Review of Income and Wealth
Series 45, Number 1, March 1999

THE AGE–WEALTH PROFILE AND THE LIFE-CYCLE HYPOTHESIS: A COHORT ANALYSIS WITH A TIME SERIES OF CROSS-SECTIONS OF ITALIAN HOUSEHOLDS

BY TULLIO JAPPELLI

CSEF, Universitá di Salerno, and CEPR

In this paper I estimate the age–wealth profile under two different identification assumptions about age, cohort and time effects. According to the life-cycle model, the two sets of assumptions should yield similar age–wealth profiles. Using the 1984–93 Italian Survey of Household Income and Wealth, the estimated average annual rate of wealth decumulation in old age is found to be between 3 and 6 percent. As in the life-cycle model, the cohort effect increases with year of birth. However, the results also uncover considerable population heterogeneity: the rates of wealth decumulation are much lower for rich households and households headed by individuals with higher education.

1. INTRODUCTION

At the individual level the life-cycle hypothesis predicts that wealth increases up to retirement, and declines smoothly thereafter. Life-span uncertainty and health hazards reduce the optimal rate of wealth decumulation during retirement but do not change the basic insight of the model. The other assumption of the life-cycle model is that growth takes place across generations but not over the lifetime of a single individual, so that any increase in growth shifts the earnings profile upwards, without affecting its shape (Modigliani, 1986; Deaton, 1999). This implies that an increase in productivity growth redistributes resources from older to younger generations, inducing an increase in the aggregate saving rate.

The prediction that the age–wealth profile is hump-shaped cannot be tested with cross-sectional data, because the individuals interviewed in any cross-section belong to different generations (cohorts) which differ in mortality rates, preferences and, most importantly, productivity. Use of out-of-sample information to impute cohort effects in cross-sectional data was proposed by King and Dicks-Mireaux (1982). While ingenious, this approach assumes that cohort effects are present in the data, and does not provide a test of the hypothesis that the magnitude of these effects is explained mainly by productivity growth, as assumed by the life-cycle hypothesis, rather than by other factors.

Panel data allow age and cohort effects to be disentangled. However panel data with information on wealth are rare; even when available, measurement errors and sample attrition pose difficult econometric problems; disentangling

Note: I thank Orazio Attanasio for very helpful suggestions; Gary Engelhardt, John Flemming, Luigi Guiso, Michalis Haliassos, Marco Pagano and seminar participants at Tilburg and Bern for comments; the Research Department of the Bank of Italy for providing the data used in this paper; the Training and Mobility of Researchers Network Program (TMR) of the European Commission DGXII, the Italian National Research Council (CNR) and the Ministry for Universities and Scientific Research (MURST) for financial support.

age and time effects requires additional identification assumptions. For countries lacking panel data on wealth, repeated cross-sectional data represent a third alternative, but these data also require strong assumptions to identify age, cohort and time effects.

The life-cycle model suggests useful (and testable) identification assumptions. The unique feature of the model is to make strong predictions about the shape of the age–wealth profile as well as about the size of the cohort effects. The elderly should run down accumulated assets (the shape of the individual age–wealth profile is concave) and the size of the cohort effect should depend positively on year of birth (the shift of the age–wealth profile induced by productivity growth across generations increases with year of birth). Only the first hypothesis concerns households behavior; the second emphasizes the mechanism by which productivity growth affects the resources of successive generations. Models with infinite horizons or altruistic bequests and buffer stocks models of saving do not have such implications.

One possibility, then, is to assume that there is no uncertainty and to regress wealth on a set of age and cohort indicators, dropping all time effects.[1] Alternatively, one can assume that cohort effects in wealth and consumption are equal. This suggests a different normalization, i.e. to regress the wealth-consumption ratio on a set of age and unrestricted time effects, dropping all cohort effects. Comparison of the wealth age-profiles obtained under the two procedures provides a convenient joint test of the validity of the life-cycle model and of the identification assumptions.

According to theoretical simulations reviewed in Section 2, the optimal yearly rate of wealth decumulation during retirement is between 3 and 8 percent, depending on the parameter assumptions and on the types of uncertainty posited (income, mortality or health). Section 3 reviews the approaches to estimating the age–wealth profile offered by the literature. The data, presented in Section 4, are drawn from the 1984–93 Italian Survey of Household Income and Wealth (SHIW).

The main findings of this paper are not at variance with the life-cycle model. Sections 5 and 6 report that the age–wealth profiles estimated using wealth and the wealth-consumption ratio as the dependent are very similar. The estimated average yearly rate of wealth decumulation after age 70 is between 5 and 6 percent. The pattern and magnitude of the estimated cohort effects are consistent with the hypothesis that they depend mainly on productivity growth. However, we also find considerable population heterogeneity: regressions by educational attainment of the household head and quantile regressions indicate the absence of wealth decumulation in old age for rich and better educated households. Section 6 summarizes the results and their implications for the explanation of the macroeconomic relation between saving and growth.

2. The Determinants of the Age–Wealth Profile

The life-cycle hypothesis posits that the main motivation for saving is to accumulate resources to be drained down for later expenditure and in particular

[1]As will be seen, a slightly less restrictive approach is to assume that time effects are orthogonal to a time trend and sum to zero (Deaton and Paxson, 1994).

during retirement. Saving should be positive for young households and negative for the retired, so that wealth should be hump-shaped. Only in special cases is it possible to derive explicitly the wealth behavior from the earnings profile and optimal consumption plan implied by the life-cycle model. The simplest case is one in which a household (individual indexes are suppressed) belonging to generation b earns a constant income y_b until retirement age N (so that lifetime resources are $H_b = y_b N$), lives T years, there is no uncertainty, the interest rate and the rate of time preference are zero and optimal consumption is constant at $c = H_b/T$. Then wealth of an individual of age a increases up to retirement, $W_{a,b} = (a/N)(1-(N/T))H_b$ for $a = 0, \ldots, N-1$, and declines afterwards, $W_{a,b} = (1-(N+a)/T)H_b$ for $a = 0, \ldots, (T-N)$. This age–wealth profile is hump-shaped and independent from lifetime resources:[2]

$$W_{a,b} = f(a)H_b \tag{1}$$

where $f(\cdot)$ is a concave function of age. For convenience $f(a)$ is specified as a function of age alone, but one should keep in mind that in more realistic examples wealth accumulation depends also on households' preferences, the interest rate, the life-cycle variation in household size and composition, and the rules governing retirement. Equation (1) implies that the shape of the wealth profile depends on age, regardless of resources, while lifetime resources, regardless of age, set the position of the profile. Introducing a positive real interest rate, or more realistic earnings profiles does not change this basic implication of the model.

Liquidity constraints early in the life cycle and uninsurable income risk modify the path of wealth accumulation. With imperfect markets households will accumulate wealth before retirement at a rate that is higher than under perfect markets. However once liquidity constraints cease to be binding and income uncertainty is resolved because the individual retires, selfish consumers will run down accumulated wealth faster than is predicted by the life-cycle hypothesis with perfect markets (Caballero, 1991). Thus, income risk and liquidity constraints increase the concavity of the $f(a)$ function.

