to equivalence scales to control for household size and composition. First, one may divide the original sample into groups corresponding to different household sizes. With the resulting separate samples of households with size 1, size 2 and so on we can repeat our estimation for each of these groups. This procedure may be interpreted as a bivariate kernel on age and family size where the smoothing parameter for family size dimension is less than one. We do not use this approach as our benchmark because of the endogeneity of household size. Individuals living alone at age 25 constitute a very different subsample of the population than individuals living alone at age 45 since the first group includes both individuals that will still live alone in 20 years and those who will form households with more than one member during the next 20 years. Despite these caveats we carry out the exercise as sensitivity analysis; figures 5.4.1, 5.4.2 and 5.4.3 plot the results for total expenditures, nondurables expenditures and durables expenditures respectively. For nearly all household sizes we observe humps in life cycle consumption expenditures of similar size and location as in our benchmark estimates.

A second alternative to correct for family size is to use a flexible specification of preferences that allows to control for demographic factors through the use of additional regressors. This approach, proposed by Attanasio *et al.* (1999) uses a linear approximation of the Euler equation for consumption, modified by the inclusion of additional demographic variables z_t (or instruments for them) to obtain

$$\Delta \log(c_{t+1}) = \text{constant} + \sigma \log(1 + r_{t+1}) + \theta \Delta z_{t+1} + \epsilon_{t+1} \quad (6)$$

Estimation of the parameters of (6) provides an implicit correction for demographic changes. There may be several problems with this approach which lets us prefer our procedure. First, a rich specification of the demographic variables to capture shifts in the utility function may

result in overparametrization and loss of efficiency. The resulting reduction in the precision of the parameter estimates may explain why the authors cannot reject the null hypothesis of correct model specification. Second, demographic variables may proxy for liquidity constraints, which makes the estimates $\hat{\theta}$ hard to interpret²⁹. Third, since labor income is hump-shaped over the life cycle and fertility choices are endogenous, households with limited access to intertemporal trade will choose a simultaneous hump in both consumption expenditure and family size, even though one does not cause the other. If the Euler equation is not estimated simultaneously with an optimality condition for fertility, (6) may spuriously pick up the hump in consumption through the estimates for the demographic parameters. Note, however, that although our methodology differs from Attanasio et al. (1999), the equivalence scale we use comes close to the one implicit in their point estimates of the demographic parameters (this scale is 1, 1.57, 1.80, 2.04, ...). Consequently, if we use their implicit estimates we obtain results similar to those from our benchmark specification.

A third approach is to use dummies, f_{it} , for different household sizes in our partial linear model in the form:

$$c_{it} = \pi_i \text{cohort}_i + \pi_t \gamma_t + \pi_{it} f_{it} + m(\text{age}_{it}) + \varepsilon_{it} \quad (7)$$

This use of dummies to correct for household size is the approach Gourinchas and Parker (2002) employ³⁰. Their results for nondurables suggest that this alternative approach yields results that are qualitatively similar to the ones presented in this paper.

Finally, an innovative alternative to controlling for household size is to estimate profiles for *individuals* directly. Deaton and Paxson (2000) report individual life cycle *saving* profiles for Taiwan. We do not follow this strategy because an analysis of profiles for individuals requires an explicit model of resource allocation within the household, and there does not seem to be widespread agreement about a “standard” model for this. The problem is especially severe for consumer durables: what is the portion of services from a TV, car or refrigerator owned by a household that each individual consumes? We do not attempt to provide an answer to these difficult questions in this paper³¹.

²⁹In the presence of liquidity constraints, the estimation of a log-linearized Euler equation may have problems on its own. See Attanasio and Low (2000), Carroll (2001) and Ludvigson and Paxson (2001).

³⁰Instead of a kernel estimator they use a set of age dummies in (7) to capture age effects (also, they deal with the multicollinearity problem by employing unemployment rates instead of time dummies to capture time effects). Using dummies is equivalent to a kernel estimator with such a small smoothing parameter that only information from households in the same cohort is used.

³¹On a technical level, implementing this strategy with U.S. data does require the integration of the member files with the family files from the CEX, not a trivial task. Future research on this issue seems particularly

5.5. Choice of the Bandwidth

An important technical point that deserves further discussion is choice of the bandwidth h of the nonparametric kernel inside the partial linear model. For example, an $h = 1$ means that the conditional expectation $m(\text{age}_{it})$ in (3) is computed using only household observations with age in the interval $[\text{age}_{it} - 1, \text{age}_{it} + 1]$, whereas an $h = 20$ uses observations with age in $[\text{age}_{it} - 20, \text{age}_{it} + 20]$. Therefore this parameter controls the degree of smoothness of the estimate of the curve m , and it determines the trade-off between small sample bias and variance (a small h gives small bias and large variance whereas a large h gives the opposite results). Since consistency of the estimates only requires that, as $n \rightarrow \infty$, h should satisfy $h \rightarrow 0$ and $nh \rightarrow \infty$, asymptotic theory does not provide much guidance for the choice of h in a small sample. Cross-validation methods³² suggest a lower value for h than the one we find preferable for our application. We search for a hump (or the absence of it) in consumption, a low frequency pattern, while cross validation tries to capture the high frequency movements. This implies that we are mostly concerned with increasing the “signal to noise” ratio of the raw data and suggests that a slightly “oversmoothed” curve is a better choice³³. We settle for a benchmark value of $h = 5$; now we check the robustness of our results by computing the regressions for $h = 1$ and $h = 10$. Figures 5.5.1-5.5.2 plot the results of this exercise, for nondurables and durables, respectively. In the first case the hump (measured as before) is basically constant, and is reduced by around 5% when $h = 10$. For expenditure on durables, the hump is increased by 11% when $h = 1$ and reduced by around 14% when $h = 10$. The fact that even for $h = 10$, where a lot of nonlocal information is used, a clear hump in consumption expenditures on both nondurables as well as durables emerges, suggests that this finding is not an artefact of the particular nonparametric procedure but rather a basic characteristic of the data.

5.6. Pseudopanel versus Pooled CEX Data

When we use all pooled observations from the CEX instead of our pseudopanel, we obtain roughly the same results as above. As a sample of our findings, in figure 5.6.1 we plot the

important.

³²Cross-validation searches for the smoothing parameter that minimizes the mean of a penalizing function (such as the Mean Square Error) when different parts of the sample are eliminated. This minimization trades off bias for variance to optimize the objective function of the researcher. Intuitively, this procedure maximizes the ability of the regression to forecast within the sample (see Härdle (1990)).

³³Other practitioners have argued in favor of this less formal approach to the choice of bandwidths. See again Härdle (1990) for examples.

age profiles from adult equivalent expenditures on nondurables both for the pseudopanel and for the pooled data. Noticeable differences occur only at the beginning of the life cycle where the pooled data set has many more observations than the pseudopanel. If anything, the use of pooled data increases the size of the hump for nondurables. However, we are skeptical about how much weight should be placed on the very few consumption observations of households with head age of 16 to build life cycle profiles. This concern and our desire to enable comparability with previous empirical studies explains our decision to take the pseudopanel estimation as our benchmark case.

