Saving and growth: another look at the cohort evidence

Angus Deaton
Christina Paxson

Research Program in Development Studies
Princeton University

February 1998

We are grateful to Bo Honoré for help in selecting an appropriate smoother, and to Vivi Alatas for research assistance. We thank Orazio Attanasio for his (anticipated) comments.

0. Introduction

In recent years, as longer time series of cross-sectional household surveys have become available, it has become possible to look at the consumption and saving behavior of birth cohorts in a number of developing and developed economies. The cohort evidence is singularly appropriate for the analysis of life-cycle models of consumption because, at least in its simpler forms, the life-cycle hypothesis (LCH) predicts that the cohort average of the logarithm of consumption in any year can be additively decomposed into a time-invariant cohort effect and an age effect, both of which can be readily recovered from the cohort data by linear regression on dummy variables. The results of these cohort level analyses have not been favorable for the LCH interpretation of the international correlation between growth and saving. According to this, higher rates of economic growth drive up rates of national saving by expanding the lifetime resources of younger generations, who are saving, relative to the lifetime resources of older generations, who are dissaving. While it is typically possible to interpret the cohort results in a way that is consistent with the LCH, it is a good deal harder to rescue the prediction that higher growth means higher national saving rates, or at least that the effect is large enough to be consistent with the international relationship in which a one percentage point increase in the rate of per capita growth is associated with a roughly two percentage point increase in the saving rate. In Section 1, we review existing cohort studies from a range of rich and poor countries, and also present new results for Indonesia. This evidence shows no strong negative relationship between saving rates and age, so that when higher growth redistributes lifetime resources towards the young, the effect on savings is modest, and in some cases even negative.

These results are largely independent of the debate over the *validity* of the LCH itself. Recent work, much of which has focussed on precautionary motives for saving, has shown that the LCH is capable in principle of accounting for a wide range of behavior that cannot be explained by earlier, simpler models. In many of these new models, consumption tracks income quite closely over the life-cycle, in which case growth will have little effect on saving, even though the LCH holds good.

There are (at least) two major unresolved problems with the cohort approach. The first is that household surveys typically do not collect comprehensive data on contributions to, or disbursement from, pension funds, whether public (social security) or private (annuities associated with previous employment or with saving). Contributions are usually missed altogether, particularly when they are made by employers, so that the saving of those in employment is understated. Pension income is frequently aggregated with other income, which misclassifies the component that is not income at all, but comes from running down assets, so that saving is overstated among the elderly. The life-cycle profiles of savings that are computed in the standard way therefore understate the hump-shape in the age-profile of saving, and therefore also understate the size of the effect of growth on saving. In this paper, we state this problem, but make no contribution to its solution; our concern here is with middle-income and poorer countries where social security is absent, and where very few people receive pensions from previous employment. Even in Taiwan (China), from where much of our evidence comes, the provision for retirement by employers is a recent phenomenon, and usually takes the form of a lump sum paid at the time of retirement.

The second problem with the cohort approach comes from the way in which cohorts are defined. Because household surveys collect consumption for households, not individuals, the usual approach is to work with household income and household consumption, and to track cohorts of households defined by the age of the household head, rather than individuals defined by their year of birth. For example, in a survey collected in 1990, we take average consumption of households whose heads were 45 years old, i.e. born in 1945, which we link to the average consumption of households whose head is 46 in 1991. That these are not *exactly* the same people does not matter provided that we are sampling from the same, unchanging cohort of households in the two dates. What matters is that the cohorts do change from one year to the next. There is emigration and immigration, and births and deaths, so that the population of *individuals* changes from one year to the next. But potentially more serious than changes in the population of individuals is changes in the population of *households*; even without migration, births, or deaths, households can form and disintegrate as people regroup through divorce, children leaving home, or the elderly moving in with their children. Once again, there is the potential for such changes to disguise the validity of the LCH, and to generate a misleading age-profile of saving. Among elderly households, survival as a household is plausibly positively related to saving, with high-saving households more likely to last as independent units. Similarly, among young adults, those who are chronically in deficit are less likely to be able to set up independent households. In consequence, the household data will tend to overstate saving among the young and the old, flattening out the hump in saving.

In life-cycle models, the effects of growth on saving are calculated by reweighting the age-profile of saving, with economic or population growth shifting weight away from the elderly toward the young. When we work on a household basis, it is the household (head's) age-profile of saving that we estimate, and that is aggregated with different weights depending on rates of growth. But this is likely to be very different from estimating individual age-profiles of saving and aggregating over individuals. Consider for example what happens if some fraction of the elderly live with their children, and that according to the individual profiles, these elderly are net dissavers. In the household data, this dissaving is hidden by the simultaneous saving that is going on in their middle-aged host households. But if the fraction of elderly is increased, or if a productivity slowdown makes their dissaving more important relative to the saving of their children, national saving will fall. The household approach does not capture the effect, because the age profile of saving is estimated on a household basis, and is inappropriately held constant when population or economic growth rates are changed. Changes in growth rates change household structure and the relative economic power of household members of different ages, and these effects need to be taken into account when assessing the effects of growth on saving.

Section 2 of the paper reviews these "selection" issues, and argues that selection through mortality and migration has different consequences from the selection that comes about through changes in household membership. It also uses the Taiwanese data to document the importance of these changes.

The main contribution of the paper is an attempt to address the selection and growth issues by working, not with cohorts of households, but with cohorts of individuals. The gain from doing so is that we dispose completely of selection associated with household formation, which compromises our estimates of the effects of growth on saving, leaving only selection through births, deaths, and migration, which does not. The cost is that we do not have data on individual consumption, and it is difficult to allocate income to individuals due to the prevalence of family enterprises, so that the standard methodology can no longer be applied. We work with an "individual" version of the LCH in which it is individuals, not households, who set their consumption so as to exhaust their lifetime resources according to an unchanging age profile of consumption. In this version, households are a veil for the individuals who live in them, and household consumption and household income are simply the sum of the consumptions and incomes of the people who live in it. This is obviously a simplification that we view as a convenient starting point; a richer model would recognize the technology of household consumption—for example the existence of economies of scale—and allow the individuals in a household to consume more in total than aggregate household consumption.

