Volume Title: International Comparisons of Household Saving

Volume Author/Editor: James M. Poterba, Editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-67621-8

Volume URL: http://www.nber.org/books/pote94-1

Conference Date: March 18-20, 1993

Publication Date: January 1994

Chapter Title: Personal Saving in the United States

Chapter Author: Orazio Attanasio

Chapter URL: http://www.nber.org/chapters/c8871

Chapter pages in book: (p. 57 - 124)

2 Personal Saving in the United States

Orazio P. Attanasio

2.1 Introduction

Saving is an intrinsically dynamic phenomenon: what is saved is to be consumed later, either by oneself or others. As such it is hard to formalize; it is difficult to assess the importance of uncertainty, expectations, changes in income and demographic variables, and so on. The most elegant theory of saving behavior is the life-cycle model of Modigliani and Brumberg (1954). According to this theory, people save to smooth consumption in the face of an uneven income profile. Its simplest version, which assumes a constant utility function, no uncertainty, no changes in the interest rate, and perfect capital markets, has very sharp implications for the life-cycle pattern of consumption, saving, and wealth. The theory can then be extended to allow for uncertainty about income and/or life length, changing discount rates, family composition, income endogeneity, and so on. In general, saving will depend on the duration of total and working life, the nature of pension arrangements, and the shape of the age profile of earnings.

While the life-cycle model can be considered a benchmark, alternative models and modifications of the original model have been proposed. In a recent paper, Deaton (1991) analyzes the implications of liquidity constraints for optimal saving under different assumptions about the dynamics of lifetime income. Kotlikoff and Summers (1981, 1988), in a lively exchange with Modig-

Orazio P. Attanasio is professor of economics at Stanford University and at the University of Bologna, a faculty research fellow of the National Bureau of Economic Research, and a research fellow of the Centre for Economic Policy Research (London) and of the Institute for Fiscal Studies (London).

The participants at the conference and the discussant, David Wise, provided helpful comments. This project was financially supported in part by NSF grant SES-9057511. The author is very grateful to Catherine C. de Fontenay for able and dedicated research and editorial assistance. Dean Maki helped in the construction of the data set on real estate wealth.

liani (1988), have argued that most wealth is not accumulated to smooth consumption over the life cycle but rather to provide bequests. The precautionary motive for saving has also received a considerable amount of attention (see, e.g., Kimball 1990; Skinner 1988).

No matter what the main motive for saving is (life-cycle, precautionary, or bequest) it is clear that aggregate savings will depend crucially on the composition of the population. In the life-cycle case, for instance, aggregate savings will depend on the relative number of young and old consumers and on the total amount of resources available to them. It is therefore essential, in order to understand aggregate savings, to analyze and model individual behavior and consider aggregation and composition effects carefully.

Unfortunately, studies of individual saving behavior are not numerous, especially for the United States, the main reason being the lack of microdata sets containing individual data on income and consumption.[1] This paper uses data from the Consumer Expenditure Survey (CEX) from the 1980s to describe and characterize individual saving behavior. The CEX is the only U.S. microdata set that contains exhaustive information on consumption, and it has been available, on a continuous basis, since 1980.[2]

The rest of the paper is organized as follows. In section 2.2 we discuss the data sources and the statistical techniques we use in what follows. The CEX contains detailed data on income, consumption, and wealth. These data are described in detail in the first part of section 2.2; in the second part, we illustrate the statistical techniques employed in the analysis. In section 2.3 we analyze the last available cross section: that for the year 1990. We estimate and tabulate the cross-sectional age profiles of some key variables: income, consumption, saving, saving rates, and wealth. In addition to age, we also control for the level of income, the years of schooling of the household head, and the number of children. The measure of location used for all the variables considered in this section is the median. This avoids the problem of dealing with top-coded observations and makes the analysis robust to the presence of outliers.

The cross-sectional analysis in section 2.3 supplies a snapshot which is easily summarized and is useful for comparison with data from other countries. However, the interpretation of the cross-sectional profiles estimated on a single year of data as age profiles can be misleading in the presence of strong cohort effects. The availability of a time series of cross sections and the use of average

1. Early evidence on individual consumption and saving behavior is contained in Friend (1954), Goldsmith (1956), Friedman (1957), Juster (1966), and in the volume edited by Friend and Jones (1960).

2. Before then various surveys were available: the first CEX was run in 1917–18. Various students have analyzed the 1960–61 and the 1972–73 surveys. In this paper we will not use them for two reasons. First, we will stress the dynamic aspects of saving behavior and use cohort techniques: surveys that are 10 years apart are therefore of limited use. Second, and more important, many important methodological aspects of the surveys are substantially different, making a comparison extremely difficult, if not impossible. Bosworth, Burtless, and Sabelhaus (1991) compare the 1972–73 CEX to those from the 1980s.

cohort techniques helps in circumventing this problem. In sections 2.4 to 2.7 we characterize the age profiles of several variables of interest using data from the 11 available CEX surveys (1980–90) and average cohort techniques.

In section 2.4 we analyze data on disposable income and pension contributions. We estimate age profiles for both total family income and its various components. In particular, income is divided into four components: labor income, capital income, transfers, and pensions. In section 2.5 we estimate age profiles for total consumption expenditure and its components. We analyze one by one those forms of expenditure that could be considered saving: durables, education, and health. In section 2.6 we construct several definitions of saving and analyze their age profiles. The relationship between the cross-sectional distribution of saving and various controls is analyzed. In section 2.7 we estimate financial wealth–age profiles, while section 2.8 concludes the paper.

2.2 Data and Methods

2.2.1 Data

The data used in this paper are primarily from the CEX 1980–90. Since 1980 the CEX has been a rotating panel of approximately 7,000 households interviewed four times over a period of one year;[3] each quarter one-fourth of the sample is replaced by new households. The sample is representative of the population of the United States:[4] each household is assigned a weight proportional to the reciprocal of the probability of its being included. For the purpose of our analysis, the information collected in the interview can be divided into three groups: expenditure information, data on income and transfers, and other variables.

The sample unit is the so-called consumer unit. The consumer unit does not necessarily coincide with a household: it may include individuals not related to the households provided they "share responsibility for at least two out of three major types of expenses—food, housing and other expenses." In this respect, the definition is similar (but not identical) to the definition of household used in the Current Population Survey (CPS).

Expenditure Information

In each interview the household's reference person is asked to report expenditure during each of the three months preceding the interview on each of about 500 different commodity categories. This level of detail is never used in the analysis that follows. As a first step, we aggregate these 500 categories into 28 expenditure categories. This level of aggregation was chosen on the basis

3. Each household is actually interviewed five times. The first interview, however, is a contact interview from which no public data are available.

4. In 1982 and 1983, nonurban households were excluded from the population of reference. For consistency, we exclude nonurban households from our sample.

of the availability of monthly price indexes at the regional level.[5] These 28 categories are then used to create monthly household-specific price indexes for the consumption aggregate we analyze (durables, nondurables, education, etc.). This is done by taking the geometric weighted average of the relevant price indexes at the regional level, with weights given by household expenditure shares.

Annual expenditure is constructed by summing over the monthly figures. Annual price indexes are given by geometric averages of monthly price indexes, with weights given by monthly expenditures. Expenditure on items for which no price index is available is deflated by the household-specific CPI.[6] In the analysis that follows we consider six categories: durables, health, education, housing (inclusive of mortgage payments), finance charges other than mortgages, and other nondurables and services. Nominal figures are converted into real figures, when necessary, using the household-specific price indexes.[7] If one attempts to estimate aggregate Personal Consumption Expenditure (PCE) by aggregating individual figures with the appropriate weights, one typically underestimates the aggregate figures. These differences exist for three reasons, the first and foremost being the fact that the CEX data come from a recall interview: it is well known that substantial omissions arise from these interviews, especially for frequently purchased items.[8]

The second reason is the presence of definitional differences between the CEX and PCE. The most important of these are in the definition of health and housing expenditures. The CEX counts as health expenditures only out-of-pocket expenditures, while the PCE includes all expenditures on health, regardless of who ultimately pays the bill. Indeed, because health insurance refunds can be given out for expenditures incurred in months previous to those covered by the interviews, in the CEX net health expenditure can be (and for some households is) negative. The PCE includes in housing expenditures the imputed rent of owner-occupied housing; no such attribution is made for the CEX. Finally, the third reason is that the universe of reference, because of the exclusion of nonurban households, is different for the CEX and the PCE.

Given these problems, it is important to assess the importance of the difference between CEX and PCE aggregates. A detailed study of this issue has been made by Gieseman (1987) and Paulin et al. (1991), who report that for

5. Our categorization of household expenditure matches closely the Bureau of Labor Statistics (BLS) classification used in the construction of the CPI. There are few exceptions: the most noticeable is the inclusion of personal computers into "entertainment goods" along with hi-fi equipment, televisions, and so on; the BLS includes them in home furniture.

6. An example is expenditure on finance charges. The household-specific CPI is constructed by excluding this expenditure from the weights.

7. Most expenditure items are not top-coded in the CEX (exceptions include boats, airplanes, etc.): we therefore decide to ignore this problem at this point.

8. The BLS runs another survey on food and other frequently purchased items based on diaries. Large differences emerged between aggregate food expenditure between the diary and the interview surveys in the early years of the survey. These differences became negligible in 1983. In what follows we do not use diary survey data.

most components of personal expenditure the differences between CEX and PCE are roughly stable.[9] A further indication of the quality of the CEX data is the evidence in Attanasio (1993b), who reports that the correlation between the rate of growth of aggregate CEX consumption and PCE is as high as .71 over the period analyzed.[10]

Income and Wealth Data

The CEX contains detailed information on total household income and various of its components. In particular, it is possible to construct four variables: labor income, capital income, transfers (including social security, food stamps, unemployment insurance, etc.), and pension income. Extensive information is also available on tax payments and refunds. Unfortunately, the tax information is not in general matched to the income source to which it refers. It is possible to construct total household after-tax income as the sum of the various income components minus tax payments net of refunds; the components, however, are before taxes.

The income questions are always asked in the first and fourth interviews. In the second and third, these questions are asked only if the employment status of some household member has changed; otherwise the figure from the first interview is repeated. The questions on income typically refer to the 12 months before the interview, and so it is not possible to construct changes in income over quarters. The timing of the income questions has implications for the construction of saving that are discussed later.[11]

Income figures are top-coded in the CEX. Until 1982 the top-coding level was $75,000; it was raised to $100,000 in 1983. In addition, in 1980 and 1981, if any of the components of income was above the top-coding level, all income variables would be top-coded. After 1981, only the components above the top-coding level are top-coded: total income for top-coded observations can therefore exceed the top-coding level. We discuss how top-coded observations are dealt with in section 2.2.2.

In addition to top-coding, income variables are plagued by another problem: that of incomplete income responses. Incomplete responses for household income are those observations that do not report all income sources. The large majority of these observations do not report any income. We decided to exclude these observations, which account for about 12 percent of the total, from the sample we analyze.[12]

9. A similar analysis has been implemented recently by Slesnick (1992), who stresses that the amount by which the CEX underestimates PCE aggregates was very different in the 1980s relative to the 1960–61 and 1972–73 surveys.

10. A further difference between CEX and PCE aggregates is the definition of a "year" used in this paper and in Attanasio (1993b). See below.

11. The CEX contains separate files on individual household members which include information on individual members' income. We do not use this information in this paper.

12. Surprisingly enough, the average level of total consumption for these observations is not statistically different from that of the households with complete income responses. The BLS routinely publishes averages for the subsample with complete income responses.

A sizeable number of households report negative values for before- and after-tax income. After consultation with the BLS statisticians, we decided to keep these observations in the sample (most of them are self-employed individuals).

If we aggregate CEX income data and compare them to the National Income and Product Accounts (NIPA) statistics, the result is much less satisfactory than for consumption. Attanasio (1993b) reports that the correlation coefficient between the aggregated CEX disposable income growth computed using the figures published by the BLS and the NIPA equivalent series is only .21. Interestingly enough, the correlation of NIPA disposable income and the CEX aggregate obtained from the sample used in this paper and in Attanasio (1993b) is much higher (around .71).[13]

Admittedly, the income data in the CEX are not of the highest quality. Other data sets, such as the CPS or the Survey of Income and Program Participation (SIPP) are much more reliable. The CEX, however, has the unique advantage of having income and consumption data simultaneously, thereby making it possible to study *saving* at the individual level. In addition, the main features of the CEX income data are not dissimilar from those of other data sets. When we estimate income-age profiles similar to those presented in section 2.4 using CPS data, we obtain very similar results, the main difference being that the CPS profile looked much smoother than the CEX. This is probably due partly to the much larger size of the CPS sample and partly to greater measurement error in the CEX.

The CEX also contains information on financial assets held by the household as of the last day of the month preceding the fourth interview. Total assets are divided into four categories: checking accounts, saving accounts, U.S. saving bonds, and other bonds and equities. In section 2.3 we describe the 1990 cross-sectional profile for total financial assets, while in section 2.4 we aggregate these four components into *liquid assets*, which include savings and checking accounts, and into *nonliquid assets*, which include the last two categories.[14] As far as top-coding is concerned, financial asset variables are treated in the same way as income variables; the top-coding level is also the same.

The quality of the CEX data on financial assets is equivalent to that of the data on income. The survey is not designed to investigate asset holding in detail. Nor can it be used to estimate aggregate household wealth: the relatively large number of nonresponses for these variables, the low level of top-coding, and the lack of oversampling among wealthy households prevent it. On the

13. The BLS published figures and the aggregates from our sample are different for three reasons. First, we exclude several groups of households from our sample (see the selection rules used). Second, the way in which we assign observations to time periods is different. Third, the BLS has access to the original data; therefore, there are no top-coded observations.

14. In addition to stock variables in the CEX, there is also a question on the change in the stock over the last year. Unfortunately, the quality of the answers to these questions is dubious: there are many missing values and a large number of zeros. We therefore decided not to use this information.

other hand, the main features of the financial asset data are similar to those found in other surveys. Attanasio (1993a) reports that the financial wealth–age profiles estimated using the CEX are similar to those typically estimated in the literature. Furthermore, while it is true that the proportion of households reporting zero assets is higher than in other surveys (about 18 percent overall, 15 percent in our sample), if one looks at the proportion of households with, say, less than $5,000, one gets figures similar to those of other data sources.[15]

The information on real estate wealth was very limited in the early years of the survey.[16] However, since 1988, the BLS has started releasing detailed information on the value of household estate property, as well as details on outstanding mortgages. In section 2.3, we use the 1990 data to estimate the cross-sectional profile of the net and gross value of real estate wealth. The gross value is the defined as the total market value of all real estate owned by the household (even if it does not live in it) and should therefore reflect capital gains and losses. The net value is obtained from the gross by subtracting total outstanding debt on real estate.

Other Variables

The "income and characteristics" file of the CEX contains information on about 500 different socioeconomic variables that range from pension contributions to family composition, region of residence, education, and so on. Detailed information on the variables available is in the CEX manual. In this section we will discuss briefly the variables that are used in the analysis below.

The information on the region of residence is not very detailed. The United States is divided into the four standard census regions: Northeast, Midwest, South, and West. In addition to the region of residence, some information is provided on the size of the city of residence.[17]

Information on pension contributions is extremely valuable to the analysis of savings. The CEX contains information on employees' contributions to both private and government pension schemes in the form of deductions from their pay. In addition, we have information on contributions to individual retirement accounts (IRAs). No information on employers' contributions to pension schemes is provided.

15. For a comparison of wealth data from the Survey of Consumer Finances, PSID, and SIPP, see Curtin, Juster, and Morgan (1987).

16. We know whether the household rents or owns (possibly with a mortgage) the house it lives in. We also know whether the consumer units live in student housing. We exclude these observations from our sample. The few observations of units that live in a nonowned house without paying rent are aggregated with the renters. Finally, a question is asked on the market value of owned homes. Unfortunately, the answers to this question are very few and extremely unreliable.

17. This information is absent for the West. Most restrictions on data availability in the CEX are motivated by confidentiality. The reason for the exclusion of city size information for the West is that there is only one city with more than 4 million people in the West.

Attrition, Sample Size, and Construction of Saving

The response rate in the CEX is reasonably good. The BLS reports that about 15 percent of the households contacted do not participate in the survey. Of the remaining 85 percent, we exclude an additional 12 percent because of incomplete income responses (see above). Various additional selection criteria were used to eliminate observations that presented apparent inconsistencies. In particular, we eliminated all observations with missing consumption or income data, those for which the age of the reference person increases by more than one during the period of interview, and those living in student housing.

In section 2.3, we focus on the 1990 sample and on 11 age groups, while in the following sections we use data from 1980–90 to analyze the behavior of 10 cohorts, defined on the basis of their year of birth. We eliminate from the sample used for the cross-sectional analysis of section 2.3 all households headed by individuals younger than age 21 or older than age 75. The sample used in the cohort analysis of sections 2.4 to 2.7 excludes all the households headed by individuals born before 1910 or after 1959. We refer to the former as the *cross-sectional* sample and to the latter as the *cohort* sample.

After all selection criteria are used, we are left with 4,623 households in the cross-sectional sample and 47,647 households in the cohort sample. The definition and size of the age groups of the cross-sectional sample are reported in table 2.1. In table 2.2 we report, for each year from 1980 to 1990, the size of the cohort sample.[18]

While, in theory, each household is interviewed four times, not all households complete the four interviews. It is common for a household to drop out of the sample and/or to miss an interview (not necessarily the last). In table 2.3, we report the number and percentage of households completing any set of interviews. As can be seen, only half the households complete four interviews. We did not eliminate households with fewer than four interviews, deciding rather to make some adjustment in the computation of saving.

Saving is defined as disposable income minus consumption. Disposable income is defined as total family income net of taxes *and* social security contributions. Deductions for and contributions to private and government retirement schemes are included in income but not in consumption. Therefore they are considered as saving. Employers' contributions to pensions are not considered as saving because of the lack of data on this item. Several definitions of saving can be constructed by using different definitions of consumption.

18. As explained below, we aggregate all the interviews of a given household: therefore we have only one observation per consumer unit. In 1986 the CEX sample was discontinued, so that it is not possible to follow into 1986 households that had their first or second interview in 1985. This explains the larger sample sizes of 1985 and 1990 (which is the last year in the sample): households that normally would have completed their cycle of interviews in 1986 and 1991 are not observed so that their "last" recorded interview is in 1985 and 1990. The size of the cohort sample in 1990 differs from the cross-sectional sample because the definition of the age groups in the latter does not coincide with the definition of the cohorts in the former.

Table 2.1 **1990 CEX: Sample Size by Age Group**

Age Group	Observations
21–25	497
26–30	639
31–35	640
36–40	572
41–45	523
46–60	375
51–55	335
56–60	251
61–65	259
66–70	303
71–75	229

Table 2.2 **CEX: Sample Size by Year**

Year	Observations
1980	4,027
1981	3,920
1982	3,954
1983	4,140
1984	4,105
1985	7,200
1986	4,315
1987	4,221
1988	3,507
1989	3,655
1990	4,603

Table 2.3 **Interviews Completed**

Interviews	Frequency	Percentage
2	4,406	9.25
3	658	1.38
3,2	3,316	6.96
4	607	1.27
4,2	88	0.18
4,3	556	1.17
4,3,2	2,923	6.13
5	3,203	6.72
5,2	71	0.15
5,3	109	0.23
5,3,2	308	0.65
5,4	2,993	6.28
5,4,2	330	0.69
5,4,3	3,558	7.47
5,4,3,2	24,521	51.46

The time span used to define the flow of consumption is the year. This is done so as to match the time spans to which consumption and income refer. Therefore, while income is taken from the last completed interview, annual consumption is defined as the sum of the monthly figures from all the interviews. For households that do not complete the four interviews, the total figure for consumption is adjusted to take into account the fact that it refers to less than 12 months. Of course this introduces some measurement error that can be particularly severe in the case of durables.

Households are interviewed every month, and, as discussed, consumption and income refer to the 12 months preceding the interviews. To construct annual aggregates, it is therefore necessary to assign households to a specific year. We chose to assign to year $n - 1$ all the households interviewed between July of year $n - 1$ and June of year n.[19]

2.2.2 Statistical Methods

All the cell statistics in sections 2.3 to 2.7 are obtained by weighting each household by the corresponding weight from the Bureau of the Census, which is proportional to the inverse of the probability of the household's being included in the sample. This procedure is used to compute percentiles (in sections 2.3 to 2.7) as well as means (estimated either by sample mean or by maximum likelihood in sections 2.4 to 2.7). The results were not dramatically affected by the weighting scheme.

The analysis of section 2.3 is based on group medians and/or other percentiles and does not present any difficulty. All the figures in that section are in current dollars.

To analyze an intrinsically dynamic phenomenon such as saving, one would like to follow the same individuals over time. If long panel data are not available, one can circumvent this difficulty by using average cohort techniques. This has the advantage, relative to the cross-sectional analysis used in section 2.3, that it controls for cohort effects. In the presence of cohort effects, the cross-sectional profile of a variable such as consumption, income, or saving, observed in a given year, might not correspond to the age profile of any individual in the population. Shorrocks (1975) constructs an example in which every individual in the population has a strictly increasing wealth-age profile and yet, because of strong cohort effects, the cross-sectional profile at a given point in time will be "hump-shaped."