5.7. Assessing the Importance of Housing

A large fraction of expenditures on consumer durables stems from housing. Since out-of-pocket expenses of owning a home are potentially significant in the first years of ownership and then decline, while consumption services from the home are roughly constant, the link between expenditure on owned dwellings and its services, the ultimate object of interest, may be particularly weak³⁴. In this section we therefore want to, at least partially, assess whether our previous results for consumer durables are primarily driven by its biggest component.

Figure 5.7.1 plots the estimation results for adult equivalent expenditure on durables, *excluding* housing and figure 5.7.2 plots the same for expenditures on housing. Both figures display a clear hump over the life cycle, suggesting that our previous results were not driven by the aggregation of expenditures on durables. It is worth noting that expenditures on housing increase more steeply over the first ten years of adult life than expenditures on other durables, so that the peak of the hump occurs earlier (mid 30's vs. 50) and is more sizeable (45% vs. 37%).

If we disaggregate one level further and separate expenditures on owned residential housing (figure 5.7.3) and expenditure on rented dwellings (figure 5.7.4) we observe that, over the life cycle, renting is substituted by home ownership (although this trend is reversed at the very end of the life cycle). This substitution pattern explains why expenditures on owned residences are hump-shaped. But even conditional on ownership, as in figure 5.7.5., or rentership, as in figure 5.7.6., we see expenditure patterns that suggest lack of consumption smoothing over the life cycle³⁵.

³⁴Our definition of housing expenditures includes expenditures on owned dwellings and rented dwellings used for residential purposes by the household. Expenditures on owned dwellings include mortgage interest, property taxes, repairs, maintenance and insurance. Expenditures on rented dwellings correspond to rent payments by the household.

³⁵We ought to interpret figures 5.7.5-6 somewhat carefully since the households that have positive renting

Housing is also the only component of durables for which the CEX contains useful information about its services, since the survey collects information about the monthly rental value of the owned residence, as estimated by the household head³⁶. Figure 5.7.7 plots the estimated life cycle profile for data, not adjusted by household equivalence scales, and figure 5.7.8 does the same for the data deflated by our equivalence scales. The first figure shows that, when controlling for quarter and cohort effects, the peak of (market valued) housing services does not occur until the mid fifties, then decreases slightly, only to mildly increase towards the end of the life cycle. Figure 5.7.7. also is one of the few instances in this paper where cohort effects play a significant role for the results³⁷, with later cohorts living in more expensive homes. The pattern in figure 5.7.7 is roughly consistent with a hypothetical model in which households face financial constraints that prevent them from obtaining their desired home at the beginning of the life cycle. As they age, these households move into better and better homes, until they reach their target house, which is kept until the end of their life cycle, to assure a smooth flow of housing services.

Figure 5.7.8, which adjusts for household size, shows a similar picture, except for the end of the life cycle. The late increase in the household-size-adjusted rental value of the home is due to the reduction in household size (usually one spouse dies) which are not associated with changes in residence (as figure 5.7.7 indicates). This empirical finding is suggestive of models with (financial or psychological) adjustment costs or models in which durables provide important collateral services (for instance, to hedge against catastrophic health expenditures), in which households at the end of the life cycle own more valuable houses than otherwise optimal³⁸.

6. 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% of which is accounted for by changes in household demographics. The other half remains to be explained

expenditures at age 40 are a very different subgroup of the population than the one that rents at age 25, raising serious sample selection issues.

³⁶Households are asked: "If someone were to rent your home today, how much do you think it would rent for monthly, unfurnished and without utilities?". Note that this question was not asked in 1980 and 1981.

³⁷Quarter effects are of minor importance for figures 5.7.7-8.

³⁸Note that the same findings as in our figures 5.7.7-8 emerge if we use the new variable of housing services defined by Fisher and Johnson (2002). They generate a series for the rental value of each households' dwelling, equal to the paid rent, equal to the imputed rent or to the sum of both.

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 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. Our limited evidence on services from the biggest durable item, residential housing, suggests that these services are not smoothed over the life cycle either.

There are a number of possible deviations from the basic life cycle model that may qualitatively account for the humps documented in this paper. First, one may relax the assumption of separability between leisure and consumption. Ghez and Becker (1973) and Heckman (1974) propose models where utility services are produced with time and goods as inputs. When time becomes more expensive (i.e. wages are higher), agents substitute goods for time in the production function of utility services, generating a correlation between labor income and consumption. A second departure, dating back to Nagatani (1972) and developed in Carroll (1992) and Attanasio *et al.* (1999), is the introduction of uninsurable uncertainty (e.g. with respect to labor income or lifetime) into a model where households are prudent. Agents wanting to self-insure against this uncertainty do not consume as much early in life as they would in the benchmark life cycle model. Finally Thurow (1969) and later Zeldes (1989), among others, have argued for the importance of liquidity constraints that prevent young households from borrowing against future (higher) labor income to finance higher current consumption.

None of these explanations is mutually exclusive. An important task in applied economics consists in measuring the quantitative importance of each of them. This paper provides an empirical target against which quantitative success these explanations may be assessed against. Since expenditures on consumer durables are an important share of total household consumption expenditures, an important asset in a households' portfolio, and an important determinant for a households' ability to borrow, an explicit incorporation of durables into models of life cycle consumption seems to be called for. In documenting life cycle profiles of expenditures on durables we want to empirically inform recent and future quantitative work on this issue, examples of which include Díaz and Luengo-Prado (2002), Fernández-Villaverde and Krueger (2002) and Laibson *et al.* (2001).

7. Appendix: Data Sources

We take our consumption data from the Consumer Expenditure Survey (CEX), as provided by the Bureau of Labor Statistics. Our sample years consist of 1980-1981 and 1984-1998, with a total of 68 longitudinal surveys.

We take each household as one observation and use the demographic information of the reference person to define cohort membership, independent of this person's gender. We select only those households with both positive income and consumption expenditure. As most of the literature we do not attempt to control for topcoding of consumption observations. The very high topcoding limits (or their nonexistence for food consumption and other items) in the CEX and the very low survey response rates of the wealthiest households in the U.S. imply that only a extremely small fraction of our sample is right-censored. As a consequence, it is unlikely that the lack of proper topcoding treatment affects the results in a significant manner.

We compute "total expenditure" using the variable with the same name in the detailed expenditure files. We divide consumption in these files into three different groups. The data on "expenditure on nondurables" include food, alcohol beverages, tobacco, utilities, personal care, household operations, public transportation, gas and motor oil, entertainment and miscellaneous expenditures. The variable "expenditure on durables" sums expenditures on owned dwelling, rented dwelling, house equipment, vehicles, books and electronic equipment. We define as ambiguous expenditure apparel, out-of-pocket health and education expenditures (unless we analyze "total expenditures" which includes all expenditures in the CEX). We account for changes in the consumption classification methodology over the sample years in the CEX, in order to assure consistency of our consumption measures.

Finally, each expenditure category is deflated using its own specific, not seasonally adjusted, Consumer Price Index (CPI) component for urban consumers. The dollar figures are adjusted to 1982-84 dollars using the "current methods" version of the CPI. This version rebuilds past CPI's with the present methodology to produce a price deflator series that is consistent over time.