Section 3 of the paper lays out this individual model, and shows how its parameters can be recovered from household level data without the need for arbitrary imputation of individual consumption from household data. We also propose a smoothing procedure that takes care of the fact that some cohorts—especially older cohorts—have relatively few representatives in the samples, and thus have noisily measured consumption. This smoothing procedure allows us to recover age and cohort effects in consumption and saving at the individual level.

Section 4 presents results based on eighteen years of cross-sectional household survey data from Taiwan. We have chosen Taiwan, not only because we have more cross-sections than for any of the other possible countries (Thailand, Indonesia, Mexico, or Pakistan), but also because it is a high saving country in which households typically report consumption much lower than income. We present estimates of the age-profile of saving on both and individual and household basis, the latter an updated version of the results in Deaton and Paxson (1994), and show that the former is much more favorable to the life-cycle story than the latter, showing a much stronger hump in saving. Under some conditions, the hump saving is large enough to generate growth-to-saving effects as large as those that come from the traditional stripped-down life-cycle models.

1. Saving and growth: evidence from cohort studies

It is useful to begin with an outline of the theoretical framework, partly to understand the results reviewed here, but also because we need the theory as a basis for the work in Section 3 below. In the standard approach, no clear distinction is drawn between a household and an individual, or perhaps more accurately, the theory is developed for individuals and then applied to households. In the simplest life-cycle model in which there is no uncertainty, consumption follows an age profile that is determined by the interaction of preferences and real interest rates, with the *level* of the profile set by the level of lifetime resources. For individual *i*, (household or person) born at date *b*, and observed at age *a* (i.e. at date $b + a$), consumption c_{iab} is given by

$$c_{iab} = f_i(a) W_{ib} \qquad (1)$$

where $f_i(a)$ is the age profile of consumption, and W_{ib} is a measure of lifetime resources. Note that while the age-profile is indexed on *i*, and so varies over individuals, it is independent of birth date *b*, so that the distribution of age profiles over individuals within each cohort is the same for all cohorts, while resources W_{ib}, although different for each cohort—with economic growth, later born cohorts have higher lifetime resources—is invariant with respect to age, which comes from the assumption of no uncertainty. (In a model with uncertainty, the age effects can be estimated following the techniques developed by MaCurdy, 1981, and Browning, Deaton and Irish, 1985). If we take logarithms of (1) and average over all members of the same cohort at the same age, we obtain

$$\overline{\ln c_{ab}} = \overline{\ln f(a)} + \overline{\ln W_b} \qquad (2)$$

where the lines over the variables denote means. Equation (2) can be estimated by regressing the average of the logarithm of consumption for those born in b and observed in $b + a$ on a set of age and cohort dummies, i.e. from the regression

$$\overline{\ln c} = \boldsymbol{D^a} \alpha_c + \boldsymbol{D^b} \gamma_c + u_c \qquad (3)$$

where $\overline{\ln c}$ is a stacked vector of log consumption levels with elements corresponding to each cohort in each year, $\boldsymbol{D^a}$ is a matrix of age dummies, and $\boldsymbol{D^b}$ is a matrix of cohort (i.e. birthyear) dummies. The coefficients α_c and γ_c are the age and cohort effects in consumption, and u_c is the sampling (or equivalently measurement) error that comes from the fact that $\overline{\ln c_{ab}}$ is a sample estimate of the average log consumption of all individuals born at b and observed at $a + b$.

Following Modigliani's lead, the standard assumption about earnings in life-cycle models is that they, like consumption, have a time invariant age profile, so that productivity growth, and variations in productivity growth, affect only the distance between the age-profiles of earnings for different cohorts, not the age profiles themselves. Income is earnings plus the interest income on accumulated wealth, so that, given that (log) consumption and (log) earnings can both be decomposed into cohort and age effects, so can income y_{iab}. Corresponding to (3), we can write:

$$\overline{\ln y} = \boldsymbol{D^a} \alpha_y + \boldsymbol{D^b} \gamma_y + u_y \qquad (4)$$

where α_y and γ_y are the age and cohort effects in income. If consumption is close to income, the difference between (3) and (4) is approximately the saving ratio, so that

$$s/y \approx \overline{\ln y} - \overline{\ln c} = \boldsymbol{D^a}(\alpha_y - \alpha_c) + \boldsymbol{D^b}(\gamma_y - \gamma_c) + (u_y - u_c). \qquad (5)$$

Under the usual assumption that there are no bequests, and that lifetime consumption exhausts lifetime resources, the cohort effects in income and consumption will be the same, so that a restricted version of (5) will have only age effects and can therefore be rewritten as

$$s/y \approx \boldsymbol{D}^{a}(\alpha_y - \alpha_c) + (u_y - u_c). \tag{6}$$

Paxson (1996) estimates models of this form for four developed and developing countries, Taiwan (China), using annual survey data from 1976 to 1990, Thailand, using five surveys, 1976, 1981, 1986, 1988, 1990, and 1992, Great Britain, using annual surveys from 1970 to 1992, and the United States, using annual surveys from 1980 to 1992. Although our main focus here is with the developing countries, the growth–savings link is found in developing and developed countries alike, and the claims of the LCH to account for that link are not confined to any subset of developed or developing countries.

Estimation of the age and cohort effects in the unrestricted saving equation (5) does not yield sensible results. Figure 1 reproduces the age and cohort effects for the two countries, the United States and Taiwan, for which the problems are most dramatic. (This information is taken from Figures 11 and 12 of Paxson (1996)). The estimates of age effects ($\alpha_y - \alpha_c$), normalized to be zero at the youngest age, are graphed in the left-hand panel, and indicate that the profile of savings rates sharply declines with age in the United States, and rises with age in Taiwan. In neither case is there low saving in *both* youth and old age, as expected under the LCH. The cohort effects ($\gamma_y - \gamma_c$) in the right panel show earlier born cohorts saving more in the US, and saving less in Taiwan. (The *x*-axis of the Figure is age in the first year of the survey, so earlier-born cohorts are farther to the right.) The presence of cohort effects contradicts the no bequest assumption and, if taken literally, implies that Americans are consuming an ever larger share, and Taiwanese an ever smaller share of their lifetime resources. The cohort effects are much too large to be plausible. For example, they indicate that Taiwanese born in 1956 will bequeath 50 percentage points more of their life-time wealth than those born in 1936.

