GROWTH AND SAVING AMONG INDIVIDUALS AND HOUSEHOLDS

Angus Deaton and Christina Paxson*

Abstract—The lifecycle theory of saving and consumption predicts that changes in an economy's rate of economic growth will affect its aggregate saving rate by changing the lifetime resources of younger people relative to older people. However, studies that track the saving behavior of cohorts of household heads over time as they age have yielded estimates of age-saving profiles that are too flat for growth to have much effect on the aggregate saving rate. One problem with the cohort approach is that multigenerational households are common in many counties, and the age-saving profiles of households may be quite different from the age-saving profiles of individuals that make up households. In this paper, we propose a method for estimating individual age-saving profiles using household data. This method is applied to data from Taiwan and Thailand. We find that the individual method yields results that are more favorable to the lifecycle model. These results imply that changes in the rate of economic growth may in some circumstances have large effects on the aggregate saving rate. However, the size and sign of these effects depends on the rate of economic growth and the rate of population growth, and in many cases the effect of growth on saving is small.

I. Introduction

THE LIFECYCLE THEORY of saving and consumption predicts that changes in an economy's rate of economic growth will affect its aggregate saving rate. In the simplest version of the model—in which young people save for retirement and old people consume their previously accumulated assets—an increase in the rate of economic growth will unambiguously increase the aggregate saving rate, because it increases the lifetime resources (and saving) of younger-age groups relative to older-age groups. More complicated and realistic versions of the model yield ambiguous conclusions about the relationship between saving and growth. For example, young people may have low current income but high lifetime wealth, and may therefore borrow to finance current consumption. If they borrow enough, then, at sufficiently high rates of economic growth, their lifetime wealth will be high enough relative to that of their elders so that further increases in the rate of growth will decrease the aggregate saving rate. Whether higher growth increases or reduces the aggregate saving rate depends on whether the age profile of saving is negatively correlated with age, which is an empirical matter.

The results of recent studies of lifecycle saving behavior have not been favorable to the lifecycle model's interpretation of the positive correlation between growth and aggregate rates of saving. (See, for example, Poterba (1994), Paxson (1996), and Deaton and Paxson (1997).) Cohort studies in both developing and developed countries have used time series of cross-sectional household surveys to trace out consumption and saving behavior of birth cohorts, and have not found evidence of the required negative relationships between saving rates and age. Estimated age-saving profiles are typically fairly flat, so that increases in economic growth that redistribute resources toward the young will have either very small positive effects (or sometimes small negative effects) on the aggregate saving rate. This micro-based evidence is not consistent with international cross-country evidence that indicates that a 1 percentage point increase in the rate of per capita growth is associated with an increase of roughly two percentage points in the saving rate, which leads to the conclusion that the aggregate relationship between growth and saving is caused by something other than growth driving saving through the lifecycle mechanism.

There are (at least) two major unresolved problems with cohort methods of estimating age-saving profiles, one of which is explored in this paper. The first problem, which we do not deal with here, is that household surveys typically do not collect comprehensive information on contributions to private and public pension funds. Contributions made by employers on behalf of employees are usually missed altogether, so that the saving of those in employment is understated. Pension income is frequently aggregated with other income, which misclassifies the component of annuity disbursement that is not income but comes from running down the underlying asset, with the result that saving is overstated among the elderly. Calculations for Italy in Jappelli and Modigliani (1998) indicate that the age-saving profile is much more "hump shaped" after adjusting for flows into and out of pension funds. When this issue is ignored and computations are done in the standard way, the effect of growth on saving may be seriously understated. In this paper, we state this problem but make no contribution to its solution; our concern here is with middle-income and poorer countries in which 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. Pensions are less common in Thailand and Indonesia, which are the other two countries which we examine.

The second problem with the cohort approach lies in the way that cohorts are defined. Household surveys collect consumption for households, not individuals. Even for income, it is often difficult to attribute earnings to individuals, especially when the household's major source of income is a family business. Because measurement at the individual level is difficult, the usual approach is to work with household income and consumption—and to track cohorts

of households defined by the age of the household head, rather than by the age of individuals. This method might be acceptable if all people (or married pairs of same-aged people) became household heads at adulthood and remained so for the rest of their lives. However, in many countries, individuals often live in multigenerational households so that household income and consumption combine data for people at different lifecycle stages, obscuring the individual age profiles. For example, a 45-year-old household head may save for retirement, but if he lives with his 20-year-old son and 70-year-old mother, both of whom save negative amounts, the net saving for the entire household could be zero or negative. More generally, the "flat" age-saving profiles that have been estimated could be the result of combining consumption and income information across different-aged individuals within households.

In this paper, we propose a method of estimating "individual" age-saving profiles using household income and consumption data. We apply this method to data from Taiwan and Thailand, and compare our results with those based on the more standard "household-based" approach. We find that the individual method yields results that are more favorable to the lifecycle model than those from the household method. In both countries, individuals appear to save in middle age and to dissave at very young and very old ages. This general hump-shaped age pattern of saving is consistent with both positive and negative effects of economic growth on the aggregate saving rate. Our results illustrate the point. For Taiwan, the results indicate that increases in growth can result in large increases in saving rates, particularly when the rate of population growth is low. By contrast, the results for Thailand indicate that, for most combinations of rates of economic and population growth, increases in economic growth raise the wealth of the very youngest individuals, who are dissavers, causing a reduction in the aggregate saving rate.