19. The BLS follows a different procedure. First it constructs monthly income figures. Then it aggregates the consumption and income figures for all households in a given year. As a consequence, observations on a given household are going to be divided between different periods, except for those households interviewed in January. Given that I am interested in matching exactly income and consumption for a given household in order to construct savings, I preferred to use the alternative procedure described in the text, even though the assignment to a given year might sound arbitrary.

The use of average cohort techniques amounts to following as they age individuals born in the same time interval.[20]

Consider a variable of interest X_t^{ch} observed for household h, belonging to cohort c at time t. It is always possible to define ε_t^{ch} by the following equation:

$$X_t^{ch} = \delta_t^c + \varepsilon_t^{ch}, \tag{1}$$

where δ_t^c is a measure of location for the cell defined by households belonging to cohort c and observed at time t. The age corresponding to cell (c,t) is given, if we identify c by the year of birth, by $t - c$.

In this paper, cohorts are defined by five-year bands. We analyze a total of 10 cohorts; data for each cohort are averaged in every year: this gives us a total of 110 cells. The cohort definition, the median age in 1980 and in 1990, and the average cell size in our sample are reported in table 2.4. The relative sizes of different cohort cells reflect approximately the composition by age of the U.S. population.[21] Changes in size across years reflect changes in the dimension of the total sample.

If δ_t^c in equation (1) is the cell mean and the data present no particular problems, a sensible and robust estimator is given by the sample mean. If, however, some observations are top-coded, as is the case both for income and financial wealth, estimation of the mean is substantially more complicated. In general, it will not be possible to use nonparametric methods. If we are willing to parametrize a density function for the cross-sectional distribution and believe that such a density fits the (unobserved) tail of the distribution, it will be possible to estimate the mean by maximum likelihood. The reliability of such an estimate hinges in a crucial way on the parametric specification used: in this respect it is important to use a flexible functional form, capable of allowing for the substantial amount of skewness and kurtosis which characterizes both income and wealth distribution.

Consistent estimates of the population quantiles can be easily obtained using the sample quantiles, as long as the top-coding level is above the quantile we are interested in estimating.[22] The sample quantiles can be compared to the quantiles of the estimated density as a specification test. When we perform this exercise in our applications, we obtain satisfactory results.

To parametrize the cross-sectional distribution of income we choose a mix-

20. Cohort techniques have been used by Browning, Deaton, and Irish (1985) and are discussed by Deaton (1985) and more recently by Moffitt (1991) and Attanasio (1993b). Particular emphasis is usually given to means; other measures of locations can however be used.

21. Of course there are sampling errors and possible biases caused by differing attrition and nonresponse rates.

22. This is true for income and wealth, where we can safely assume that the top-coded observations are on the right tail of the distribution. The same is not necessarily true for savings. The bias introduced is however much less than for the estimation of the mean.

Table 2.4 Cohort Definition and Cell Size

Cohort	Year of Birth	Age in 1980	Age in 1990	Average Cell Size
1	1955–59	23	33	759
2	1950–54	28	38	672
3	1945–49	33	43	580
4	1940–44	38	48	432
5	1935–39	43	53	350
6	1930–34	48	58	325
7	1925–29	53	63	334
8	1920–24	58	68	340
9	1915–19	63	73	295
10	1910–14	68	78	243

ture of two normal densities with different means and variances. Therefore we fit five parameters for each cell.

The unconditional distribution of financial asset holdings is parametrized as a distribution with mass p at zero (where p is the proportion of observations with zero assets) and the remaining mass distributed on the positive axis as a mixture of two log-normal distributions, which results in six parameters to be fitted for each cell.

Finally, these techniques can also be modified to control for within-cell heterogeneity. One can either define new cells, interacting year cohort dummies with other dummies (education, race, and sex of the household head), or allow these variables to affect the within-cell conditional mean, possibly imposing across-cell restrictions.

To analyze data from several years it is necessary to transform current into constant dollars. Given that we construct household-specific price indexes (see section 2.2.1), we have a choice between estimating cell means in nominal terms and deflating them by cell-specific price indexes[23] and estimating cell means of quantities deflated at the individual level. We chose the former alternative to avoid the possibility that measurement error in consumption would affect income indirectly through the individual price level, while preserving the heterogeneity in price indexes determined by the different expenditure patterns of different cohorts. In practice, the two procedures give extremely similar results.

2.3 Cross-Sectional Analysis

In this section we use the 1990 CEX survey, which includes all households interviewed during 1990. In what follows, we tabulate and plot the cross-sectional profiles for income, consumption, saving, and wealth. Because in-

23. The cell-specific price indexes are constructed by averaging individual price indexes in each cell.

come and consumption refer to the 12 months preceding the interview, for most households the period of reference includes some months of 1989. All figures are obtained using current dollars: no attempt was made to correct for inflation. Given the short length of the period considered, we do not think this is a serious issue. The sample is divided into the 11 age groups described in section 2.2.1.

2.3.1 Income

In the first column of table 2.5 we tabulate median disposable income[24] by age group for the total sample. In the other four columns we divide each age group on the basis of income quartiles and compute the median within each group.[25] The numbers in parentheses are cell sizes.[26] The figures in table 2.5 are plotted against age in figure 2.1. The evidence is not particularly surprising: the cross-sectional profile for income presents the usual hump shape. Family income peaks between ages 41 and 55 (it is reasonably flat over that interval) and declines afterward. The difference among the percentiles plotted in figure 2.1 evidentiates the degree of inequality of the cross-sectional distribution of income at various ages. Inequality is more pronounced for the central age groups as the cross-sectional profile is steeper for higher percentiles than for lower percentiles.

For reasons discussed below, we divide the sample not only by income quartiles but also on the basis of the educational attainment of the household head. The three groups we consider are high school dropouts, high school graduates, and college graduates. The cross-sectional profile of median household after-tax income of these three groups is reported in table 2.6 and plotted in figure 2.2. Household income peaks slightly earlier for high school dropouts than for high school and college graduates. In addition, the hump is much more pronounced for more highly educated households: this indicates that the returns to education,[27] and especially college education, increase strongly with age, at least until age 50. Median income for college graduates in the 51–55 age group is 1.67 times that of high school graduates in the same age group and 2.84 times that of high school dropouts. The same figures in the 26–30 age group are 1.38 and 1.88, respectively. The median income of high school graduates is (not surprisingly) just around the overall median.

Some of the cells (especially for college graduates) are very small: for instance we have only 19 households whose head is in the 71–75 age group and

24. As stressed in section 2.2, we subtract Social Security contributions from the BLS definition of "after-tax family income."

25. This is done only for comparability with the subsequent tables that report median consumption and saving by income quartiles. The second to fifth columns table 2.5 are the 12.5, 37.5, 62.5, and 77.5 percentiles of each age group.

26. These are not the same in the second to fifth columns because income quartiles (and medians) are computed on the basis of the Bureau of Census weights.

27. The figures reported in table 2.6 refer to total household income, not wages. As a consequence they can vary because of variations both in wages and in household labor supply.

Table 2.5 **Median Disposable Income by Age Group and Income Quartile**

Age Group	Total Sample	Below First Quartile	Between First and Second Quartiles	Between Second and Third Quartiles	Above Third Quartile
21–25	12,855	4,486	10,375	15,699	29,890
	(497)	(120)	(120)	(124)	(133)
26–30	20,663	8,428	15,689	25,738	41,557
	(639)	(152)	(157)	(160)	(170)
31–35	24,137	8,609	18,847	30,377	49,909
	(640)	(151)	(148)	(163)	(178)
36–40	28,057	8,726	21,920	33,482	56,099
	(572)	(141)	(138)	(146)	(147)
41–45	31,573	11,040	24,980	39,836	64,280
	(523)	(120)	(132)	(135)	(136)
46–50	29,719	8,920	24,728	40,626	65,886
	(375)	(83)	(96)	(92)	(104)
51–55	30,127	9,251	22,003	37,852	62,435
	(335)	(79)	(85)	(85)	(86)
56–60	25,511	8,409	19,521	29,452	55,252
	(251)	(63)	(56)	(60)	(72)
61–65	20,617	6,300	16,923	28,871	56,112
	(259)	(64)	(62)	(68)	(65)
66–70	17,718	6,608	12,759	22,178	45,137
	(303)	(75)	(78)	(78)	(72)
71–75	13,996	6,535	11,331	17,879	31,666
	(229)	(52)	(65)	(53)	(59)

Note: Numbers in parentheses are cell sizes.

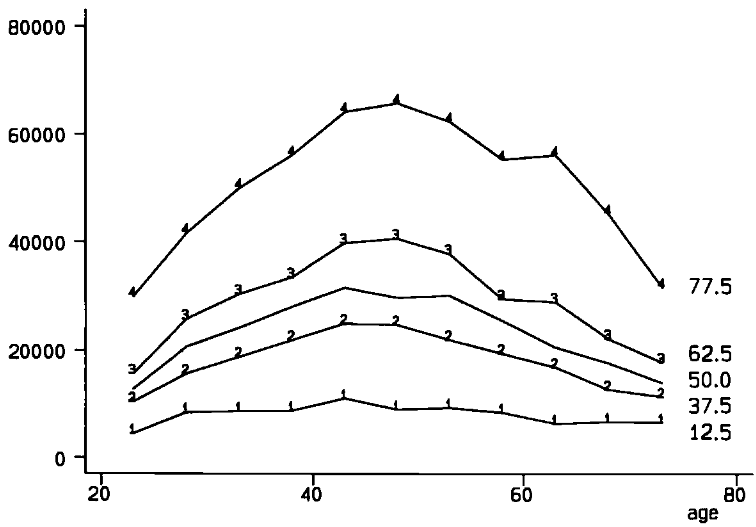

Fig. 2.1 Median and various percentiles of disposable income

Table 2.6 **Median Disposable Income by Age and Education Groups**

Age Group	Total Sample	High School Dropouts	High School Graduates	College Graduates
21–25	12,855	12,649	12,235	16,566
	(497)	(64)	(342)	(91)
26–30	20,663	13,911	18,937	26,108
	(639)	(63)	(388)	(188)
31–35	24,137	12,688	22,138	36,083
	(640)	(65)	(399)	(176)
36–40	28,057	18,050	25,660	38,111
	(572)	(59)	(316)	(197)
41–45	31,573	19,000	26,464	46,772
	(523)	(60)	(288)	(175)
46–50	29,719	13,341	28,307	47,900
	(375)	(69)	(200)	(106)
51–55	30,127	16,748	28,520	47,572
	(335)	(74)	(193)	(68)
56–60	25,511	16,095	25,894	39,653
	(251)	(84)	(105)	(62)
61–65	20,617	16,176	18,184	37,686
	(259)	(79)	(120)	(60)
66–70	17,718	13,618	16,610	30,993
	(303)	(101)	(154)	(48)
71–75	13,996	11,648	15,244	28,597
	(229)	(113)	(97)	(19)

Note: Numbers in parentheses are cell sizes.

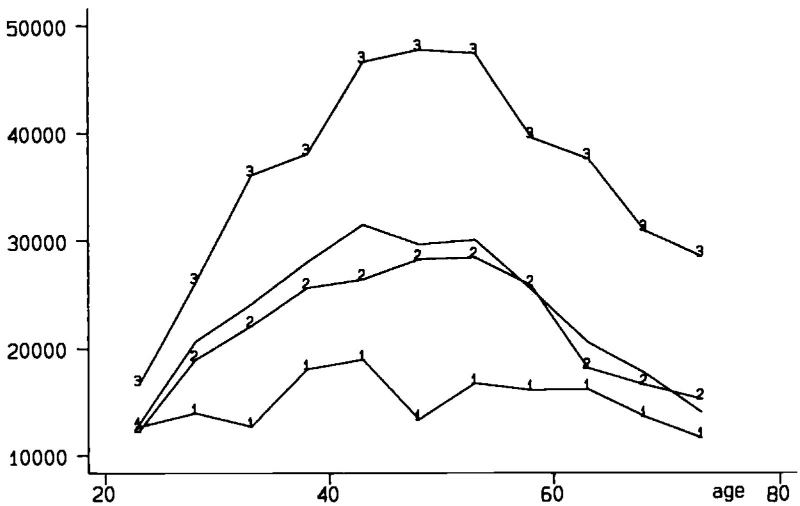

Fig. 2.2 Median income cross-sectional profile by education group

is a college graduate. Therefore we need to use some caution in interpreting these and other statistics computed on such small cells.

2.3.2 Consumption

The consumption definition we use in this section is the closest we can get to the NIPA. Therefore it includes all consumption expenditures reported in the CEX. As discussed in section 2.2, the main differences are in that it includes only out-of-pocket health expenditure and that it excludes imputed rents on owner-occupied housing. Pension contributions are not included in consumption.

Median consumption cross-sectional profiles for the total sample and by income quartile are reported in table 2.7 and plotted in figure 2.3. Median consumption by age-education group is tabulated in table 2.8 and plotted against age in figure 2.4. The cross-sectional profile of consumption is similar to that of income in that it presents a pronounced hump. However, the consumption profile peaks slightly earlier and is flatter than the income profile.

The differences in consumption cross-sectional profiles across education groups mirror those observed in income profiles. This evidence is consistent with that reported, for instance, by Carroll and Summers (1991), who show that differences in cross-sectional age profiles for consumption across different occupational groups parallel corresponding differences in income profiles.[28] Median consumption of college graduates in the 51–55 age group is 1.5 times that of the high school graduates and 2.23 times that of the high school dropouts in the same age group. The same figures are 1.27 and 1.83 respectively for the 26–30 age group.

2.3.3 Saving

In table 2.9 we report median saving by age group and by age-income quartile group and plot them in figure 2.5. Saving is always negative for the lowest income group. Median saving is also negative for the two youngest and the oldest groups of households whose disposable income is between the first and second income quartiles (for their age group). In general, saving increases with income and, with the exception of households with income below the first quartile, increases until ages 51–55 and declines afterward. This hump shape is more pronounced for the households with income in the highest income quartile.

Median saving rates for the whole sample and by disposable income quartile are tabulated and plotted in table 2.10 and figure 2.6, respectively.[29] Saving rates are very flat for the two groups above the median, exhibit a substantial

28. Attanasio and Browning (1992) argue that this is not necessarily inconsistent with the life-cycle model. It should be remembered, for instance, that the profiles considered here do not control for either cohort effects or changes in family composition.

29. We do not plot the saving rates for the first income group because their variability would swamp that of the other groups.

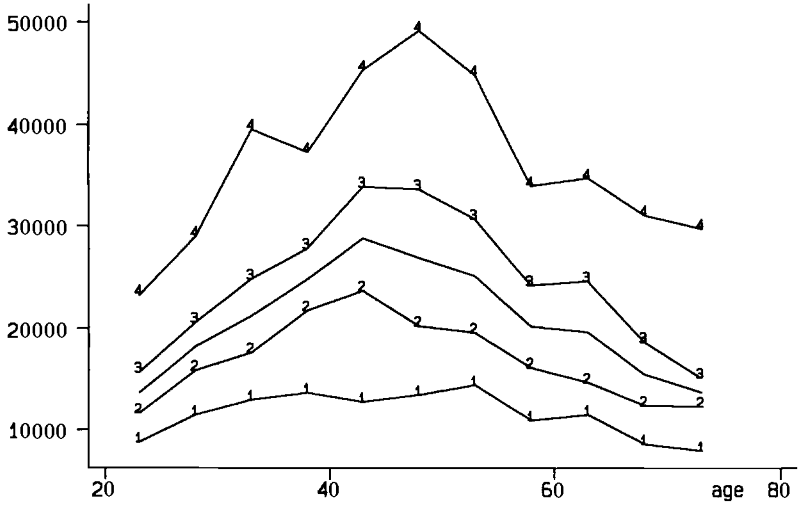

Fig. 2.3 Median consumption cross-sectional profile by income quartile

Table 2.7 Median Total Consumption by Age Group and Income Quartile

Age Group	Total Sample	Below First Quartile	Between First and Second Quartiles	Between Second and Third Quartiles	Above Third Quartile
21–25	14,359	8,771	11,644	15,728	23,279
26–30	18,902	11,503	15,927	20,618	29,005
31–35	21,987	12,992	17,651	24,845	39,564
36–40	24,820	13,693	21,758	27,839	37,299
41–45	28,041	12,785	23,713	33,868	45,294
46–50	24,467	13,489	20,245	33,624	49,109
51–55	25,704	14,455	19,581	30,670	44,776
56–60	19,653	10,934	16,199	24,217	33,934
61–65	21,042	11,513	14,705	24,608	34,696
66–70	16,823	8,563	12,423	18,672	31,018
71–75	14,362	7,899	12,323	15,177	29,705

Note: Cell sizes are the same as in table 2.5.

hump for the households with disposable income between the first and second quartiles, and are negative and extremely variable for households with income below the first quartile.

The fact that, beside their life-cycle dynamics, both saving and saving rates are an increasing function of disposable income is not inconsistent with the life-cycle permanent-income theory. Households in the lowest income quartiles are affected, on average, by lower transitory income shocks than households in higher income quartiles. According to the theory, transitory

Table 2.8 Median Total Consumption by Age and Education Groups

Age Group	Total Sample	High School Dropouts	High School Graduates	College Graduates
21–25	14,359	13,343	13,453	16,543
26–30	18,902	12,372	17,896	22,699
31–35	21,987	15,246	20,112	29,732
36–40	24,820	19,260	23,160	32,681
41–45	28,041	16,790	25,170	37,925
46–50	24,467	15,907	24,285	41,645
51–55	25,704	17,276	25,704	38,952
56–60	19,653	13,484	19,653	32,669
61–65	21,043	15,267	19,972	32,115
66–70	16,823	12,301	16,273	24,929
71–75	16,362	11,686	15,245	35,280

Note: Cell sizes are the same as in table 2.6.

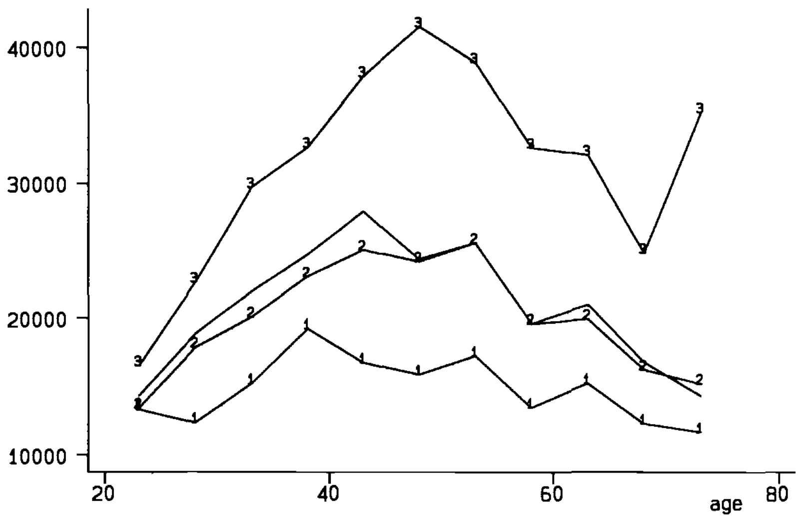

Fig. 2.4 Median consumption cross-sectional profile by education group

shocks to income should be smoothed by saving. Households receiving large negative shocks should dissave to smooth consumption.[30]

To characterize the relationship between saving and income we should therefore divide the households in the sample according to their permanent rather than disposable current income. An instrumental variable approach is to divide the sample on the basis of a variable which is correlated with permanent

30. If measurement error in income is uncorrelated with that in consumption, it introduces a similar bias: measurement error in income is fully reflected in saving.

Table 2.9 Median Saving by Age Group and Income Quartile

Age Group	Total Sample	Below First Quartile	Between First and Second Quartile	Between Second and Third Quartile	Above Third Quartile
21–25	−1,431	−4,620	−1,857	723	5,673
26–30	1,414	−2,951	−58	4,905	13,910
31–35	1,570	−3,925	1,065	4,648	11,299
36–40	2,341	−4,329	1,395	5,657	18,922
41–45	2,895	−3,732	1,359	6,249	15,635
46–50	2,461	−4,060	2,681	7,278	19,789
51–55	3,411	−5,111	2,805	8,293	22,331
56–60	2,574	−2,862	4,523	6,525	21,606
61–65	1,059	−7,148	1,766	4,531	18,124
66–70	914	−2,017	465	3,383	11,930
71–75	139	−600	−1,535	2,840	6,413

Note: Cell sizes are the same as in table 2.5.

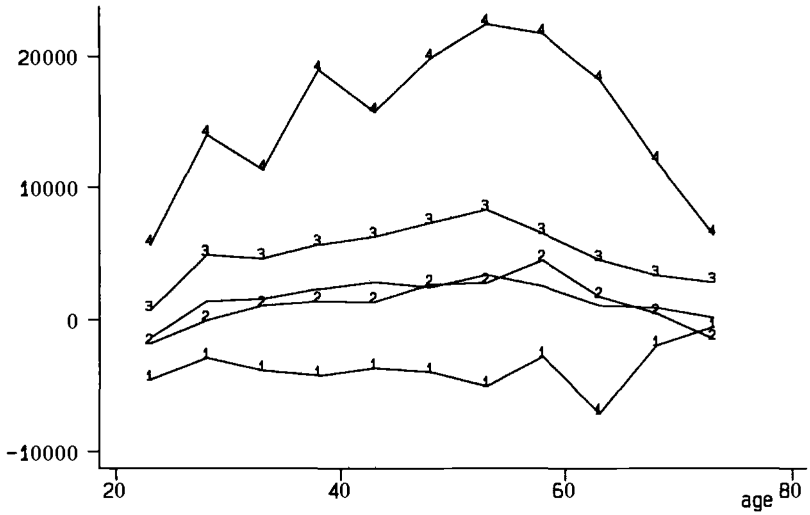

Fig. 2.5 Median saving cross-sectional profile by income quartile

income and uncorrelated with transitory shocks. For such a purpose we use the educational attainment of the household head.[31]

In tables 2.11 and 2.12 we report median saving and saving rates by age-education group. The same figures (with the exception of saving rates for the lowest education group) are plotted in figures 2.7 and 2.8.