Lifetime uncertainty and non-insurable health hazards, both increasing with age, have an opposite effect on the age–wealth profile: they induce the elderly to hold assets for precautionary purposes, thus reducing the rate of wealth decumulation during retirement. Davies (1981) points out that in the absence of annuity markets the elderly do not run down assets completely unless there is zero probability of surviving any longer. Davies' analysis further suggests that even with life uncertainty wealth must decline at some age—depending on retirement age and the form of the utility function—and that after this age wealth should continue to decline smoothly. Using standard parameter assumptions, he shows that life-span uncertainty reduces the average annual rate of wealth decumulation between ages 65 and 85 from 7 to 3 percent with respect to the standard life-cycle model.[3]

[2] In the presence of an upward sloping earnings profile, without liquidity constraints or income uncertainty, the standard model also predicts dissaving in the early part of the life-cycle.

[3] He assumes an isoelastic utility function with a coefficient of relative risk aversion of 4, a rate of time preferences of 0.015 and an interest rate of 0.03. When the coefficient of risk aversion is increased to 5 the implied rate of wealth decumulation is 2.1 percent.

Hubbard, Skinner, and Zeldes (1994) perform a simulation for a representative household facing life-span, health and income uncertainty. In the absence of all sourcès of uncertainty, the model predicts decumulation rates of about 8 percent per year during retirement; when length of life is the only source of uncertainty, the average decumulation rate slows to about 3 percent, as in Davies. When all sources of uncertainty are introduced, the rate of wealth decumulation is about 5 percent per year.

Altruistic behavior also affects the age–wealth profile. If the primary purpose of accumulation is to leave a bequest to one's children, households behave as if their horizon were infinite and wealth does not decline in old age (Barro, 1974). Even if bequests are not motivated by altruism, as in Barro, but rather by strategic motives, by the joy of giving or by deriving utility from terminal assets, wealth decumulation is smaller than predicted by the life-cycle model under uncertainty (Hurd, 1989). The prediction is not as sharp when intergenerational transfers take the form of *inter vivos* gifts, rather than bequests. Altruistic parents should transfer resources when they are most needed by their heirs. For instance, if the young are liquidity constrained this will occur around the retirement age of the donors. Significant wealth decumulation after retirement may then reflect gifts transferred to liquidity-constrained children, rather than selfish behavior.

In principle, almost any shape of the age–wealth profile can be made consistent with altruism, and evidence on decumulation rates does not provide a conclusive test for the existence and nature of the bequest motive. However, in Italy, which is the focus of the present analysis, gifts are not a major component of wealth. A special section of the 1991 SHIW indicates that only 5.4 percent of households report to have ever received a sizable gift (in the form of real estate or financial assets) from parents or relatives, and that gifts represent about 4 percent of households' net worth and 20 percent of total intergenerational transfers (Guiso and Jappelli, 1998). Thus, in countries in which gifts and other *inter vivos* transfers do not represent a large share of intergenerational transfers, the shape of the age–wealth profile is a significant indicator of the importance of life-cycle and precautionary motives for saving as compared with the bequest motive.

The shape of the age–wealth profile also has important implications for the debate about the source of the well-established macroeconomic relation between saving and growth. The life-cycle model derives such relation assuming that productivity growth is generation-specific.[4] The growth-saving link results purely from the aggregation mechanism because growth redistributes resources towards younger cohorts (Modigliani, 1986). Note, however, that a hump-shaped wealth profile is by no means a necessary condition for the existence of a positive relation between growth and aggregate saving.[5]

[4]When productivity-growth is individual-specific, higher expected income growth reduces individual saving. With both types of growth, the aggregate correlation between saving and growth is *a priori* ambiguous.

[5]Other consumption models, such as habit persistence, predict a positive relation between aggregate saving and growth. Neoclassical and endogenous growth models imply a link running from growth to saving, rather than from saving to growth.

3. What Kind of Data, and What Kind of Identification Assumptions?

In order to illustrate the problems encountered in estimating the age–wealth profile empirically, it is convenient to express generation-specific resources as $H_b = H_0 e^{\rho b}$, where H_0 denotes resources common to all generations, regardless of year of birth and ρ the growth rate of productivity across generations, and write equation (1) in logs:[6]

$$\ln (W_{a,b}) = \ln f(a) + \ln (H_0) + \rho b. \tag{2}$$

Use of cross-sectional data to estimate (2) can be highly misleading (Shorrocks, 1975; Mirer, 1979). The individuals interviewed in any cross-section belong to different generations that differ in mortality rates, preferences and institutional arrangements; and most importantly, from the point of view of this paper, productivity differs between generations. For instance, the productivity of an individual entering the labor force in the fifties is much less than that of an individual born in the seventies and just now entering the labor force. Since the elderly are considerably poorer than the young over the lifetime, they also have lower lifetime resources and wealth. Thus, a finding that wealth declines with age in a cross-section may derive from the fact that older generations are less productive than younger generations, and tells little about households' behavior. In short, in a cross-section one cannot identify both age and cohort effects (in year t, the difference in wealth between a 20 and a 21 years old is equivalent to the difference between somebody born in year $t-20$ and somebody born in $t-21$).

There are three ways to control for the presence of cohort effects: panel data, out-of-sample information, and repeated cross-sectional data. Ideally, the best approach is the first. Wealth panel data allow the econometrician to measure decumulation rates of retired people of one particular cohort b according to the length of retirement (rather than age). For instance, Diamond and Hausman (1984), find rates of dissaving after retirement of about 5 percent per year in the National Longitudinal Survey of Mature Men. Hurd (1987), using the Retirement History Survey, finds decumulation rates of about 1.5 percent per year (3 percent excluding housing in the definition of wealth) and that couples with independent children dissave more during retirement than childless couples.

In Italy and many other countries this approach is not practicable for lack of long panel data. The ingenious method devised by King and Dicks-Mireaux (1982) is to construct a proxy for generation-specific resources by estimating an earnings function and to use out-of-sample information on ρ to scale the earnings of the different cohorts according to their productivity growth rates. The (log) of the ratio of wealth to permanent income is then regressed on a function of age and other demographic controls. This is equivalent to assuming that cohort effects depend only on generation-specific productivity growth. This approach has some unique advantages: it allows interpretation of the cross-sectional variation in the age-profile as the pure age effect; by introducing permanent income as a separate regressor it allows a straightforward test of homotheticity; by adding

[6] Deaton and Paxson (1997) propose the same approach to model consumption and saving profiles.

a measure of social security wealth it provides an estimate of the degree of substitution between private and pension wealth. King and Dicks-Mireaux estimates, obtained on Canadian data, imply a rate of wealth decumulation between ages 65 and 85 of 0.7 to 1.5 percent per year. Brugiavini (1987) applied this method to 1984 Italian data and finds that the average rate of wealth decumulation during retirement ranges from 2 to 8 percent (depending on the specification). One limitation of this approach is that cohort effects are assumed to affect the wealth profile, i.e. one cannot test if they depend mainly on productivity growth, as in the life-cycle model.