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] Auerbach, A. and L. Kotlikoff (1987), *Dynamic Fiscal Policy*, Cambridge University Press.
- [3] Atkinson, A., Rainwater, L. and T.M. Smeeding (1995), "Income Distribution in OECD Countries: Evidence from the Luxembourg Income Study (LIS)". *Social Policy Studies* 18, OECD.
- [4] Attanasio, O. P. (1998), "Cohort Analysis of Saving Behavior by U.S. Households". *Journal of Human Resources* 33, 575-609.
- [5] Attanasio, O.P., 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. and M. Browning (1995), "Consumption over the Life Cycle and over the Business Cycle". *American Economic Review* 85, 1118-1137.
- [7] Attanasio, O., P.K. Goldberg and E. Kyriazidou (2000), "Credit Constraints in the Market for Consumer Durables; Evidence from Microdata on Car Loans". *NBER Working Paper* 7694.
- [8] Attanasio, O. and H. Low (2000), "Estimating Euler Equations", *NBER Technical Working Paper* 253.
- [9] 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.
- [10] Barrow, L. and L. McGranahan (2000), "The Effects of the Earned Income Credit on the Seasonality of Household Expenditures". *National Tax Journal*, forthcoming.
- [11] 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.
- [12] Browning, M. and T. Crossley (1999), "Shocks, Stocks and Shocks: Consumption Smoothing and the Replacement of Durables during an Unemployment Spell". *Mimeo*, Institute of Economics, University of Copenhagen.
- [13] Browning, M., A. Deaton and M. Irish (1985), "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle". *Econometrica*, 53, 503-543.
- [14] Browning, M. and M. Ejnæs (2002), "Consumption and Children". *Mimeo*, Institute of Economics, University of Copenhagen.

- [15] Buhmann, B., Rainwater, L., Schmauss G. and T.M. Smeeding (1988), "Equivalence Scales, Well-Being, Inequality and Poverty: Sensitivity Estimates across Ten Countries Using the Luxembourg Income Study Database". *Review of Income and Wealth*, 34, 115-142.
- [16] Cardia, E. and S. Ng (2000), "How Important are Intergenerational Transfers of Time? A Macroeconomic Analysis". *Mimeo*, John Hopkins University.
- [17] Carroll, C.D. and L.H. Summers (1991), "Consumption Growth Parallels Income Growth: Some New Evidence" in B.D Bernheim and J.B Shoven (eds), *National Saving and Economic Performance*, University of Chicago Press for the NBER.
- [18] Carroll, C.D. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence". *Brookings Papers on Economic Activity*, 23, 61-156.
- [19] Carroll, C.D. (2001), "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)". *Advances in Macroeconomics*, 1, article 6.
- [20] Citro, C.F. and R. T. Michael (1995), *Measuring Poverty: A New Approach*. National Academy Press.
- [21] 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.
- [22] Cutler, D. and L. Katz (1991), "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's". *American Economic Review*, 82, 546-551.
- [23] Datzinger S, J. van der Gaag, M.K. 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.
- [24] Deaton, A. (1985), "Panel Data from Time Series of Cross-Sections". *Journal of Econometrics* 30, 109-126.
- [25] Deaton, A. (1992), *Understanding Consumption*. Oxford University Press.
- [26] Deaton. A. (1997), *The Analysis of Households Surveys*. John Hopkins University.
- [27] Deaton, A. and C. Paxson (2000), "Growth and Saving Among Individuals and Households". *Review of Economics and Statistics*, 82, 215-225.
- [28] Díaz, A. and M. Luengo-Prado (2002), "Durable Goods and the Wealth Distribution". *Mimeo*, Universidad Carlos III.
- [29] Eberly, J. (1994), "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases". *Journal of Political Economy* 102, 403-436.
- [30] Federal Register (1991), 56, n 34.

- [31] 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
- [32] Ferreira, M.L., R.C. Buse and J.P. Chavas (1998), “Is There Bias in Computing Household Equivalence Measures”. *Review of Income and Wealth*, 44, 183-198.
- [33] Fisher, J.D. and D.S. Johnson (2002), “Consumption Mobility in the United States: Evidence from Two Panel Data Sets”. *Mimeo*, Bureau of Labor Statistics.
- [34] Garner, T.I. and K. de Vos (1995), “Income Sufficiency vs. Poverty: Results from the United States and The Netherlands”. *Journal of Population Economics* 8, 117-134.
- [35] Garner, T.I. and K.S. Short (2001), “Owner-Occupied Shelter in Experimental Poverty Measurement with a ‘Look’ at Inequality and Poverty Rates”. *Mimeo*, Bureau of Labor Statistics.
- [36] Ghez, G.R. and G. Becker, (1975), *The Allocation of Time and Goods over the Life Cycle*. Columbia University Press for the NBER.
- [37] Gourinchas, P.O. and J. A. Parker (2002), “Consumption over the Life Cycle”. *Econometrica* 70, 47-89.
- [38] Hall, P. (1992a), *The Bootstrap and Edgeworth Expansion*. Springer-Verlag.
- [39] Hall, P. (1992b), “Effect of Bias Estimation on Coverage Accuracy of Bootstrap Confidence Intervals for a Probability Density”. *Annals of Statistics* 20, 675-694.
- [40] Härdle, W. (1990), *Applied Nonparametric Regression*. Cambridge University Press.
- [41] Heckman, J.J. (1974), “Life Cycle Consumption and Labor Supply: An Explanation of the Relationship Between Income and Consumption over the Life Cycle”. *American Economic Review* 64, 184-194.
- [42] Heckman, J.J. and R. Robb (1985), “Alternative Methods for Evaluating the Impact of Interventions” in J.J. Heckman and B. Singer (eds), *Longitudinal Analysis of Labor Market Data*. Cambridge University Press.
- [43] Heckman, J.J. and E. Vytlacil (2001), “Identifying the Role of Cognitive Ability in Explaining the Level and Change in the Return to Schooling”. *Review of Economics and Statistics* 83, 1-12.
- [44] Horowitz, J. (2001), “The Bootstrap” in J.J. Heckman and E. Leamer (eds), *Handbook of Econometrics*, vol. 5.
- [45] Johnson, D.S. and T.I. Garner (1995), “Unique Equivalence Scales: Estimation and Implications for Distributional Analysis”. *Journal of Income Distribution* 4, 215-234.

- [46] Johnson, D.S. and T.M. Smeeding (1998), "Intergenerational Equity in the United States: the Changing Well-being of the Old and the Young, 1960-1995". *Mimeo*, Bureau of Labor Statistics.
- [47] Jorgenson, D. and D. Slesnick (1987), "Aggregate Consumer Behavior and Household Equivalence Scales". *Journal of Business Economics and Statistics* 5, 219-232.
- [48] Kotlikoff, L.J. (2001), *Essays on Saving, Bequest, Altruism and Life-Cycle Planning*. MIT Press.
- [49] Laibson, D., A. Repetto and J. Tobacman (2001), "A Debt Puzzle". *Mimeo*, Harvard University.
- [50] Lazear, E. and Michael, R.T. (1980), "Family Size and the Distribution of Real Per Capita Income". *American Economic Review* 70, 91-107.
- [51] Lindert, P. H. (1978), *Fertility and Scarcity in America*. Princeton University Press.
- [52] Lindert, P.H. (1980), "Child Cost and Economic Development" in R. Easterlin (ed), *Population and Economic Change in Developing Countries*. Chicago University Press for NBER.
- [53] Ludvigson, S. and C. H. Paxson (2001), "Approximation Bias in Linearized Euler Equations". *Review of Economics and Statistics* 83, 242-256.
- [54] Mace, B. (1991), "Full Insurance in the Presence of Aggregate Uncertainty". *Journal of Political Economy* 99, 928-56.
- [55] Nagatani. K. (1972), "Life Cycle Savings: Theory and Fact". *American Economic Review* 62, 344-353.
- [56] Nelson, J. (1988), "Household Economies of Scale in Consumption: Theory and Evidence," *Econometrica* 56, 1301-1314.
- [57] Nelson, J. (1993), "Independent of a Base Equivalence Scales Estimation Using United States Micro-Data". *Annales d'Économie et de Statistique* 29, 43-62.
- [58] Organization for Economic Cooperation and Development (1982). *The OECD List of Social Indicators*.
- [59] Pendakur, K. (1999), "Semiparametric Estimates and Test of Base-Independent Equivalence Scales". *Journal of Econometrics* 88, 1-40.
- [60] Phipps, S. and T.I. Garner (1994), "Are Equivalent Scales the Same for the United States and Canada?". *Review of Income and Wealth* 40, 1-17.
- [61] Pollak, R.A. and T.J. Wales (1979), "Welfare Comparisons and Equivalence Scales (in Equality: The Individual vs. the Family)". *American Economic Review* 69, 216-221.