Our method relies on the assumption that households serve as "veils" for individuals, each of whom consumes and saves according to his or her own lifetime wealth; household membership does not alter individual consumption or income, but only obstructs our ability to measure these quantities. This is obviously a simplification of the role of households in individual behavior. 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. It would also recognize that household formation is endogenous and that the behavior of those who choose to live in extended families is likely to be systematically different from those who live by themselves or with partners of similar ages. We do not estimate models that handle these extensions, and they remain important priorities for future work.

The paper is organized as follows. Section II briefly presents the basic theory and reviews existing (household-based) cohort studies from a range of rich and poor countries, and also presents new results for Indonesia. Section III presents a simple model that illustrates how using household-level aggregates can bias estimates of age-saving profiles and lead to incorrect inferences about the relationship between growth and saving. We also document the problems with defining cohorts by age of household head in countries such as Taiwan and Thailand. Section IV lays out the individual model and shows how its parameters can be recovered from household-level data. The model is estimated using data from Taiwan and Thailand and results are presented in section V. Section VI turns to growth accounting and examines the implications of our estimates of age-saving profiles for the relationship between economic growth and the aggregate household saving rate. Section VII concludes.

II. Saving and Growth: Basic Theory and Evidence to Date

We begin with an outline of the theory. In the standard approach, no clear distinction is drawn between a household and an individual (or, more precisely, the theory is developed for individuals and then applied to households). We begin with that individual theory. The literature on lifecycle saving and growth works with a model in which uncertainty is ignored, and in which consumption follows an age-profile determined by the interaction of preferences and real interest rates, with the level of the profile set by the level of lifetime resources. We use the index i to denote an individual, and a to denote age. Lifetime resources, W_i, are the sum of assets at birth (inheritance) and the discounted present value at birth of future labor earnings:

$$W_i = A_i^0 + \sum_0^L y_{ia}^l (1 + r)^{-a} \tag{1}$$

where L is the length of life,
r is the constant real interest rate, and
y_{ia}^l is labor income or earnings by i at age a.

According to the lifecycle hypothesis (LCH), consumption at age a is proportional to lifetime resources, where the constant of proportionality depends only on age and the (constant) real interest rate. Suppressing the latter, we write

$$c_{ia} = f_i(a) W_i. \tag{2}$$

The effects of economic growth on saving depend on how labor income is affected by growth. In the version we adopt in this paper, which follows Modigliani's original lead, labor income has an age profile that is invariant to changes in the growth rate of the economy, so that productivity growth affects only the distance between the age profiles of earnings for different cohorts and not the age profiles themselves. The assumption is in principle testable given cross-sectional data

on earnings profiles from spans of years in which growth rates are different, although finding such periods and enough data is likely to be difficult. For example, there is little difference in the slope of Taiwanese age-earnings profiles between the 1976–1985 and 1986–1995 periods, but growth rates during these two time periods were also similar. As has long been recognized, if individual earnings growth over life mirrors aggregate growth, with steeper age profiles at higher growth rates, and if young consumers are free to borrow, higher growth will increase borrowing among the young and must ultimate decrease aggregate saving. Given the Modigliani specification, we can conveniently write earnings as proportional to lifetime resources,

$$y_{ia}^l = g_i(a) W_i \tag{3}$$

where the shape of $g_i(a)$ is set by the age profile of earnings and is scaled so as to make equation (3) hold. If we define α_i^0 so that we can write initial assets (inheritances) as

$$A_i^0 = \alpha_i^0 W_i, \tag{4}$$

then the lifetime budget constraint, which requires that the present value of consumption equal the present value of earnings and assets, can be written as

$$\alpha_i^0 + \sum_{a=0}^{L} (1 + r)^{-a} [g_i(a) - f_i(a)] = 0 \tag{5}$$

which is independent of the scale factor W_i provided that α_i^0 is independent of W_i, which we will argue below. Assets evolve over the lifecycle in order to allow equation (2) and (3) to hold simultaneously; the asset evolution equation is

$$A_{ia} = (1 + r)(A_{ia-1} + y_{ia}^l - c_{ia}), \tag{6}$$

so that, cumulating, we have

$$A_{ia} = (1 + r)^a A_i^0 + \sum_{k=0}^{a} (1 + r)^k (y_{ik}^l - c_{ik}). \tag{7}$$

Total income, y_{ia}, is the sum of asset income, rA_{ia}, and labor income, so that we can write

$$y_{ia} = h_i(a) W_i \tag{8}$$

where the scale factor $h_i(a)$ is given from equation (7), (4), and (3) as

$$h_i(a) = g_i(a) + r(1 + r)^a \alpha_i^0$$
$$+ r \sum_{k=0}^{a} (1 + r)^k [g_i(a) - f_i(a)]. \tag{9}$$