The key to interpreting these findings is the result that time trends in saving rates produce offsetting age and cohort profiles. A linearly decreasing (increasing) age profile and increasing (decreasing) cohort profile of the same slope is equivalent to a time trend in which the saving rate of all age groups is falling (rising) through time. The offsetting age and cohort effects in Figure 1 come from the decline in saving rates in the US data and the upward trend in saving rates in Taiwan. It is important to note that these offsetting age and cohort profiles should not be present if the LCH is giving a complete account of the data; the trends are supposed to be an aggregation effect associated with slower growth in the US and faster growth in Taiwan. Instead, Paxson's results show that, each American (Taiwanese) cohort is saving less (more) at each age than did its predecessors at the same age. If growth is causing saving, the mechanism is not macroeconomic, working through differential weighting of different age groups, but microeconomic in some way that we do not understand. Even without further analysis, these results cast considerable doubt on the proposition that changes in saving rates in the US and Taiwan are attributable to changes in growth through the mechanism postulated by the LCH.

Paxson shows that it is possible to obtain more sensible age profiles for saving rates by allowing for a time trend, i.e. by estimating (6) with the addition of a set of year dummies. While this procedure essentially concedes defeat in the attempt to explain the major features of the saving ratio by the LCH, it allows investigation of the LCH as a secondary cause. Indeed, the age profiles of saving in both the US and Taiwan now have the general shape predicted by the theory, with saving rates among the elderly up to 20 percentage points lower than those at younger ages. For Britain and Thailand, by contrast, there is little evidence of any well-defined age profile of saving rates. But even those estimated for Taiwan and the United States are insufficiently variable with age to allow growth rates to have much effect on aggregate saving. An increase in the per capita growth rate of income from 2 percent to 4 percent a year is predicted to increase the US household saving rate from 6.2 percent to 6.5 percent, the British saving rate by less than 0.1 percentage points, the Taiwanese rate from 15.9 percent to 16.2 percent, and to *decrease* the Thai saving rate from 17.6 percent to 17.5 percent. Even after allowing for the time trends, which are necessary to obtain sensible age profiles of saving at all, the variation of saving rates with age is simply not large enough—or even in the right direction—to allow growth to have much effect on savings through the aggregation effects postulated by the LCH.

The results reported by Paxson can be extended to any country where there are a sufficiently large number of household surveys. In principle, age and cohort effects can be identified from only two surveys, provided year effects are excluded, but because age, cohort, and year effects are only identified up to an arbitrary time trend, the use of two surveys makes the estimated age and cohort effects hostage to any abnormality in either of the two years. To eliminate such effects, three survey years are the theoretical minimum, though more are required for safety. The only other country where we have so far been able to obtain such data is Indonesia, where the consumption and income data were collected in the SUSENAS household surveys for 1980, 1984, 1987, 1990 and 1993, with about 50,000 households in each sample. For reasons that we have yet to unravel, the 1980 income figures are not usable, and are reported as non-missing for only about one eighth of the sample; as a result, we work with the four surveys from 1984 to 1993.

Figure 2, which corresponds to Figure 4 of Paxson (1996), shows averages of the logarithm of real income and real consumption for 17 cohorts defined by head's age in terms of three year age bands. Each connected plot follows one cohort for four observations over the nine year period, on the left for the logarithm of income, and on the right for the logarithm of consumption. This figure shows the standard patterns of cohort effects—the younger cohorts have higher real income and consumption—and age effects, wherby consumption and income first rise with age and then fall. There is also some evidence for year effects which show up as common patterns within the traces for the different cohorts. Figure 3 shows the age effects from regressions of the logarithms of income and of expenditure on a set of age and cohort effects without any constraint that the two sets of cohort effects γ_c and γ_y be the same. The left and right panels correspond to Figure 5 (age effects in consumption and income) and Figure 11 (age effects in the saving ratio) in Paxson (1996). In the latter, the savings rates monotonically decline with age, but as in the US, there are offsetting cohort effects that, together with the estimated age effects, mimic a declining time trend. The saving rates in the survey decline over time, although not linearly—the aggregate saving rate using the survey data is 11.7 in 1984, 9.8 in 1987, 8.6 in 1990, and 9.7 in 1993.

As with Taiwan and the United States, we concede the model's inability to explain the time trend, and attempt to estimate a more sensible age-profile for saving by regressing the approximate saving ratios on age dummies and year dummies, with cohort dummies excluded. The resulting age profiles are shown in the left panel of Figure 4. These decidedly strange profiles show the exact opposite of a savings hump; instead there is savings dip in middle age. However, the differences across ages in the savings rates are not very large, and the right hand panel shows the same graph plotted on the same scale as in Figure 14 of Paxson (1996), which gives the impression of an essentially flat age-profile of saving. These Indonesian results are therefore very close to those reported earlier for Thailand, as are the calculated effects of growth on saving. Higher growth in Thailand or Indonesia, working through the life-cycle mechanism, is estimated to *decrease* the saving ratio, albeit very modestly.

These Indonesian results, as well as those obtained earlier by Paxson (1996) are consistent with other studies from developed economies. For the US, Bosworth, Burtless and Sabelhaus (1991) had earlier found the result that the decline in the saving ratio is in the microeconomic data, and is not accounted for by the aggregation effects postulated by the LCH. Other international evidence comes from the NBER's project on international comparisons of household saving, Poterba (1994), which covers Canada, Italy, Japan, Germany, the United Kingdom, and the US. The results for the US and UK are consistent with those reported above, while those for Canada show both the matching cohort and age effects identified by Paxson for the US and Taiwan, as well as a generally rising age-profile of saving rates conditional on the time trends. The Japanese, German, and Italian studies do not have enough survey years to estimate reliable cohort and age effects, but all three sets of authors remark on high saving rates among the elderly. All of these studies, like Paxson's, are based on household data, and on the identification of cohorts by the age of the household head, and all appear to include as pension income the component that should rightfully be regarded as dissaving.