- [62] Speckman, P. (1988), “Kernel Smoothing in Partial Linear Models”. *Journal of the Royal Statistical Society B* 50, 413-436.
- [63] Thurow, L. (1969), “The Optimal Lifetime Distribution of Consumption Expenditures”. *American Economic Review* 59, 324-330.
- [64] U.S. Department of Commerce (1991). *Trends in Relative Income: 1964 to 1989*. Series P60-177.
- [65] Zeldes, S. (1989), “Consumption and Liquidity Constraints: An Empirical Investigation”. *Journal of Political Economy* 97, 305-346.

Figure 2.2.1: Total Expenditure




Figure 2.2.2: Expenditures non Durables




Figure 2.2.3: Expenditures Durables




Figure 3.2.1: Family Size




Figure 4.1.1: Total Expenditure




Figure 4.1.2: Expenditures non Durables

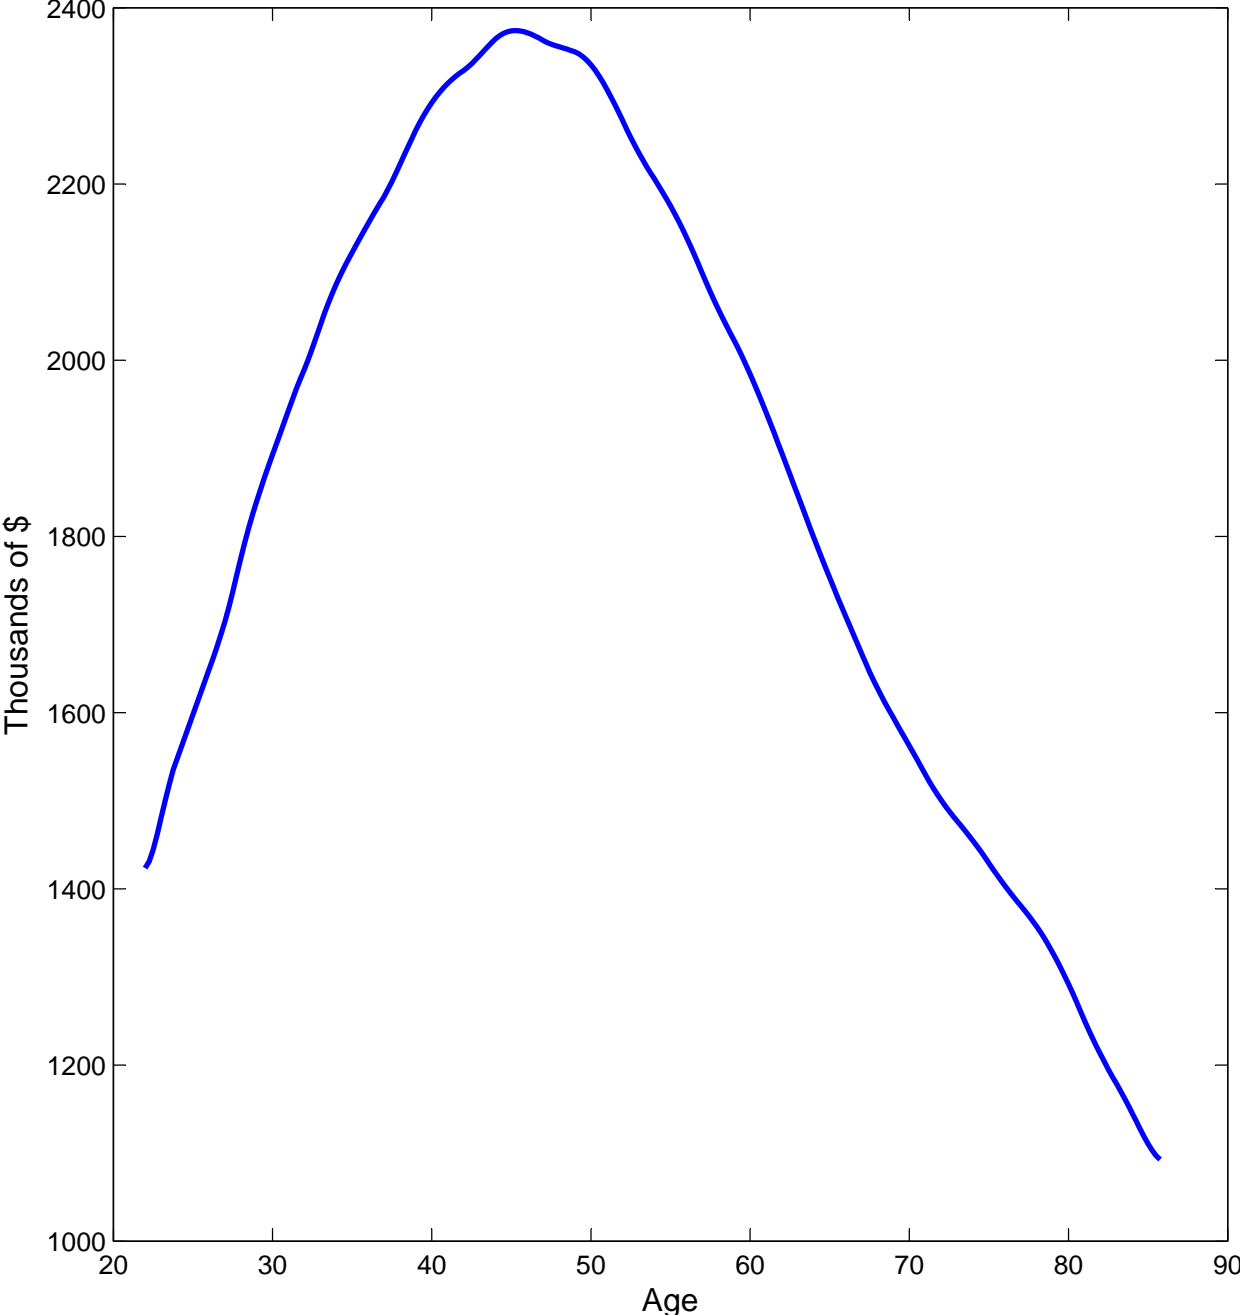


Figure 4.1.3: Expenditures Durables




Figure 4.1.4: Total Expenditure adult equivalent




Figure 4.1.5: Expenditures non Durables adult equivalent




Figure 4.1.6: Expenditures Durables Adult Equivalent




Figure 5.1.1 a): 95% confidence interval




Figure 5.1.1 b): Widest confidence interval




Figure 5.1.1 c): 95% confidence band




