

Copyright ? 1994. University of Chicago Press. All rights reserved. May not be reproduced in any form without permission from the publisher, except fair uses permitted under U.S. or applicable copyright law.

---

# Studies in the Economics of Aging

Edited by **David A. Wise**



The University of Chicago Press

*Chicago and London*

---

DAVID A. WISE is the John F. Stambaugh Professor of Political Economy at the John F. Kennedy School of Government, Harvard University, and the area director for Health and Retirement Programs at the National Bureau of Economic Research.

The University of Chicago Press, Chicago 60637  
The University of Chicago Press, Ltd., London

© 1994 by the National Bureau of Economic Research  
All rights reserved. Published 1994  
Printed in the United States of America  
03 02 01 00 99 98 97 96 95 94 1 2 3 4 5  
ISBN: 0-226-90294-3 (cloth)

Library of Congress Cataloging-in-Publication Data

Studies in the economics of aging / edited by David A. Wise.  
p. cm.—(A National Bureau of Economic Research project report)

Includes bibliographical references (p. ) and index.  
1. Old age—Economic aspects—Congresses. 2. Aging—Economic aspects—Congresses. 3. Retirement—Economic aspects—Congresses.

I. Wise, David A. II. Series.

HQ1061.S844 1994

305.26—dc20

94-16196  
CIP

⊗ The paper used in this publication meets the minimum requirements of the American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

---

## 9 Saving, Growth, and Aging in Taiwan

Angus S. Deaton and Christina H. Paxson

This paper examines issues of life-cycle saving, growth, and aging in Taiwan. We are mainly concerned with standard issues of life-cycle saving and their implications for the living standards of the elderly. We investigate whether saving appears to be motivated by life-cycle factors, how income growth has affected the profiles of income, consumption, and saving, and how changes in the demographic structure of Taiwan have influenced saving behavior. We use data from 15 consecutive household income and expenditure surveys, from 1976 through 1990. Although there is no panel element to these data, the large number of households surveyed (approximately 14,000 per year) and the large number of cross sections allows us to track cohorts of people through time and to observe the evolution of their levels of income, consumption, and saving.

Taiwan provides an excellent laboratory in which to study the determinants of household saving. First, as table 9.1 shows, the last four decades have seen both very rapid growth and very high saving rates. The annual growth of GNP per capita averaged 4.3 percent in the 1950s, 7.0 percent in the 1970s, and 6.7 percent in the 1980s, and although the oil price shocks of the early 1970s and 1980s were associated with periods of relatively slow growth, there is no evidence that per capita growth rates are beginning to slow. Indeed, growth rates in the latter half of the 1980s exceeded any in the previous three decades. The high saving rates are also documented in table 9.1. The national accounts data show that gross national saving as a fraction of GNP has increased over the past three decades, from an average of 18.7 percent in 1961–65 to over 32

Angus S. Deaton is the William Church Osborn Professor of Public Affairs and professor of economics and international affairs at Princeton University and a research associate of the National Bureau of Economic Research. Christina H. Paxson is associate professor of economics and public affairs at Princeton University and a research associate of the National Bureau of Economic Research.

The authors are grateful to Jonathan Skinner for useful comments, to Alec Levenson for assistance with the data, and to David Card and Gregory Chow for helpful discussions.

**Table 9.1**                      **Statistics on Income, Growth, Saving, and Population**

Year	GNP per Capita (base = 1986)	Growth in GNP per Capita	Income per Household (base = 1986)		Gross Saving as a Percentage of GNP	Household S Rate National Accounts
			National Accounts	Survey		
1961	25,408	-	-	-	18.4	-
1965	33,386	6.83	-	-	20.7	-
1970	47,710	7.14	143,088	-	25.7	-
1975	63,428	5.69	193,062	-	26.7	20.0
1978	84,316	9.49	250,100	248,854	34.4	20.7
1979	89,605	6.08	271,363	279,332	33.4	19.8
1980	92,098	2.74	288,934	288,173	32.3	17.9
1981	94,374	2.44	294,759	286,102	31.3	20.2
1982	96,944	2.69	294,322	287,190	30.1	20.1
1983	104,130	7.15	310,414	308,659	32.1	21.4
1984	114,903	9.84	326,692	331,271	33.8	22.9
1985	119,581	3.99	331,193	338,155	33.6	23.9
1986	137,992	14.32	341,728	359,387	38.5	28.5
1987	154,838	11.52	364,700	380,111	38.5	28.3
1988	164,229	5.89	404,138	423,080	34.5	23.6
1989	174,407	6.01	443,866	465,652	30.8	19.5
1990	180,053	3.19	478,296	507,242	29.2	20.0

*Notes:* All numbers are from National Accounts statistics (from Taiwanese Statistical Office) unless otherwise specified. All money figures are in N.T. dollars, base year = 1986.

percent in 1976–90. Household saving represents a large fraction of national saving, and the national accounts data show that roughly 50 percent of national saving is done by households, even if we exclude the savings of private corporations. What is more, the household survey data that we use yield savings figures that are of comparable magnitude to those in the national accounts. This situation stands in sharp contrast to that in many developing countries, where survey data yield savings estimates that are much lower than those in the national accounts.

The second feature of Taiwan that is relevant to life-cycle saving is its rapid transition over the last several decades from high to low population growth; see again table 9.1. In the 1950s, Taiwan's rate of population growth averaged 4.13 percent per year, reflecting the influx of immigrants from mainland China and then a postwar baby boom. The rate declined slowly through the early and middle 1960s and fell sharply in the late 1960s and early 1970s. The overall decline in the rate of population growth masks two offsetting factors, both of which are important to our topic. Life expectancy has increased, and there has been a sharp drop in fertility. In 1961, a 60-year-old male had a life expectation of 13.92 years. By 1990, this expectation had increased to 17.93 years (Republic of China 1991). For females, comparable numbers are 17.85 and 20.51 years. The decline in population growth rates despite these increases reflect the behavior of fertility, which began to fall in the late 1960s. The total fertility rate fell from 6.1 in 1958 to 1.9 in 1985. An important implication of this decline is that Taiwanese who are currently young will, when old, have far fewer adult children to provide potential support than do those who are old now. In fact, the fraction of the elderly who live with their children is already beginning to decline (see Lo 1987), although this may be due more to an increase in wealth that makes independent living possible, rather than to a decline in the number of children who can potentially provide care.

High growth, declining fertility, and increasing life expectancy all have consequences for saving that are predictable by standard life-cycle theory, and it is these predictions that we examine in this paper. We first investigate whether the observed patterns in consumption and saving across different households can be fit into the life-cycle story. In particular, we examine the basic implication of a life-cycle model that savings should (at least eventually) decline with age. We also examine cohort effects in consumption and in saving. Those who are currently young in Taiwan have lifetime earnings that are many times larger than those of their parents and grandparents, and we test whether their consumption profiles are shifted in the appropriate way. We also look for the increased saving rates among the younger cohorts that might be expected in the face of rising life expectancy and falling numbers of children. We find reasonable conformity between the Taiwanese facts and the life-cycle theory; it is certainly possible to interpret the broad features of the data in life-cycle terms.

We also look at the familiar implications of the life-cycle model for the relationship between savings and productivity growth. High income growth

across cohorts should yield high national savings rates, because younger savers have a much higher lifetime wealth level than do older dissavers, a mechanism that links growth to saving at the aggregate, but not at the individual, level. We find strong evidence of a link between growth and saving at the individual level, something that is not predicted by the theory. We also run a battery of standard excess sensitivity tests, which do not provide any evidence against the life-cycle story. Even so, we find a very marked coherence between the age profiles of consumption and income for different cohorts, a coherence that is consistent with the life-cycle but strongly suggests a more direct link between consumption and income. We find it hard to rationalize such a relationship in the Taiwanese context, so that the life-cycle model remains perhaps the most satisfactory account of our evidence.