2. Selection effects on saving in household data

There is a literature on the effects of various kinds of selection on the measurement of wealth and saving. Most of it is concerned with the relationship between wealth and age among the elderly, and with how to test the prediction of the LCH that the elderly run down their assets during retirement. In an important paper, Shorrocks (1975) identified two sources of bias in using a single cross-section to examine whether or not assets decline with age. The first, labeled the "productivity" bias, comes from looking across ages and interpreting the result as if it were the behavior of a single individual or group of individuals. In particular, in a growing economy, the older people in a cross-section have had lower lifetime resources, and would be expected to have lower wealth than younger people, even if no decumulation is taking place. Indeed, Shorrocks constructed examples in which people accumulate steadily throughout life, but where, because of growth, any single cross-section shows a negative relationship between wealth and age. The second source of bias comes from mortality differentials that are correlated with wealth. For reasons that are far from being fully understood, mortality is higher among low wealth individuals; in the US, age specific mortality rates in late middle age can vary by as much as threefold between the bottom and top deciles of the wealth distribution, depending on what other factors are controlled for. As a result, an older sample tends to be one with higher levels of wealth. The mortality effect biases upwards the age-profile of wealth, while the productivity effect biases the age-profile downwards, at least in an economy with positive growth in per capita resources.

In the cohort studies discussed in Section 2, there is no productivity bias because cross-sections are not used directly, but only to follow cohorts through successive surveys. Indeed, this is one of the main advantages of the cohort approach, and the results typically confirm Shorrocks suspicions, that inferences about the age profiles of wealth and saving from single cross-sections are so badly biased as to be useless. However, mortality remains a problem. Cohorts become smaller as they age, and mortality almost certainly selects an increasingly wealthy subsample from the original birth cohort. The extent of this effect can be estimated from panel data that tracks the same households or individuals as they age and die. For the US, Jianakoplos. Menchik and Irvine (1989) and Attanasio and Hoynes (1996) have carried out such studies using the National Longitudinal Study of Mature Men (NLSO) and the Survey of Income and Program Participation (SIPP) respectively. Both studies find that differential mortality by wealth causes an understatement of the decline in wealth among the elderly, but there are other effects at work too. Jianakoplos et al find that panel attrition is as important as mortality, and that it is the richer households who are more likely to attrit, which imparts a bias that is in the opposite direction to that imparted by differential mortality. It is unclear in this study how the authors deal with changes in household membership, but Attanasio and Hoynes do so by working only with married couples, and by defining mortality to be the death of either spouse, at which point the household ceases to exist, at least in its original form.

In the US, where the elderly tend to live alone or with other elderly, mortality is relatively important as a cause of household dissolution. In developing countries, few elderly live alone, and relatively few as isolated elderly couples, so that mortality plays a relatively small part in explaining why cohort sizes fall with the age of the head among the elderly. Older men and women can cease to be household heads in many ways other than by dying, something that is also true in the US and other developed economies. Even when people are heads in one year and heads in the next, they may not be heads of the same household—for example if a divorce has taken place—and the probability of such outcomes is almost certainly conditioned by income, wealth, and saving.

When we are concerned with the truth of the life-cycle hypothesis, and in particular with testing its prediction that retirees run down their assets, it is clearly important to recognize and correct for mortality. However, this is not the case for assessing the effects of growth on saving, where what counts is the *actual* age-profile of saving, with the effects of mortality included, and not the *hypothetical* profile that would be the case if there were no wealth differentials in mortality. The dead make no contribution to aggregate saving. The situation is quite different when a household is dissolved because an elderly person moves in with a child, or more generally when the people who were originally in the household move to other households. These people are still alive, still receive income, and still consume, and their behavior will have an impact on the saving behavior of the receiving household. As with mortality, household dissolution of this nature is almost certainly non-random—for example, divorce is less likely at high income, and the viability of a household depends on its asset position—and should be taken into account when testing life-cycle theory. But unlike the case of mortality, we must keep track of the members of dissolved households and allow for their effect on aggregate saving.

In the remainder of this section, we document some of the characteristics of the Taiwanese data, focussing on the issues discussed above and in particular on the relationship between household and individual data. We have 20 years of data, from 1976 through to 1995, which is five years more than in our previous work, or in Paxson (1996). Figure 5 starts by following *individuals*, not households, through the various surveys. As usual in these cohort pictures, we connect up points for the same birth cohort where, in this case, the points show the estimates of the total number of people in each cohort as it ages. These are obtained by summing the sampling weights for all individuals of each age in each year, and following through the birth cohort. The changes from year to year are partly genuine changes in the cohort size, through death, emigration, and immigration, but also show sampling variability since the surveys do not track members of the cohort, but only sample from it. The most obvious feature of the graph is the dip in the population around the age of 20. This is not genuine, but is due to the fact that the military and other institutionalized population are not covered by the surveys: Taiwan has universal military service for males, and large numbers of young people live in dormitories for college students and workers. The postwar baby boom shows up as a marked expansion in the distance between the cohort traces in the middle of the diagram. Mortality is apparent among the older cohorts; for those born earlier than about 1930, the traces slope downwards as the number of survivors falls with age.