Figure 5.1.1 d): All simulations




Figure 5.1.2 a): 95% confidence interval




Figure 5.1.2 b): Widest confidence interval




Figure 5.1.2 c): 95% confidence band




Figure 5.1.2 d): All simulations




Figure 5.1.3 a): 95% confidence interval




Figure 5.1.3 b): Widest confidence interval




Figure 5.1.3 c): 95% confidence band




Figure 5.1.3 d): All simulations




Figure 5.1.4 a): 95% confidence interval




Figure 5.1.4 b): Widest confidence interval




Figure 5.1.4 c): 95% confidence band




Figure 5.1.4 d): All simulations




Figure 5.2.1: Comparison of Different Equivalence Scales, nondurable Consumption




Figure 5.3.1: Age-Changing Equivalence Scales




Figure 5.3.2: Comparison of Consumption Profiles




Figure 5.4.1: Total Consumption Expenditure




Figure 5.4.2: Nondurable Consumption Expenditure




Figure 5.4.3: Durable Consumption Expenditure




Figure 5.5.1: Expenditures non Durables adult equivalent, different smoothing parameters

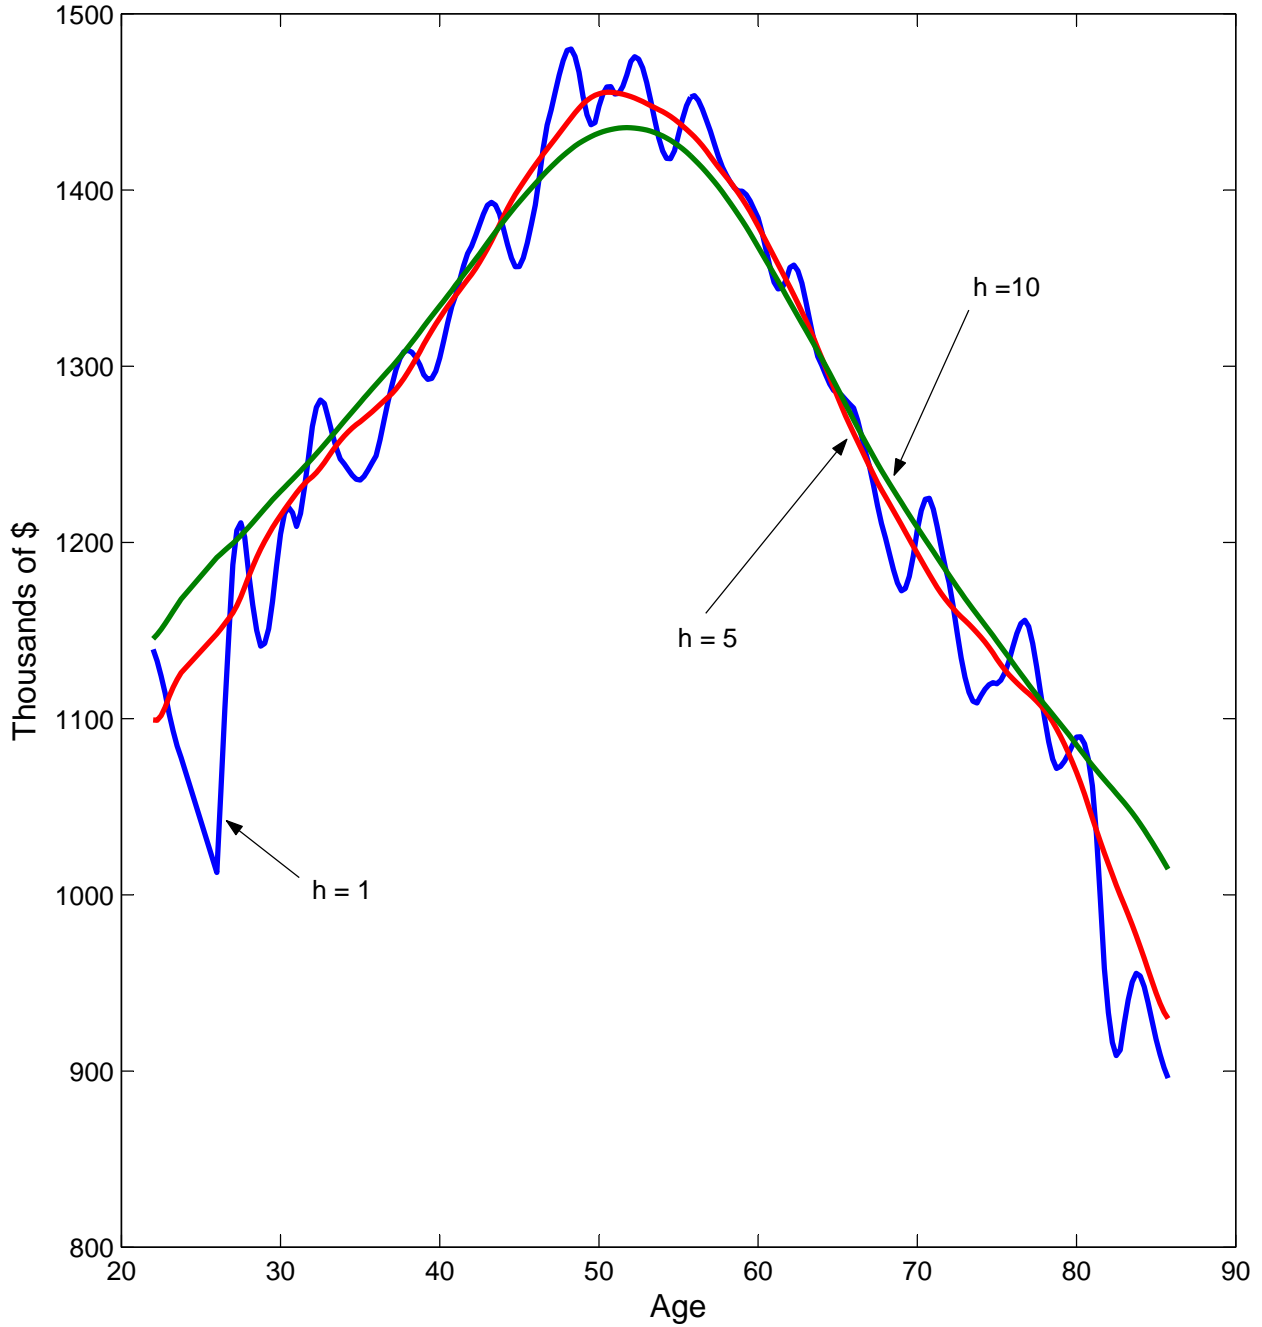


Figure 5.5.2: Expenditures Durables adult equivalent, different smoothing parameters