The third approach, pioneered by Shorrocks (1975) and Masson (1986), is to control for differences in productivity and preferences between generations using a time-series of cross-sectional data.[7] Repeated cross-sections allow the econometrician to track cohorts over time. Although the same individual is only observed once, a sample from the same cohort is observed in a later survey. Empirically, one can specify a flexible functional form for wealth:

$$\ln (W_{a,b}) = g(a) + X\beta + h(b) + d_{a+b}\delta + \varepsilon \tag{3}$$

where $g(a)$ is a polynomial in age, X is a matrix of variables that affect households resources H_0 (such as sex, region of residence and schooling) regardless of year-of-birth, $h(b)$ is a cohort polynomial, ε an idiosyncratic component and d_{a+b} a set of time effects to be discussed below. In principle, this equation allows one to estimate not only if the $g(a)$ function is concave, but also if the $h(b)$ function is related to productivity growth. However, these estimates do not come for free.

Equation (3) was derived under special assumptions and in the absence of uncertainty. With uncertainty, macroeconomic shocks lead to revisions in households resources, and therefore in assets accumulation. Measurement errors can also generate disturbances to the wealth equation. To incorporate time effects, one can add to equation (3) a set of calendar year fixed effects d_{a+b} $(a+b=m, m+1, \ldots, M$, where m and M are the first and last available cross-sections). The subscript highlights that the year in which each household is sampled equals age plus year of birth. The separate effect of a, b and $a+b$ is therefore not identified unless one is willing to make additional identification assumptions.

One possibility is to rule out uncertainty and measurement errors and eliminate all year dummies. Deaton and Paxson (1994) adopt a slightly less restrictive approach and assume that $\sum_{a+b=m}^{M} d_{a+b} = 0$ and $\sum_{a+b=m}^{M} (a+b)\, d_{a+b} = 0$, so that the year dummies sum to zero and are orthogonal to the time trend $(a+b)$. This is equivalent to assuming that all trends in the data can be interpreted as a combination of age and cohort effects and are therefore, by definition, predictable. The time effects then reflect additive macroeconomic shocks or the residual influence of non-systematic measurement error.

This approach is useful but restrictive. In the model with uncertainty the impact of macroeconomic shocks on saving and wealth depends on age (for

[7]Shorrocks (1975) used 60 years of estate-duty statistics, concluding that wealth is an increasing function of age. These statistics over-represent the most affluent households. Masson (1986) constructed cohort-adjusted age–wealth profiles using four cross-sections of French data. He found annual rates of decumulation ranging from 0.7 percent for wealthy self-employed persons to 3–4 percent for wage-earners.

instance, young households should react more to a given shock than individuals close to retirement).[8] Such interaction terms are ruled out in equation (3). Furthermore, since wealth is equal to accumulated saving, past macro shocks may affect wealth at time $(a+b)$. An alternative normalization that would allow past macro shocks to be reflected in the time effects would be to redefine the time dummies as $d_m = 1$ for $(a+b) = m, \; m+1, \ldots, M$; $d_{m+1} = 1$ for $(a+b) = m+1, \ldots, M$; and so on up to $d_M = 1$ for $(a+b) = M$ and impose the Deaton–Paxson normalization on this different set of dummies.[9] This alternative normalization of the time effects does not affect the fit of the model and the age–wealth profile estimated in Sections 4 and 5.

Since equation (3) imposes a life-cycle interpretation to the data (with a rather limited role for uncertainty), it is a useful framework to check if the estimated age–wealth profile and the cohort effect are consistent with the theory or violate some of its basic requirements. However another set of identification assumptions can be used. Homothetic preferences imply that consumption is proportional to lifetime resources. This implies that cohort effects in wealth are the same as in consumption, and one can subtract the log of consumption from both sides of equation (3), drop the cohort polynomial and regress the log of the wealth-consumption ratio on the age polynomial and unrestricted time effects:

$$\ln(W_{a,b}/C_{a,b}) = s(a) + X\beta + d_{a+b}\delta + v \tag{4}$$

where $s(a)$ is the shape of the wealth-consumption ratio. Having dropped $h(b)$, equation (4) allows unrestricted time dummies to affect the wealth-consumption ratio. Comparison of equations (3) and (4) highlights advantages and costs of the two specifications. If time effects are constrained, as in equation (3), one can estimate cohort effects in wealth. If time effects are left unconstrained, one must assume that cohort effects in wealth and consumption are equal, as in equation (4), but one can test if the unconstrained time dummies sum to zero and are orthogonal to a trend. If one augments equation (4) with cohort effects, but restricts the time effects, one can test if the cohort effect in wealth equals that in consumption. Finally, if the life-cycle model provides a good description of the data, so that consumption is smooth through life, the shape of the estimated age–wealth profiles obtained under the different identification assumptions should be similar, i.e. $g(a) \cong s(a)$.

4. The Data

The primary purpose of the Bank of Italy Survey of Household Income and Wealth (SHIW) is to collect detailed data on demographics, households' consumption, income and balance sheets (Brandolini and Cannari, 1994). The data set used in this study includes six independent cross-sections of Italian households (1984, 1986, 1987, 1989, 1991 and 1993), a total of 44,792 observations. Net worth is the sum of household's financial assets and net real assets. The major weakness

[8] Attanasio and Weber (1994) test this implication of the model for the U.K. and find that the consumption of the young increased in anticipation of the economic boom of the late 1980s.

[9] I thank the referee for pointing out this alternative normalization of the time effects.

of the SHIW is that data on financial wealth for some years are not available. Scaling up the flow of financial income allows one to impute financial wealth. Missing data for financial income are imputed using the estimated coefficients of a regression that takes into account the probability of non-responses. The imputation scheme, which is described in an Appendix available upon request, does not necessarily solve the missing value problem if the modeling of non-responses is subject to systematic bias arising from the omission of relevant variables.[10] The regressions presented in Sections 5 and 6 were therefore replicated dropping observations with missing financial income; since the differences with the full-sample coefficients are of minor importance (in particular for decumulation rates and cohort effects, which are the main focus here), these results are not reported for brevity.

Households headed by persons born before 1910 and after 1959 are excluded. These exclusions are motivated by concern over two sources of potential sample bias. The first arises because survival probabilities may be positively correlated with wealth, implying that rich households are over-represented in the oldest cohorts. This correlation implies that one may find a low rate of decumulation after retirement simply because the poor tend to disappear from the sample earlier than the rich. We thus drop 1,758 households born before 1910 who would be over 83 years old in 1993. Any residual correlation between wealth and mortality and between wealth or household headship should not seriously affect the estimates. However even if it does, the residual presence of old households will bias the coefficients against the life-cycle model, rather than in its favor, due to the correlation between wealth and mortality and the finding that the poor tend to decumulate more rapidly than the rich (see below).