### 9.1 The Survey Data

The data used in this paper come from a time series of household surveys, the Personal Income Distribution Surveys, collected in the Taiwan area of the Republic of China (ROC), or Taiwan for short, in the 15 calendar years 1976–90. In 1976 and 1977, the sample sizes are a little over 9,000 households, but from 1978 there are over 14,000 households in each survey. The number of persons covered varies from around 50,000 in the first two years to around 75,000 later. New samples are drawn each year, so that it is not possible to track individual households over time; instead, we shall track individual age cohorts through the successive surveys.

The survey design is described in Republic of China (1989). For income and consumption, data are collected both from interviews and from diaries. At the single interview, questions are asked about major items of income and expenditure in the past year, while the diaries, maintained throughout the year and regularly inspected by field workers, keep track of all items of income and expenditure. Only a “small number of households” keep the diaries; these households are also interviewed, and the results of comparing the two methods are “used to check and/or correct results of all interviews in the survey.” The survey is a comprehensive one and collects information on household structure, on socioeconomic characteristics (including industry, occupation, and education), on household expenditures in some detail, and on the incomes by source of the household and of each of its members. All money amounts in this paper are in real N.T. dollars, base year = 1986.

Our impression from the work that we have done and from the results reported below is that the data are of good quality. One of the attractions of working with data from a high saving country is the ability to study saving at the household level. For many developing countries, it is difficult to find much reported saving in household survey data, even in those cases where national accounts data indicate that households as a whole are saving substantial amounts; see, for example, Visaria (1980) for several Asian countries, Paxson


(1992) for Thailand, and Deaton (1992a) for Côte d'Ivoire. The Taiwanese survey data do not have this problem, and the high national saving rates are reflected in the behavior of the individual households. Table 9.1 shows that, according to the national accounts, saving by households has varied around 20 percent of household income since 1975, while the figures from the surveys, although far from identical, are of the same order of magnitude. Given that saving is the difference between two large magnitudes, each estimated with error, given that the coverage of the surveys and the national accounts is different, and given the difficulties that are typically encountered in reconciling household and national accounts, we view the correspondence as remarkably close. For household incomes, which are also shown in table 9.1, the survey data are again close to those from the national accounts, with a maximum discrepancy of 5 percent.

Since we shall be concerned with saving within a life-cycle context, it is useful to look first at the demographic information in the surveys. Figure 9.1 shows the age-sex pyramids for each of the survey years; as with all of the summary statistics presented in the paper, numbers are weighted by the appropriate inflation factors so as to provide estimates of the underlying population. No inflation factors are available for the first two years, 1976 and 1977, for which the results are presented in unweighted form; evidence from the other years suggests that the unweighted data are unlikely to be misleading.

Note first that the surveys evidently fail to capture a substantial number of young males aged 20–25, presumably those in college or engaged in military service. In the last few years, 1988–90, a similar phenomenon is beginning to be evident for young females over a somewhat broader age range. This may again be related to college, and indeed more than half of college students in Taiwan are women; it may also reflect employment patterns among young women. In any case, these people, both male and female, are “missing” as a result of the survey design, and not because of any similar feature in the population.

Another feature of these survey demographics reflects a genuine peculiarity of the Taiwanese population. In all years, there is an excess of middle-aged to older men over women in the same age group; see the sex ratios by age in figure 9.2. In the first year, 1976, the number of males per 100 females is greater than 100, over the range of ages 40–65, reaching a peak of 140 at age 50. This peak ages one year per year, and by 1990, there are somewhat less than 140 65-year-old males for each 65-year-old female. These numbers are a consequence of the cohort of predominantly male “mainlanders” who came to Taiwan in 1949, and an unusually large fraction of whom have never married. These people differ in a number of other respects from the rest of the population; for example, many were soldiers, and a high proportion have government pensions (or pension rights), and so they provide a group whose life-cycle consumption and saving behavior can be expected to differ from that of other Taiwanese.

The changing shapes of the pyramids also shows very clearly the marked



**Fig. 9.1** Age-sex pyramids

Note: Positive values represent males. Negative values represent females.

(and remarkable) aging of the population over the 15-year period of our sample. Figure 9.3 shows, for each year, the fractions of the population older than each age. In 1976, .046 of the population were over 60; in 1990, the fraction was .092. The ratio of the elderly to their surviving children, although still high because of older fertility rates, will fall rapidly over time. For example, Hermalin, Ofstedal, and Li (1991) indicate that women aged 61–65 in 1990 had, on average, 5.1 children ever born in their lifetimes. The number of children ever born is expected to drop to 3.6 for those aged 46–50 in 1990 and to 2.4 for those aged 31–35 in 1990. Although numbers of children ever born is



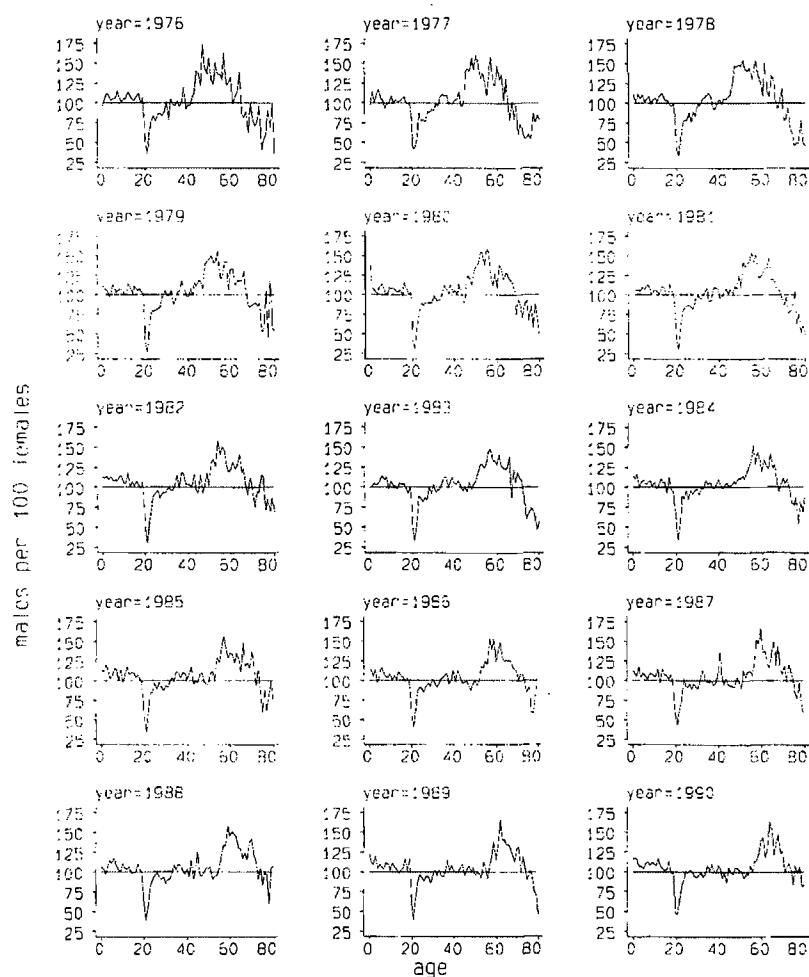


Fig. 9.2 Sex ratios

not the same as children surviving into their parent's old age, the evidence strongly suggests that many of the current generation of young adults cannot expect to avail themselves of the traditional Chinese pattern whereby elderly parents move into their son's home. This idea is supported by survey evidence from Chang (1987), which indicates that a growing fraction of currently young Taiwanese women do not expect to live with their sons when they are old. Absent state intervention in the form of social security (and although social security is being expanded in Taiwan, the coverage and amounts are still very small), the current generation of young workers has a strong incentive to save a larger fraction of their incomes than did their parents. One of our aims in this




Fig. 9.3 Age distribution of the population, by year

paper is to see if we can find evidence of such an increase in life-cycle saving among the young and whether it can plausibly be attributed to the effects of the demographic transition.