For high school dropouts, median saving and saving rates are negative for 4

31. Of course this approach would not solve the problem if different education groups have been affected by different transitory shocks in the year considered.

Table 2.10 Median Saving Rates by Age Group and Income Quartile

Age Group	Total Sample	Below First Quartile	Between First and Second Quartiles	Between Second and Third Quartiles	Above Third Quartile
21–25	−11.8	−111.8	−16.7	4.3	20.0
26–30	7.4	−28.8	−0.4	19.2	33.8
31–35	7.1	−44.0	5.4	17.2	24.2
36–40	9.4	−55.3	5.4	17.1	35.6
41–45	9.8	−35.3	5.1	13.8	28.2
46–50	11.2	−33.5	12.4	20.3	32.1
51–55	13.9	−65.4	12.4	22.7	33.9
56–60	16.6	−34.0	26.6	23.8	41.4
61–65	8.6	−148.3	11.0	14.4	34.7
66–70	7.1	−28.4	4.2	16.9	24.4
71–75	1.1	−8.1	−15.8	14.5	21.0

Note: Cell sizes are slightly different from those in table 2.5 because of zero-income observations.

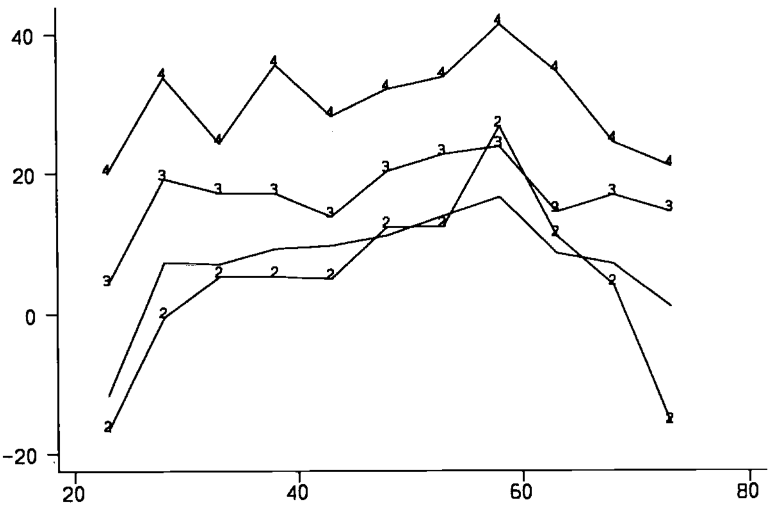

Fig. 2.6 Median saving rates cross-sectional profile by income quartile

of the 11 age groups, and not as variable as for the households in the lowest income quartile. For this group, both saving levels and saving rates do not have a distinctive pattern over age. For the other two groups, the level of saving exhibits a distinctive hump which is more pronounced for college graduates. Saving rates are substantially flat for college graduates, while for high school graduates they are highest before retirement.

Saving rates are not defined for households with zero or negative income. It is therefore useful to consider the ratio of saving to consumption rather than

Table 2.11 Median Saving by Age and Education Groups

Age Group	Total Sample	High School Dropouts	High School Graduates	College Graduates
21–25	−1,431	60	−2,245	298
26–30	1,414	112	1,013	2,960
31–35	1,570	−651	1,157	4,074
36–40	2,341	−426	1,843	4,879
41–45	2,895	−50	1,378	7,695
46–50	2,461	−1,338	3,374	5,209
51–55	3,411	1,537	3,488	5,762
56–60	2,574	1,725	4,904	2,174
61–65	1,059	909	−198	5,294
66–70	914	900	1,236	1,524
71–75	139	139	584	−3,287

Note: Cell sizes are the same as in table 2.6.

Table 2.12 Median Saving Rates by Age and Education Groups

Age Group	Total Sample	High School Dropouts	High School Graduates	College Graduates
21–25	−11.8	0.7	−18.8	2.2
26–30	7.4	1.6	6.9	14.8
31–35	7.1	−7.2	6.0	11.5
36–40	9.4	−2.0	7.7	16.1
41–45	9.8	−0.2	7.0	19.0
46–50	11.2	−13.4	13.6	14.0
51–55	13.9	8.8	14.4	17.7
56–60	16.6	16.6	21.5	10.6
61–65	8.6	10.1	−2.0	15.2
66–70	7.1	4.2	6.0	7.7
71–75	1.1	1.7	5.3	−10.0

Note: Cell sizes are slightly different from those in table 2.6 because of zero-income observations.

to income. This variable, besides being defined at zero income, has several advantages. First, it is a monotonic transformation of saving rates (for those values of income when they are defined). Second, consumption might be a more appropriate denominator because, in theory, it reflects variations to permanent income and is therefore less affected by transitory shocks. Third, the monotonic transformation that maps saving rates into the variable we consider has the effect of damping extreme observations, just as a log transform or a Box-Cox transform would do.[32] This can be useful given the enormous variability of individual saving rates.

32. Think, for instance, what happens when income goes to zero. Saving rates diverge to minus infinity, whereas our measure goes to minus one.

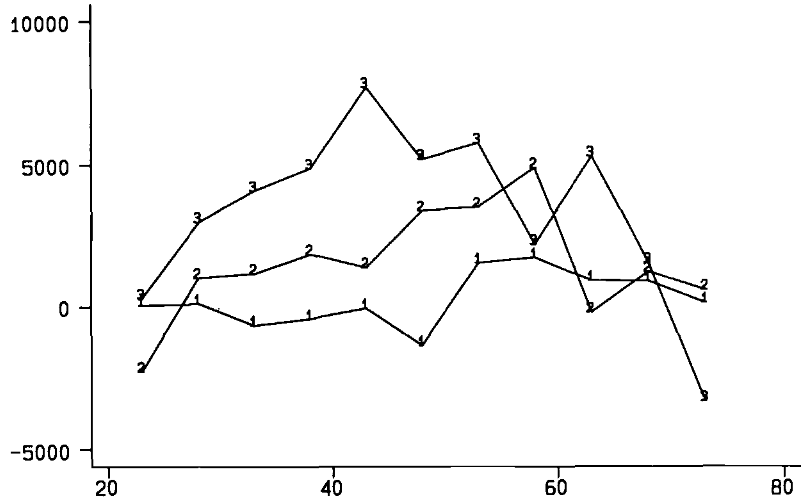

Fig. 2.7 Median saving cross-sectional profile by education group

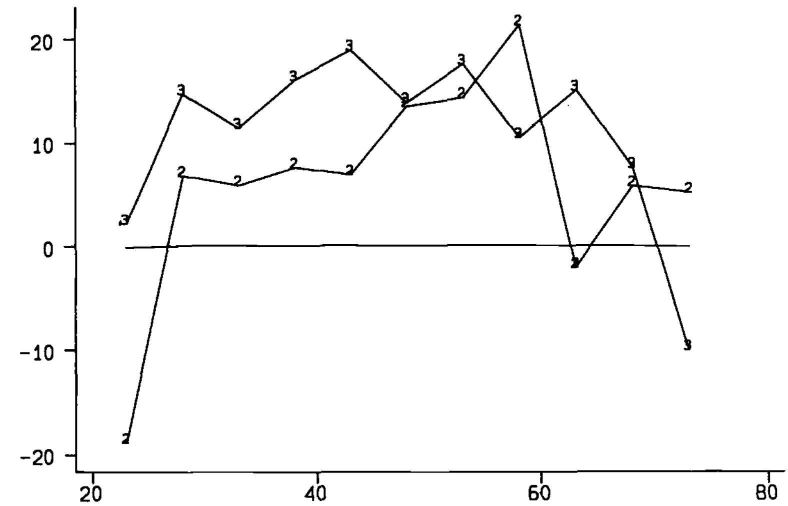

Fig. 2.8 Median saving rates cross-sectional profile by education group

Ratios of median saving to consumption by age group for the whole sample and by income quartile and education group are reported in tables 2.13 and 2.14 and the latter are plotted against age in figure 2.9.

The picture that emerges is substantially similar to that in tables 2.10 and 2.12, and it does not deserve further comment except to note that the corresponding figure 2.9 is, as expected, substantially smoother than figures 2.6 and 2.8 for saving rates.

Table 2.13 **Median Saving to Consumption Ratio by Age Group and Income Quartile**

Age Group	Total Sample	Below First Quartile	Between First and Second Quartile	Between Second and Third Quartile	Above Third Quartile
21–25	−10.8	−53.2	−14.3	4.5	24.9
36–30	7.5	−26.8	−0.4	23.7	51.0
31–35	7.5	−31.2	5.8	20.8	32.0
36–40	9.8	−38.9	5.7	20.7	55.3
41–45	10.8	−26.1	5.4	16.1	39.2
46–50	11.1	−31.0	14.1	25.5	47.2
51–55	16.0	−39.8	14.2	29.3	51.2
56–60	18.6	−40.2	36.3	31.2	70.6
61–65	9.5	−59.7	12.4	16.9	53.0
66–70	6.3	−25.7	4.3	20.3	32.2
71–75	1.1	−7.5	−13.6	17.0	26.5

Note: Cell sizes are the same as in table 2.5.

Table 2.14 **Median Saving to Consumption Ratio by Age and Education Groups**

Age Group	Total Sample	High School Dropouts	High School Graduates	College Graduates
21–25	−10.8	0.7	−16.3	2.3
26–30	7.5	1.7	6.2	17.3
31–35	7.5	−6.7	6.1	12.9
36–40	9.8	−2.2	8.0	19.2
41–45	10.8	−0.2	7.6	23.5
46–50	11.1	−11.8	14.9	16.2
51–55	16.0	6.8	16.9	21.5
56–60	18.6	19.9	27.3	9.6
61–65	9.5	11.2	−2.0	17.9
66–70	6.3	4.3	6.3	7.7
71–75	1.1	1.8	5.6	−9.1

Note: Cell sizes are the same as in table 2.6.

Controlling for demographic changes in saving behavior is not an easy task. In table 2.15, however, we report a very simple attempt to control for the number of children in households. We will not give any interpretation to these results. In particular, we tabulate ratios of median saving to consumption by age and number of children. With the exclusion of some cells which are extremely small (and should be ignored), these results do not show strong effects of children on saving behavior. The only possible exception is that of households with no children (*at that point in time*), which exhibit consistently higher saving than the households with children. This probably reflects the greater consumption needs of larger households.

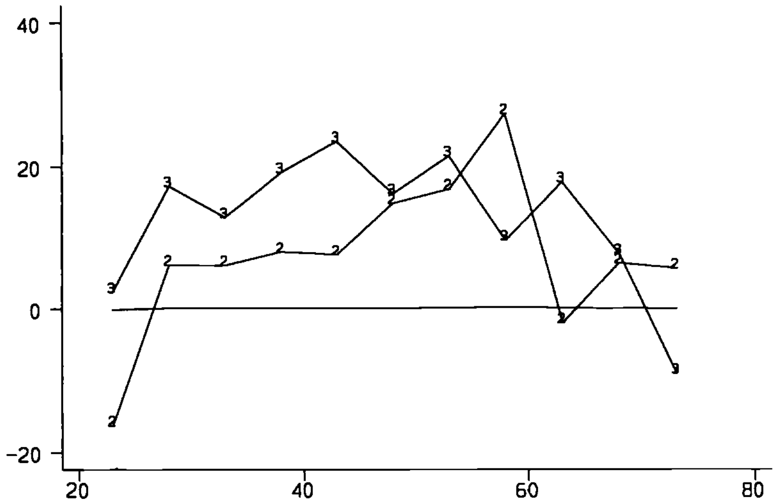

Fig. 2.9 Ratio of median saving to consumption cross-sectional profile by education group

Table 2.15 Median Saving to Consumption Ratio by Age Group and Number of Children

Age Group	No Children	One Child	Two Children	Three or More Children
21–25	−16.3	−1.0	−3.1	6.7
	(351)	(77)	(46)	(23)
26–30	16.8	5.6	3.2	−7.3
	(344)	(135)	(96)	(64)
31–35	12.2	9.6	−5.6	−0.7
	(227)	(135)	(165)	(113)
36–40	19.3	2.9	7.5	7.8
	(183)	(130)	(156)	(103)
41–45	12.1	20.6	9.4	6.3
	(196)	(132)	(114)	(81)
46–50	10.3	14.4	13.1	8.1
	(221)	(86)	(37)	(31)
51–55	17.8	−0.3	0.8	60.9
	(255)	(43)	(24)	(13)
56–60	19.9	2.8	21.2	−4.6
	(219)	(20)	(8)	(4)
61–65	11.2	−11.2	10.2	38.2
	(238)	(16)	(4)	(1)
66–70	7.7	5.0	14.6	2.3
	(289)	(7)	(3)	(4)
71–75	1.1	23.6	−49.5	62.4
	(217)	(8)	(2)	(2)

Note: Numbers in parentheses are cell sizes.

2.3.4 Wealth

Because the life-cycle model (at least in some simple versions) has strong implications for the pattern of wealth accumulation (and decumulation) the latter has been studied extensively.[33] In spite of the large volume of the literature on this issue, no firm answer on the shape of a typical age profile for wealth (and in particular on the issue of asset decumulation by the elderly) or on the main motivation for savings has emerged. This is both because of the scarcity and poor quality of data sets and because of difficult conceptual issues.[34]

We conclude this section describing the cross-sectional age profile of real and financial wealth (see section 2.2 for definitions). These two components do not exhaust household net wealth. The main exclusions are, on the asset side, pension wealth and durable commodities and, on the liability side, loans and debts other than mortgages.

In table 2.16 we report mean, median, and standard deviation for financial wealth (in the first two columns) and for gross and net real estate wealth (in the fourth to sixth columns). In the last column we report the percentage of home owners (with or without mortgages).

Mean and median financial wealth are plotted by age group in figure 2.10. Two considerations are in order as far as financial wealth is concerned. First, the median level of financial wealth is very low: for all age groups it is below $7,000. For most groups it is around one-tenth and for all groups is well below one-half of median annual disposable income. Second, there is no tendency for either the mean or the median to decline in the last part of the life cycle. This could be due, of course, to a variety of reasons and does not necessarily contradict the life-cycle model.[35]

Real wealth is substantially higher than financial wealth, confirming that real estate constitutes a very important part of households' portfolios. Both the mean and the median (plotted in fig. 2.11) increase very rapidly in the first part of the life cycle and show a slight tendency to decline in the last part. The difference between net and gross wealth (which roughly corresponds to mortgage debt) tends to decline in the last part of the life cycle as households repay their mortgage debts.[36]

A pattern similar to that of the stock of real estate wealth is followed by the percentage of home owners by age. The percentage of home owners is as low

33. The papers on this topic are too numerous to be cited here. Some interesting studies are those by Shorrocks (1975), King and Dicks-Mireau (1982), Hurd (1989), and Jianakoplos, Menchik, and Irvine (1987).

34. Wealth information (especially for rich households) is very difficult to obtain and no long panels exist. For a comparison of different data sets containing wealth information see Juster et al. (1987).

35. See Attanasio (1993b) for a discussion of these issues and Attanasio and Hoynes (1993) for a discussion of differential mortality by wealth class.

36. The difference between median gross and net real estate wealth is obviously *not equal* to median mortgage debt.

Table 2.16 **Real and Financial Wealth Holdings by Age Group**

Age Group	Financial Assets Mean	Financial Assets Median	Gross Real Estate Mean	Gross Real Estate Median	Net Real Estate Mean	Net Real Estate Median	Percentage of Home Owners
21–25	1,801 (6,466)	500	7,020 (26,898)	0	3,844 (20,126)	0	0.105
26–30	4,233 (9,583)	850	36,913 (61,175)	0	25,163 (48,876)	0	0.397
31–35	7,763 (16,928)	2,527	59,732 (81,523)	35,000	38,079 (63,630)	5,000	0.570
36–40	10,579 (24,198)	2,000	66,380 (103,064)	33,000	43,842 (87,517)	6,803	0.556
41–45	14,060 (28,860)	2,286	83,515 (97,421)	45,000	59,488 (80,193)	37,660	0.725
46–50	14,151 (27,928)	2,275	97,657 (104,619)	70,000	71,224 (89,113)	45,000	0.729
51–55	14,177 (31,313)	1,575	98,098 (131,458)	60,000	79,136 (116,350)	45,000	0.831
56–60	21,795 (41,267)	2,819	91,620 (102,945)	62,500	80,133 (94,759)	57,500	0.773
61–65	34,777 (50,136)	6,950	107,269 (106,318)	79,000	95,391 (98,490)	65,000	0.816
66–70	27,107 (45,744)	5,750	105,707 (107,420)	75,000	100,984 (104,955)	67,350	0.811
71–75	33,663 (56,815)	6,093	81,539 (98,192)	55,000	77,042 (95,949)	52,500	0.775

Note: Numbers in parentheses are standard deviations.

as 10 percent for the first age group, and it is already equal to 57 percent for the third group. It peaks at 83 percent for the 51–55 age group. This indicates that most of the sharp increase in average real estate wealth in the first part of the life cycle is explained by an increase in home ownership rather than by an increase in the stock owned.[37]

The same words of caution used for the interpretation of the financial asset age profile are in order here. The fact that the cross-sectional profile for total assets declines in the last part of the life cycle it is not necessarily an indication of asset decumulation: it could be due to cohort effects or to biases introduced by differential mortality. For instance, the decline in mean and median real estate wealth, as well as in the percentage of home owners in the last age group could be explained either by some of the elderly liquidating their real estate wealth or by the fact that the percentage of home owners (and their average real estate wealth) for that particular cohort had always been lower.

37. The variable reported in the survey should, at least in theory, be equal to the market value of the house and therefore reflect capital gains and losses as well as "active" additions to the household's estate wealth.

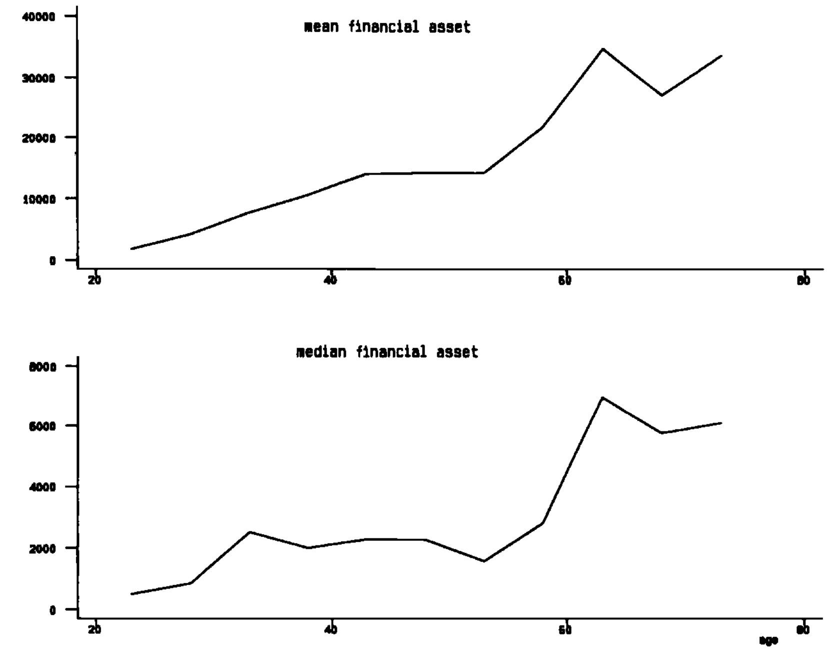

Fig. 2.10 Mean and median financial assets cross-sectional profile

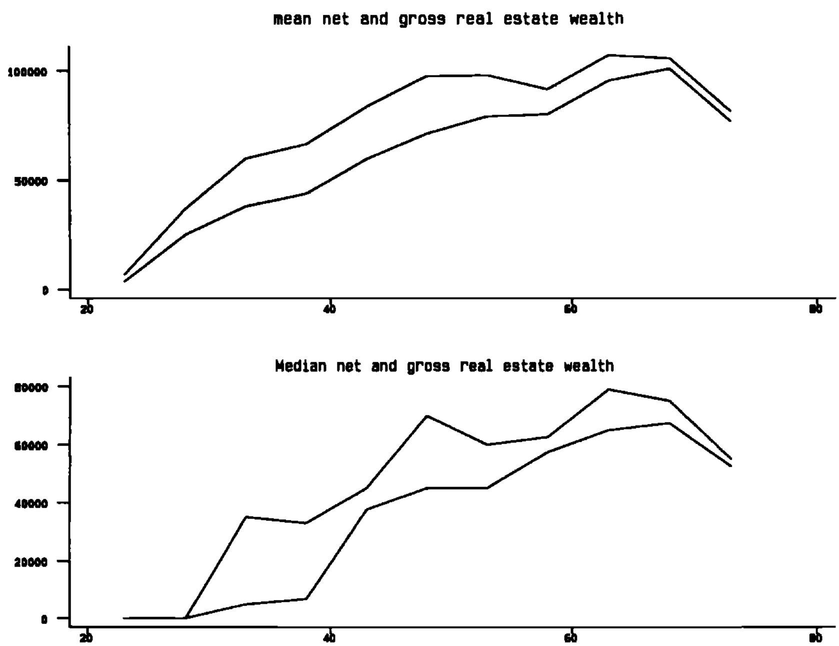

Fig. 2.11 Mean and median net and gross real estate wealth cross-sectional profile

2.4 Income and Pension Contribution Age Profiles

As stressed in the introduction, the analysis of a single cross section can provide only a useful snapshot of a given phenomenon, but it is of very limited use in describing the behavior of a dynamic variable such as saving. In the presence of strong cohort effects, the interpretation of the cross-sectional profile as the life-cycle profile of a given variable can introduce serious biases. In the absence of panel data and repeated cross sections there is not much one can do to control for cohort effects. In this and in the following sections we use the 11 CEX surveys available since 1980 to construct synthetic cohort averages in the attempt to measure life-cycle and cohort effects on the variables of interest.