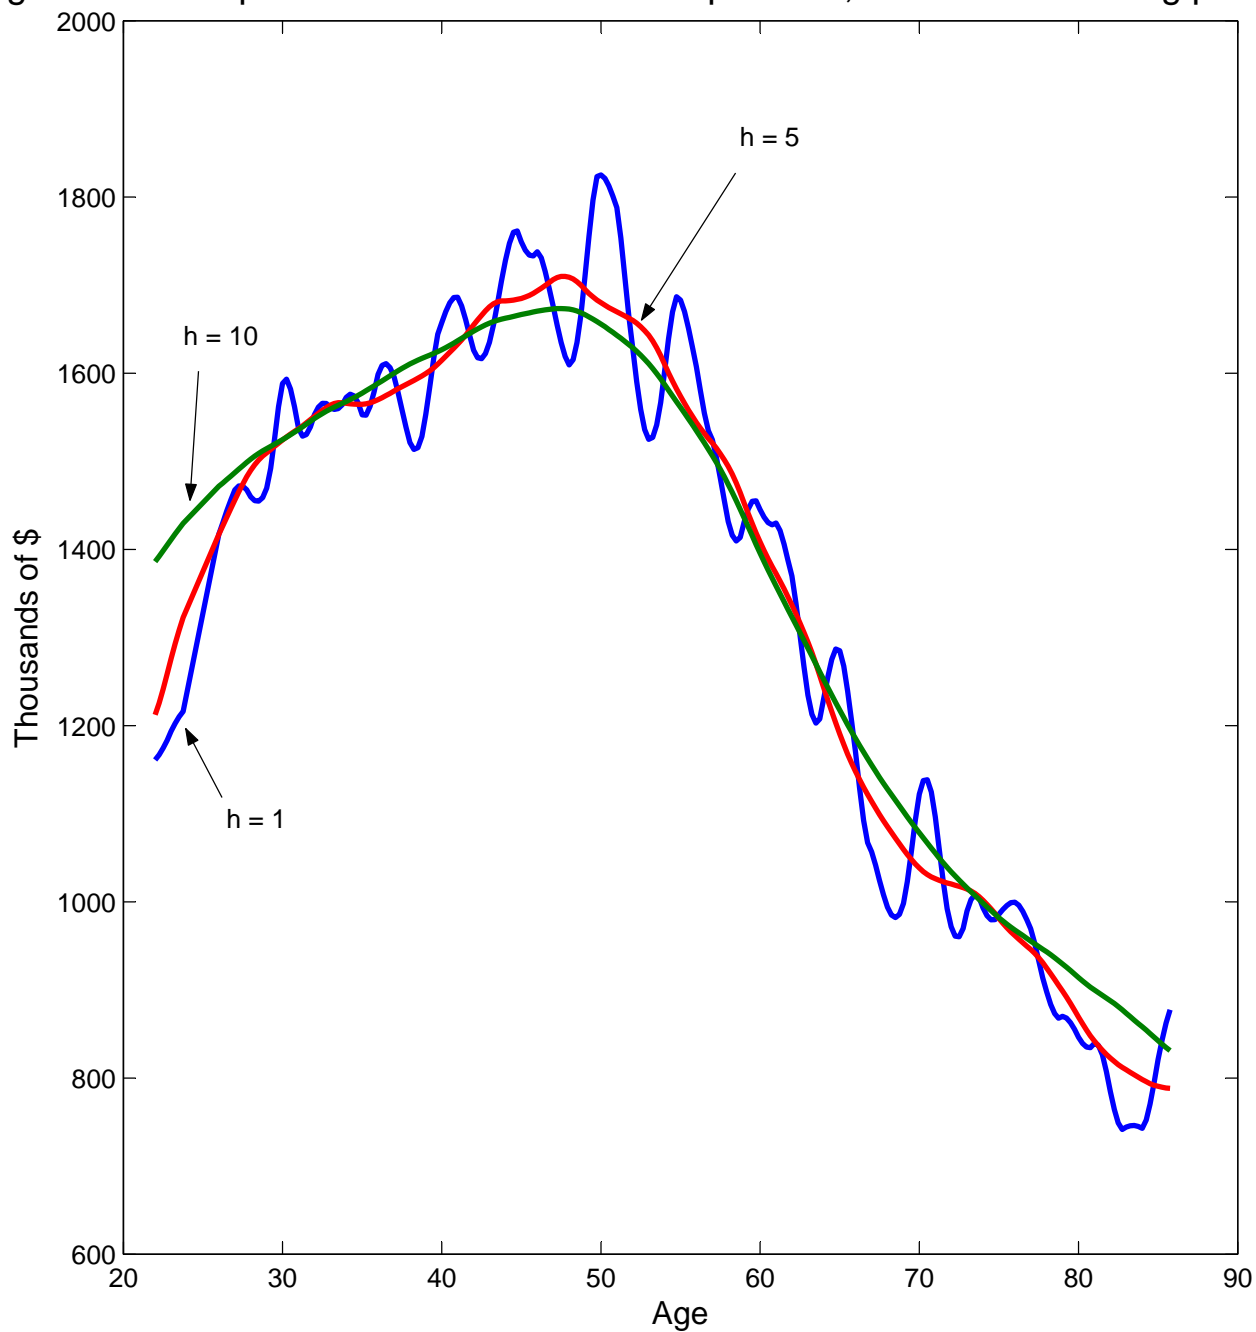


Figure 5.6.1: Expenditures non Durables adult equivalent




Figure 5.7.1: Expenditures Durables, Non Housing, adult equivalent

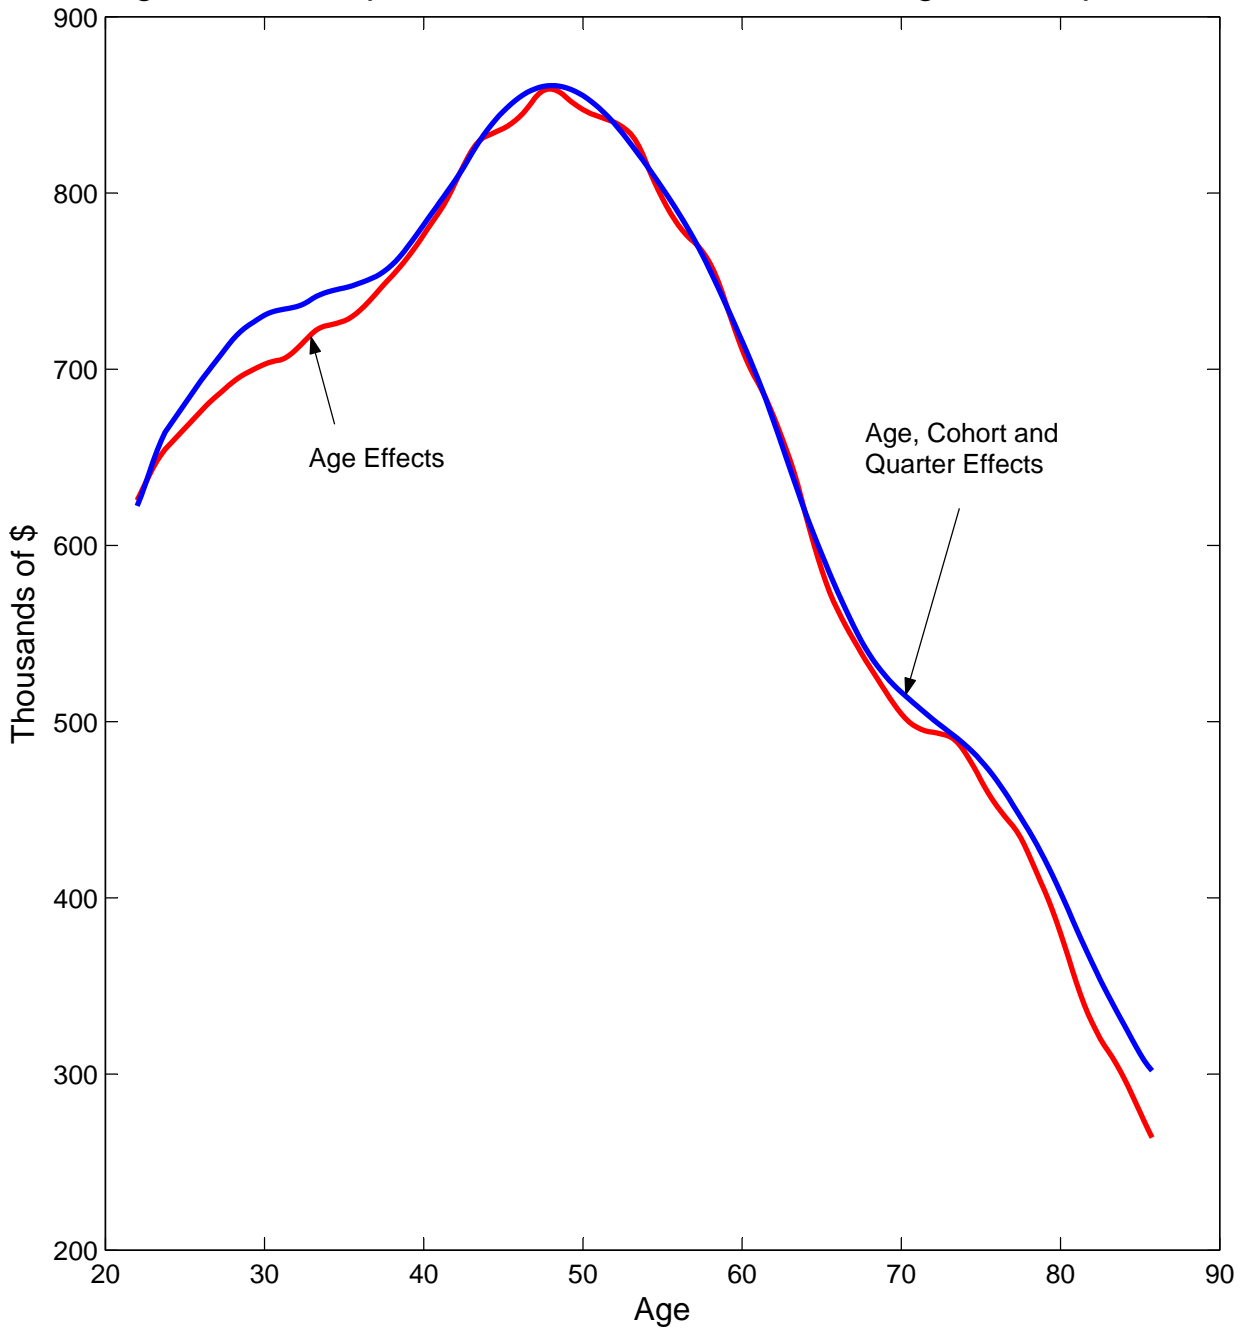


Figure 5.7.2: Expenditures Durables, Housing, adult equivalent