Figure 6 shows the fractions of people in each cohort at each age who are household heads, and shows that, as people age, they cease to be heads of household, even though they are still alive; in these Taiwanese data, mortality plays a relatively minor role in terminating household headship. Among those alive and in the survey, the fraction who are heads of households falls from a maximum of over a half around age 45 to about a third by age 60, and is less than 10 percent at around age 80. (The variability of the plots at high ages comes from the small numbers of survivors.) It should be noted that household headship is notoriously difficult to define in survey practice, and that, in the Taiwanese case, headship tends to be associated with being the main earner, as defined by the enumerator early in the interview. In consequence, the age pattern shown here is not the same in other countries; headship *rises* with age in the US, the UK, and Thailand, see Paxson (1996, Fig. 2). Previous applications of cohort techniques to the LCH in Taiwan and elsewhere have identified cohorts of households by the age of the household head, and Figure 6 makes it clear that such an identification does not accomplish its objective, which is to provide labels that enable us to identify the same unchanging cohort of households from one survey to the next. For example, a household headed by a 50 year old male could remain unchanged from one year to the next, say 1990 to 1991, but if a 31 year old son becomes the main earner in 1991, the household which was in the cohort of 1940 in 1990 would be allocated to the 1960 cohort in 1991. Even without the instability introduced by taking the main earner to be the household head, headship can change for many other reasons, so that cohorts defined by head's age are not in fact true cohorts at all, and changes in their behavior can have as much to do with changes in the population included in the cohort as with it does with changes in the behavior of a fixed subset of the population.

A major concern for life-cycle analysis with household cohort data is that survival, as a household head, not just as an individual, is conditioned by factors that also condition saving and wealth. One way to look at this is to look at the education of household heads, and to see whether those with higher education are more likely to survive as heads. Figure 7 shows the cohort tracks of years of schooling, and we have shown only every fifth cohort in order to clarify the presentation. The traces are lower at higher ages because the earlier born cohorts are less well educated. Of more interest here is what happens within cohorts, and in particular, whether schooling levels increase with age as the less educated fail to survive as heads. In fact, at higher ages, there is some evidence of exactly the opposite, of education levels falling. However, this comes from higher mortality of men, the consequential increase in the proportion of female heads at high ages, and the lower education of women. Figure 8 is drawn for males only, and shows the average education of heads relative to the average education of all males of the same age. Rather than show the raw cohort plots, we have regressed the ratios on age and cohort dummies (the latter are not important in this case), and show the estimated age effects. This gives what we might have expected, which is evidence of increasing selection with age toward the more educated members of each cohort. This is consistent both with an inverse relationship between mortality and education, and with a higher propensity of men with lower education to give up headship and live with their children.

This rather preliminary examination of headship in the Taiwanese data casts a good deal of doubt on the validity of using household cohorts as in our previous work. The probability of being a household head is linked to both age and education so that by constructing cohorts based on head's age, we are not sampling from an unchanging population, and our results are affected by selection into the population as well as by the behavior of the population itself. It is possible that the (non-standard) definition of headship in the Taiwanese surveys makes the data particularly susceptible to these problems, but headship is always difficult to define in survey practice —for example, it is often left to respondents to interpret—and we suspect that household cohorts in other countries are likely to be subject to similar problems.

3. Saving and age on an individual basis: methods

To construct a personal, not household, version of the life-cycle hypothesis, we suppose that each individual follows his or her own life-cycle trajectory, in which both consumption and income at each age is the product of an age-effect and a life-cycle wealth effect. Individuals are grouped into households in some way into which we do not inquire, and we observe, not the individual consumption and income levels, but household consumption and income, defined as the sums over the individual members. Households operate as "veils" that mask the life-cycle behavior of their individual members.

This version of the LCH has a number of unfamiliar features, and some genuine problems. That we treat household formation as an exogenous and largely irrelevant process is clearly unsatisfactory; as far as we are concerned, the household plays a role only in the measurement process. But the standard approach, which takes the household as the atomic unit, is no more satisfactory, and has equally little to say about how households form, who is head, and what happens when households form or disintegrate. And as we have seen in the previous section, we cannot define a constant population of households from which to sample, which is not a problem for individuals. A more remediable problem with our approach is the assumption that household consumption is the sum of the consumption of household members. While this is an obvious starting point, it makes no allowance for economies of scale, or other interactions between household members in consumption. The unfamiliar part of our story is the assumption that life-cycle trajectories are set from birth and persist to death, rather than being tied to household formation. Siblings may be differentially successful in life, and we model their consumption in their parental household as being determined by their individual lifetime resources, and not by that of their parents. This is a different way of thinking about life-cycles than the usual one—which typically has very little to say about children and their life-cycles—and it remains to be seen whether it is a useful one.

Our econometric procedures are inspired by the work of Chesher (1996a, b), who is also concerned with recovering individual behavior from household data, though we use different models and econometric techniques. We begin in the same way as before, with equation (1), although now the *i* subscript is clearly taken to refer to an individual, not a household. It is convenient to rewrite (1) in the form

$$c_{iab} \; = \; c_{ab} + \epsilon_{iab} \; = \; f(a) W_b + \epsilon_{iab} \tag{7}$$

where ϵ_{iab} is a mean zero error, and we are now decomposing the mean cohort consumption into an age effect $f(a)$ and a cohort effect W_b, interpretable as cohort average lifetime resources. (This is not identical to the original formulation, but is equivalent if, within each cohort, individual deviations from the age profile are independent of individual deviations from lifetime resources.) For a household *h*, included in the survey at time *t*, we observe household consumption c_{ht} which is the sum of individual consumption, so that

$$c_{ht} \; = \; \sum_{a=1}^{A} n_{aht} f(a) W_{t-a} + \sum_{i \in h} \epsilon_{iat-a} \tag{8}$$

where n_{aht} is the number of people aged *a* in household *h* at time *t*, where *A* is the maximum age in the population, and where we have used the fact that someone aged *a* and observed in *t* was born in $t - a$.