## 9.2 Constructing Cohorts

For most of the interesting questions about saving and the life cycle, it is necessary to track individuals over time and to observe the changes in consumption, income, and saving as people age. Although we cannot track individual households in the data, we can track “cohorts” of households, with cohorts defined according to their year of birth. For each year of data, we average the variables of interest by age of individuals or age of household head and then track, not the same individuals, but the sample from the same cohort one year older in the next survey. In this way, we can follow cohort means for 15 years of each age cohort of households or individuals. We do this both at the level of the household, with age defined by the age of the head, and at the level of the person, again using age as the grouping variable. Where we have individual data, for example on earnings, the (sample weighted) cohort means of individuals by age provide unbiased estimates of mean earnings for the underlying population of that age. For other variables, which can be collected only at the household level, matters are a good deal less clear. Averaging by the age of the household head has the inevitable effect of confounding genuine

changes in stable households with changes in both household formation and in headship. For example, since earnings (eventually) decline with age and since older men cease (by definition in these data) to be heads when they cease to be the main earner, households with older heads will be an increasingly selected sample of the population and will therefore display behavior that increasingly disguises actual patterns of household change over time. For example, households with older heads will be those whose head has unusually high earnings compared with others of his or her cohort, or those who have an unusual amount of wealth that enables the household to survive as an independent unit. As a consequence, and if, for argument's sake, savings rates are positively correlated with earnings, or are positively correlated with wealth (households who are predisposed to save are typically richer than those who are not), we may observe that, at high ages, household saving rates increase with the age of the head even though assets are typically being run down in old age.

The alternative, of averaging over individuals by age, although minimizing selectivity, is difficult, because so many variables are reported only at the household level, so that we have to impute individual consumption or incomes based on those of the household to which they belong. Old people who live with their children are therefore attributed some fraction of household income or consumption, an attribution that is essentially arbitrary when there is no information about intrahousehold allocation. These problems are not soluble, and they have to be kept continuously in mind when interpreting the results presented below.

Table 9.2 reports the numbers of households in five selected household cohorts, for each year of the survey data; table 9.3 records the same information for the numbers of people in the "individual" cohorts. To save space, we show numbers for only five cohorts, but we have constructed (and will use) the cohorts at each age, although in some cases it will be necessary to eliminate the youngest and the oldest groups. In total, the data allow us to calculate 1,031 household head age/year pairs for the household data and 1,161 age/year pairs for the individual data; only 75 of these are shown in each table. Cohorts are defined by year of birth, or more conveniently, by age in 1976. Hence, the first column of table 9.2 shows the sample representation of households headed by those born in 1956, who were age 20 in 1976 and 34 in 1990. The increase in numbers, from 26 households in 1976 to 640 households in 1990, shows not that the cohort of those born in 1956 is increasing, nor that the sample is incorrectly drawn, but that 34-year-olds are much more likely to be household heads than 20-year-olds. For the older cohorts, the variation in sample size is much less, and it is these middle years of household headship for which the averages are likely to be most accurate, both because there are more households in each average, and because selection problems are less severe.

Figure 9.4 shows how the cohort grouping can be used to show both the life-cycle pattern of family formation and the cohort effects of falling fertility. The age of the head of the household is plotted on the horizontal axis, while the

Table 9.2 Numbers of Households in Selected Household Cohorts (indicated by age of head in 1976)


Year	Cohort 1 (age 20)	Cohort 11 (age 30)	Cohort 21 (age 40)	Cohort 31 (age 50)	Cohort 41 (age 60)
1976	26	203	304	275	113
1977	35	274	280	237	71
1978	70	401	380	344	106
1979	133	352	383	300	88
1980	188	393	359	329	82
1981	261	375	372	318	87
1982	363	372	370	327	74
1983	394	364	376	278	71
1984	463	410	347	237	70
1985	519	400	296	242	48
1986	543	398	303	222	53
1987	563	367	281	191	65
1988	604	390	260	208	49
1989	636	364	256	184	41
1990	640	299	228	196	59

vertical axis shows the average number of children in the household; children are here defined according to the official survey definition, which is 20 years old or younger. The plotted points are connected when we are following the same cohort through time, but different cohorts are left unconnected. To avoid complete clutter, we show only every fifth cohort, so that the first line segment on the left-hand side of the figure shows the number of children in households headed by those who were age 20 in 1976; by the time these people are age 35 in 1990, they are well launched into their child-rearing years and have a little under two children per household. The second cohort, the 25-year-olds in 1976, overlap with the first cohort for 10 years, but take us another five years into the life cycle, since these people are age 40 in the last, 1990 survey. To the extent that these two cohorts and the next one overlap, the profiles of children by age have similar shapes, with the maximum number of children attained around age 40 and falling thereafter, as the oldest leave home and new children cease to be born. However, the falling fertility shows up in the profiles as pronounced vertical shifts, or cohort effects, as we move from one cohort to the next.

One way of turning these pictures into numbers is to fit to these means a polynomial in age together with a series of cohort dummies, one for each date of birth. For figure 9.4, we do this only for those older than age 19 and younger than age 61, so as to exclude those cohorts where there are few individuals in the means. The estimated cohort effects show that, apart from the first five groups, aged 20–24 in 1976, which we know to be unreliable both because of the sample design and because of selectivity, there is an average decline of 0.06 children per year, so over the 35-year age span, from heads aged 25 to

**Table 9.3** Numbers of Persons in Selected Person Cohorts (indicated by age in 1976)

Year	Cohort 1 (age 20)	Cohort 11 (age 30)	Cohort 21 (age 40)	Cohort 31 (age 50)	Cohort 41 (age 60)
1976	788	521	608	461	249
1977	648	602	535	422	186
1978	1008	854	738	629	287
1979	1126	796	708	574	269
1980	1238	834	723	625	268
1981	1261	794	720	624	293
1982	1401	771	695	655	273
1983	1368	737	718	597	292
1984	1392	825	711	541	291
1985	1426	766	651	596	238
1986	1305	725	659	549	263
1987	1347	634	632	513	254
1988	1366	674	617	548	196
1989	1331	672	600	519	202
1990	1330	601	575	508	219

**Fig. 9.4** Number of children in the household

heads aged 60, there has been a decline in number of children per household of 2.3, from 4.71 to 2.44 at the maximum of the profile. Of course, this measure is neither a measure of the decline in children ever born, nor of children available to parents in old age, but it nevertheless seems likely that for those in their mid-thirties now, there will be only half as many sons to look after them in old-age as there are now for those who are 70 years old.

The demographic transition is one remarkable fact about modern Taiwan; another is its rate of economic growth and the effects on the earnings of the different cohorts. Figure 9.5, which is constructed for individuals and not for households, shows the cohort earnings patterns for every fifth cohort, again starting with those who were 20 years old in 1976. Again we can make out the usual pattern of earnings with age, but superimposed is very rapid within-cohort growth, at least for some of the cohorts. For those aged 25 in 1976, real earnings have grown at the astonishing rate of 12 percent per year over the 15 years since. For older cohorts, the amount is less, falling with age and eventually becoming negative for the cohort aged 50 in 1976. The bias of earnings growth toward the young is plausibly attributable to the much higher levels of education among young cohorts and is consistent with evidence from the United States, where earnings functions frequently show a positive interaction between education and experience. In the United States, much of the growth in earnings among the young is associated with job changes and with the greater tendency of young workers to move from one job to another. The same phenomenon might well hold true in Taiwan, with young people moving into the new jobs created by the rapid rate of economic growth.