Figure 5.7.3: Expenditures Durables, Owned Housing, adult equivalent




Figure 5.7.4: Expenditures Durables, Rented Housing, adult equivalent

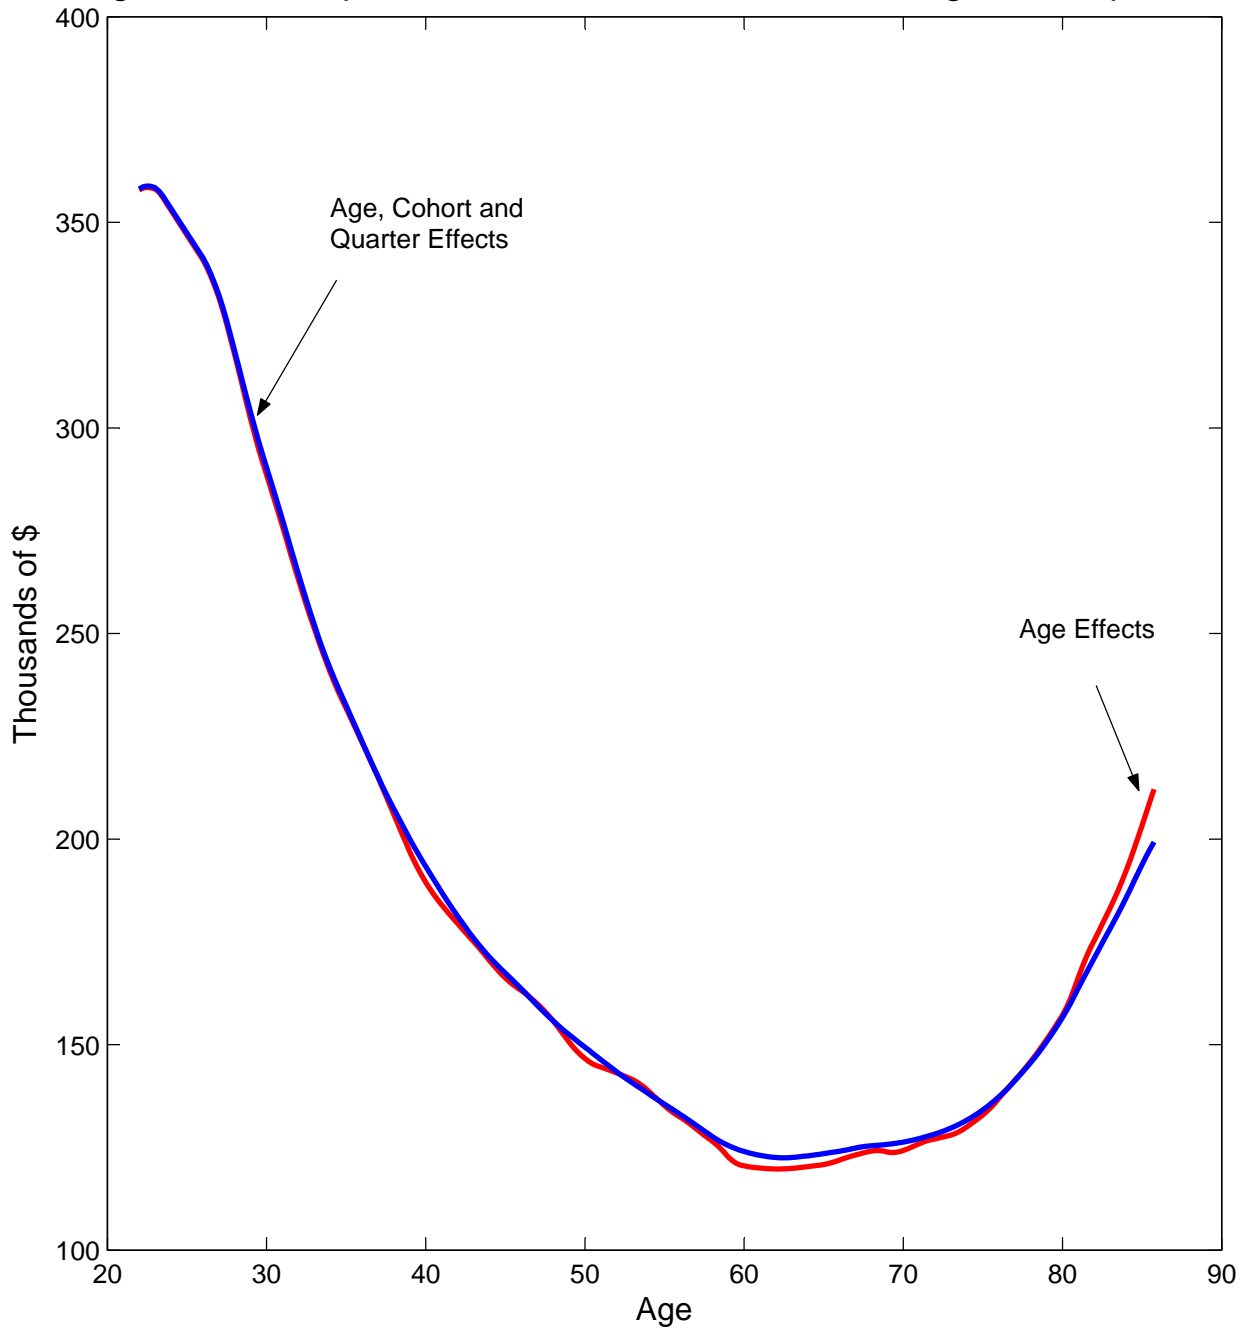