Before proceeding with the analysis, however, it might be of some interest to assess the importance and the magnitude of the potential bias introduced by interpreting cross-sectional profiles as life-cycle profiles. One can use the techniques discussed below to identify a (smoothed) age profile and compare it to a cross-sectional profile. As a representative variable we chose total household consumption expenditure.

In figure 2.12 the line labeled "cross-section profile" is constructed using the 1984 survey to compute consumption averages for different age groups. The line labeled "age profile" is the same as that in figure 2.26 (in section 2.5 below) with cohort effects removed. It is constructed by regressing average cohort data on a fifth-order polynomial in age and cohort-specific intercepts. What we plot is the age polynomial with the intercept of the first cohort. As can be seen, because of the presence of sizeable and positive cohort effects, the cross-sectional profile exaggerates the hump in consumption. A similar picture can be obtained choosing different cross sections and different variables (provided that there are sizeable cohort effects). The method used to remove cohort effects from the consumption profile is crude and based on some strong identification assumptions. The picture makes clear, however, the potential importance of these effects.

2.4.1 Total After-Tax Household Income

In this subsection we use the cohort techniques outlined above to estimate total household age-income profiles. We correct the CEX definition of total after-tax income by subtracting from it Social Security contributions which will be considered (here and in the section on saving) as a tax rather than as a form of saving.

As discussed in section 2.2.2, we parametrize the cross-sectional distribution of disposable household income within each cell as a mixture of two normal densities.[38] The estimated means are, as expected, higher than the simple

38. It is not possible to fit a log-normal distribution (or the mixture of two log-normals) because of the presence of negative and zero income. These densities are fitted to the CES definition of disposable income. The corrected means are obtained by subtracting from the maximum likelihood means the sample means of Social Security contributions.

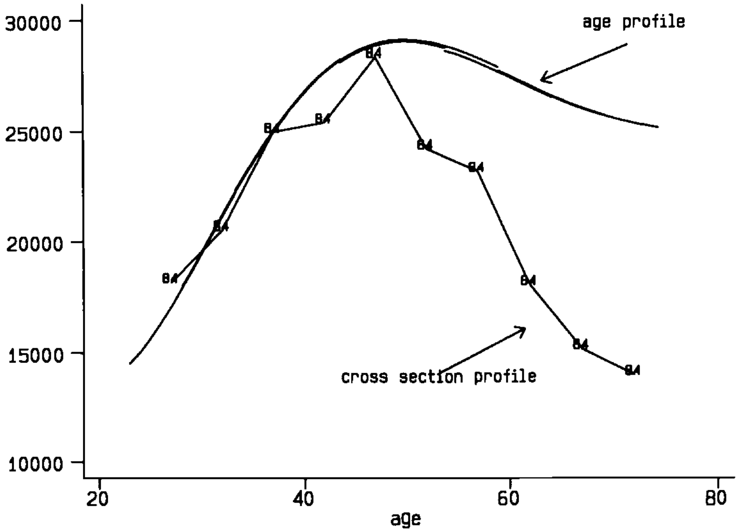

Fig. 2.12 Cross-sectional vs. cohort consumption profile

averages one obtains setting the top-coded observations at the top-coding value. The magnitude of the difference between these two estimates obviously depends on the number of top-coded observations.[39] Because the proportion of top-coded observations varies systematically with age and cohort, the distortion introduced by ignoring top-coded observations will also have a systematic pattern.

In the top panèl of figure 2.13 we plot the estimated means for year-cohort cells against age. Because this kind of graph will be used extensively throughout the rest of the paper, it is worthwhile spending a few moments explaining it in detail. Each connected segment represents the behavior of a cohort over the 11 years of our sample. For instance, the first segment on the left is average household income for the first cohort—i.e., for households headed by a person born between 1955 and 1959—in each year from 1980 to 1990. These individuals were, on average, 23 years old in 1980, 24 in 1981, and so on until 1990 when they were 33. Because a cohort is defined by a five-year interval and we have 11 years of data, each cohort overlaps at six ages with the following cohort: for instance, cohort 2 is observed between ages 28 and 38, while cohort 3 is observed between ages 33 and 43.

In the bottom panel, the same data points are smoothed by regressing them on a fifth-order polynomial in age, cohort-specific intercepts, and year dummies whose coefficients are constrained to sum to zero and to be orthogonal to a linear trend.[40] The smooth profiles in the graph are given by the polynomial

39. Without top-coded observations the maximum likelihood estimator is the sample mean.

40. Deaton and Paxson (1992) use a similar procedure. Estimates of the coefficients can be obtained either by OLS or by weighted least squares, using as weights the standard errors of the cell means. The results are almost identical.

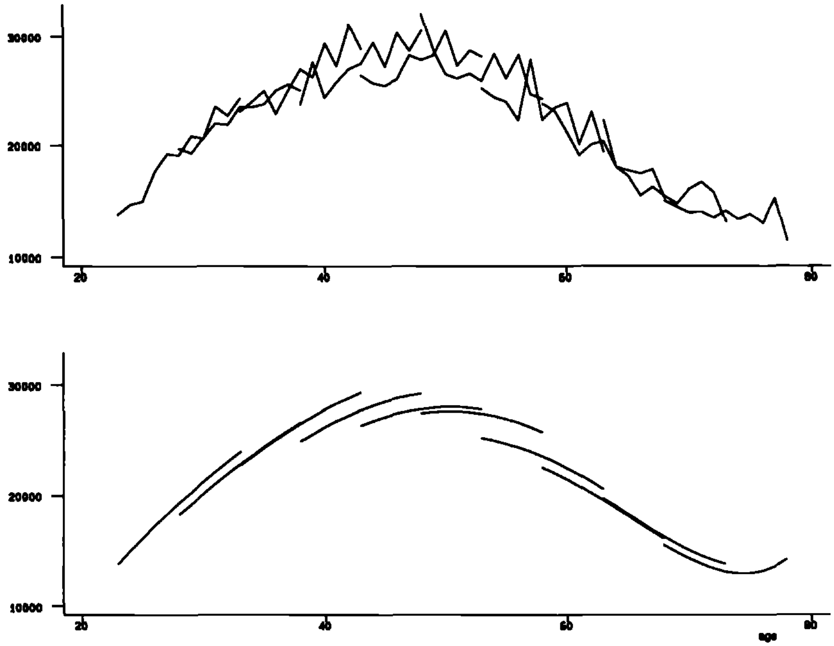

Fig. 2.13 Mean disposable income

in age with the cohort-specific intercepts. One should not give any structural interpretation to this graph: without a structural model or additional information it is not possible to identify cohort, age, and time effects separately because of the linear relationship that links them.[41] The only aim of the bottom panel is to smooth the estimated means; to interpret them as pure age profiles we would have to assume that time effects are common for all cohorts, sum to zero, and are orthogonal to a linear trend.

Figure 2.14 is similar to 2.13 except that we plot cohort medians instead of means against age.[42] Because of the skewness of income distribution, medians are lower than means; otherwise the picture that emerges from this figure is similar to that from figure 2.13.

Several elements of interest emerge from figures 2.13 and 2.14. First, household disposable income has the typical hump-shaped profile that is often found in the literature: a similar profile emerges from CPS data. The smoothed age profile peaks at age 51 for the means and at age 48 for the medians.

Cohort effects are also quite evident: for all cohorts but one, the smoothed profiles lie above that of the next older cohorts. In table 2.17 we report the

41. For a discussion of identification issues in this framework see Heckman and Robb (1987), MaCurdy and Mroz (1990), and Attanasio (1993b).

42. In this figure and in figure 2.15 I use sample quantiles. The figures obtained using the quantiles of the density estimated by maximum likelihood are very similar.

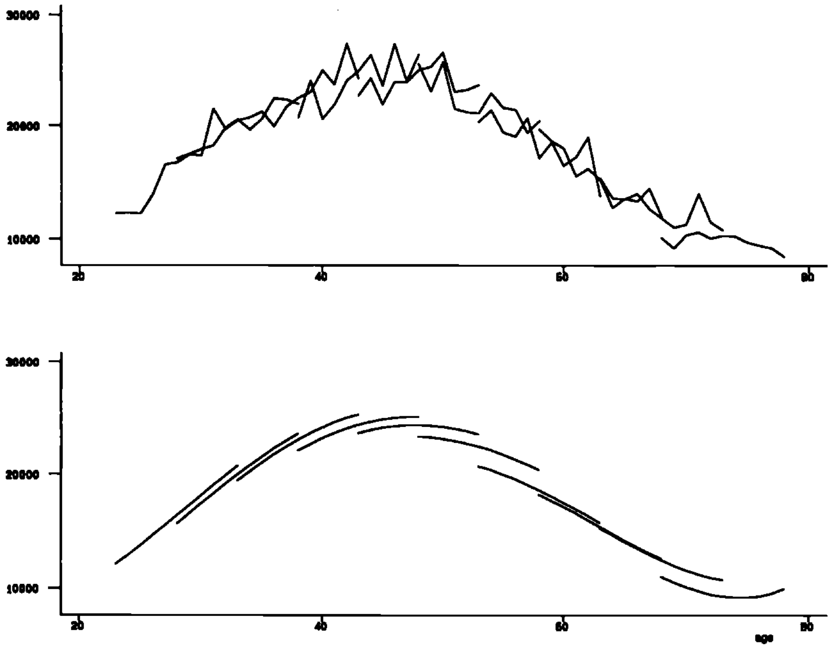

Fig. 2.14 Median disposable income

Table 2.17 Cohort Intercepts and Their Changes for Income Profiles

Cohort	Intercept for Medians	Percentage Increase	Intercept for Means	Percentage Increase
1	25,336.2	2.99	28,944.3	3.92
2	24,599.7	2.16	27,853.6	0.34
3	24,079.0	4.05	27,758.6	6.06
4	23,142.7	3.38	26,172.1	5.63
5	22,386.6	4.86	24,776.4	1.77
6	21,350.0	8.97	24,344.9	9.94
7	19,592.6	1.95	22,143.4	4.82
8	19,218.1	0.66	21,124.7	−0.70
9	19,093.0	8.23	21,273.3	3.57
10	17,641.5		20,539.1	

intercepts of the profiles in the bottom panels of figures 2.13 and 2.14 and their percentage increase relative to the intercept of the next older cohort. The "gains" from one cohort to the next seem to be higher (both in means and medians) for the middle cohorts than for the youngest and the oldest. For instance, for the two youngest cohorts, the intercept for the means rises on average by 2.1 percent, while for the next four cohorts it rises by 5.8 percent. The figures for the medians are 2.6 percent and 5.3 percent. For the three oldest cohorts, the average increase in the intercept is 2.6 percent for the means profiles and 3.6 percent for the medians.[43]

Finally, it is possible to detect some business-cycle effects:[44] for most cohorts the raw profiles increase considerably more in the years after 1982. In figure 2.15 we plot the coefficients on the time dummies estimated for the median income profile.[45]

So far we have not controlled for within-cell heterogeneity. In the next graph we look at the income-age profiles for different education groups. Mean and median income-age profiles for college graduates, high school graduates, and high school dropouts are plotted in figure 2.16. The size of the cells on which these estimates are based is relatively small, especially for college graduates and high school dropouts: this is reflected in the larger variability of these estimates, as is most evident in the graph for high school dropouts. As expected, the profiles are higher and steeper for higher education groups (notice the different scale). These profiles are going to be compared with analogous profiles for consumption and savings.

2.4.2 Components of Household Income

In this section we analyze the components of household income. We decompose total before-tax household income into interest income, labor income, pensions, and transfers. These four components exclude some minor items because their economic significance is minor and because they are not easily classifiable. Labor income is defined as wages and salaries, plus income from own business.[46] Interest income includes interest, dividends, and royalties. Interest income does not include capital gains or income for the sale of assets. Pensions are all payments of pensions. Transfers include unemployment compensation, welfare payments, food stamps, Social Security payments, alimony for child support, and other transfer income. Unlike the BLS definition, this one leaves out income or loss from taking in boarders.

In the four panels of figure 2.17, the percentage of households with positive quantities of each of the four components of income is plotted in turn against age, for each year-cohort cell.

The top two panels are not surprising. The percentage of households with

43. Similar results are obtained if, in smoothing the means, we do not include year dummies at all.

44. Most of the period under study is characterized by a boom: only the first few years were affected by the 1981–82 recession. Unfortunately, as we mentioned in section 2.3, the quality of the 1980–81 survey is doubtful: therefore we should be careful in comparing the means of those years with those for subsequent years. Analyzing the data for 1991 and 1992, when they become available, will be extremely interesting.

45. One should remember that these coefficients are constrained to sum to zero and to be orthogonal to a time trend. Also notice that what we call 1982 includes data referring to 1981 (see section 2.2); this is why figure 2.15 starts in 1979.

46. The inclusion of this last category is questionable. The results, however, do not change substantially.

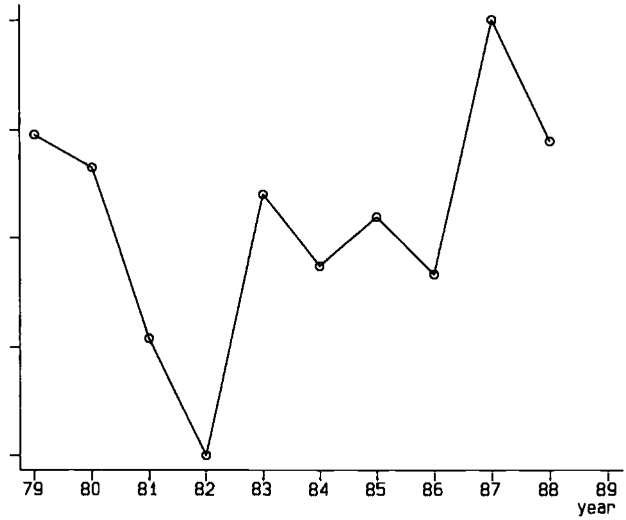

Fig. 2.15 Year effects for median cohort income

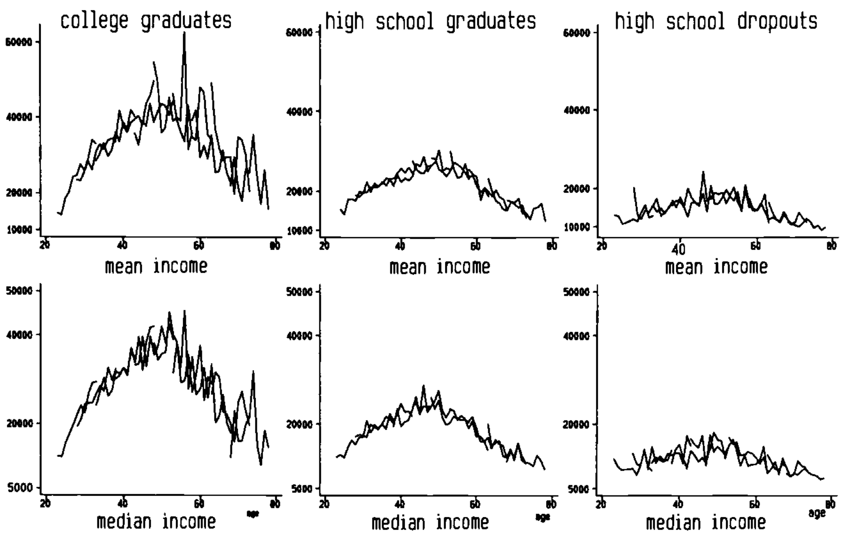

Fig. 2.16 Mean and median income by education group

positive values for labor income is very close to unity until age 50, and declines after that, when an increasing proportion of earners retire.

The evidence in the third panel (lower left) shows that the proportion of households receiving interest income increases until around age 65 and flattens out after that. At the beginning of the life cycle, this proportion is very low (around 30 percent) but reaches almost 70 percent by retirement age. Given

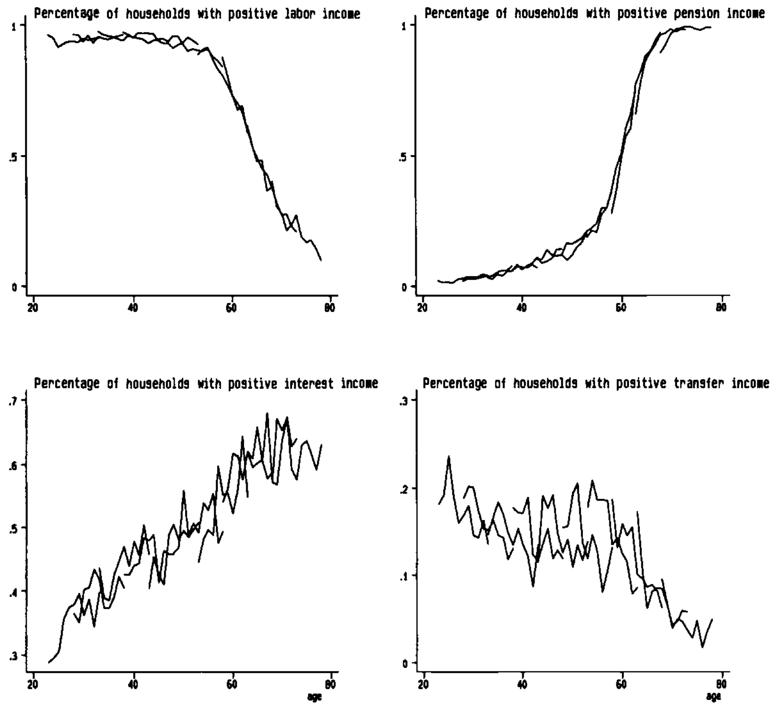

Fig. 2.17 Income composition

the amount of nonreporting that plagues this measure, this evidence indicates that a substantial proportion of households hold assets that provide income by the time the household head reaches retirement.

The last panel of figure 2.17, with the proportion of households receiving some form of transfer income, shows that this proportion is much lower and much more variable than for the other income sources. As expected, a substantial amount of fluctuation over the business cycle is evident in this graph (the figures for the early years of the sample are relatively higher than the others). In terms of life-cycle fluctuation, the proportion seems to decrease until age 40, to increase slightly after that, and to decline again.

In figure 2.18 we plot the age profile for median labor income.[47] The profile increases until around age 50 and declines steeply at retirement. There are no strong cohort effects: the profiles for overlapping cohorts are very close to each other.

Figure 2.19 plots the profile for the median and the third quartile of interest income. Both graphs, and especially that of the median, are very noisy: some

47. In this subsection, we will only look at the quantiles of income components, thus avoiding maximum likelihood estimation.