Figure 9.6 shows nonlabor income by age and by cohort. In a country with savings rates as high as those in Taiwan, total income behaves very differently from earnings, and the figure shows how asset income and transfers replace earnings with age, so that total income does not collapse with age nearly as rapidly as does earnings. Lo (1987) provides more detail on the composition of income of the elderly and on how it has changed over time.

### 9.3 A Life-Cycle Interpretation of Consumption and Saving

With the demographic and earnings environment described, we turn to consumption behavior and the extent to which it fits the standard life-cycle framework. We are particularly concerned with whether Taiwan's high saving ratio can be attributed to its high rate of growth through the traditional life-cycle mechanism, with younger, richer cohorts being responsible for the bulk of the saving. We also want to look for traces of the demographic transition in the behavior of younger cohorts and, in particular, to investigate whether they are saving not just large amounts because their earnings are high, but are actually saving at a higher rate than their predecessors.

The cohort data are ideally suited for examining these issues. There are a large number of cohorts, with younger cohorts very much richer in lifetime

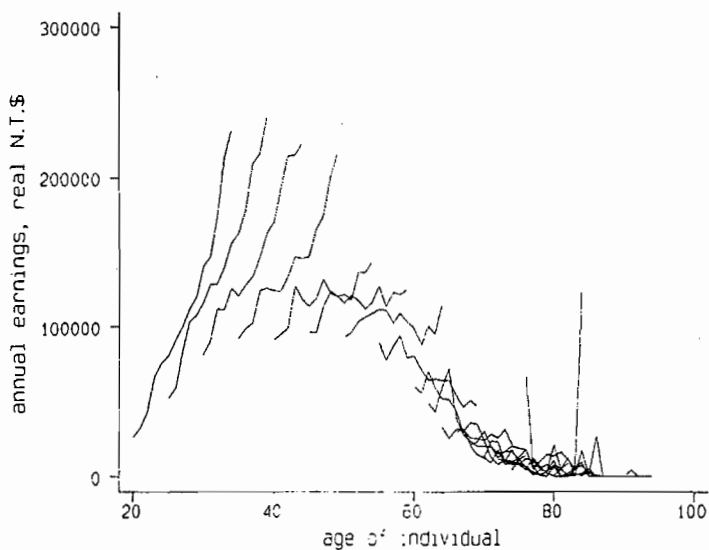


Fig. 9.5 Individual earnings, by age




Fig. 9.6 Individual nonlabor income, by age

Copyright ? 1994. University of Chicago Press. All rights reserved. May not be reproduced in any form without permission from the publisher, except fair uses permitted under U.S. or applicable copyright law.

resources than their parents or grandparents. Even at 10 percent per annum real earnings growth, which is less than the 12 percent average experienced by the younger cohorts, 25-year-olds in Taiwan can expect real lifetime earnings that are nearly 11 times greater than those of 50-year-olds. At the same time, since the data run for 15 consecutive years, we observe behavior at the same age for a range of different cohorts, so that it is possible to separate cohort (wealth) effects from age (preference) effects. Indeed, as we shall show below, it is even possible to say something about the effects of common macroeconomic shocks, as represented by year effects.

With no income uncertainty, the life-cycle model predicts that consumption is a function of lifetime resources (earnings plus inherited assets), with the fraction of resources consumed being a function of age, as dictated by preferences and the life-cycle variation in household size and composition. We write

$$(1) \quad c = g(a)W,$$

where  $W$  is the sum of assets and the discounted present value of current and expected future labor income, and  $g(a)$  is some function of age  $a$ . Taking logs, we have

$$(2) \quad \ln c = \ln g(a) + \ln W.$$

This equation holds at the level of the individual, but given its additive structure, it can be averaged over all households of the same age (as defined by head's age) in each year, so that the average of the logarithms of consumption for each age/year combination should be additively decomposable into a wealth term, which is constant within cohorts, and an age term. Equation (2) can then be estimated nonparametrically using the cohort data, by regressing the cohort averages of the logarithm of consumption against cohort dummies and age dummies.

The raw data on total household consumption are plotted in figure 9.7 in the same format as the earlier cohort diagrams; the figure shows the average of consumption, although it is the average of log consumption that will be used in the regressions. (Similar pictures for consumption per head, or consumption per adult equivalent, for various definitions of the latter, give very similar results.) Figure 9.7 shows a remarkable resemblance to figure 9.5, for earnings, a resemblance that will be explored further in the next section. Note again that, while there is a distinctly visible life-cycle pattern to consumption, rising with age and then falling, there is also a great deal of within-cohort consumption growth, especially for the younger cohorts. Clearly, old and young households are not sufficiently altruistically linked for their consumption to move in lockstep.

The results can be given an explicit life-cycle interpretation by estimating equation (2), and the results are given, not in numerical form, since there are 66 cohort effects and 51 age effects, but graphically in figure 9.8, where cohort






Fig. 9.7 Total household consumption, by age of household head

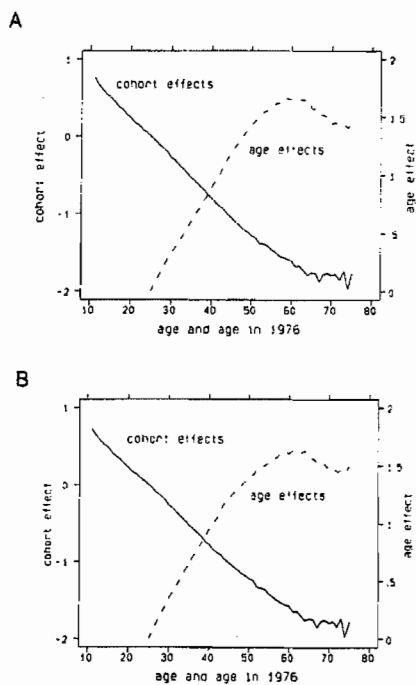


Fig. 9.8 Age and cohort effects in consumption

Note: Controls for numbers of adults and children are excluded from panel A and included in panel B.

Copyright © 1994. University of Chicago Press. All rights reserved. May not be reproduced in any form without permission from the publisher, except fair uses permitted under U.S. or applicable copyright law.

effects (age in 1976) and age effects are shown together. In panel *A* equation (2) is estimated with age and cohort dummies only, while in panel *B* we show the age and cohort effects when the (average) numbers of adults and children are added as additional regressors. (Additional adults markedly increase consumption; additional children have little effect.) As predicted by the life-cycle model for a growing economy, the cohort effects decline steadily from younger to older cohorts; both graphs show that, at the same age, consumption increases at about 4 percent per cohort, so that, for example, the consumption at age 40 of someone born in 1950 is on average 3.4 times larger than the consumption at age 40 of someone born in 1925. As is to be expected, the age effects depend to some extent on whether the household composition variables are explicitly included in the regressions. In either case, the age profile is much steeper than is typically seen in the United States, or even in other LDCs, flattening out only very gradually with age. In both panels, consumption falls with age only after age 60, although the decline is less pronounced in the bottom panel, where we control for the numbers of household members. Note again, however, that the selection effects should make us very skeptical about conclusions for households with elderly heads. While steep age-consumption profiles are unusual, they are certainly not inconsistent with the life-cycle model. In an economy growing as rapidly as Taiwan's, real interest rates are very high, especially in the large and rapidly developing small-business sector, so that there are strong intertemporal incentives to postpone consumption and for consumption to grow rapidly over time. Alternatively, it is possible that the Taiwanese are more patient than most other people, or that the Chinese veneration for the elderly extends to a perceived high marginal utility of consumption in old age.