The second source of potential bias is a correlation between wealth and young household heads peculiar to our sample. In Italy young working adults with independent living arrangements tend to be wealthier than average, because most young working adults live with their parents.[11] For instance, in 1989 the fraction of income recipients below 30 years of age was 19.8 percent, while the fraction of household heads in that age bracket was a tiny 7.6 percent. Households whose head was born after 1959 (who would be less than 24 years old in 1984) are therefore excluded (2,435 households).

Also excluded are households with missing disposable income, disposable income or consumption less than 1 million lire or missing information for the variables used in the estimation (146 households). Since the dependent variable is the logarithm of net worth, we exclude households with negative or zero wealth. The sample truncation may lead to biased estimates of the age–wealth profile. However, Italian households borrow very little, so the number of these

[10]The Appendix also compares the SHIW measures of wealth and income with the aggregate national accounts and describes the imputation method for financial wealth. I have also tried to match the microeconomic data with the aggregate data blowing up the survey data by applying separate scaling factors for real and financial assets so that aggregate values are increased to the balance sheet figures. The estimated age–wealth profile that I obtain is similar to the one plotted in Figure 2.

[11]The reasons for such behavior includes mortgage market imperfections, which prevent young households from borrowing, and imperfections in the rental market for housing.

TABLE 1

CELL SIZE AND WEALTH STATISTICS BY YEAR OF BIRTH OF THE HOUSEHOLD HEAD

Year of Birth (1)	Age of Cohort in 1984 (2)	Age of Cohort in 1993 (3)	Average Cell Size (4)	Mean (5)	25th Percentile (6)	Median (7)	75th Percentile (8)
1910–14	70–74	79–83	372	114.3	11.6	39.6	113.6
1915–19	65–69	74–78	426	142.0	16.9	51.6	142.5
1920–24	60–64	69–73	689	169.2	21.8	62.2	166.7
1925–29	55–59	64–68	762	196.7	26.5	76.5	199.7
1930–34	50–54	59–63	806	203.5	26.6	75.1	204.6
1935–39	45–49	54–58	892	209.7	27.0	79.9	212.6
1940–44	40–44	49–53	809	212.1	23.8	74.3	215.2
1945–49	35–39	44–48	879	172.6	17.8	61.1	173.9
1950–54	30–34	39–43	706	146.1	14.5	50.6	155.1
1955–59	25–29	34–38	629	126.9	8.7	39.2	126.8
All cohorts	25–74	34–83	738	176.4	20.6	63.9	178.4

Note: Wealth is expressed in millions of 1991 lire. The averages and percentiles in columns (5)–(8) are computed using sample weights.

households is tiny (530 observations, or 1.2 percent of the sample).[12] The final sample covers 39,939 households.

The first four columns in Table 1 report the year-of-birth intervals, the range over which the age of each cohort is observed in 1984 and in 1993, and the average cell size in each survey. Columns (5) to (8) display the sample means and the 25th, 50th and 75th percentiles of wealth for the 10 cohorts. All cohorts hold substantial amounts of wealth: even the 25th wealth percentile has as much as 20 million lire. The concave pattern of wealth neither supports or contradicts the life-cycle model, because it reflects a mixture of cohort, age and time effects. The wealth distribution is highly skewed: average wealth is much closer to the 75th percentile than to the median. The wealth skewness and the likely presence of influential values suggest that OLS regressions may not adequately characterize the age–wealth profile. We therefore supplement the statistical analysis by quantile regressions.

Figure 1 offers fundamental insights into the process of wealth accumulation, plotting the average wealth by age of 10 cohorts: cohort 1 includes all households whose head was born between 1955 and 1959, cohort 2 those born between 1950 and 1954, and so on up to cohort 10, those born between 1910 and 1914. Each cohort is observed at six different times, one for each cross-section. The figure shows that the young and the middle-aged do most wealth accumulation. The wealth profile is relatively flat between age 60 and 70, and then declines for all cohorts, with the exception of cohort 9. As will be seen in the next section, the econometric results broadly corroborate the descriptive analysis. Common shocks also clearly affect the data in Figure 1. For instance, several cohorts show that wealth declines between 1984 and 1986 (the first segment of each broken line).

[12]The potential selection bias in other studies that use the logarithm of wealth as the dependent variable is more severe. For instance, King and Dicks-Mireaux (1982) exclude 18 percent of the sample with net worth lower than $2,500.

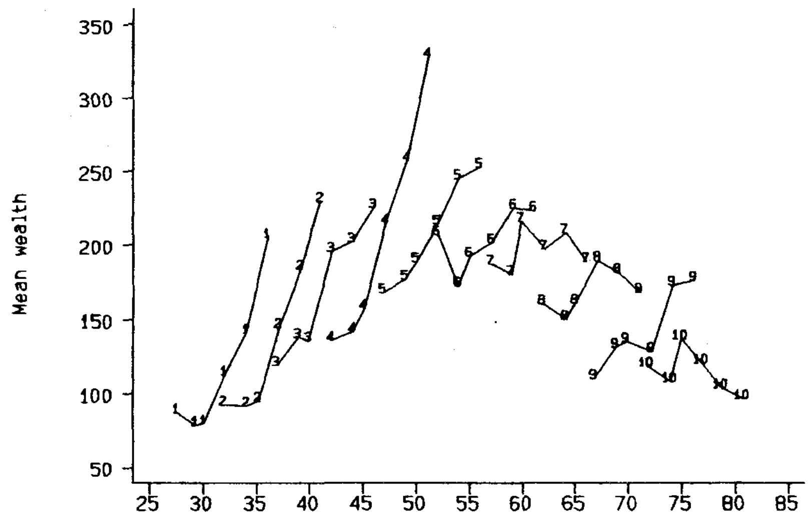

Figure 1. Mean Wealth by Age and Cohort

Either cohort or age does not account for these features of the data (reflecting either measurement errors, macroeconomic shocks or preference shifts). Time effects are therefore potentially important.

5. Regressions for Wealth Levels

In the basic specification of Table 2 I drop the X variables from equation (3) and regress the log of wealth on two fifth-order age and year-of-birth polynomials and a set of restricted time dummies.[13] The results are most clearly illustrated by plotting the implied age–wealth profile in Figure 2. In order to highlight the importance of controlling for cohort effects, I also plot the profile obtained by a regression that treats all observations in the six cross-sections as a large sample (i.e. dropping the time dummies and the cohort polynomial). The cross-sectional profile shows the typical humped shape predicted by the life-cycle hypothesis: wealth peaks at age 55 and declines steadily thereafter. At age 80 wealth has the same level as for the 30-years old. However, cohort effects contaminate this shape. In fact, the cohort-adjusted profile peaks about 10 years later than the cross-sectional profile. The decumulation of wealth during retirement is less but still substantial.

The effect of the cohort-polynomial is plotted against year-of-birth in Figure 3. Recall that the shape of the age–wealth profile (the age polynomial) is constrained to be the same for all generations, which thus differ only in the level of the profile. On average, wealth grows by 3.43 percent per year for the generations born between 1910 and 1950. The cohort effect then declines for generations born

[13]Results using age or cohort dummies are similar to those obtained with the age and cohort polynomial. Furthermore, the age profile in regressions with interacted time and age dummies do not differ much from those displayed in Table 2 and are not reported.