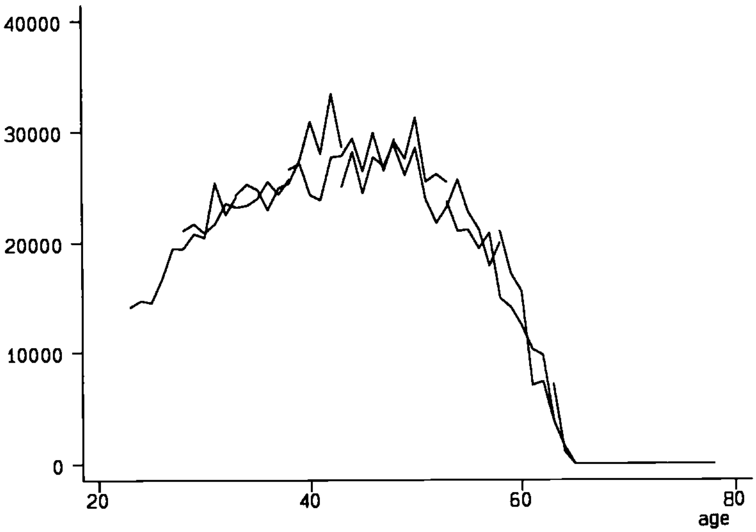

Fig. 2.18 Median labor income

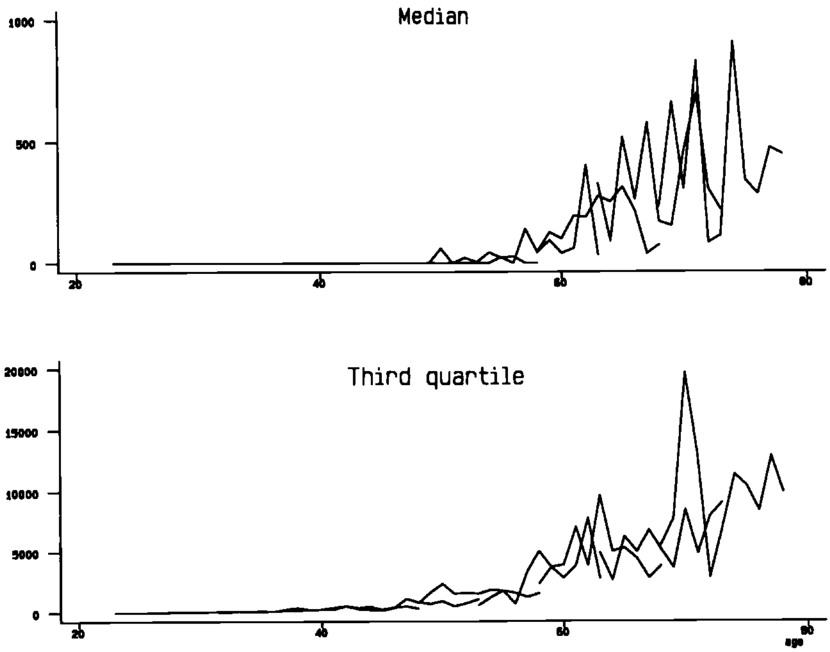

Fig. 2.19 Interest income

interesting features, however, emerge clearly. Interest income becomes somewhat important only late in the life cycle. Median interest income, for instance, is below 1 percent of median labor income at age 60 and is above 10 percent for only one of the cells. The median of the ratio of interest to labor income has a similar pattern: only for two cells it is above 10 percent, and it is always below 20 percent. On the right tail of the distribution, however, interest income is much more important. At age 60 the third quartile of interest income is above 10 percent of the third quartile of labor income, and by age 66 is 36 percent for cohort 8 and 71 percent for cohort 9.[48]

When compared to the quantiles of pension income, the quantiles of interest income exhibit a similar pattern. If we take the ratio of medians (or compute the median of the ratios), we notice that interest income is a small fraction of pension income, while, when we consider the third quartiles, it is much more important. The ratio of the medians averages 5 percent for ages above 65, while the median of the ratio averages just under 1 percent for ages above 65. The same numbers for the third quartiles are 66 percent and 18 percent.[49]

These patterns are a reflection of the fact that the inequality in the distribution of interest income is, as expected, much higher than that of labor or pension income. At age 66, the median of interest income averages at just over $300, while the third quartile averages at almost $5,000. With an average interest rate of 5 percent these figures imply assets worth $6,000 and $100,000 for the median and the third quartile of financial wealth.[50]

For the sake of completeness, in figure 2.20 we plot the age profile for median pension income and for the 90th percentile of transfer income. The 90th percentile for transfer income was chosen because the proportion of households receiving transfer income is often just above 10 percent.

2.4.3 Pension Contributions

In this section we analyze Social Security and pension contributions. The latter are divided into contributions to government pensions (including railroad retirement schemes), private pensions, and individual retirement accounts.[51] In this paper we consider pension contributions as saving, while Social Security contributions are considered as taxes.

We start with the latter. In the top left panel of figure 2.21 we plot average

48. Obviously the household earning the third quartile (or the median) labor income is not the same one that earns the third quartile (or the median) interest income. The third quartile of the ratio of interest income to labor income is on average 0.1 for cells with ages between 60 and 65 and 0.27 for ages between 65 and 70.

49. By age 66 almost 90 percent of the households in the sample receive pension income.

50. It should be remembered that these figures do not include capital gains and/or income from the sale of assets.

51. In 1980 and 1981 the data on Social Security and pension contributions were top-coded when any other income variable was top-coded. In subsequent years these variables were not top-coded at all. In this section we ignore the problem completely. All the figures are deflated by a cohort-specific CPI.

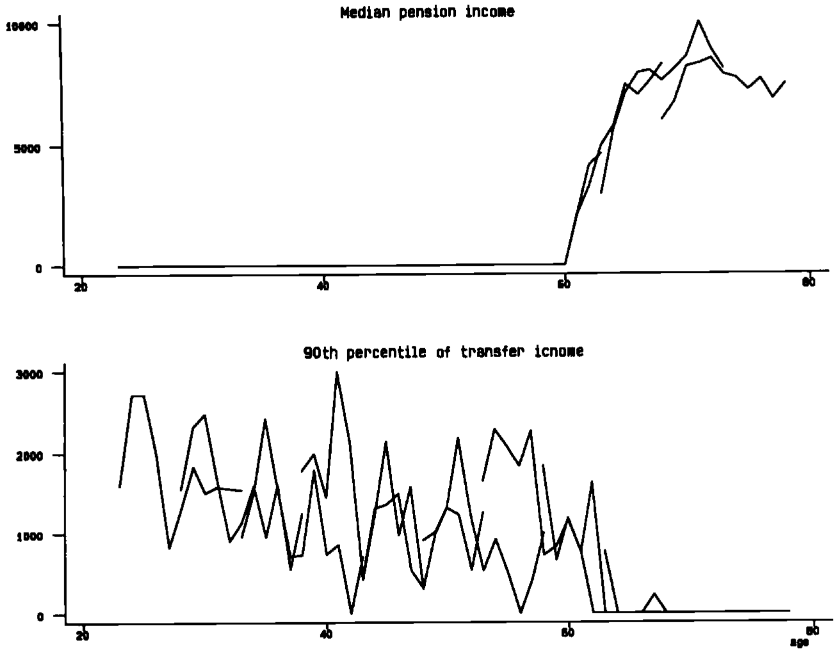

Fig. 2.20 Pension income

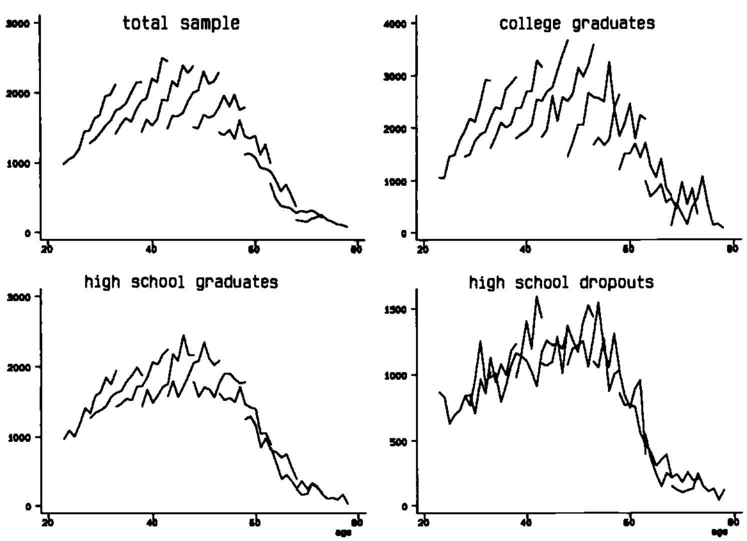

Fig. 2.21 Contributions to social security by education group

contributions to Social Security for the whole sample. The impressive feature of this picture is the steady and strong increase (for most cohorts) of the average contribution. The decline in the last part of the life cycle is explained by the decline in the percentage of households contributing to Social Security as an increasing number of individuals retire. The strong increase in the average contribution is reflected in a strong increase in the ratio of contribution to gross income.

The increase in Social Security during the 1980s is explained in part by an increase in the tax rate of Social Security, but even more important, by several increases in the taxable base. This is apparent in the remaining three panels of the figure 2.21 which plot average contributions for the three education groups considered above. Notice that the increase over time is strongest for college graduates (who perceive the highest income) and weakest for high school dropouts (note the difference in scale).

In the top panel of figure 2.22 we plot the proportion of households contributing to government pension schemes. This proportion is, as expected, fairly low. In the bottom panel we plot the average level of contributions to government pensions, conditional on positive amounts. This profile is slightly hump-shaped and peaks around $2,000.

In figure 2.23 we look at contributions to private pension schemes. As stressed in section 2.2 these include only employees' contributions (deducted from the pay) but not employers' contributions. In the top panel, we plot the proportion of households with positive contributions to private pension schemes, while in the bottom panel we plot the average annual contribution for those households with positive contributions. The top profile rises sharply in the early part of the life cycle and peaks at around .15 after age 40. Notice the sharp increase for some cohorts in the last years of the sample. The average level shows substantial increases over time for most cohorts. At the end of the sample, most middle cohorts are at a level of about $2,000.

In the top panel of figure 2.24 we plot the percentage of households contributing to IRAs, while in the bottom panel we plot the average contributions for those households with a positive contribution. While the average contribution is mostly stationary around the $2,000 level, we see that there are two large fluctuations, corresponding to 1982 and 1986, for the percentage of households contributing to IRAs.

There is now a voluminous literature on the effects that the two tax acts of 1982 and 1986 have had on IRAs and, more generally, on savings. In 1982 IRAs were given strong fiscal incentives: contributions became tax-deductible, while a limit of $2,000 ($2,500 for a couple with one earner, $4,000 for a couple with two earners) was kept on the total amount of the contribution. Most of the fiscal incentives were subsequently removed with the tax act of 1986. From the two panels of figure 2.24 it is evident that the tax incentives had a large effect on participation: the percentage of households participating jumps from below 5 percent to almost 30 percent around the tax act of 1982.

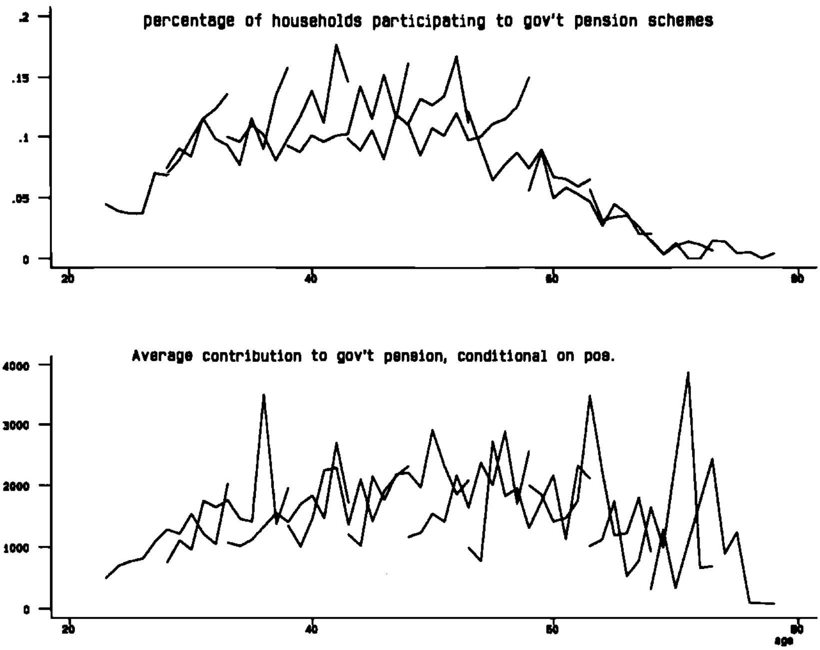

Fig. 2.22 Contribution to government pension schemes

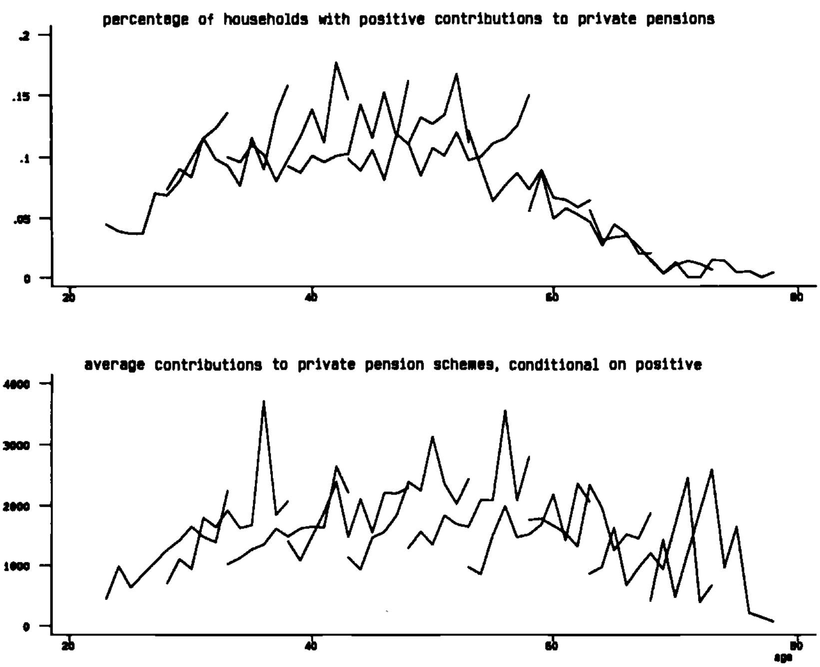

Fig. 2.23 Contribution to private pensions schemes

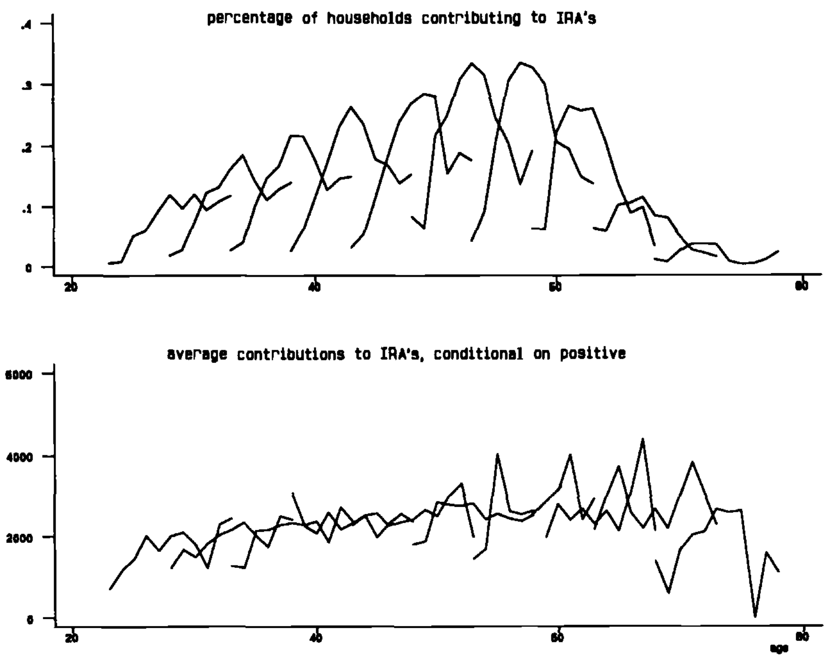

Fig. 2.24 Contribution to IRAs

This percentage drops following the tax act of 1986, but not by the same amount. Most contributions, as documented by Feenberg and Skinner (1989) were at the statutory limit of $2,000 (or $4,000). This explains the fact that the average level of the contributions in the bottom panel of the figure is stable around $2,000.

It is beyond the scope of this paper to address the main issue in the literature on IRAs, i.e., whether the tax incentives given to IRAs in 1982 increased aggregate savings. It is interesting to notice, however, that the decline in IRA participation in 1987 occurs simultaneously with an increase (for many cohorts) in the average level of private pension contributions, perhaps indicating that, after the fiscal incentives to IRAs were removed in 1987, some households moved some funds (back?) from IRAs to other pension schemes. If the movement had been symmetric, this would indicate that the IRA legislation caused a reshuffling of existing savings rather than the creation of "new" savings. Further investigation of this issue is needed.[52]

Whether Social Security contributions should be considered as a tax or as a form of saving, as we have done for pension contributions, is questionable. To give an idea of the order of magnitude of these variables, in the top panel of

52. A nonexhaustive list of papers on the effects of IRA legislation includes Venti and Wise (1990), Feenberg and Skinner (1989), and the survey by Gravelle (1991).

figure 2.25 we plot the age profile for the ratio of average total contributions (pensions and Social Security) to consumption. In the bottom panel we exclude Social Security contributions.

2.5 Consumption

2.5.1 Total Consumption Expenditure

In this section we estimate age-consumption profiles following the same steps that we used for household disposable income. The only difference is that, because consumption is not top-coded, we can use sample averages to estimate cell means.

In figures 2.26 and 2.27 we plot cohort means and medians for total consumption expenditure against age, with the same method for smoothing.

Age-consumption profiles present features similar to those of disposable income: the characteristic hump shape is, if anything, even more apparent than for income. As in the case of income, mean consumption age profiles are higher for younger cohorts. The pattern of the increase in the smoothed profiles, however, is slightly different. In table 2.18 we see that, while it is still true that the percentage increase in the intercept is smaller for the two youngest cohorts, cohorts 8 and 9 present percentage increases comparable to those of the other cohorts. The profiles for the medians, instead, are much flatter and the "cohort effects" do not present any distinguishable pattern. The difference in the dynamics of mean and median consumption indicates that there have been changes in, the cross-sectional distribution of consumption, discussed below.

The similarity between income and consumption age profiles has been interpreted as evidence against the life-cycle model.[53] In these simple graphs, however, we are ignoring family composition and labor supply behavior, not to mention investment in human capital. While the specification of a structural model including flexible adult-equivalent schemes and endogenous labor supply in a dynamic framework is necessary to test the life-cycle hypothesis, it is interesting to look at the effects of some extremely simple adult-equivalent scheme. In figure 2.28 we plot mean per capita consumption (i.e., consumption divided by family size) and mean per adult-equivalent consumption (which is constructed by considering children under age 15 to be equivalent to half an adult). From this figure is evident that, while some life-cycle movements are still apparent (such as the decline in consumption corresponding to retirement), these consumption-age profiles look much flatter than in figure 2.26.

The variability of consumption and income in figures 2.13, 2.14, 2.26, and

53. See Carroll and Summers (1991) and the discussion on these issues in Attanasio and Browning (1992). It is interesting to notice that this same evidence has also been interpreted as evidence in *favor* (Ghez and Becker 1975) of the life-cycle model.

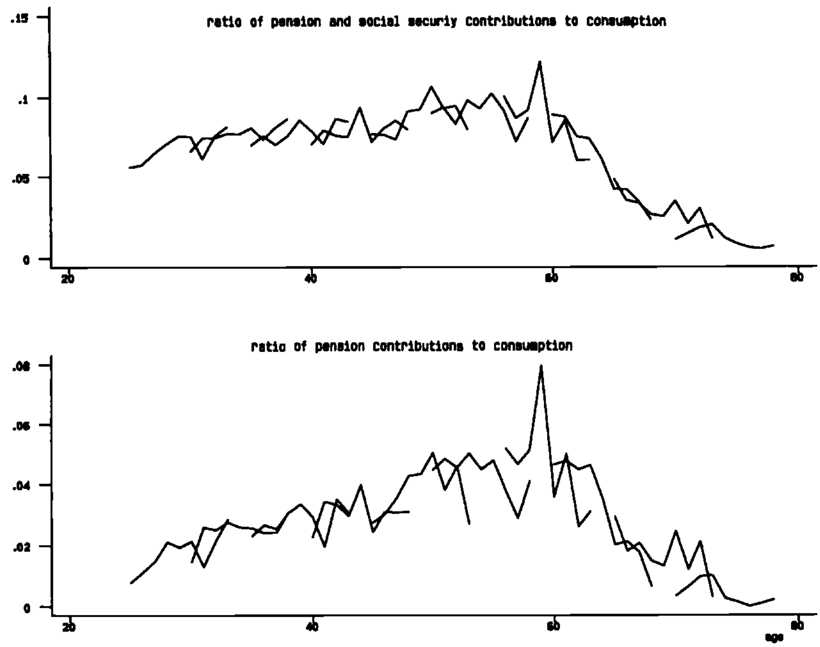

Fig. 2.25 Ratio of pension and social security to consumption

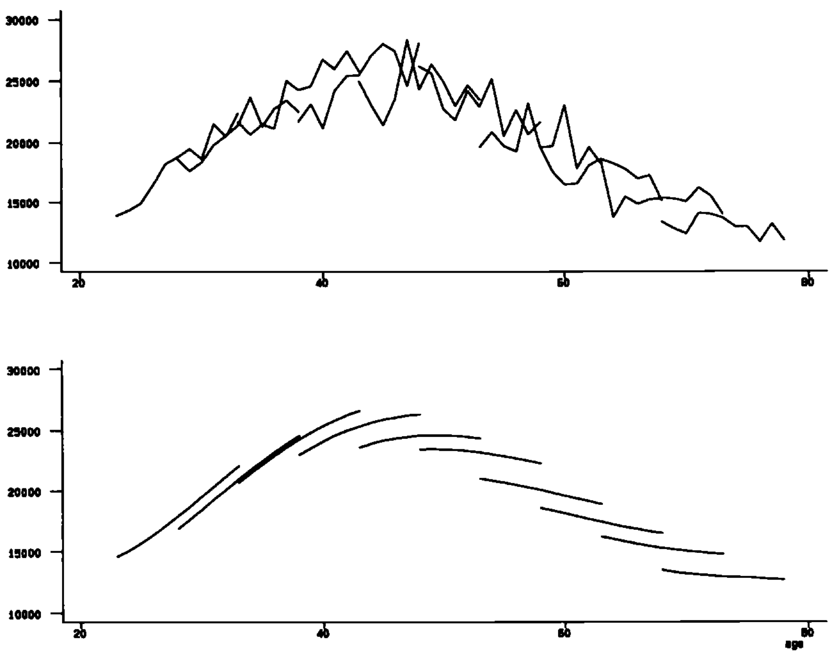

Fig. 2.26 Mean total consumption expenditure

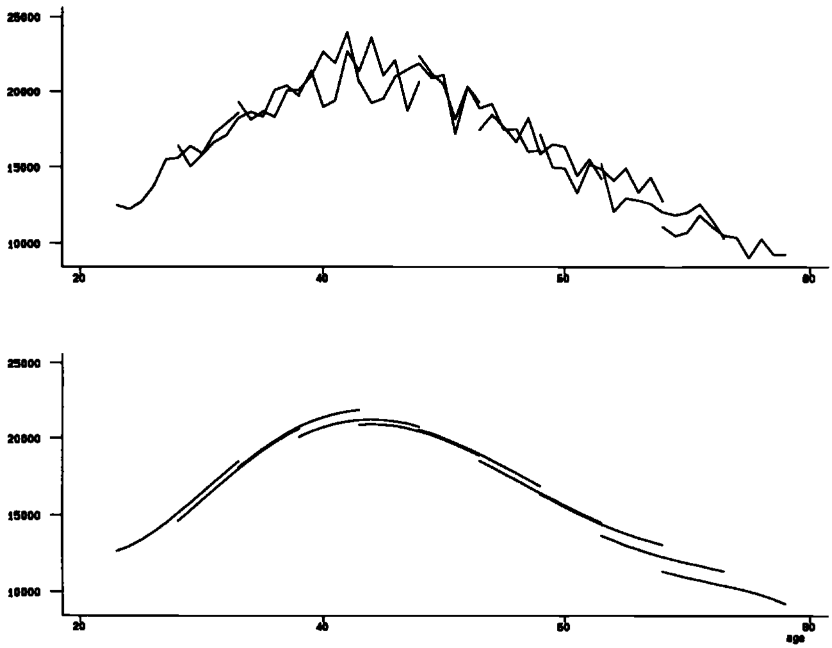

Fig. 2.27 Median total consumption expenditure

Table 2.18 Cohort Intercepts and Their Changes for Consumption Profiles

Cohort	Intercept for Medians	Percentage Increase	Intercept for Means	Percentage Increase
1	21,750.95	2.309	26,765.93	4.174
2	21,260.13	−0.501	25,693.38	1.158
3	21,367.19	3.164	25,399.34	5.214
4	20,711.78	1.507	24,140.68	7.681
5	20,404.27	−0.458	22,418.60	5.364
6	20,498.21	2.200	21,277.34	11.558
7	20,056.91	0.475	19,072.93	8.470
8	19,962.07	3.991	17,583.60	7.309
9	19,196.02	5.195	16,386.01	12.486
10	18,247.99		14,567.19	

2.27 reflects both life-cycle and business-cycle effects. In an attempt to remove life-cycle variability and cohort effects, we consider the standard error of the deviations of the cell means and medians from the smoothed profiles. According to this measure, income is more variable than consumption: in the case of means, the standard deviation of consumption around the smoothed profile is $1,277 and $1,340 for income; in the case of medians, these two figures are $926 and $1,293. This is not inconsistent with the prediction of the life-cycle model, but it might also reflect greater measurement error in income.