The age profile $f(a)$ and the cohort lifetime wealth levels W_b are, as before, estimated non-parametrically using dummy variables, but we must now proceed in two steps. At the first stage, define

$$\beta_{at} \; = \; f(a) W_{t-a} \tag{9}$$

the product of the age a age effect and the wealth effect for those of age a in year t. For a single cross section from year t, we estimate the regression, from (8),

$$c_{ht} = \sum_{a=1}^{N} n_{aht}\beta_{at} + v_{ht} \tag{10}$$

so that the β's are recovered by a cross-sectional regression of household consumption levels on the numbers of people in the household in each age from 1 to 99. This is done for each of the cross-sectional surveys from 1976 to 1995, so that we have 20 sets of β's, each with 99 elements. At the second stage, these β's are decomposed according to (9) by taking logarithms, pooling all 1,980 estimates, and regressing on a set of age and cohort dummies, which corresponds to the first (and only) stage when we are working with household level data. The procedure is repeated in exactly the same way for income, although we impose that the age effects in income are zero for those aged 15 or younger and age 80 and older. The reason for imposing these restrictions is that, without them, the age effects for very young and old people are close to zero (as they are expected to be), and are sometimes negative. Taking logarithms would require eliminating the negative values and bias our results. The cohort and age effects for the savings ratio are computed as the difference between the logarithmic income and consumption effects, again exactly paralleling the calculations in the household case.

Somewhat to our surprise, the β's estimated from the first stage are typically sensible, for both income and consumption. The (unsmoothed) lines in Figure 9 shows the estimated consumption and income β's for 1995; these estimates conflate age and cohort effects, so that the graph should not be interpreted as showing genuine age profiles. The consumption profile is low for young children, rises to age 45, and then falls with age. Among the elderly, there are increasingly few observations, and the estimates become quite imprecise, and include several negative values. At the other end, for children, the estimates are much more precise, but they are close to zero so that, even with small sampling variability, there are occasional negatives. For income, where we allow non-zero effects only for those older than 15, we get negative estimates only for the youngest people. (If income profiles are estimated for all ages, we get larger negative effects for the youngest children, presumably because of reductions in mother's labor supply. A fuller model of household consumption and income would permit us to handle such effects.)

In order to be able to implement the second stage of the estimation, which requires taking logarithms of the estimated β's, we have smoothed the raw estimates as shown in Figure 9, to eliminate unnecessary variance, and in the hope that, when the variance is reduced, the estimates will be consistently positive. The smoothed β's are defined by

$$\tilde{\beta}_{at} = \sum_{\alpha=1}^{99} K\left(\frac{\alpha - a}{h_a}\right)\hat{\beta}_{at} \Big/ \sum_{\alpha=1}^{99} K\left(\frac{\alpha - a}{h_a}\right) \tag{11}$$

where α runs over age from 1 to 99, $K(.)$ is a kernel—in this case a Gaussian unit normal density—and h_a is an age specific bandwidth. The optimal bandwidth in the sense of minimizing integrated mean square error, can be shown to be

$$h_a^{opt} = \left[\frac{B\sigma_\beta^2(a)}{[\hat{\beta}^{//}(a)]^2}\right]^{\frac{1}{5}} \tag{12}$$

where B is the integrated squared kernel

$$B = \int K(\eta)^2 \, d\eta \tag{13}$$

which, in the Gaussian case is $1/2\sqrt{5}$, $\sigma_\beta^2(a)$ is the variance of the estimated β for age a, and $\beta^{//}(a)$ is the second derivative of β regarded as a function of age a. The denominator of (12) makes the bandwidth smaller where the slope of $\beta(a)$ is changing rapidly, so as not to smooth over nonlinearities, while the numerator makes the bandwidth larger when the estimates are imprecise. Here, we are not much concerned with nonlinearities, so that we replace the denominator by a single number for all a, calculated by fitting a cubic in age to the estimated β's, and then evaluating the estimated second derivative, which is linear in age, at $a = 50$. The $\sigma^2_\beta(a)$ in the numerator is replaced by the estimated variances of the parameters in the first stage regressions. These variances are functions of the number of people in each cohort, so the smoother can be thought of as a weighted moving average of the unsmoothed points, with the number of points in the average increasing with age.

The smoothed estimates of the β's for 1995 are shown in Figure 9 together with the original estimates. It is important to note that the smoothing does not make large alterations to the pattern of the original parameters, the major effect being the elimination of the variability at high ages. Even so, the smoothing, by averaging over nearby ages, especially at young and old ages, eliminates all 21 negative β's for income and all but 9 of the 167 negative estimated β's for consumption, so that, after dropping 9 estimates, we can take logarithms.

4. Saving and age on an individual basis: results

Figure 10 shows for each of the 20 years the (smoothed) (log) consumption profiles for the individual method together with the consumption profiles estimated on a household basis. (The household profiles show the average of the log of consumption by the age of the household head.) Once again, note that these are cross-section profiles, not age profiles, to which we shall turn below. The individual profiles are shown only from age 20 to age 75 so that they can be compared with the household profiles, which are calculated only over that range because of the scarcity of very young and very old household heads. Because household consumption levels are the sum of individual consumption levels, the household profiles are above the individual profiles; we have not attempted to deflate by household size to put them on a comparable basis, if only because the graphs would then be difficult to distinguish. The main point of this figure is not the levels, but the differences in shapes of the household and individual profiles, and in particular, that the latter are much more hump shaped than the former. The individual profiles rise from age 20 to age 45 or so, and fall thereafter. By contrast, the household profiles are almost flat until age 50, and the subsequent fall is less. Households headed by 25 year olds (say) are not representative of the consumption levels of 25 year olds in general and, as we might expect, the 25 year olds who are heads have relatively higher consumption levels than the 25 year olds who are not, artificially raising the profile among the young, and flattening out the life-cycle profile. The same effect operates at a somewhat smaller scale among the elderly, flattening out the profile at the upper end. In our earlier work, we attributed the flatness of the cross-sectional consumption profiles to economic growth, the “productivity effect,” which moves up the consumption of the young relative to that of the middle-aged and elderly. To this effect we must now add the effect of selection into headship among the youngest households.