One feature of figure 9.7 is not accounted for by either the age or the cohort effects in figure 9.8. This is the fact that the individual cohort "tracks" have similar shapes, so that, for example, almost all show a decline in consumption with age between the fourth and fifth observations. These are clearly caused by the presence of macroeconomic effects that impinge on all cohorts to a greater or lesser degree, but which are located in real time and which cannot be explained by cohort or age effects. The obvious extension to equation (2) is to allow for fixed-year effects, so that the consumption equation becomes

$$(3) \quad \ln c_{at} = \ln g(a) + \ln W_b + \theta_t,$$

where  $\theta_t$  is a year fixed effect, and the subscripts  $a$ ,  $b$ , and  $t$  denote age, cohort, and time, respectively. Note that given age and time, cohort is determined; indeed we have been measuring  $b$  as age in 1976, which is  $a - t + 1976$ .

Equation (2) is an implication of the life-cycle model of consumption when there is no uncertainty, so that cohort wealth levels are never revised. Once uncertainty is admitted, wealth levels will be revised in response to macroeconomic shocks, so that the life-cycle model with uncertainty provides at least some basis for equation (3). However, the link between equation (3) and the

model under uncertainty is not as clean as was the link with equation (2) under certainty, since the effects of a common macroeconomic shock on wealth levels ought to vary with age. While it is possible in principle to include interaction terms between year and cohort effects, unrestricted estimation would not be possible even with the current data. Nevertheless, the year effects are certainly present in the data, and it seems useful to estimate equation (3) as an approximation to the more general model, if only to ensure that the failure to accommodate the year effects does not contaminate the estimates of the age and cohort effects, with which the year dummies are strongly correlated.

The estimation of equation (3) requires some thought about the relationship between age, cohort, and year dummies. Write  $C$ ,  $Y$ , and  $A$  for matrices of dummy variables of cohort, year, and age dummies; each matrix has 1,031 rows, the number of year/head's-age pairs, while the numbers of columns are 65 for cohorts (for those aged 11 in 1976 to those aged 75 in 1976), 15 for years, and 51 for ages (from age 25 to age 75). Note that we have truncated on age of head, eliminating those below 25 to avoid the "missing" males below that age group and those above 75 to avoid the very imprecisely estimated cohort means from households with very old heads. Of course, there are no heads of households in 1976 who were then 11 years old, but we do include one observation on that cohort, when they are 25-year-old heads of household in 1990. Equation (3) can be rewritten in terms of the dummy variable matrices as

$$(4) \quad \ln c_{it} = \iota\beta + C\gamma + A\alpha + Y\psi + \varepsilon,$$

where  $\iota$  is a vector of units, and the vectors  $\alpha$ ,  $\gamma$ , and  $\psi$  are parameters of age, cohort, and year effects. As usual, one category from each set of dummies must be excluded, and doing so presents no nonstandard issues of interpretation. However, the dependency between age, cohort, and year introduces a slightly less standard complication.

Let  $\sigma_n$  denote the (transpose of the) vector  $(1, 2, 3, \dots, n)$ . Then, since cohort is age minus year plus a constant, the matrices of dummy variables satisfy the exact linear relationship

$$(5) \quad C\sigma_{n_c} = A\sigma_{n_a} - Y\sigma_{n_y} + n_y\iota,$$

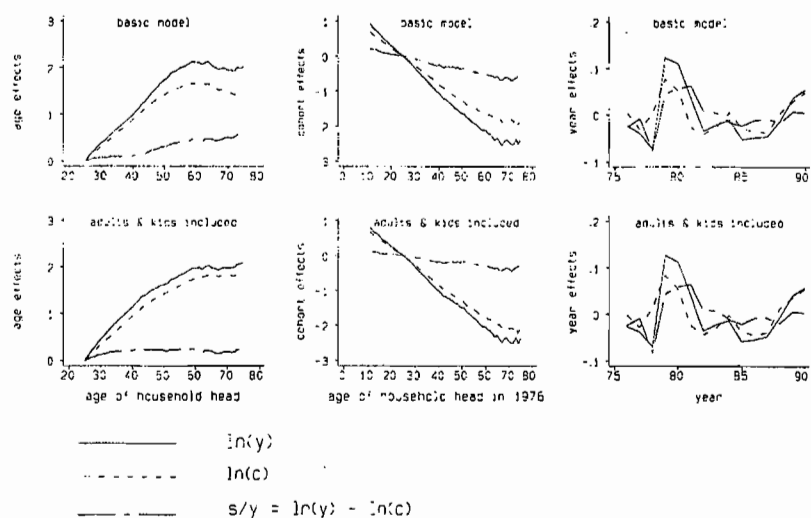
where  $n_c$ ,  $n_a$ , and  $n_y$  are the numbers of cohorts, ages, and years. To see that the constant term in equation (5) equals  $n_y$ , note that the youngest cohort of people (aged 11 in 1976) will be the youngest age of 25 in the  $n_y^{\text{th}}$  year. As a result of the identity expressed by equation (5), the parameters in equation (4) are not identified, even after one category has been dropped from each set of dummies. In particular, if the vectors  $\alpha$ ,  $\gamma$ ,  $\beta$ , and  $\psi$  are replaced by  $\tilde{\alpha}$ ,  $\tilde{\beta}$ ,  $\tilde{\gamma}$ , and  $\tilde{\psi}$ , where

$$(6) \quad \begin{aligned} \tilde{\alpha} &= \alpha - \kappa\sigma_{n_a}, & \tilde{\beta} &= \beta - n_y\kappa, \\ \tilde{\gamma} &= \gamma + \kappa\sigma_{n_c}, & \tilde{\psi} &= \psi + \kappa\sigma_{n_y}, \end{aligned}$$

and  $\kappa$  is some arbitrary nonzero scalar, then, as is easily checked by substitution into equation (4), there is no consequence for the predicted value of consumption. The way that this works can readily be seen from the original cohort-consumption plots in figure 9.7. For the first cohort, suppose that there is no cohort effect, so that the first trace is the sum of age and year effects. Since the cohort ages one year at a time, adding equal amounts to year and age dummies will leave the trace unchanged. The trace for the second cohort starts from the same year but one year of age younger, so that the second trace will also be left unchanged if the second cohort effect is incremented by one, and so on through the cohorts according to equation (6). In effect, any trend in the data can be arbitrarily reinterpreted as a year trend or (since year equals age minus cohort plus a constant) as trends in ages and cohorts that are equal but of opposite sign. However, it is clear that the appropriate normalization is to require that the year effects sum to zero. As an example, consider a hypothetical case in which there are apparently no cohort or age effects, but where the logarithm of consumption increases by  $\psi$  per annum for all age groups. Consider what the cohort diagram figure 9.6 would look like in such a case. Each cohort trace would increase by  $\psi$  per year of age, but each cohort trace would start at the same horizontal level as the previous one, or equivalently, at a vertical shift of  $\psi$  below it. A steady growth in year effects simply means that consumption is growing with age and declining with cohort, and it is appropriate to attribute the effects to age and cohort, not time.

In the light of this discussion, we estimate equation (4) with the first age group, and the fifteenth cohort omitted, so that the reference group is that for a household headed by a 25-year-old in 1976. The 15-year dummies are constrained to be orthogonal to a time trend and to add to zero. The "base year" is thus a timeless average of all years, and any time trend is attributed to cohorts and ages, not to time.