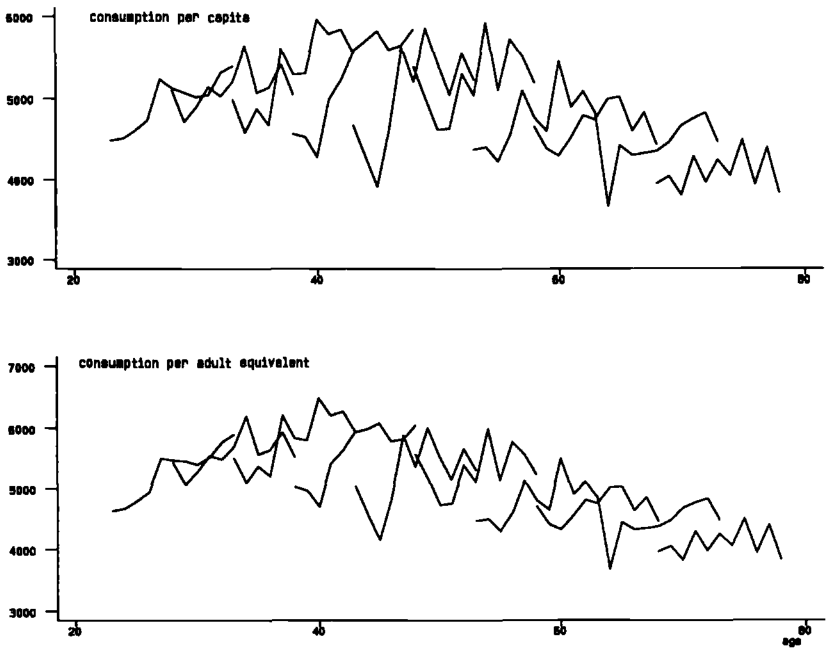

Fig. 2.28 Consumption per capita and adult equivalent

In studying inequality, consumption might be a better measure to look at than income for two reasons: first, if some version of the life-cycle permanent-income hypothesis holds, consumption should not be affected by transitory shocks that increase both the time-series and the cross-sectional variability of income; second, consumption gives a more direct measure of the resources available to a household and therefore is a better indicator of inequality.

In figure 2.29 we plot cell mean and median consumption for the three education groups considered above. The profiles of the three groups are, after rescaling, qualitatively similar. However, as in the case of income, age-consumption profiles are higher and steeper for better educated households: inequality in consumption across education groups increases with age. This similarity in the patterns of income and consumption has been interpreted by Carroll and Summers (1991) as a failure of the life-cycle model. However, as we saw with figure 2.28, these patterns could be explained by changing family composition and/or by nonseparabilities between leisure and consumption.

2.5.2 The Components of Consumption

In this section we look at various components of expenditure. There are several reasons to do so. Ideally, one would like to observe *consumption* rather than *expenditure*. Unfortunately, this is virtually impossible with the data sources currently available. It is therefore interesting to focus on components

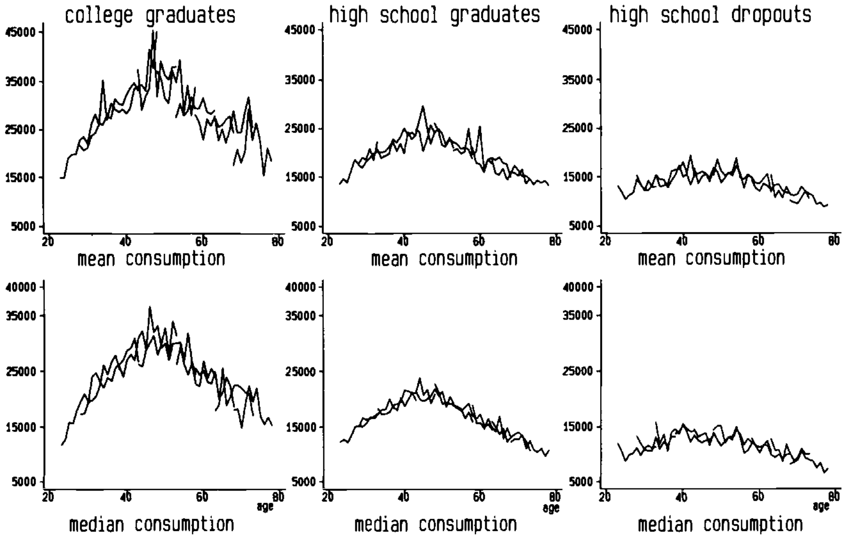

Fig. 2.29 Median and mean consumption by education group

of expenditure that are directly related to consumption. This means excluding all components of expenditure that have an element of durability, i.e., both expenditure on durable commodities and on some services whose effect (it is hoped) lasts in time, such as education and health care. Within the framework of the life-cycle model this is important because durability automatically introduces intertemporal nonseparability in the utility function.

Furthermore, we want to exclude from consumption items that are all but unmeasured for a substantial part of the sample. The obvious example here is housing: for home owners we should impute housing services from owner-occupied houses, a quantity which is very difficult to estimate with the data source available.

Finally, it might be interesting to focus on some items whose relative cost has increased considerably in the last decade, such as education and health care.

In figures 2.30 and 2.31 we plot mean and median expenditure on nondurable and service consumption. This excludes all expenditures on durables, housing, education, and health. While the overall shape of this profile is similar to that in figures 2.26 and 2.27, some important differences emerge. In particular, note that in figure 2.30 the "cohort effects" discussed for income and consumption disappear almost completely, and that in figure 2.31 they are, if anything, reversed.

The relative price of education increased tremendously during the 1980s. In addition, education expenses can be considered, to a certain extent, a form of savings. In the two panels of figure 2.32 we plot average expenditure on per-

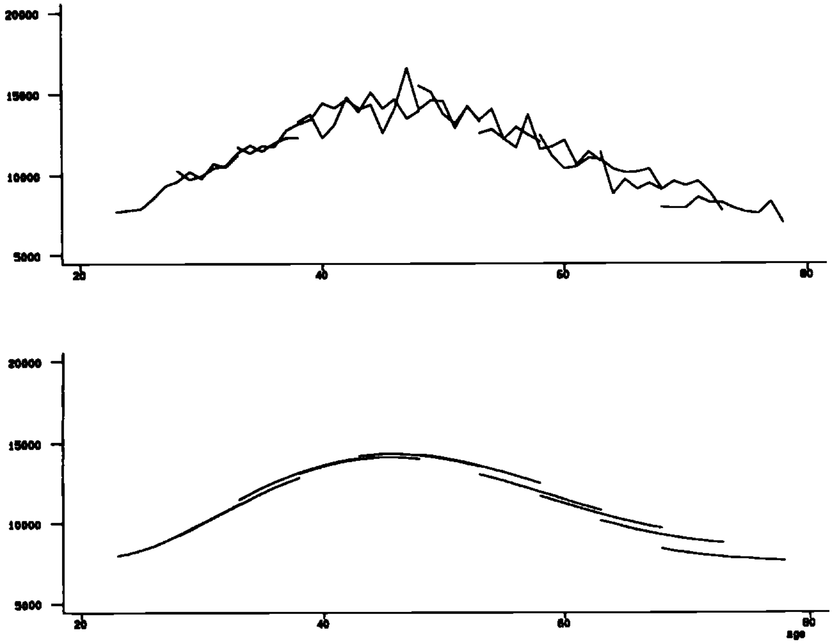

Fig. 2.30 Mean nondurable consumption

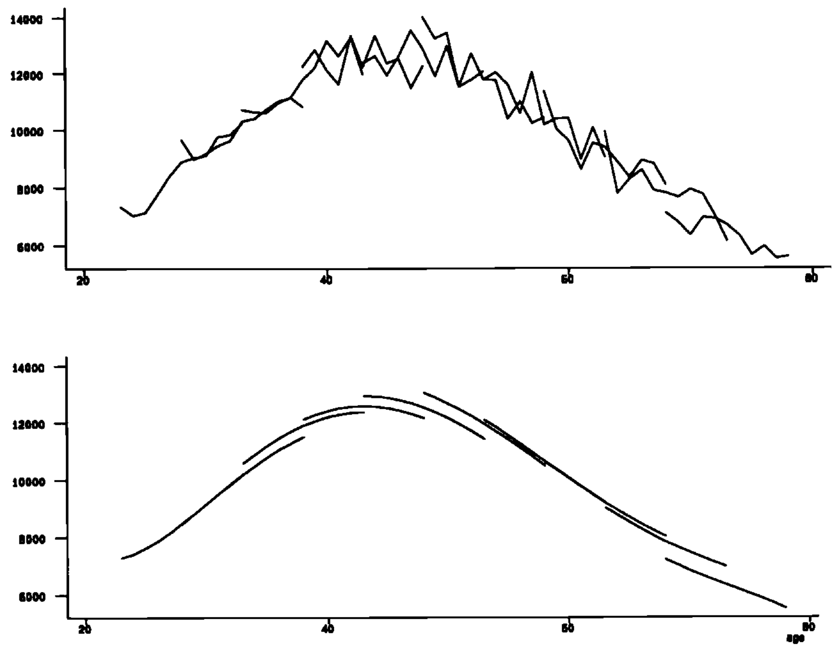

Fig. 2.31 Median nondurable consumption

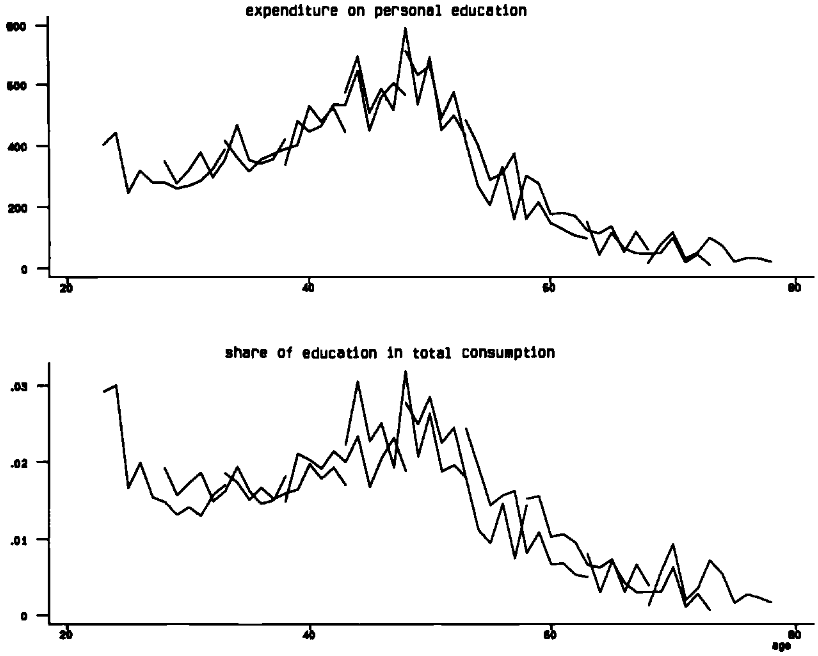

Fig. 2.32 Personal education expenditure

sonal education and its share in average total household consumption. The shape of these profiles is, as expected, correlated with the average number of children: the estimated profiles increase until around age 50 and decline rapidly afterward. Expenditure on education as a share of total consumption is much flatter than the profile of the *level* of education. In several instances, older cohorts are above younger ones, probably reflecting both the fact that education was relatively cheaper in the early 1980s than at the end of the decade and the fact that some of the middle cohorts had, on average, a higher number of children or potential students.

The other item whose relative price has increased tremendously is health care. In the two panels of figure 2.33 we plot the age profiles for health expenditure. The difference between the two panels is in the deflator used to convert current into constant dollars. In the top panel we deflate average health expenditure by a cohort-specific CPI, while in the bottom one we use a health price index. It should be remembered that the CEX data on health expenditure are out-of-pocket figures. Therefore they do not include those items that are covered by health insurance (see section 2.2.1).

The profiles increase monotonically with age. In the top panel we also notice that there are strong time effects, evident in the large spikes corresponding to the last years of the sample for most cohorts. As a consequence of the large increase in the relative price of health care, the top profile, deflated by the CPI, increases over time at a much faster rate than the bottom one.

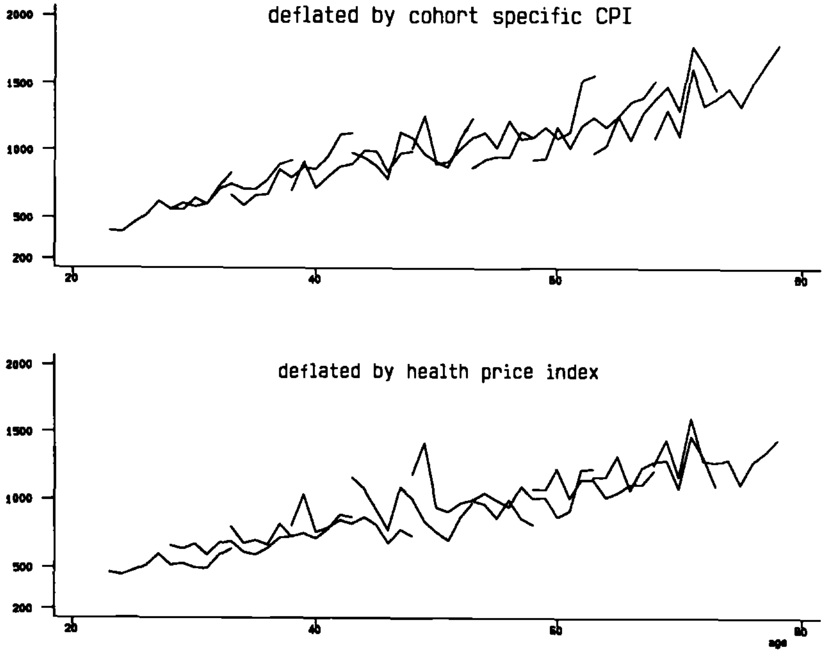

Fig. 2.33 Health expenditure

In figure 2.34 we plot the ratio of out-of-pocket health expenditure to total consumption. The ratio increases monotonically with age, reaching a level around 15 percent toward the end of the life cycle.

2.6 Saving

In this section we characterize the saving behavior of U.S. households. The analysis is similar to that in Attanasio (1993b), with two main differences: the definition of disposable income and the analysis of several consumption measures. The section is divided into three subsections. In subsection 2.6.1 we look at the age profiles for the *level* of saving. In subsection 2.6.2 we look at saving levels across education groups. Finally in subsection 2.6.3 we analyze individual *saving rates* conditional on various observable variables.

2.6.1 Saving Levels

As in section 2.3 saving is defined as disposable income minus consumption expenditure. Unlike in Attanasio (1993b), we do not consider as disposable income the BLS definition, but subtract from it contributions to Social Security. Employees' pension contributions are considered as saving because they

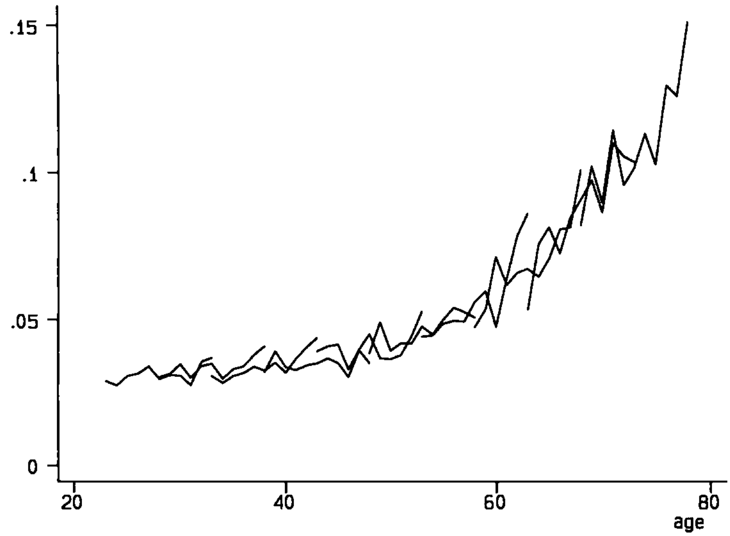

Fig. 2.34 Share of health in consumption

are not subtracted from disposable income.[54] That Social Security contributions should be considered as a tax is questionable: even if compulsory, they give the right to a stream of future income, just as pensions do.

Consumption can also be defined in different ways, since some expenditure items can be legitimately considered a form of investment. The obvious example is expenditure on durables: one should exclude it from total consumption expenditure and add instead the service flow from the stock of existing durables. Another example is expenditure on education and possibly health, which can be considered investment in human capital. Finally, mortgage payments include both service of the debt and repayment of the principal: the latter part should be considered saving.

Ideally, one would like to measure consumption, while only expenditure data are available. Unfortunately, the information available in the CEX is not sufficient to construct all the variables of interest. There is not enough information to estimate the service flow from the stock of existing durables, or the part of mortgage payments that repays the principal. Therefore we use different definitions of consumption that make, in turn, opposite and extreme assumptions: for instance, we consider cases where durable expenditure and mortgage payments are considered first as saving and then as consumption. In total, we consider five different definitions of consumption and therefore of saving. The first is the closest to the National Account definition and includes both expenditure on durables and mortgage payments. The second and third exclude ex-

54. As stressed in section 2.2 the CEX does not contain any information on employers' contributions to pensions.

penditure on education and health, respectively. The fourth and fifth exclude mortgage payments and durable expenditure. In a sense, in this subsection we put together the evidence presented in sections 2.4 and 2.5.

In figure 2.35 we plot the age profile for our benchmark definition of saving.[55] For each year-cohort cell we compute average disposable income by maximum likelihood and subtract from it average contributions to Social Security (as in section 2.4.1). Finally, we subtract average consumption (as in figure 2.26) to obtain saving.[56]

As in figures 2.13 and 2.26, we plot the raw cell means in the top panel, and in the bottom, smoothed profiles obtained using the same method as before. Several considerations are in order. First, it seems that saving increases in the first part of the life cycle, is highest just before retirement, and declines afterward. While these features are roughly consistent with the implications of the life cycle, given that we have not controlled for family composition, labor supply behavior, or any other variables, this evidence cannot be used in support of the model.

The smoothed age-saving profiles peak later than the corresponding profiles for disposable income: the polynomial in the bottom panel of figure 2.35 peaks around age 56, while that in figure 2.13 peaks at 50.[57]

A possible interpretation of figure 2.35, supported by the analysis of the smoothed profiles in the bottom panels of the two figures, is that the middle cohorts (cohorts 4 to 8) saved less, given their age, than the other cohorts in the sample. The only difference among the smoothed profiles of different cohorts is in the intercepts: in the second column of table 2.19 we report the ratio of the intercepts of each cohort's saving-age profile (in figure 2.35) to that of cohort 3. As a term of reference, in the first column of the same table we report the same ratio for the consumption-age profiles of figure 2.26.