TABLE 2

WEALTH REGRESSIONS

	Dependent variable: ln W		Dependent variable: ln (W/C)	
	(1)	(2)	(3)	(4)
Age	0.504	0.762	0.215	0.310
	(0.044)	(0.041)	(0.016)	(0.016)
Age^2	–0.226	–0.166	–0.094	–0.067
	(0.029)	(0.027)	(0.015)	(0.015)
Age^3	0.018	0.019	0.017	0.022
	(0.015)	(0.014)	(0.009)	(0.009)
Age^4	0.011	0.007	–0.001	–0.002
	(0.005)	(0.005)	(0.004)	(0.003)
Age^5	–0.004	–0.004	–0.003	–0.003
	(0.002)	(0.002)	(0.001)	(0.001)
Female head		–0.470		–0.139
		(0.020)		(0.019)
Resident in the North		–0.013		–0.056
		(0.019)		(0.017)
Resident in the South		–0.263		–0.129
		(0.019)		(0.018)
Self-employed head		0.945		0.787
		(0.018)		(0.016)
Education (in years)		0.012		0.055
		(0.002)		(0.001)
R^2	0.050	0.184	0.051	0.114
Implied pattern of wealth accumulation				
25 < age < 40	9.09	10.46	5.63	5.77
40 < age < 60	5.19	7.76	2.30	3.28
60 < age < 70	0.08	3.75	–0.37	1.31
70 < age < 83	–4.43	–2.81	–6.84	–6.34

Notes: The standard errors reported in parentheses are computed using White's heteroskedasticity consistent covariance matrix estimator. Regressions (1) and (2) also include a fifth-order polynomial in year-of-birth and a set of time dummies, constrained to sum to zero and to be orthogonal to a time trend. Regressions (3) and (4) also include a set of unrestricted time dummies. Excluded attributes in columns (2) and (4) are: households headed by a male, households living in the Center, households whose head is not self-employed. The age variable is expressed in deviations from 50. The number of observations is 39,923.

after 1950, possibly reflecting the post-1973 slowdown in productivity (2.24 percent per year).[14]

It is useful to compare the estimated cohort effect with an out-of-sample index of productivity growth. As in Alessie, Lusardi, and Kapteyn (1996), I replace the cohort-polynomial with an index of generation-specific productivity growth in the wealth regression and plot the effect of productivity growth in Figure 3 (the coefficient of productivity growth is 0.76 with a standard error of

[14]It is hard to reconcile the pattern of the cohort effect with preference-based explanations. For instance, suppose that the shape in Figure 3 arises from different rate of time preferences or attitudes towards bequests. For this to be true, however, younger generations should be much thriftier and more altruistic than older ones, and that there is absoultely no evidence that this is the case.

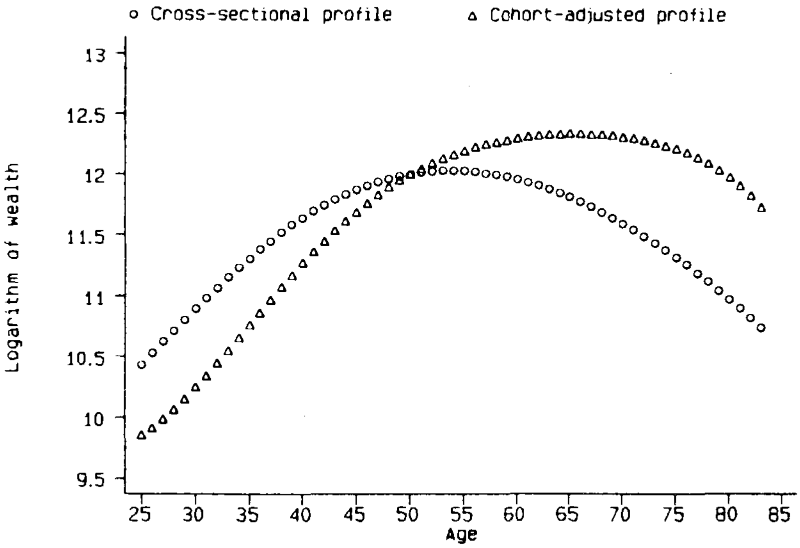

Figure 2. Age-wealth Profile

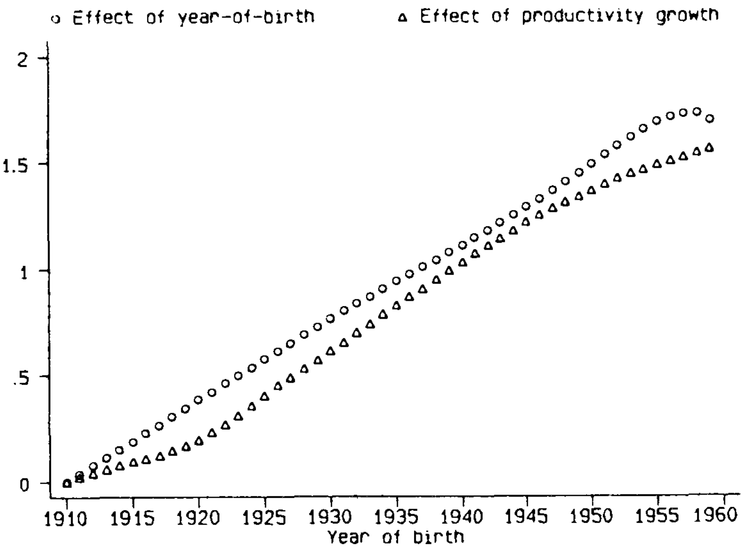

Figure 3. Cohort Effect

0.059, the age profile is unaffected).[15] The high correlation coefficient between the two lines in Figure 3 (0.98) is highly suggestive of the hypothesis that the amplitude of the cohort effect is explained mainly by productivity growth, rather than by preferences or attitudes towards bequests.[16] If one adds the index of productivity growth to the regression, the hypothesis that the coefficients of the

[15]Each generation is assumed to enter the labor market at age 25 (for instance, individuals born in 1959 are assumed to enter the labor market in 1984). Average productivity growth is estimated on the basis of a 9-year moving average of per capita GDP (highly volatile and non-reliable data for 1944–46 are excluded). The source for GDP is Rossi, Sorgato, and Toniolo (1992).

[16]In the stylized version of the life-cycle model predictable variations in generation-specific productivity growth are the only sources of variation in the cohort effect.

cohort-polynomial are jointly equal to zero yields an F-test statistic of 2.59, which is statistically different from zero at the 5 percent level (but not at the 1 percent level). This should perhaps not be surprising given the flexible parametrization of the effect of year-of-birth and the rough index used to measure the growth in generation-specific resources.

In column (2) of Table 2 I add to the basic specification a set of demographic variables to control for the determinants of households' resources: sex of the household head, two regional dummies, a dummy for self-employment, and schooling measured in years. Even though these variables are meant to measure H_0, inspection of the data suggests that each of these variables varies over the life cycle, and cohort effects (particularly education) also effect some of them. For instance, "female head" increases in old age, given the longer life expectancy for women; "self-employed" declines sharply after retirement, due to the lower participation rate of the elderly; the "North" dummy is higher at younger ages, owing to the positive correlation between age, wealth and headship discussed in Section 3; for "South" the opposite is true.