Figure 9.9 shows the resulting age, cohort, and year effects, with (*bottom panels*) and without (*top panels*) the addition of the average adult and children variables. The figure also shows the same decomposition for the average of the logarithms of total income (earnings plus asset and other income) and for the "saving ratio," measured here as the difference of the logarithms of income and consumption. (For Taiwanese saving ratios, the approximation is not particularly accurate, but the logarithmic form is convenient, given that we are working with the logarithms of income and consumption.) The consumption profiles are not markedly different from those in the figure 9.8, even though the year effects are quite significant, with an associated  $F$ -value of 59.8 in the basic model and 81.5 in the model with adults and children. The corresponding  $F$ -values for age and cohort effects are, respectively, 408.8 and 373.1 in the basic model, and 437.0 and 60.4 in the extended model, and it is these effects that are our primary concern in examining life-cycle behavior. Conditional on cohort and age effects, adults increase income by about as much as they in-



**Fig. 9.9** Age, cohort, and year effects in consumption, income, and saving rates, for household heads

crease consumption and have little effect on household saving, while children decrease income with little effect on consumption and so decrease saving.

The cohort effects are larger the younger the cohort for consumption, for income, and for the *saving ratio*, so that, at the same age, those born later in calendar time are saving a higher fraction of their incomes. This is exactly what we might expect from the life-cycle model in the Taiwanese context in the face of rapidly declining fertility and rapidly rising life expectancy. Although the cohort differences in saving ratios are too large to be plausible in the model without the addition of numbers of adults and children, the extended model shows a flatter profile of savings rates across cohorts. For example, the saving cohort effects in the second plot in the bottom panel of figure 9.9 show a difference of  $-0.23$  in  $\ln(y/c)$  between the cohort aged 11 in 1976 (aged 25 in 1990) and the cohort 35 years older, aged 60 in 1990. The average value of  $\ln(y/c)$  in 1990 for households headed by 25-year-olds was 0.276, a saving rate of 24 percent, so that the model predicts a saving rate of 4.5 percent for the older cohort at the same age, i.e., in 1955. As we saw above, this is exactly the comparison group for which the number of children has fallen by a half, and the increase in the saving rate is perhaps of the order of magnitude required to replace family by autarkic old-age insurance.

The age effects in the left-hand panels are not obviously consistent with the life-cycle story. The rising age-profile of consumption is matched by an even more rapidly rising age pattern of income, so that the saving rate, instead of being positive at young ages and negative later, shows a steady increase with age. If such a result is correct, an increase in the rate of growth of productivity

would *decrease*, not increase, aggregate saving since it would redistribute income toward the young, who have the lowest saving propensity, at least provided cohort effects are eliminated, as would be expected when the demographic transition is complete. However, it is quite likely that the age effects on saving are contaminated by the selection processes for headship, certainly at old ages, and perhaps also at young ages. As already discussed, it is entirely plausible that those households headed by older people are those with an unusually high propensity to save, so that the right-hand panels in figure 9.9 are revealing only the selection effects, and not the true patterns of saving ratios. Nor can this problem be resolved by switching from a household to an individual basis, as is done in figure 9.10, which repeats figure 9.9, but with averages taken over individuals, not households. The problem is that here we are still looking at *household* incomes, consumption, and saving, only the household magnitudes have been attributed to individuals. The age effects on saving are now even more extreme, because the old individuals on the right of the diagram are now predominately located in households headed by younger adults, so that we lose the flattening of income, consumption, and savings profiles with age that is observed in the household-level data.

Our results can thus be interpreted in terms of a more-or-less standard life-cycle model, at least provided we are allowed to choose a desired consumption profile that rises rapidly with age, which indeed is theoretically plausible in the circumstances. Saving *rates* are systematically higher for the younger cohorts, and the effect is consistent with the increased need to provide for old age. And if we fail to observe any life-cycle profile in saving rates by age, there are statistical grounds that can perhaps explain what we see.

### 9.5 Alternative Explanations: Tracking and Growth

The ability to give a life-cycle interpretation to the data does not mean that the life-cycle model is in fact correct, and in this section we look at some more negative evidence. In particular, we consider whether there is evidence in these data that consumption “tracks” income by more than would be expected if the life-cycle model is correct, and we look more closely at the relationship between saving and growth.

Tracking has been found in other data; see particularly Carroll and Summers (1991) for the United States and Deaton (1992b) for Thailand and Côte d’Ivoire. However, in all these examples, the large mass of households accumulate little or nothing, and saving is used to buffer short-term fluctuations in income; in the absence of accumulation, consumption and income must move together in the long run. However, there is certainly no lack of accumulation by households in Taiwan, so that we should expect to observe a greater decoupling of consumption and income, even over long life-cycle spans. Taiwanese households have demonstrated their ability to accumulate enough assets to accomplish even “low-frequency” consumption smoothing. Although the evi-

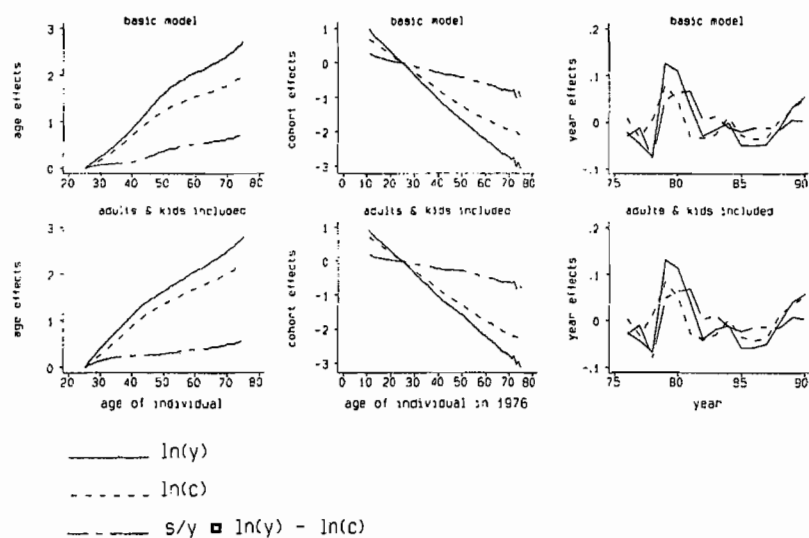


Fig. 9.10 Age, cohort, and year effects in consumption, income, and saving rates, for individuals

dence in the previous section is consistent with the life-cycle story, it does not immediately follow that much long-term smoothing is being done, since the preferred profile of consumption may not call for it. It is therefore necessary to look at the data in a different way.