Figure 5.7.5: Expenditures Durables, Owned Housing, adult equivalent




Figure 5.7.6: Expenditures Durables, Rented Housing, adult equivalent

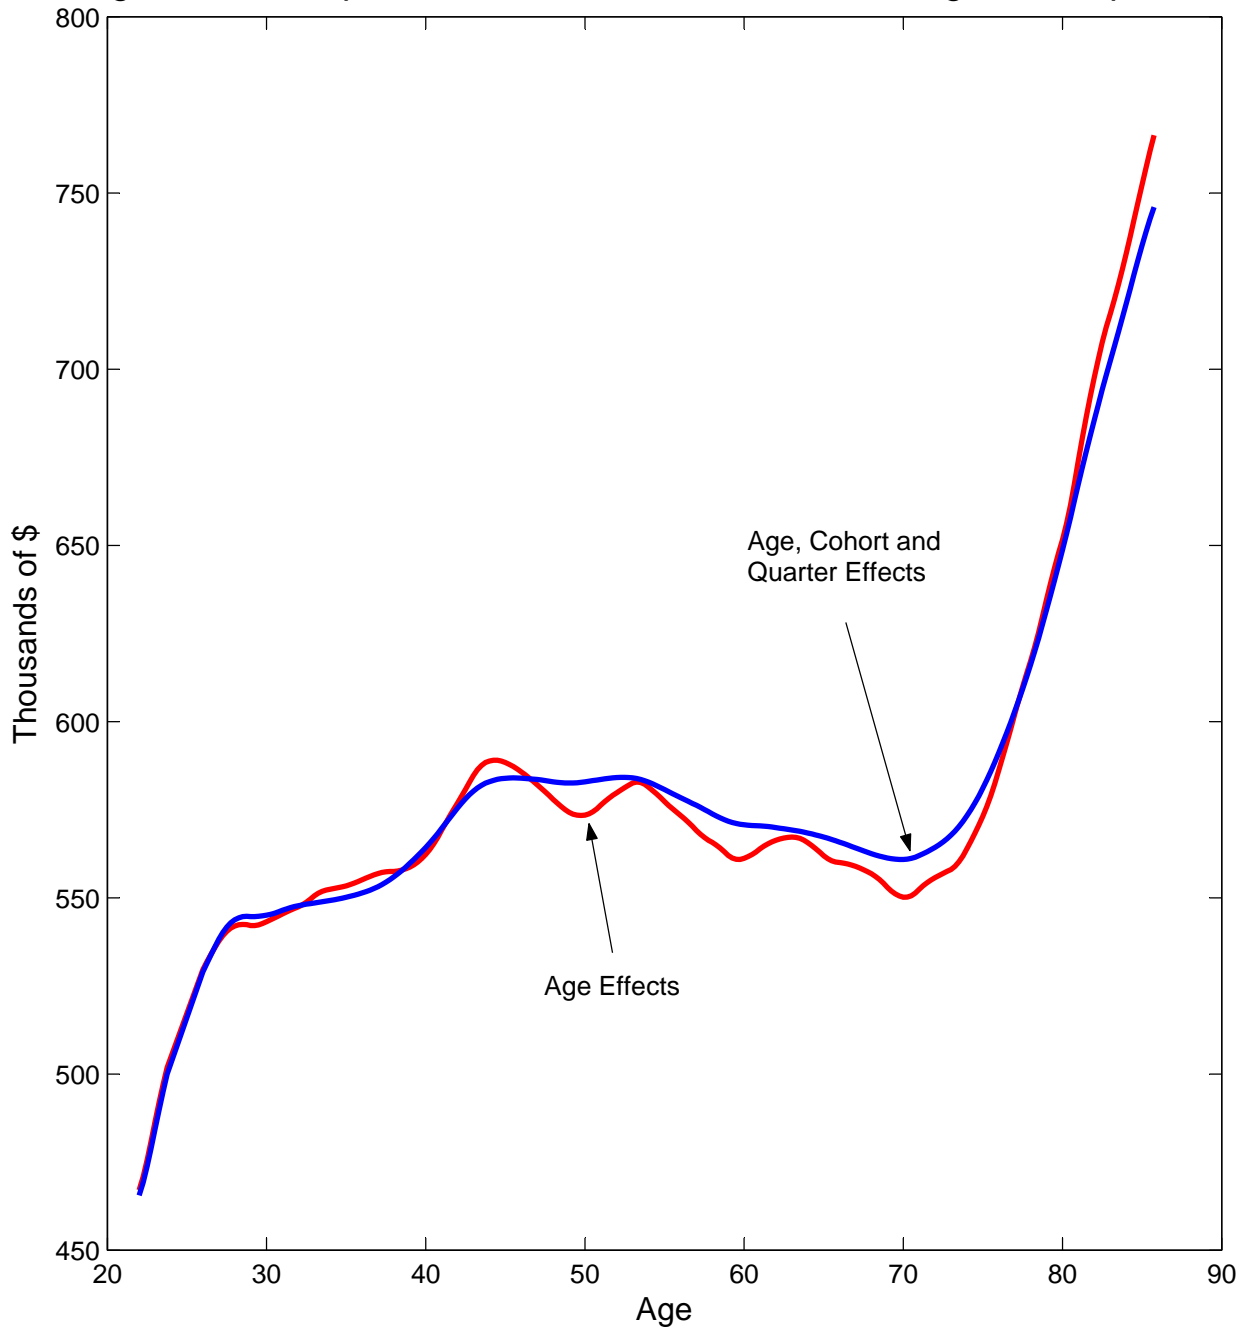


Figure 5.7.7: Equivalent Rental Value




Figure 5.7.8: Equivalent Rental Value, adult equivalent