It is therefore difficult to identify the pure age–wealth profile in the augmented regressions because the age effect is now captured not only by the age polynomial but also by some of the additional control variables.[17] On the other hand, these variables may help correcting for systematic sample bias. For instance, if wealth and mortality are correlated, education, which might be correlated with both, can attenuate the bias arising from the changing composition of the sample (Attanasio, 1994).

Column (2) of Table 2 reports the results of the augmented specification. Residence in the South and less schooling signal lower permanent income and reduce considerably asset accumulation.[18] The dummy for self-employment is large and positive. To the extent that self-employment proxies for income volatility, the results confirm the theoretical work of Caballero (1991) showing that the age–wealth profile shifts upwards in the presence of uninsurable income risk.

Both the unadjusted and the cohort-adjusted age profiles plotted in Figure 4 show less wealth decumulation in old age than the profiles in Figure 2. The demographic variables affect the size but not the general pattern of the cohort polynomial, which again indicates that the log of wealth increases almost linearly with year of birth.

So far I constrain the age–wealth profile to be the same for all population groups. I thus divide the sample according to the educational attainment of the head, a good proxy for initial resources. Since the earnings profile, family composition, retirement age or social security replacement rates differ across education groups, the age–wealth profiles may also differ. However, it is still an implication of the theory that there should be wealth decumulation in old age for each education group.

[17]This problem is even more serious for family-related variables—such as number of children or number of adults, both hump-shaped over the life-cycle—which are therefore omitted from the estimation. The coefficients of the variable "number of adults" is positive and significant if the variable is included in the estimation. As expected, the estimated age profile is flatter; however, the decumulation rates during retirement are not affected. Instead, the coefficient of the variable "number of children" is not significantly different from zero.

[18]An alternative explanation is greater under-reporting of wealth in the South.

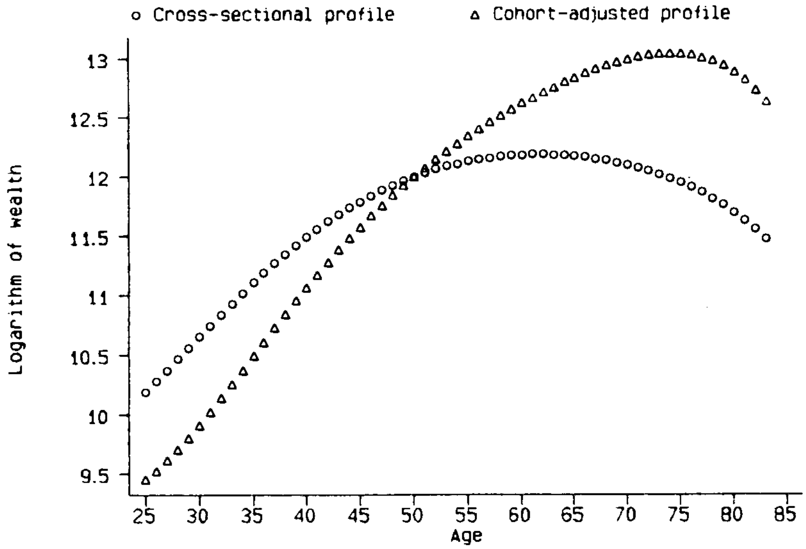

Figure 4. Demographic Controls

For brevity, I do not report the estimated coefficients and plot in Figure 5 the three cohort-adjusted age–wealth profiles. The profiles of households with elementary education or junior high school (8 years of schooling or less) and high school education (9 to 13 years) are similar to the full sample estimates. These groups represent the bulk of the population (70 percent and 22 percent, respectively). However households with college education (more than 13 years,

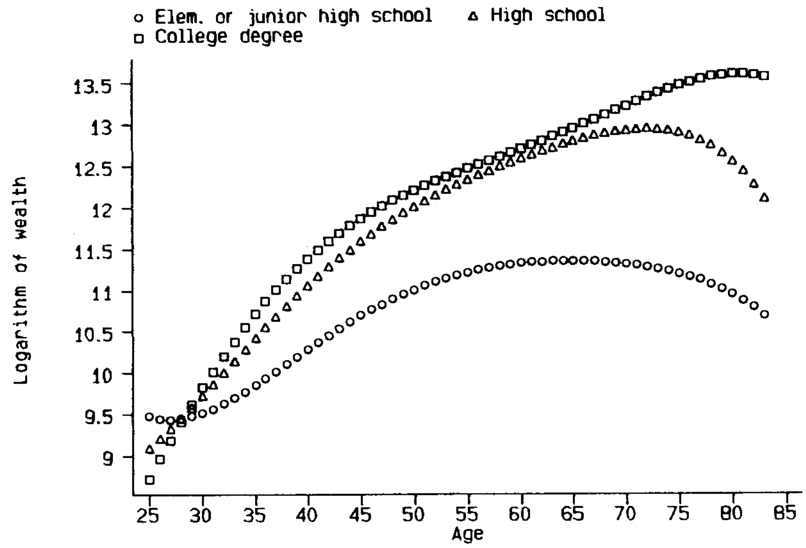

Figure 5. Age–Wealth Profiles by Education

representing approximately 8 percent of the population) do not exhibit any tendency to decumulate wealth after retirement.[19]

There are several possible explanations for the different pattern of wealth accumulation of the college-educated, not mutually exclusive. The rate of time preference of the college-educated may be lower than that of the non-college-educated. Non-homothetic utility functions with consumption floors also predict that wealth is an increasing function of lifetime earnings. Another possibility is that everyone has the same concern for their heirs' consumption but that the bequest motive is not operative for the poor. Since negative bequests are not allowed, poor households may choose a corner solution for bequests. Flemming (1979) shows that if earning power regresses towards the mean, those with high wealth relative to earning power and those at the top of the ability range will plan to leave larger bequests. Even if earnings are positively correlated across generations, the better educated should receive greater bequests, also inducing skewness in the wealth distribution. Given the reduced form of the estimates, and the lack of information on the resources or ability of the potential heirs, it is not possible to distinguish between these possibilities.