The fact that the cohort-specific intercepts for the middle cohorts are lower than those for the younger ones does not necessarily reflect a behavioral change. It could conceivably be explained by the fact that we are considering saving *levels:* the intercepts of the smoothed profiles depend in an obvious way

55. In figures 2.35–2.38 we exclude data from 1980 and 1981. This was motivated by the fact that in 1980 and 1981 top-coded observations had all components of income (including Social Security contributions) top-coded. It turns out that the treatment of Social Security contributions and of top-coded observations in these two years affects substantially the shape of the estimated saving-age profiles. As a consequence we prefer to exclude these data from the analysis at this point. When considering Social Security as saving, the inclusion of 1980 and 1981 has no strong effects on the shape of the estimated profiles. The fact that cohorts are defined on a five-year band and that we have (after the exclusion of 1980 and 1981) nine years of data still leaves us with four ages of overlap between adjacent cohorts. After 1981 Social Security contributions were never top-coded.

56. Alternatively, we could have fitted a density to the cross-sectional distribution of savings with similar results.

57. If we consider Social Security contributions as a form of saving (as in Attanasio 1993b), we obtain a picture very similar to figure 2.35, except that the raw means are slightly smoother and the smoothed profiles peak earlier (around age 53).

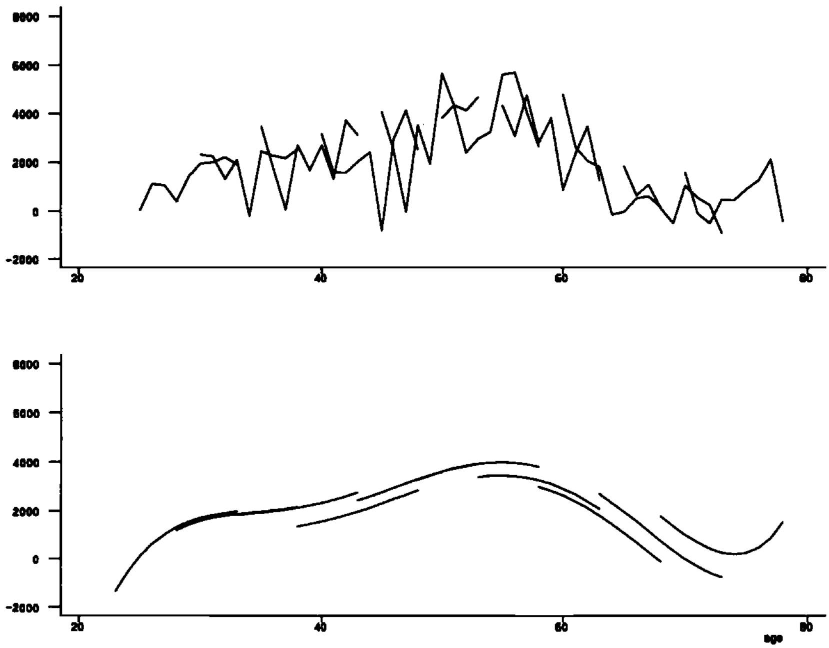

Fig. 2.35 Mean saving (benchmark)

on the amount of resources available to a given cohort. However, it is much harder to use this explanation for the fact that the intercepts for the middle cohorts are lower than those for the older cohorts, if, as we saw in the analysis of income and consumption data, cohort effects are positive and reasonably strong. Furthermore, the presence of negative cohort effects on the saving behavior of the middle cohorts is confirmed by the analysis of saving *rates* in section 2.6.3 below.

This interpretation is not uncontroversial, and it is not the only possible: an alternative would be to say that a typical saving-age profile is flat rather than bell-shaped in the middle part of the life cycle and that, therefore, the middle cohorts have *not* saved relatively less than the others. The problem is, of course, one of identification. As stressed above, it is not possible, without additional information, to identify separately age, cohort, and time effects and therefore a "pure" age profile. In this sense, the smoothing procedure used to obtain the bottom panel of figures 2.13, 2.26, and 2.35 is not neutral: to interpret the smoothed profiles as pure age profiles, one has to assume that the year effects sum up to zero over the sample period and that they are orthogonal to a linear trend. All trends in saving would then be interpreted as age effects. In Attanasio (1993b) we provide another, and maybe stronger, justification for this interpretation based on evidence from financial asset accumulation.

Table 2.19 Ratio of Cohort-Specific Age Profile Intercepts to Cohort-3 Intercept

Cohort	Consumption (1)	Saving (benchmark) (2)	Saving with Social Security (3)	Consumption Except on Education (4)	Consumption Except on Health (5)	Consumption Except on Housing (6)	Consumption Escept on Durables (7)
1	1.06	1.07	1.11	1.13	1.09	1.14	1.10
2	1.01	1.02	1.03	1.05	1.04	1.06	1.03
3	1.00	1.00	1.00	1.00	1.00	1.00	1.00
4	0.94	0.67	0.82	0.72	0.71	0.86	0.79
5	0.85	0.86	0.80	0.87	0.83	0.71	0.76
6	0.77	0.85	0.86	0.82	0.78	0.59	0.66
7	0.70	0.62	0.77	0.60	0.60	0.44	0.48
8	0.63	0.50	0.79	0.51	0.46	0.31	0.41
9	0.58	0.88	1.05	0.82	0.68	0.29	0.53
10	0.51	1.32	1.29	1.18	0.93	0.28	0.68

This reading of the evidence in figures 2.35 and 2.36, i.e., that the middle cohorts saved relatively "less," could explain the decline in aggregate saving observed during the 1980s. The cohorts that were observed in that part of their life cycle when saving is highest, saved "less." This is the main explanation for the decline in aggregate personal saving given in Attanasio (1993b).

Using alternative definitions of saving we obtain figures very similar to figure 2.35, which are not reported for the sake of brevity. In table 2.19, however, we report the ratio of cohort-specific intercepts to that of cohort 3 for the definitions of saving derived by adding Social Security contributions (column [3]), and excluding from consumption expenditure on education (column [4]), health (column [5]), housing, and mortgage payments. There are no substantial changes in the pattern of these intercepts.

Figure 2.36 is analogous to figure 2.35 except in the definition of saving, which is constructed excluding from consumption expenditure on durable commodities. Unlike the other definition of saving we experimented with, we observe a quite different pattern in the shape and relative position of the age profiles. This is particularly evident if we consider the smoothed profiles in the bottom panels. The only feature that they have in common with figure 2.35 is the characteristic hump shape. The increase in the early part of the life cycle, however, is steeper: young households are more likely to buy durables than old households. In fact, the share of durables over total consumption declines monotonically with age.

What is more interesting, however, is that the pattern of cohort effects, as represented by the cohort-specific intercepts, is very different. The ratio of the cohort-specific intercepts to the intercept of cohort 3 are reported in column (7) of table 2.19. We observe that the cohort-specific intercepts increase monotonically with the year of birth of the cohort. This is consistent with the pattern of cohort effects that we saw both for income and consumption in sections 2.4 and 2.5.

2.6.2 Saving by Education Group

In section 2.6.1 we ignored within-cohort heterogeneity. Yet we know from the analysis in sections 2.4 and 2.5 that the dynamics of disposable income and consumption is very different across education groups and that most of the changes in inequality occurred across rather than within groups.

It is therefore interesting to recast the analysis of the previous section conditioning on education. The drawback of this, of course, is that we will be looking at much smaller cells than in the unconditional case.

In figure 2.37 we plot age profiles for the basic definition of saving for the three education groups. As expected, the profiles are much noisier.[58]

58. To obtain the smoothed profiles we used weighted OLS on the cohort means, using as weights an estimate of the standard errors of mean income. For a few cells the maximum likelihood algorithm to estimate mean income did not converge: in these cases we consider a missing value for savings for that cell.

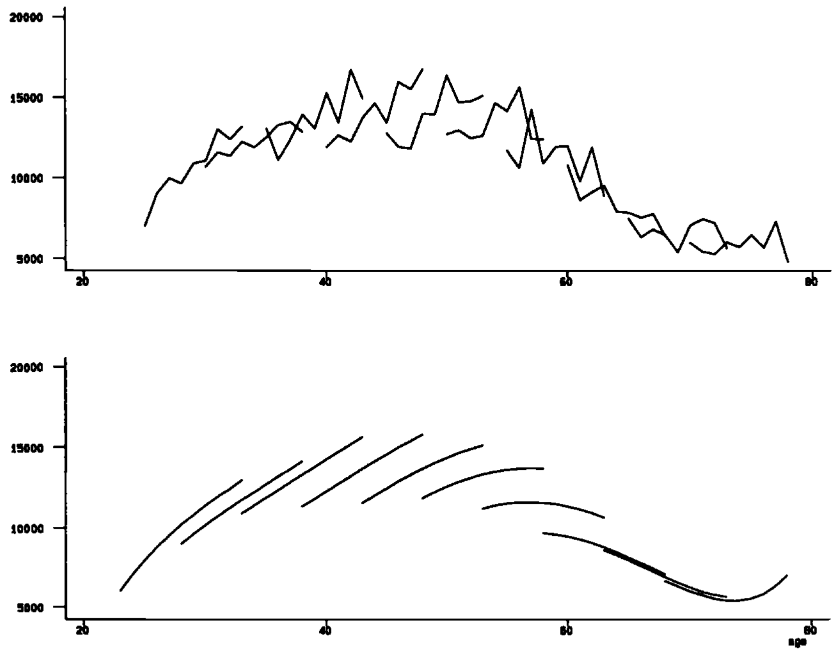

Fig. 2.36 Mean saving (excluding durable expenditure from consumption)

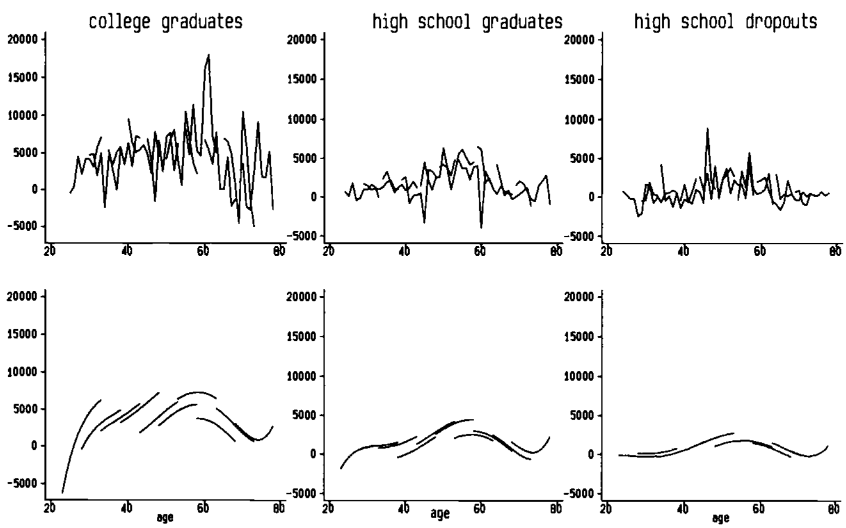

Fig. 2.37 Mean saving by eduction group

Some features, however, emerge pretty clearly. Saving is obviously higher for richer households: this could be explained simply by the larger amount of lifetime resources available to wealthier households. In the next section, where we look at saving *rates* rather than *levels*, we will come back to this problem.

The three smoothed profiles peak approximately at the same age group (55–57), but are quite different in shape, as one would expect, given the difference in the shape of income profiles.

The decline for the middle cohorts, discussed in section 2.6.1, while less evident because of the noisiness in the data, is still visible. The decline for college graduates obviously has larger aggregate effects, since they save more on average.

2.6.3 Individual Saving Rates

While the characterization of saving *levels* carried out in sections 2.6.1 and 2.6.2 can be translated directly into significant information about aggregate saving, the microeconomic behavior underlying these aggregates deserves to be analyzed further. To do so it might be more interesting to consider saving *rates*. Furthermore, both for statistical and economic reasons it might be more useful to consider measures of location other than means, in particular medians and various quantiles. The drawback, of course, is that neither ratios nor quantiles aggregate; they might provide a more interesting description of individual behavior, however.

Individual saving rates are very noisy almost by construction, as is evident by looking at their large variance and kurtosis. To deal with this problem we use two devices. First, we consider measures of location that are relatively robust to the presence of large outliers. Furthermore, instead of considering the traditional definition of saving rates (saving over income) we consider, as in section 2.3, saving over consumption.

In this section we will estimate several versions of the following equation:

$$SC_t^i = \theta' X^c + f(\text{age}_t^i) + \gamma' D_t + \beta' Z_t^i + u_t^i, \tag{2}$$

where SC_t^i is the ratio of saving to consumption for household i observed at time t, X^c and D_t are sets of cohort and year dummies, respectively, f is a polynomial in age, the Z_t^i are household-specific variables we want to control for, the u_t^i are residuals, and β, θ, and γ are parameter vectors; the γs are constrained to have zero mean and to be orthogonal to a time trend. The first three terms are equivalent to the terms used to obtain the smoothed age profiles in sections 2.4, 2.5, 2.6.1, and 2.6.2; the Zs control for various observable variables.[59] To construct the dependent variable we use the benchmark definition of saving.

59. As an alternative we also substituted the cohort and year dummies and the polynomial in age with fully interacted year-cohort dummies. The results were extremely similar and are available on request.

It is important to distinguish between Zs that are time invariant (such as race, education, and region) and Zs that may vary over the life cycle (such as the number of children, home ownership, and so on). When we condition on Zs that do not vary over time, the interpretation of the results is straightforward. If we were to fully interact the year-cohort dummies with these Zs, we would get entirely different profiles for different groups; this is similar to the analysis conducted in sections 2.4, 2.5, and 2.6 for the three education groups. In equation (2) we impose more structure in that we assume that the effects of the Z are the same across cohorts and time; this kind of restriction is motivated only by the lack of a large enough data set.

When we consider Zs that vary over time, the interpretation of the results is more complicated because of two related reasons. First, it is often the case that the variables we condition on are endogenous and are planned in advance by most households (home ownership or children). Second, most of these variables have a very distinctive life-cycle pattern which is going to interact with the age polynomial and the cohort dummies. Both of these problems make it very difficult to interpret the estimated coefficients. Nonetheless, as long as we do not intend to give the parameters any structural interpretation, they can be interesting.

Finally, notice that we never condition on income. This is because income is obviously endogenous and because it includes transitory shocks that, if the life-cycle model holds, are correlated with saving rate innovations. Instead we condition on variables such as education or race, which are likely to be correlated with permanent income.

Before turning to the regression results, we plot ratios of median saving to consumption in figure 2.38.[60] As with the other graphs of this kind we also plot a smoothed version of the same graph obtained using the same technique as above. The evidence in this graph confirms the interpretation of the evidence given in section 2.6.1: the middle cohorts seem to have saved, *given age, less* than other cohorts.

Table 2.20 contains estimates of θ and β in equation (2), and of the cohort dummies, obtained by least absolute deviations.[61] This is equivalent to estimating conditional medians. The reference group is households headed by a nonblack, high school dropout, residing in the West and belonging to cohort 10.

In column (1) we condition on variables that do not vary over the life cycle: the race and education of the household head and the region of residence. Education has a positive effect on saving rates: both the dummies for college grad-

60. The computation of the medians is equivalent to estimating eq. (2) with only year-cohort dummies on the right-hand side.

61. In this section, we completely ignore the problem of top-coding. Given that we estimate conditional *medians*, the bias introduced should not be large. The standard errors are computed using standard formulas for least absolute deviation estimation which assume homoskedasticity and gaussian residuals.

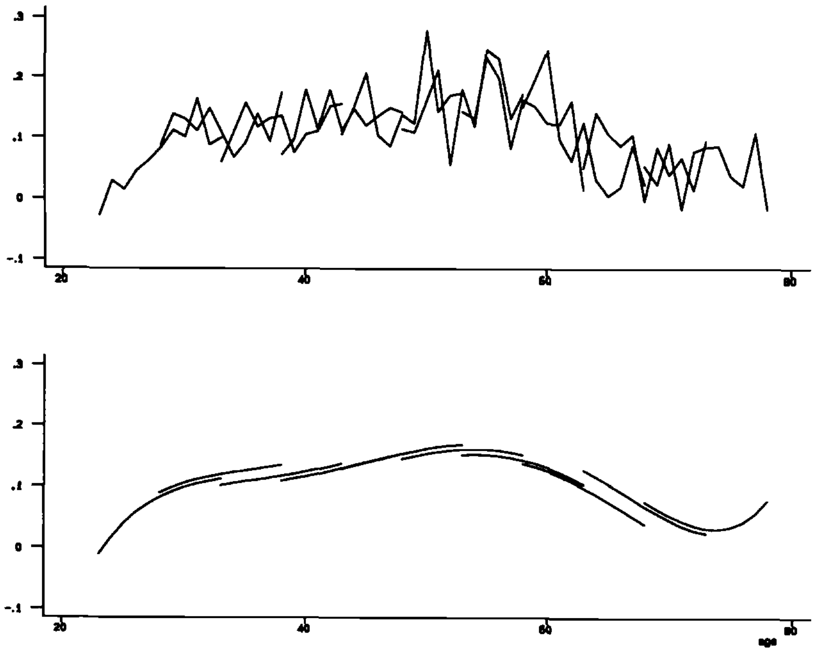

Fig. 2.38 **Median saving rates**

Table 2.20 **Saving Rate Regressions Using Least Absolute Deviations**

Variable	(1)	(2)	(3)	(4)	(5)
College graduate	0.125	0.128	0.091	0.124	0.084
	(0.008)	(0.009)	(0.009)	(0.009)	(0.008)
High school graduate	0.057	0.054	0.030	0.055	0.030
	(0.007)	(0.008)	(0.008)	(0.007)	(0.007)
Northeast	0.035	0.036	0.016	0.028	0.037
	(0.008)	(0.009)	(0.009)	(0.008)	(0.008)
Midwest	0.056	0.059	0.034	0.054	0.053
	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)
South	0.022	0.025	0.011	0.021	0.014
	(0.008)	(0.008)	(0.010)	(0.008)	(0.007)
Black	−0.021	−0.029	0.002	−0.017	0.020
	(0.009)	(0.010)	(0.010)	(0.009)	(0.008)
Self-employed head		−0.127	−0.142	−0.139	−0.188
		(0.012)	(0.012)	(0.011)	(0.010)
Self-employed spouse		−0.050	−0.067	−0.057	−0.088
		(0.018)	(0.017)	(0.016)	(0.015)
Home owner			0.115		
			(0.007)		
Home owner with mortgage			0.105		
			(0.009)		
Earners					0.095
					(0.004)

(continued)

Table 2.20 (continued)

Variable	(1)	(2)	(3)	(4)	(5)
Retired head					−0.049
					(0.012)
Female head					−0.119
					(0.006)
Children aged 3–15				−0.031	−0.030
				(0.003)	(0.003)
Children aged 16–17				−0.051	−0.097
				(0.009)	(0.009)
Infants				−0.024	−0.017
				(0.010)	(0.009)
Adults				0.064	−0.031
				(0.003)	(0.005)
Cohort 1	−0.040	−0.026	−0.055	−0.025	0.000
	(0.046)	(0.050)	(0.050)	(0.046)	(0.043)
Cohort 2	−0.040	−0.041	−0.052	−0.021	0.003
	(0.044)	(0.049)	(0.048)	(0.045)	(0.042)
Cohort 3	−0.042	−0.039	−0.063	−0.026	−0.009
	(0.043)	(0.047)	(0.046)	(0.043)	(0.040)
Cohort 4	−0.059	−0.058	−0.071	−0.052	−0.028
	(0.040)	(0.044)	(0.044)	(0.041)	(0.038)
Cohort 5	−0.064	−0.057	−0.089	−0.069	−0.047
	(0.037)	(0.041)	(0.041)	(0.038)	(0.035)
Cohort 6	−0.046	−0.043	−0.074	−0.042	−0.023
	(0.034)	(0.037)	(0.037)	(0.034)	(0.032)
Cohort 7	−0.070	−0.065	−0.100	−0.060	−0.056
	(0.030)	(0.033)	(0.032)	(0.030)	(0.028)
Cohort 8	−0.074	−0.082	−0.094	−0.070	−0.048
	(0.024)	(0.026)	(0.026)	(0.024)	(0.023)
Cohort 9	−0.021	−0.018	−0.048	−0.027	−0.012
	(0.019)	(0.021)	(0.021)	(0.019)	(0.018)

Notes: The dependent variable is the ratio of saving to consumption. The regression also includes a fifth-order polynomial in age and year dummies constrained to sum to zero and to be orthogonal to a linear trend. Numbers in parentheses are standard errors. The regression uses 47,647 observations. Year dummies are excluded from the specification in col. (4) because of convergence problems.

uates and high school graduates are very strongly significant and significantly different from one another. The higher saving rates for more educated households might indicate differences in tastes and/or economic opportunities, probably related to the higher level of permanent income of better educated households.

The regional dummies are statistically significant and positive, indicating that households living in the West have the lowest saving rates. The highest saving rates prevail in the Midwest.

The dummy for black household heads is negative and significant. Again, a variable likely to be related to households' permanent income takes a sign that implies a positive relationship between permanent income and saving rates.[62]

Finally, the pattern of the cohort dummies is consistent with the evidence presented in section 2.6.1: cohorts 5 to 8 seem to be those with the lowest conditional median saving rates. This is true for all the columns in table 2.20.