Figure 9.11 shows the same data on consumption and income by age as in figure 9.5, but with a separate diagram for each cohort, and with income and consumption shown together. Once again, we see the marked differences in the shape of the profile between the younger and older cohorts, with much more rapid growth for younger households. We also see the growing saving ratio as a widening between consumption and income over time. But the most marked feature of figure 9.11 is the extraordinary coherence between the patterns of consumption and income for each of the cohorts. The shape of each changes with age, but they change together, with consumption tracking income, year by year, and age by age. Such a picture suggests a crude Keynesian consumption function, with consumption strongly linked to income. And yet we have seen in the previous section that the data, which are the same as those used here, are consistent with a life-cycle interpretation. The key to the reconciliation is the fact that in the earlier analysis, in figure 9.9, the cohort effects, the age effects, and the year effects have very similar shapes, across cohorts, age, and time, for both consumption and income. While this is consistent with the life-cycle model, only the coherence between income shocks and consumption shocks is directly predicted by it, although it could also be argued that in a rapidly growing economy, real returns will be high, so that optimal intertem-

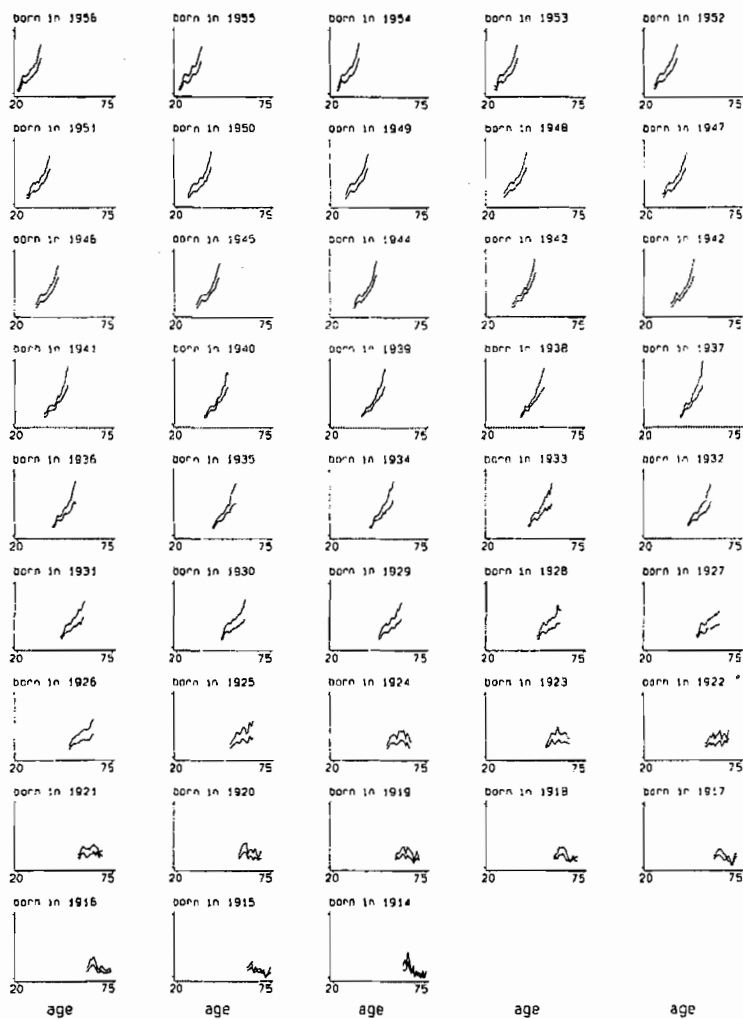


Fig. 9.11 Consumption and income, by age of household head

poral planning should generate rapid consumption growth to match rapid income growth. Even so, figure 9.11 is strongly suggestive of some simpler model embodying a much more direct link between consumption and income. That said, it is no simple matter to think of a coherent alternative. Liquidity constraints make very little sense in an environment in which so many people are saving so much. Nor do buffering models seem relevant here, whether with borrowing constraints as in Deaton (1991) or with precautionary saving and large shocks as in Carroll (1991). Both of these models work by supposing that people generally wish to accumulate little or nothing, but use saving to buffer



consumption against short-term fluctuations in income. For such people, consumption tracks income because, in the long run, consumption is identical to income. That is not the case here. Perhaps the life-cycle model still provides the best theoretically consistent account of the facts.

Before accepting that position, consider the positive relationship between saving and growth, perhaps the most celebrated prediction of the life-cycle model. We have already seen that the estimated models do not clearly support this prediction in the Taiwanese case, since the saving rate appears to increase with age, although we cannot use this to reject the model because of the increasing selection of household heads with age. Even so, the *fact* that there is a connection between saving and growth seems hard to dispute. Taiwan, like Korea, Japan, Hong Kong, and Singapore, is among the most rapidly growing and highest saving economies in the world, and the cross-country relationship between saving and growth has been repeatedly documented. Furthermore, the productivity slowdown in the OECD countries has been accompanied by a reduction in national saving ratios, essentially without exception (see Modigliani 1990). However, there is still far from general acceptance that the *mechanism* is that of the life-cycle model, whereby slower growth redistributes national income away from younger, saving cohorts, to their lower-saving or dissaving elders. Indeed, for the United States, Bosworth, Burtless, and Sabelhaus (1991) have undertaken the very difficult task of interpreting the time-series evidence from the various Consumer Expenditure Surveys and have provided strong evidence that the decline in saving in the United States has been a decline that has affected all age groups; it is not the aggregation phenomenon that the life-cycle model requires. This evidence suggests that there is a mechanism linking growth and saving *within* cohorts, and it is this possibility in Taiwan to which we now turn.

The top panel of table 9.4 shows regressions of the "saving rate" (defined as before by the logarithm of the ratio of total income to consumption) against growth rates of earnings. These are run on the same averaged data displayed in the household diagrams above, with each observation corresponding to an age/year pair, and once again include only ages 25–75. As appropriate, averages are averages of the logarithms of the underlying household data, not the logarithms of averages, and all regressions are weighted by (the square root of) the number of households in the cohort average. The regression of the saving rate on the current rate of change of earnings, i.e., the growth of current earnings over last period's earnings, is *not* shown; it produces a strong positive coefficient, but such a result tells us little, since measurement error will induce a positive correlation between current saving and current income. Instead, the first line shows the regression of the saving rate on the first four lags of earnings growth; all the coefficients are positive and significantly different from zero. The second line explores the possibility that the relationship is between saving and current earnings growth, but using instrumental variables estimation to purge the measurement error. Since measurement error will induce a

Table 9.4 Saving Rates, Consumption Growth, and Earnings Growth

Effects Included	$d\ln e$	$d\ln e_{-1}$	$d\ln e_{-2}$	$d\ln e_{-3}$	$d\ln e_{-4}$	F-Test for Overidentification
Saving Rate Regressions: $\ln y - \ln c$						
1. None	0.6995 (1.6)	0.0793 (8.3)	0.1049 (9.2)	0.0888 (7.7)	0.0632 (5.8)	10.75
2. Cohorts	0.2150 (0.9)	0.0302 (4.1)	0.0483 (5.5)	0.0431 (4.9)	0.0321 (3.9)	10.36
3. Cohorts and ages	0.2881 (1.5)	0.0187 (2.3)	0.0421 (4.1)	0.0430 (4.2)	0.0329 (3.5)	3.09
4. Cohorts, ages, and years	0.0318 (0.7)	-0.0136 (1.7)	-0.0114 (1.0)	-0.0099 (0.9)	0.0075 (0.8)	0.56
Consumption Growth Regressions: $d\ln c$						
5. None	-0.0354 (0.3)		-0.0188 (1.8)	-0.0116 (1.0)	0.0064 (0.6)	2.37
6. Cohorts	0.0552 (0.3)		-0.0307 (2.8)	-0.0260 (2.2)	-0.0062 (0.5)	4.80
7. Cohorts and ages	0.1510 (1.1)		-0.0158 (1.4)	-0.0033 (0.3)	0.0260 (2.1)	4.16
8. Cohorts, ages, and years	0.0258 (0.4)		-0.0005 (0.1)	0.0024 (0.2)	0.0071 (0.6)	0.03

Notes: The regressions in the top panel have the cohort average of the individual "saving rates"  $\ln(y/c)$  as the dependent variable, where  $y$  is total income and  $c$  is consumption. In the first regression,  $\ln(y/c)$  is regressed on four lags of earnings growth  $d\ln e$ ; in the second, saving growth is regressed on current earnings growth by instrumental variables, with the second through fourth lags of earnings growth as instruments. The regressions are repeated with the inclusion of cohort effects, cohort and age effects, and cohort, age, and year effects. The bottom panel repeats the regressions with consumption growth as the dependent variable. All consumption regressions, OLS and IVE, exclude the first lag of earnings growth. All regressions include the average numbers of adults and children, and all are weighted by the square root of the number of observations in each cohort. When computing the cohort averages of log earnings (used to construct earnings growth), the few households with no earnings are assigned a value of earnings equal to one.