I next explore other dimensions of the conditional distribution of wealth by quantile regressions. The age–wealth profiles implied by the 25th, 50th and 75th

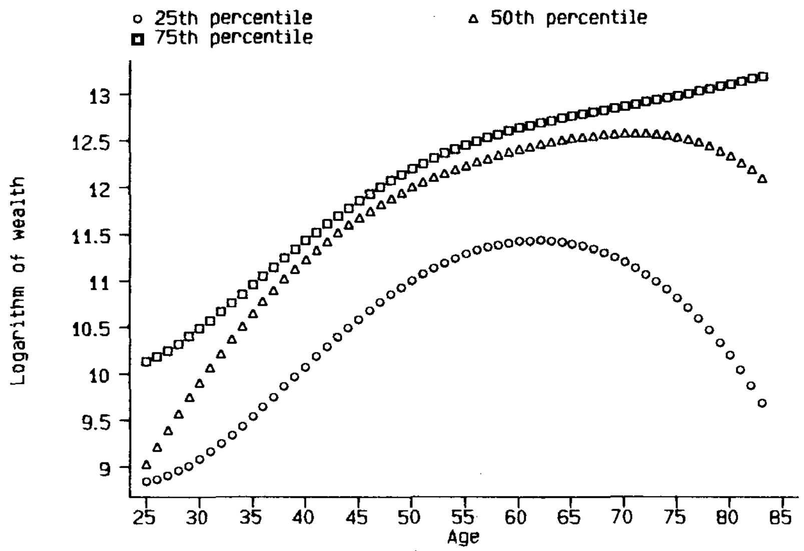

Figure 6. Quantile Regressions

quantile regressions are displayed in Figure 6. The figure suggests considerable heterogeneity in the data. For the bottom and medium part of the wealth distribution (up to the 75th quantile) there is substantial decumulation in old age. Wealthy households, however, do not seem to conform to the model: the wealth profile shows no tendency to decline with age after retirement for the 75th wealth percentile. These results are in agreement with those by educational groups, because of the positive correlation between schooling and the level of wealth. For

[19]Including demographic variables does not change the general picture.

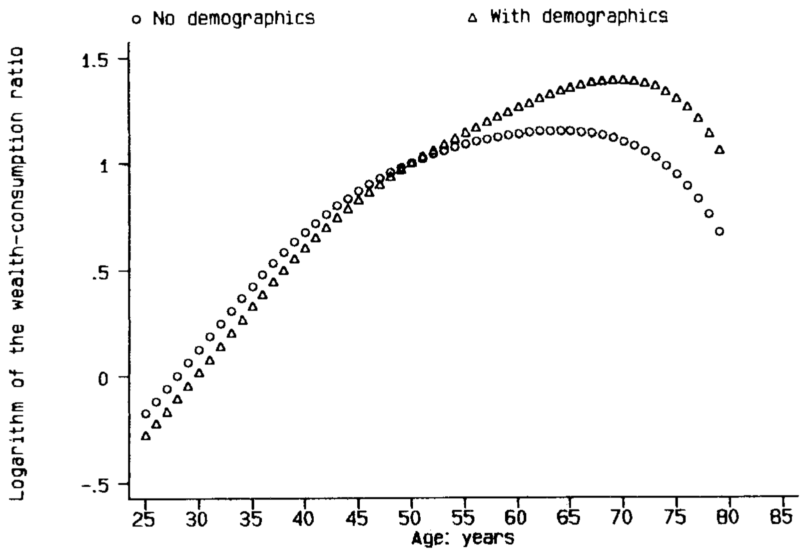

Figure 7. Age–Profile of Wealth/Consumption

instance, average wealth in the group with medium education (8 to 13 years) is 205 million lire, a value that is between the 75th and 90th wealth percentile; in the group with higher education (more than 13 years) average assets exceed the 85th wealth percentile.

6. Regressions for the Wealth-Consumption Ratio

Table 2 also reports the coefficients of the regressions that use the logarithm of the ratio between wealth and total consumption expenditure as the dependent variable (equation (4)).[20] As in the previous section, I report the coefficients of the basic model (column 3), and of the model augmented by demographic variables (column 4). Each regression contains a full set of unrestricted time dummies. Figure 7 plots the growth rates of wealth implied by the age effects estimated in Table 2 (results using the augmented specification are similar). Since in old age wealth declines at a much larger speed than consumption, there is substantial wealth decumulation after age 70, in agreement with the results obtained for wealth levels.

When I add the cohort polynomial to the specification in Table 2 (with restricted time dummies), I cannot reject the hypothesis of absence of cohort effects in the wealth-consumption ratio (the F-test statistic is 2.12, which given 5 and 39,908 degrees of freedom is not significant at the 5 percent level). A proper test of the life-cycle hypothesis, however, is to confront a model with unrestricted age, cohort and time effects against an alternative including only the age polynomial and restricted time dummies (constrained to sum to zero and to be orthogonal to a time trend, as explained in Section 3). In the basic specification, this

[20]In principle, one should use the sum of non-durable consumption and the consumption of the durable stock. However, the latter is not available in 1986.

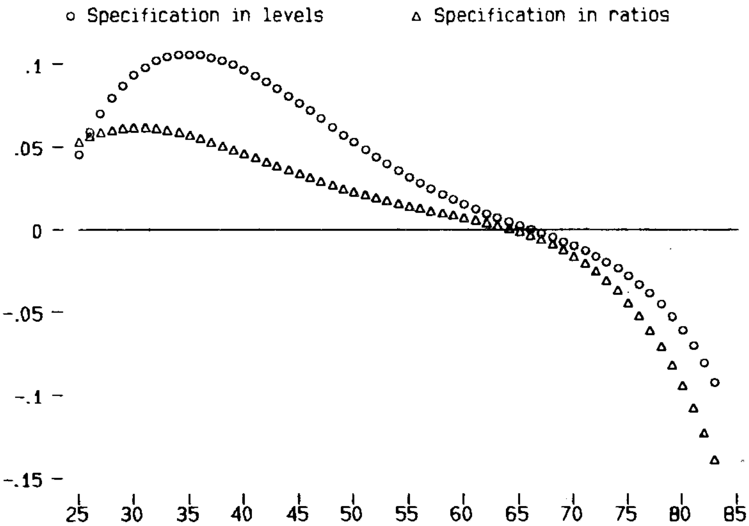

Figure 8. Growth Rate of Wealth

test yields an F-statistic of 5.39, which is significantly different from zero at the 1 percent level given 7 and 39,908 degrees of freedom. Thus, the restrictions implied by the life-cycle model are rejected on statistical grounds. However, this test is very demanding. Given the large sample size, the flexible parametrization of cohort effects and the particular form of the life-cycle model that is fully consistent with both sets of identification assumptions, one would perhaps be surprised not to reject the null hypothesis.

It is useful at this stage to compare the age–wealth profiles obtained under the two identifying assumptions, i.e. unrestricted cohort effects and constrained time effects in the regressions for $\ln W$ and equality of cohort effects in consumption and wealth and unrestricted time effects in the regressions for $\ln (W/C)$. The bottom panel of Table 2 displays the average annual rates of wealth accumulation for younger households (age 25–40), the middle-aged (41–60), the old (61–70) and the very old (age 71–83) implied by the regressions of Table 2. The rates of growth of wealth are also plotted in Figure 8. In the case of the regressions for $\ln (W/C)$ the numbers can be interpreted as average wealth growth rates only if one makes the strong assumption that consumption is constant over the life-cycle, so that the $g(a)$ polynomial in equation (3) equals the $s(a)$ polynomial in equation (4).