In column (2) we add two dummies to the previous specification, dummies that equal one if the household head or the spouse is self-employed. These variables have received some attention in the precautionary saving literature (see, e.g., Skinner 1988) because it is believed that self-employed individuals face riskier income. Consistent with that literature, the measured effect of self-employed status on the conditional median of saving rates is negative and strongly significant. One should treat these results with care: as we saw in section 2.3, there are strong variations in the proportion of self-employed individuals over both the life cycle and the business cycle. This indicates that assuming self-employed status is probably correlated with various economic variables and therefore with variations in labor income. In addition, it is likely that less risk-averse individuals will select into riskier occupations. Finally, income is likely to be underreported and consumption overreported for self-employed individuals. The coefficients on the other variables do not change substantially.

In column (3) we control for the effect of home ownership by introducing two dummies that equal one if the household owns its place of residence without and with a mortgage, respectively. Both dummies are positive and strongly significant. The interpretation of these results is difficult because of the life-cycle pattern that characterizes home ownership. It is likely that the positive sign of the coefficients reflects a correlation with individuals' permanent income.

Column (4) adds to the specification of column (2) the number of children between ages 3 and 15, the number between ages 16 and 17, the number of infants, and the number of adults. Not surprisingly, all the children variables are negative and strongly significant, probably indicating that a given level of consumption will produce different levels of utility depending on the number of children present, while the same children have, for the most part, no effects on income. The coefficient on the number of adults is positive and strongly significant.

Finally, in column (5) we add to the specification of column (4) the number of earners, a dummy for retired individuals, and a dummy for a female household head. While the first variable is strongly positive, the other two, perhaps not surprisingly, are negative and significant. The other noticeable features of

62. It might be interesting to note that one of the facts that was discussed in the early literature on the permanent-income hypothesis, namely, that black households save *more* for each income level than nonblack households, holds true in the CEX sample.

this specification are that the dummy for black household heads is no longer significant and that the coefficient on the number of adults is now negative, probably indicating that the positive sign in column (4) picked up the correlation between the number of adults and the number of earners.

2.7 Financial Wealth Accumulation

In this section we look briefly at one of the most investigated issues in the life-cycle literature: the accumulation of wealth over the life cycle. In what follows, we analyze the main features of the wealth data in the CEX: the focus, however, is not the issue of asset decumulation by the elderly. Most of the results in this section are taken from Attanasio (1993a).

As described in section 2.2, the CEX provides some information on financial asset holdings. This information is only collected in the last interview: as a consequence of the fact that some households do not reach the fourth interview and that some households that have valid information on consumption and income do not respond to the questions on assets, in this section of the paper we use a reduced sample of 32,050 observations.

We divide financial wealth into two components: the first, defined as liquid assets, includes savings and checking accounts; the second includes U.S. savings bonds and other bonds and equities. In figure 2.39 we plot mean and median total financial wealth age profiles.[63] These means and medians are calculated by fitting a distribution that has a mass point at zero (to take into account households reporting zero assets, who constitute approximately 15 percent of the sample) and is given, for households with positive assets, by the mixture of two log-normal densities.

The main features of these graphs are by and large consistent with what is found in the literature: financial wealth increases with age at decreasing rates until retirement. After retirement average and median wealth seem to have little tendency to decline. In Attanasio (1993a) we discuss why this does not necessarily contradict the life-cycle model: failure to observe a decline in assets at the end of the life cycle could be caused by changes in preferences, the failure to control properly for retirement (see Hurd 1989), the correlation between mortality and wealth (see Shorrocks 1975; Attanasio and Hoynes 1993), or uncertainty.

In figure 2.40 we plot the proportion of households with positive quantities of nonliquid assets. A first observation is that this proportion is extremely low. In the whole sample, only 25 percent of households report positive quantities of nonliquid assets. If we look at this proportion in the year-cohort cells we see that it increases rapidly until age 40 and then flattens out well below 40 percent. In the last part of the life cycle it declines, even though it is not clear

63. These graphs are reproduced from Attanasio (1993a).

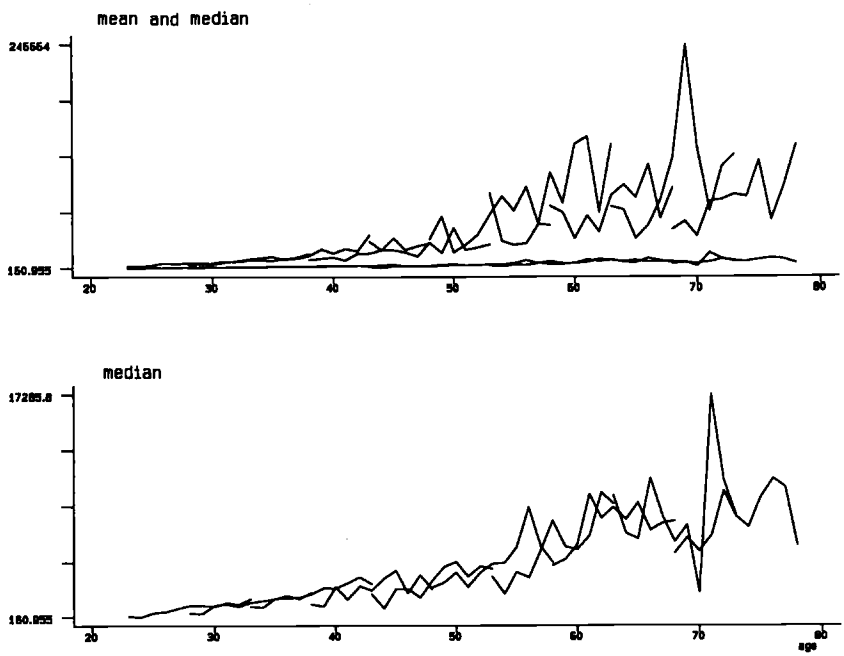

Fig. 2.39 Mean and median total financial assets

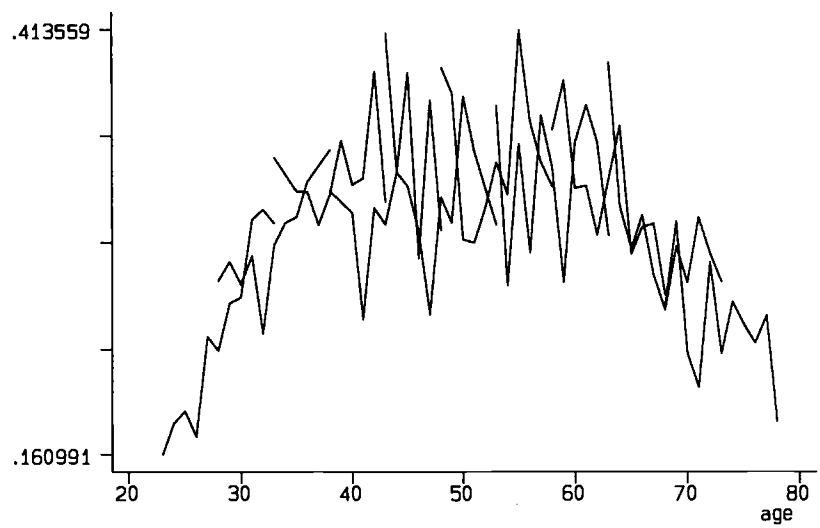

Fig. 2.40 Percentage of households with positive nonliquid assets

whether this is a sign of asset decumulation or a reflection of the fact that older cohorts were less likely to hold nonliquid financial assets.

Given that the sample we use to study financial assets is considerably smaller than the one used in the rest of the paper and that less than 30 percent of it holds positive quantities of nonliquid assets, the maximum likelihood estimates of the density's parameters exhibited a great deal of variability in many cells and often take implausible values. We only report the sample 90th percentile for nonliquid assets, which should give us consistent estimates of the corresponding parameters in the population as long as the top-coding level is above that percentile. These estimates are plotted in figure 2.41, and tell a story similar to that of the median in figure 2.39: nonliquid financial assets keep growing until around age 60. In the last part of the life cycle, asset levels stabilize and possibly decline slightly: unfortunately the estimates are much noisier in the last part of the sample, partly because of the smaller size of those cells.

Probably the most important feature that emerges from these data is that a large number of households hold (at least directly) very small quantities of financial assets. The ratio of the stock of financial assets to annual total consumption expenditure averages only 0.57 for the whole sample.[64] This average, however, is somewhat misleading because of the strong skewness of the ratio: its median is only 0.08, and even its third quartile is below the mean (0.43).[65]

These figures hide large differences across the education groups we considered in the previous sections: the median of the ratio of financial assets to consumption is 0.21 for college graduates, 0.07 for high school graduates, and below 0.01 for high school dropouts.

The low level of assets for most households does not necessarily mean that their behavior is irrational or that they are liquidity constrained.[66] It should be remembered that the assets we are considering do not include two of the most important assets in the portfolio of U.S. households: real estate and pension wealth. This low level is, however, consistent with the evidence on interest income presented in section 2.4.2. In that section we saw, however, that interest income becomes somewhat important only around retirement age. It is therefore interesting to look at the life-cycle profile of the ratio of financial assets to consumption, as in figure 2.42. The most noticeable feature of the figure is probably the sharp increase just before retirement: the median ratio goes from

64. We prefer to use the ratio of financial assets to consumption rather than to income for two reasons. First, under the life-cycle permanent-income hypothesis, consumption is less affected than income by temporary fluctuations: consumption should be related to permanent rather than current income. Second, in the present data set, consumption is better measured than income. Furthermore, consumption unlike income is not top-coded; this allows us to determine at least the sign of the bias introduced by the observations which have top-coded financial wealth. The results are not affected dramatically, however, if income rather than consumption stands in the denominator.

65. The ratio of average financial assets to average consumption in the sample is about 0.6.

66. Deaton (1991) constructs a model of optimizing behavior with liquidity constraints in which agents hold (optimally) very small amounts of assets.

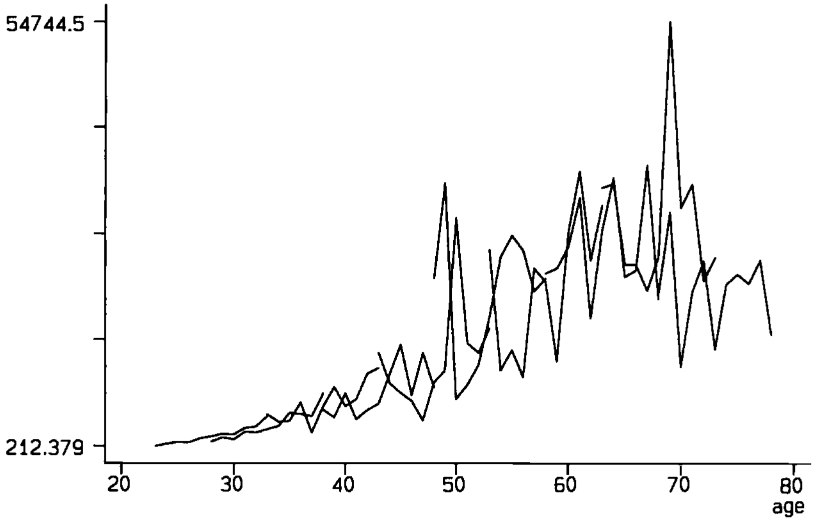

Fig. 2.41 90th percentile of nonliquid financial assets

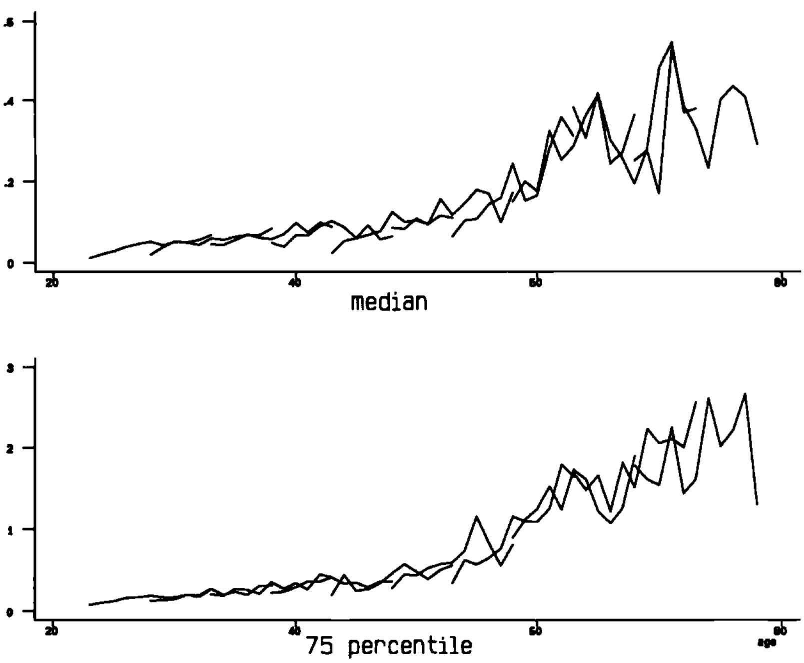

Fig. 2.42 Ratio of financial wealth to consumption

around 0.17 at age 60, to around 0.4 at age 65, to around above 0.5 at age 71; the third quartile goes from around 1.1 at age 60 to around 1.45 at age 65.

2.8 Conclusions

This paper presents a detailed analysis of the only U.S. microdata set that contains complete and exhaustive data on consumption, as well as information on income, wealth, labor supply, and a variety of other demographic and socioeconomic variables. While the quality of the data is far from perfect, the data set gives a unique opportunity to study savings at the individual level, as is essential for understanding the recent dynamics of aggregate saving rates.

The model we use as a benchmark and frame of reference is the life-cycle theory of consumption. The paper, however, is descriptive in nature: we estimate and characterize age profiles for a variety of variables directly or indirectly relevant to savings. We think of this as an essential first step which should be followed, in the future, by a more structural analysis. The descriptive statistics are very suggestive but, without additional structure, are consistent with numerous and diverse interpretations.

The main features that emerge from the analysis are as follows:

1. Both the income and consumption age profiles are hump-shaped, though without a model of household production and of joint consumption and labor supply decisions, this cannot be taken as evidence against the life-cycle model. There are fairly strong cohort effects, probably reflecting productivity growth, as younger cohorts appear "wealthier" than older ones.

2. Profiles of savings by age are hump-shaped. This is roughly consistent with the life-cycle model, and it is true across education groups, but especially for more highly educated individuals.

3. A possible interpretation of the evidence presented on savings is that the middle cohorts (roughly, the households headed by individuals born between 1925 and 1940) saved less than the other cohorts, keeping age constant. This is consistent with the decline in aggregate savings during the 1980s, because those cohorts were in that part of the life cycle when savings are the highest.

4. When we consider alternative definitions of savings to take into account the fact that some expenditure items have an important element of savings, a different pattern emerges for the definition that excludes durable consumption from consumption.

5. A regression analysis of saving rates shows that savings are higher for better educated individuals, nonblacks, and households residing in the Midwest.

6. Very few households hold financial assets directly (we do not have information on pension wealth and how it is invested); furthermore, the median level of financial assets is very low. Very few households, especially in the low

tail of the distribution, receive substantial amounts of interest income, at least until retirement age. Interest income is much more skewed than labor income.

7. There is little evidence of asset decumulation by the elderly.

8. Real estate is one of the principal means by which U.S. households hold wealth.

9. There are large differences in the dynamics of consumption, income and savings across education groups.

Several issues are left for future research. We think that there are four areas of research that are particularly promising. First, we think it necessary to develop, estimate, and test a structural model which considers consumption and labor supply decisions jointly. It seems obvious that saving and consumption choices cannot be considered separately from labor supply decisions; however no systematic evidence yet exists on this issue. Second, it is important to understand the effects of changing family composition on consumption and saving. This is only possible with the help of a structural model which allows for flexible adult-equivalent schemes. Third, it is important to gather more data on real estate wealth, as it is clear that it is one of the most important assets in the portfolio of U.S. households. The new data sets released by the BLS since 1988 seem the right ones to use. Fourth, it is important to study the implications of the dynamics of income and consumption inequality for aggregate savings.

References

Attanasio, O. P. 1993a. An analysis of life-cycle accumulation of financial assets. *Ricerche economiche* 47.

———. 1993b. A cohort analysis of saving behavior by US households. NBER Working Paper No. 4454. Cambridge, Mass.: National Bureau of Economic Research.

Attanasio, O. P., and M. Browning. 1992. Consumption over the life cycle and over the business cycle. NBER Working Paper No. 4453. Cambridge, Mass.: National Bureau of Economic Research.

Attanasio, O. P., and H. W. Hoynes. 1993. Differential mortality and wealth accumulation. In preparation.

Bosworth, B., G. Burtless, and J. Sabelhaus. 1991. The decline in saving: Evidence from household surveys. *Brookings Papers on Economic Activity*, no. 2.

Browning, M. J., A. S. Deaton, and M. Irish. 1985. A profitable approach to labor supply and commodity demands over the life-cycle. *Econometrica* 53:503–44.

Carroll, C. D., and L. H. Summers. 1991. Consumption growth parallels income growth: Some new evidence. In *National savings and economic performance*, ed. B. D. Bernheim and J. B. Shoven. Chicago: University of Chicago Press.

Curtin, R. T., F. T. Juster, and J. N. Morgan. 1987. Survey estimates of wealth: An assessment of quality. In *The measurement of saving, investment and wealth*, ed. R. E. Lipsey and H. S. Tice. Chicago: University of Chicago Press.

Deaton, A. 1985. Panel data from time series of cross sections. *Journal of Econometrics* 30:109–26.

———. 1991. Saving and liquidity constraints. *Econometrica* 59:1221–48.

Deaton, A., and C. Paxson. 1992. Aging and consumption in Taiwan. Princeton University. Mimeograph.

Feenberg, D., and J. Skinner. 1989. Sources of IRA savings. In *Tax policy and the economy 1989*, ed. L. H. Summers. Cambridge: MIT Press.

Friedman, M. 1957. *A theory of the consumption function.* Princeton, N.J.: Princeton University Press.

Friend, I. 1954. *Individuals' saving.* New York: Wiley.

Friend, I., and R. Jones, eds. 1960. *Proceedings of the conference on consumption and saving.* Philadelphia: University of Pennsylvania.

Ghez, G. R., and G. S. Becker. 1975. *The allocation of time and goods over the life cycle,* New York: Columbia University Press.

Gieseman, R. 1987. The Consumer Expenditure Survey: Quality control by comparative analysis. *Monthly Labor Review* 108.

Goldsmith, E. 1956. *A study of saving in the US.* Princeton, N.J.: Princeton University Press.

Gravelle, J. G. 1991. Do individual retirement accounts increase savings? *Journal of Economic Perspectives* 5:133–48.

Heckman, J. J., and R. Robb. 1987. Using longitudinal data to estimate age, period and cohort effects in earnings equations. In *Cohort analysis in social research,* ed. W. M. Mason and S. E. Fienberg. New York: Springer-Verlag.

Hurd, M. 1989. Mortality risk and bequest. *Econometrica* 57:779–813.

Jianakoplos, N. A., P. L. Menchik, and F. O. Irvine. 1987. Using panel data to assess the bias in cross-sectional inferences of life-cycle changes in the level and composition of household wealth. In *The measurement of saving, investment and wealth,* ed. R. E. Lipsey and H. S. Tice. Chicago: University of Chicago Press.

Juster, T. F. 1966. *Household capital formation and financing.* New York: Columbia University Press.

Kimball, M. 1990. Precautionary saving in the small and in the large. *Econometrica* 58:53–73.

King, M. A., and L. D. L. Dicks-Mireau. 1982. Asset holdings and the life cycle. *Economic Journal* 92:247–67.

Kotlikoff, L., and L. Summers. 1981. The role of intergenerational transfers in aggregate capital accumulation. *Journal of Political Economy* 86:706–32.

———. 1988. The contribution of intergenerational transfers to total wealth: A reply. In *Modelling the accumulation and distribution of wealth,* ed. D. Kessler and A. Masson. Oxford: Clarendon.

MaCurdy, T. E., and T. Mroz. 1990. Measuring macroeconomic shifts in wages from cohort specifications. Stanford University. Mimeograph.

Modigliani, F. 1988. Measuring the contribution of intergenerational transfers to total wealth: Conceptual issues and empirical findings. In *Modelling the accumulation and distribution of wealth,* ed. D. Kessler and A. Masson. Oxford: Clarendon.

Modigliani, F., and R. Brumberg. 1954. Utility analysis and the consumption function: An interpretation of cross section data. In *Post Keynesian economics,* ed. K. K. Kurihara. George Allen and Unwin.

Moffitt, R. 1991. Identification and estimation of dynamic models with a time series of repeated cross sections. Brown University. Mimeograph.

Paulin, G., M. Boyle, R. Branch, and R. Cage. 1991. Comparison of 1987–1988 CEX integrated survey and PCE annual estimates. Washington, D.C.: Bureau of Labor Statistics. Mimeograph.

Shorrocks, A. F. 1975. The age-wealth relationship: A cross-section and cohort analysis. *Review of Economics and Statistics* 57:155–63.

Skinner, J. 1988. Risky Income, life cycle consumption and precautionary savings. *Journal of Monetary Economics* 22:237–55.

Slesnick, D. T. 1992. Aggregate consumption and saving in the postwar United States. *Review of Economics and Statistics* 74:585–97.

Venti, S. F., and D. A. Wise. 1990. Have IRA's increased US saving? Evidence from consumer expenditure surveys. *Quarterly Journal of Economics* 105:661–98.

This Page Intentionally Left Blank