Numbers in parentheses are  $t$ -values.

spurious correlation between successive lags of earnings growth, we use only the second and third lags as instruments. This regression shows no relationship between the saving rate and predictable earnings growth, and the overidentification test rejects the hypothesis that lagged growth rates affect saving rates only through their ability to predict current growth. (Indeed they have little ability to predict current earning growth.)

The rest of the top panel explores the same regressions but with the addition first of cohort dummies, then of cohort and age dummies, and finally of cohort, age, and year dummies. Given the relationships that exist between earnings growth, age, and cohorts, we would argue for including these dummies, but for excluding the year dummies, since the latter should be subsumed in the effects for which we are testing. Although the coefficients are smaller than in the first row, the positive association between lagged growth of earnings and saving rates remains, at least until the year effects are introduced. None of the results suggest that the effect is working through anticipated earnings growth. These results show a positive relationship between saving rates and growth, a relationship that is not an aggregation effect, but holds for specific cohorts over time. Tracking the same group of people as they age, we find that their saving rate is higher in the years following rapid earnings growth. Because this effect is from *lagged* earnings growth, it has nothing to do with the standard response of saving to transitory income, and it is not an effect that is predicted by life-cycle theory.

These effects of within-cohort growth on within-cohort savings are a good deal smaller than those that are typically predicted from the aggregation effects of life-cycle saving, or that are typically found in the cross-country or long time-span evidence. Those findings, and again see Modigliani (1990), typically show an effect of growth on savings of about two, so that a shift in earnings growth from 5 to 10 percent will generate an increase of 10 percentage points in the saving rate. If we take the results from the third row of table 9.4, the effect is 0.137, which is less than a tenth of the typical cross-country estimates. Even the first row estimates, which condition on neither cohort nor age effects, add up to only 0.33. Hence, while the within-cohort growth effects are present and while they cast doubt on the life-cycle model, they probably cannot account for a major share of the relationship between saving and growth in Taiwan.

The second panel of table 9.4 reports similar regressions, but with the dependent variable the rate of growth of consumption. These correspond closely to the standard "excess sensitivity" tests in the macroeconomic consumption literature; see, for example, Campbell and Mankiw (1991), who look for a relationship between consumption growth and anticipated growth in labor income. The only difference between the top and bottom panels, apart from the dependent variable, is that the first lag of earnings growth is excluded from the unrestricted regression; if consumption growth is correlated with unanticipated growth in earnings, as the life-cycle theory supposes, and if earnings are mis-

measured, there will be a spurious correlation between consumption growth and lagged earnings growth that does not contradict the model. The results provide no evidence of excess sensitivity in Taiwan, although there are good reasons to question the power of the test. Lagged earnings growth has little predictive power for consumption growth in any of the regressions, but it also has little predictive power for current earnings growth. With little ability to anticipate earnings growth, we cannot tell whether the component of earnings growth that might be anticipated is or is not correlated with changes in consumption.

### 9.6 Summary and Conclusions

Where do these results leave us, and what is their implication for saving in Taiwan, particularly for retirement saving in a rapidly aging population? We have found it a good deal harder to find fault with the life-cycle theory than we supposed would be the case. Although consumption seems to move very closely with income, the Taiwanese save a great deal, and the comovement can be explained in terms of earnings shocks, which induce a short-run correlation between consumption and income, together with a taste for rapid consumption growth with age, which itself might be attributed to the intertemporal incentives generated by very high rates of return in a rapidly growing economy. We also find that younger cohorts have higher saving rates, as is to be expected given falling fertility and rising life expectancy. If we accept such a picture, we can attribute at least part of Taiwan's high saving ratio to farsighted young consumers preparing for a "modern" old age in which they will be thrown on their own resources, rather than on those of their married sons.

We are not entirely convinced that the picture is quite so simple, although we have no really coherent alternative to offer. The life-cycle explanation for the coherence between consumption and income seems farfetched compared with some simpler and more parsimonious story in which consumption is directly linked to income, although we also recognize that it is difficult to explain why people should behave in this way, especially in a high-saving economy where liquidity constraints are unlikely to be a problem. There is certainly no evidence of excess sensitivity in our results, though better and more powerful tests might yield different results. We also found evidence that saving responds to growth within cohorts, something that is consistent with the importance of habits in consumption. Even so, the effect is not nearly as large as that delivered by simple life-cycle stories and is small compared with other empirical evidence on the relationship. Our data do not show dissaving or even decreased saving among older households; indeed the estimated age effects in saving behavior increase steadily with age. Of course we recognize that this evidence is severely contaminated by the processes of household dissolution and combination as individuals age. Nevertheless, such results are hardly overwhelming evidence in favor of the life-cycle model. We therefore hesitate to endorse the

position that the life-cycle model is alive and well and living in Taiwan, especially if such a position is taken to imply that there is no need for public action to supplement social security among those who will be old 25 years from now.

## References

- Bosworth, Barry, Gary Burtless, and John Sabelhaus. 1991. The decline in saving: Some microeconomic evidence. *Brookings Papers on Economic Activity* 1:183–256.
- Campbell, John Y., and N. Greg Mankiw. 1991. The response of consumption to income: A cross-country investigation. *European Economic Review* 35:715–21.
- Carroll, Christopher D. 1991. Buffer stock saving and the permanent income hypothesis. Washington, D.C.: Board of Governors of the Federal Reserve System. Circulated manuscript.
- Carroll, Christopher D., and Lawrence H. Summers. 1991. Consumption growth parallels income growth: Some new evidence. In *National saving and economic performance*, ed. B. Douglas Bernheim and John Shoven. Chicago: University of Chicago Press.
- Chang, Ming-cheng. 1987. Changing familial network and social welfare in Taiwan. In *Conference on economic development and social welfare in Taiwan*, vol. 2. Taipei: Institute of Economics, Academia Sinica.
- Deaton, Angus S. 1991. Saving and liquidity constraints. *Econometrica* 59:1221–48.
- . 1992a. Saving and income smoothing in Côte d'Ivoire. *Journal of African Economies* 1:1–24.
- . 1992b. *Understanding consumption*. Oxford and New York: Oxford University Press.
- Hermalin, Albert, Mary Beth Ofstedal, and Chi Li. 1991. The kin availability of the elderly in Taiwan: Who's available and where are they? Population Studies Center, University of Michigan. Circulated manuscript.
- Lo, Joan Chi-chiung. 1987. The changing patterns of household structure and economic status of the elderly: 1976 to 1985. In *Conference on economic development and social welfare in Taiwan*, vol. 1. Taipei: Institute of Economics, Academia Sinica.
- Modigliani, Franco. 1990. Recent declines in the savings rate: A life cycle perspective. Frisch Lecture, Sixth World Congress of the Econometric Society, Barcelona, August.
- Paxson, Christina H. 1992. Using weather variability to estimate the response of savings to transitory income in Thailand. *American Economic Review* 82:15–33.
- Republic of China. 1989. *Report on the survey of personal income distribution in Taiwan area of the Republic of China*. Taipei. Directorate-General of Budget, Accounting and Statistics, Executive Yuan.
- . 1991. *Statistical yearbook of the republic of China*. Taipei: Directorate-General of Budget, Accounting and Statistics, Executive Yuan.
- Visaria, Pravin. 1980. Poverty and living standards in Asia. Living Standards Measurement Study Working Paper. Washington, D.C.: World Bank. Circulated manuscript.