The main difference between the two profiles is in the earlier part of the life cycle. In fact, the regressions for the log of wealth predict higher accumulation for younger households, especially up to age 40. However the similarity between the two profiles are more striking, also considering that the estimated age-profile for $\ln (W/C)$ truly reflects the growth of the wealth-consumption ratio, rather than of wealth alone: both indicate dissaving after 60, and substantial decumulation rates between age 70 and 83 (an average of 6 percent per year).[21] If demographic variables are introduced in the regressions, the growth rate of wealth is

[21]The median decumulation rates do not differ much from the averages and are not reported for brevity.

uniformly higher for the level specification. However even in this case, most of the difference between age effects is for younger households. In sum, the equality of the two specifications is rejected on statistical grounds, but both models explain wealth accumulation reasonably well, and the predictions of the alternative identification assumptions for the age–wealth profile are similar. Indeed, the correlation coefficient between the growth rates of wealth in Figure 8 is 95.1 percent (97.1 percent including demographic controls).

7. Summary and Implications

The pattern of wealth accumulation of Italian households conforms reasonably well with the predictions of the life-cycle model. Most wealth accumulation is done by the young and the middle-aged; the wealth profile flattens between age 60 and 70; the average annual decumulation rates in old age are between 3 and 6 percent, as predicted by many simulations of life-cycle models without bequest motives.[22] Even if rejected by a statistical test, the estimated age–wealth profile is quite similar under two very different identification assumptions, i.e. that all trends in wealth are explained either by age or cohort effects, or that cohort effects in wealth and consumption are equal.

The cohort effect in the regressions for the (log) of wealth increases almost linearly with year of birth for the generations born between 1910 and 1950; for younger generations the cohort effect declines considerably. This is consistent with the hypothesis that the cohort polynominal captures the effect of generation-specific productivity growth. If the young save and the old dissave (as in Figure 8), and if the differences in the resources that belong to each generation reflect cohort-specific productivity growth, rather than attitudes towards bequests, the aggregation mechanism of the life-cycle model leads to a positive relation between growth and saving. However, it would be premature to draw such a strong conclusion from the results.

First of all, a relatively small number of wealthy households and households with higher educational attainment show no tendency to decumulate wealth during retirement. This implies that one cannot predict the aggregate saving rate from the behavior of the average household. Even if the absolute number of wealthy households is not large, the rich control a substantial share of aggregate wealth. For instance, 8 percent of college-educated households controls 18 percent of aggregate wealth, and the richest 5 percent in the sample controls over 30 percent. Understanding the behavior of these households becomes crucial for an understanding of the saving-growth link.

Second, decumulation during retirement implies absence of a bequest motive only if gifts are not a major component of households' wealth. If most transfers occur *inter vivos* one would also observe wealth decumulation in old age, but the implication for the aggregate saving rate would be rather different than the predictions of the life-cycle model. Available empirical evidence for the Italian

[22] These rates also broadly confirm the findings with panel data for the U.S. by Diamond and Hausman (1984) and Hurd (1987) and by Brugiavini (1987) with the 1984 SHIW.

economy suggests that this is not the case, but more evidence on the exact timing and amount of *inter vivos* transfers is clearly needed.

Finally, annuities have a potentially strong impact on asset accumulation. If pension benefits are considered as part of dissaving during retirement, decumulation of total wealth in old age would be even greater than that found in the data. However measuring this impact is not easy, because it requires an estimate of pension wealth (including social security) and of the degree of substitution between private and pension wealth.

References

Alessie, R., A. Lusardi, and A. Kapteyn, Saving and Wealth Accumulation of Different Cohorts: A Panel Analysis, Tilburg University, mimeo, 1996.

Attanasio, O., Personal Saving in the United States, in Poterba, J. (ed.),*International Comparison of Personal Saving*, University of Chicago Press, Chicago, 1994.

Attanasio, O. and G. Weber, The U.K. Consumption Boom of the Late 1980s: Aggregate Implications of Microeconomic Evidence, *The Economic Journal*, 104, 1269–302, 1994.

Barro, R. E., Are Government Bonds Net Wealth?, *Journal of Political Economy*, 81, 1095–117, 1974.

Brandolini, A. and L. Cannari, Methodological Appendix: The Bank of Italy's Survey of Household Income and Wealth, in Ando, A., Guiso, L., and Visco, I. (eds.),*Saving and the Accumulation of Wealth. Essays on Italian Household and Government Saving Behavior*, Cambridge University Press, Cambridge, 1994.

Brugiavini, A., Empirical Evidence on Wealth Accumulation and the Effects of Pension Wealth: An Application to Italian Cross-Section Data, Financial Markets Group, Discussion Paper no. 20, London School of Economics, 1987.

Caballero, R., Earnings Uncertainty and Aggregate Wealth Accumulation, *American Economic Review*, 76, 676–91, 1991.

Davies, J. D., Uncertain Lifetime, Consumption and Dissaving in Retirement, *Journal of Political Economy*, 89, 561-77, 1981.

Deaton, A., Saving and Growth, in K. Schmidt-Hebbel and K. Serven (eds.), *The Economics of Saving and Growth*, Cambridge University Press, Cambridge, 1999.

Deaton, A. and C. H. Paxson, Saving, Aging and Growth in Taiwan, in Wise, D. (ed.), *Studies in the Economics of Aging*, University of Chicago Press, Chicago, 1994.

———, The Effects of Economic and Population Growth on National Saving and Inequality, *Demography*, 34, 97–114, 1997.

Diamond, P. A. and J. A. Hausman, Individual Retirement and Savings Behavior, *Journal of Public Economics*, 23, 81-114, 1984.

Flemming, J. S., The Effects of Earnings Inequality, Imperfect Capital Markets, and Dynastic Altruism on the Distribution of Wealth in the Life-Cycle Model, *Economica*, 46, 363-80, 1979.

Guiso, L. and T. Jappelli, Private Transfers, Borrowing Constraints and the Timing of Home Ownership, CEPR Discussion Paper no. 2050, 1998

Hubbard, G. R., J. Skinner, and S. P. Zeldes, The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving, *Carnegie-Rochester Conference Series on Public Policy*, 40, 59-126, 1994.

Hurd, M. D., Savings of the Elderly and Desired Bequests, *American Economic Review*, 77, 298-312, 1987.

———, Mortality Risk and Bequests, *Econometrica*, 57, 779-813, 1989.

King, M. and L. Dicks-Mireaux, Asset Holdings and the Life-Cycle, *Economic Journal*, 92, 247-67, 1982.

Masson, A., A Cohort Analysis of Wealth–age Profiles Generated by a Simulation Model in France (1949–75), *Economic Journal*, 96, 173-90, 1986.

Mirer, T. W., The Wealth-Age Relation Among the Aged, *American Economic Review*, 69, 435–43, 1975.

Modigliani, F., Life Cycle, Individual Thrift, and the Wealth of Nations, *American Economic Review*, 76, 297-313, 1986.

Rossi, N., A. Sorgato, and G. Toniolo, Italian Historical Statistics, University of Venice, Discussion Paper no. 92–18, 1992.

Shorrocks, A. F., The Age–Wealth relationship: A Cross-Section and Cohort Analysis, *Review of Economics and Statistics*, 57, 155-63, 1975.