Figure 11 shows the estimated age effects in log income and log consumption, in the left-hand panel from the second-stage estimation from the individual data, and in the right-hand panel for the estimates obtained directly from the household data. In both sets of estimations, the cohort effects are restricted to be identical for log income and log consumption, so that there are no cohort effects in the saving rate. As always, the household age effects run from 20 to 75, while the individual effects are estimated from 1 to 99 for consumption, and from 16 to 79 for income. Both sets of plots show age-profiles of consumption and income that rise with age. For the household data, the income and consumption profiles are very closely related. This is the "tracking" phenomenon; since the cohort effects in income and consumption are (constrained to be) equal, the similarity of the age profiles is a consequence of the fact that, cohort by cohort, consumption and income are very closely related on a year to year basis. In the individual data, tracking is somewhat reduced; although the profiles still rise together, the two shapes are more clearly distinct, and there are periods both early and late in life when income is lower than consumption.

Figure 12 shows age profiles of the (approximate) saving rate and contains the key results of our analysis. In this figure, the results from the individual and household analyses are shown in both panels; for the former, age is the age of the individual, for the latter age is age of the household head. In the left-hand panel, the cohort effects are allowed to differ between consumption and income, so that there are cohort effects in the saving rates, while in the right-hand panel, the cohort effects are constrained, and there are no cohort effects in saving, corresponding to the results in Figure 11. The cohort effects are statistically significant, and the data strongly reject the hypothesis that there are no cohort effects in saving. In both individual and household data, later born cohorts save a larger share of their lifetime resources, and it is this effect—essentially an increase in bequests—that "explains" the increase in the saving rate in the aggregate data. This was the finding from the household data, and it is not altered by moving to an individual basis. Furthermore, for the cohort effects to generate a time trend—which is what appears to be in the data—they must be accompanied by a linear increasing trend in the age-profile. In consequence, when we eliminate the cohort effects, the age profile of saving is raised among the young, and lowered among the elderly, bringing it more into accord with the LCH, see the difference between the left-hand and right-hand panels of Figure 12.

Once we impose the restriction that there are no cohort effects—in violation of the data—we obtain an interpretation of the evidence that is very much in the spirit of the LCH—the rising saving rate is not explained by the rising importance of bequests (in fact it is not explained at all), and the age-profile of saving is hump-shaped—and, what is most important for our current purpose, the individual-based age profile of saving is a good deal more hump-shaped than is the household-based age-profile of saving. Of the two profiles in the right-hand panel, the individual profile rises from a sharply negative saving rate among teenagers to nearly 30 percent from the later 20s to 60 years of age falling back to negative values after about age 75. In contrast, the household profile is essentially flat over its range of estimation. The differences between the two profiles come more from the extended age range of the individual profile than from differences in estimates over the common range. Because the individual profiles are extended into childhood and old age, where there are no or few heads, and where saving is negative, and because both profiles must fit the data, the individual profile must be higher in middle ages than is the household profile. But otherwise, the shapes are similar over most of the range of the overlap. Nevertheless, this apparently minor change has major consequences when we come to do the growth or demographic accounting. If the individual profile is correct, and if we change the relative numbers or resources of different age groups, the implied household age-profile must change, and this effect is omitted if we do the accounting on a household basis, changing only the numbers and resources of households at different ages, holding the age-profile of household saving constant.

As we shall see below, accounting for this additional effect has a large effect on the predictions about the effects of growth on saving.

Before moving on to the growth accounting, it is worth looking at the age profiles of saving in another way which also provides a cross-check on whether they make sense. Figure 13 shows the wealth profiles with age at a range of real interest rates. These assume that each individual starts out with nothing and then accumulates according to the difference between consumption and income. These profiles actually start out by going negative, although it is hard to see because the amounts are small since both consumption and income are zero or close to zero for young children. At a real interest rate of 5 percent, wealth is exhausted at age 99; above or below that, someone who dies at 100 will leave either positive bequests or debts. Since few people live to such an age, and since 5 percent is a conservative interest rate for Taiwan, our results predict that there will be substantial bequests. Note that, even with the cohort restrictions imposed, so that we are making substantial concessions in favor of the LCH, there is little decumulation of wealth among the elderly. Indeed, according to these profiles, and depending on the interest rate, decumulation does not begin before about age 80.

The saving and growth results are computed using the stable population proportions at different rates of population growth, rather than the actual population in Taiwan now. Suppose that the rate of per capita economic growth is g, so that each cohort has $(1+g)$ times as much lifetime resources as its immediate predecessor. Hence

$$W_b = W_0(1+g)^b = W_0(1+g)^{(t-a)} \tag{14}$$

so that, for the corresponding income expression, we have

$$y_{ab} = (1+g)^{(t-a)} W_0 \exp(\alpha_{ay}) \tag{15}$$

where α_{ay} denotes the age effect in income for individuals (or, for the household model, for household heads) aged *a*. In the individual model, income is set to zero for those younger than 16 and older than 79. (We shall not complicate the notation by noting this exception again, but it is taken into account in all the calculations.) Consumption for cohort *b* at age *a* is

$$c_{ab} = (1+g)^{(t-a)} W_0 \exp(\alpha_{ac}) \tag{16}$$

The aggregate saving ratio at time *t* is therefore given by

$$\left(\frac{S}{Y}\right)_t = \frac{\sum_{a=1}^{A} \eta_{at}(1+g)^{-a}[\exp(\alpha_{ay}) - \exp(\alpha_{ac})]}{\sum_{a=1}^{A} \eta_{at}(1+g)^{-a}\exp(\alpha_{ay})} \tag{17}$$

where *A* is the maximal age in the population, and η_{at} is the fraction of households or of individuals aged *a* in the population at time *t*, depending on whether we are working with the household or individual version of the model. In the individual case,

$$\eta_{at} = \frac{(1+n)^{-a} p_a}{\sum_{\alpha} (1+n)^{-\alpha} p_{\alpha}} \tag{18}$$

where *n* is the rate of population growth and p_{α} is the probability of survival to age α, taken from Keyfitz and Flieger (1990). In the household case,

$$\eta_{at} = \frac{(1+n)^{-a} p_a^h}{\sum_{\alpha} (1+n)^{-\alpha} p_{\alpha}^h} \tag{19}$$

where p_{α}^h is the probability that someone survives to age *a* and is a household head at that age, and is calculated as described in Paxson (1996).

We are also interested in the derivative of the aggregate saving ratio with respect to the rate of growth g, since it is this number that must be compared with the slope of the line linking saving and growth in the international comparisons. From (17), we have

$$\frac{\partial\left(\frac{S}{Y}\right)_t}{\partial g} = \left(\frac{1}{1+g}\right)\frac{\sum_a a\left[\left(\frac{S}{Y}\right)_t - [\exp(\alpha_{ay}) - \exp(\alpha_{ac})]\right](1+g)^{-a}\eta_{at}}{\sum_a \eta_{at}\exp(\alpha_{ay})(1+g)^{-a}}. \tag{20}$$

The results of the growth calculations are shown in Figure 14 and Table 1. In the former, the four panels show the relationship between saving and growth for four different rates of population growth. There are two main points. First, there are strong interaction effects, so that the effects of economic growth depend on the rate of population; these are the “variable rate of growth” effects that have been emphasized by Fry and Mason (1982) and Mason (1988). When the rate of population growth is high, there are relatively large numbers of children, who are not saving, so that making them relatively rich has only a small or even negative effect on the aggregate saving rate. When population growth is slow, the traditional effect dominates, with higher growth of income redistributing income towards the high savers and causing saving rates to rise. The second point is that, at low rates of population growth, where the traditional growth to saving effect operates, the effect is large, much larger than the essentially zero effect in the household based model, and large enough to be a serious candidate for explaining the correlation between growth and saving in the aggregate data. The size of the effect is documented in the second column of Table 1; for example, with a growth rate of population of one percent a year, an increase in the rate of per capita economic growth from three percent to four percent would increase the aggregate saving rate by about two percentage points, which is close to the numbers that come out of the simplest stripped-down life-cycle model.

5. Saving and growth: conclusions

The results of this paper are a surprise to us, and have caused us to think more favorably about the possibility that the life-cycle hypothesis could be responsible for the relationship between growth and saving in the international aggregate data. Although we were aware that previous tests of the growth effect in the LCH were biased both by the omission of pension contributions and withdrawals, and by selection effects in household formation and dissolution, we are surprised by the difference that is made by our corrections for the latter. As it turns out, the biggest changes are brought about, not by correction to the shape of the age profiles of consumption, but by the recognition that the household age-profiles change along with changes in the age structure of the population—and thus of households—and with changes in the relative economic power of different age groups. If this is modeled as we do it here, and at least on these Taiwanese data, we get the possibility of large effects of growth on saving, at least at some rates of population growth.

As in any out-of-sample simulation study of this kind, the results are a mixture of assumption and evidence. We draw attention to our extreme—but extremely useful—assumption that each individual follows his or her own life-cycle path from birth to death, and that the level of that path is set by lifetime resources. The assumption has some (perhaps) implausible consequences. For example, children consume but do not earn. Adding children to a household will decrease household saving from what it otherwise would have been, something that is standard enough (although not necessarily well-documented empirically) in these sorts of models. But the individual life-cycle model also implies that the negative effect on household saving will be larger the larger is the rate of growth, because higher growth enriches children relative to their parents within the same family, and thus redistributes household resources towards the non-savers. The same is true when elderly people move in with their children; their consumption levels are determined by their own lifetime resources and not by those of their host family, an assumption that is perhaps more palatable than the corresponding assumption about children, and is in fact much more important in generating the result that growth can have a large effect on saving. Nevertheless, we would be a good deal more confident in our conclusions if there were any direct empirical corroboration of these effects, whether from Taiwan or elsewhere.

6. List of works cited

Attanasio, Orazio P., and Hillary W. Hoynes, 1996, "Differential mortality and wealth accumulation," NBER Working Paper No. 5126, May.

Bosworth, Barry, Gary Burtless, and John Sabelhaus, 1991, "The decline in saving: evidence from household surveys," *Brookings Papers on Economic Activity*, 183–241

Browning, Martin, Angus Deaton, and Margaret Irish, 1985, "A profitable approach to labor supply and commodity demands over the life-cycle, *Econometrica*, **53**, 503–43

Chesher, Andrew, 1996a, "Diet revealed? Semiparametric estimation of nutrient intake—age relationships," Department of Economics, University of Bristol, processed, Feb.

——, 1996b, "Individual demands from household aggregates: time and age variation in the composition of diet," Department of Economics, University of Bristol, processed, Apr.

Deaton, Angus and Christina Paxson, 1994, "Saving, growth, and aging in Taiwan," in David A. Wise, ed., *Studies in the Economics of Aging*, Chicago and London. Chicago University Press.

Fry, Maxwell J., and Andrew Mason, 1982, “The variable rate-of-growth effect in the life-cycle saving model: children, capital inflows, interest and growth in a new specification of the life-cycle model applied to seven Asian developing countries.” *Economic Inquiry*, 20, 426–42

Jianakoplos, Nancy A., Paul L. Menchik, and F. Owen Irvine, 1989, “Using panel data to assess the bias in cross-sectional inferences of life-cycle changes in the level and composition of household wealth,” in Robert E. Lipsey and Helen S. Tice, eds., *Measurement of savings, investment and wealth*, Chicago. Chicago University Press.

Keyfitz, Nathan and Flieger, Wilhelm, 1990, *World population growth and aging: Demographic trends in the late twentieth century*, Chicago and London, University of Chicago Press.

MaCurdy, Thomas, 1981, “An empirical model of labor supply in a life-cycle setting,” *Journal of Political Economy*, **89**, 1059–85.

Mason, Andrew, 1988, “Saving, economic growth, and demographic change,” *Population and Development Review*, **14**, 113–44

Paxson, Christina, 1996, “Saving and growth: evidence from micro data,” *European Economic Review*, **40**, 255-88

Poterba, James M., ed., 1994, *International comparisons of household saving*, National Bureau of Economic Research Project Report series, Chicago and London: University of Chicago Press.

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