Savings, s_{ia}, is the difference between total income and consumption. We work, not with savings itself, but with the difference between the logarithm of total income and the logarithm of consumption which, when the saving rate is small, is approximately equal to the saving rate. (We shall rely on this approximation only in the discussion; all calculations, in the estimation and in the calculations of the effects of growth on saving, are done without approximation.) From equation (3) and (8), the approximate saving rate is

$$s_{ia}/y_{ia} \approx \ln y_{ia} - \ln c_{ia} = \ln h_i(a) - \ln f_i(a), \tag{10}$$

which depends on age, the interest rate, and on idiosyncratic variation in tastes, but not on lifetime resources, nor on the growth of lifetime resources.

As developed so far, we have made no allowances for bequests, so that initial assets and α_i^0 should logically be zero. But it is straightforward to incorporate a bequest motive in which bequests, like consumption, are a fixed proportion of lifetime resources, so that initial assets will also be a fixed proportion of the resources of the donor, and therefore also a fixed (although different) proportion of resources of the beneficiary.

A time series of cross-sectional household surveys gives us the ability to track over time the average consumption and average income for different birth cohorts. If we denote year of birth by b, and average the logarithm of equation (2) over all individuals born in b, we can write

$$\overline{\ln c_{ab}} = \overline{\ln f(a)} + \overline{\ln W_b} \tag{11}$$

where the lines over the variables denote means. Equation (11) 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:

$$\overline{\ln c} = D^a \alpha_c + D^b \gamma_c + u_c, \tag{12}$$

where $\overline{\ln c}$ is a stacked vector of log consumption levels with elements corresponding to each cohort in each year,

D^a is a matrix of age dummies,

D^b is a matrix of cohort (year of birth) dummies,

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 $a + b$.

Corresponding to equation (12) and (8), we have an income regression

$$\overline{\ln y} = D^a \alpha_y + D^b \gamma_y + u_y, \tag{13}$$

where α_y and γ_y are the age and cohort effects in income. Subtracting equation (12) from (13) yields

$$s/y \approx \overline{\ln y} - \overline{\ln c} = D^a(\alpha_y - \alpha_c)$$
$$+ D^b(\gamma_y - \gamma_c) + (u_y - u_c). \qquad (14)$$

Under the assumption that bequests are a fixed proportion of lifetime resources and that lifetime consumption and bequests together exhaust lifetime resources, the cohort effects in income and consumption will be the same, so that a restricted version of equation (14), which corresponds to equation (10), will have only age effects and can be rewritten as

$$s/y \approx D^a(\alpha_y - \alpha_c) + (u_y - u_c). \qquad (15)$$

Models of the form of equation (12) through (15) have been estimated by a number of authors, always on the (forced) assumption that the individuals of the theory can be replaced by households, and individual age by the age of the household head. Paxson (1996) uses data from four developed and developing countries:

1. Taiwan (China), with annual survey data from 1976 to 1990;
2. Thailand, with six surveys from 1976, 1981, 1986, 1988, 1990, and 1992;
3. Great Britain, with annual surveys from 1970 to 1992; and
4. 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 (14) does not yield sensible results. Figures 1 and 2 show versions of Paxson's results (updated using data to 1995) for Taiwan, Thailand, and the United States, together with new results from Indonesia using the SUSENAS household surveys for 1984, 1987, 1990, and 1993, with approximately 50,000 households in each sample.

Figure 1 shows the estimated age effects ($\alpha_y - \alpha_c$). The age profile of savings rates sharply declines with age in the United States, by almost 70 percentage points from age 25 to age 70, and equally dramatically rises with age in Taiwan, by nearly 50 percentage points. Both Indonesia and Thailand have declining age effects, although the decline is more modest (10 points in Indonesia and 15 points in Thailand). In no case is there low saving in both youth and old age, as might be expected under the LCH. The cohort effects ($\gamma_y - \gamma_c$) in figure 2 show older cohorts saving more in the

FIGURE 1.—AGE EFFECTS IN LN (Y)-LN (C), COHORT EFFECTS INCLUDED (GRAPHED IN FIGURE 2)
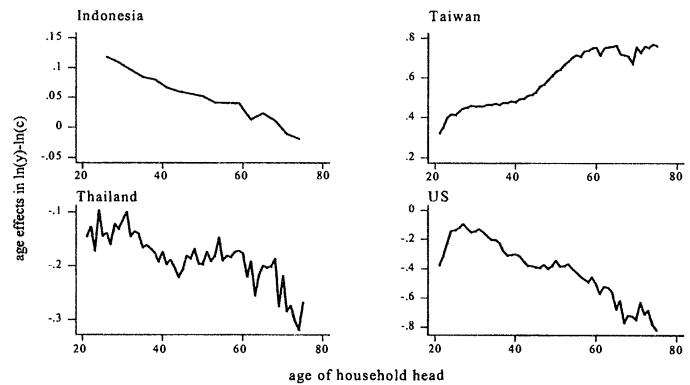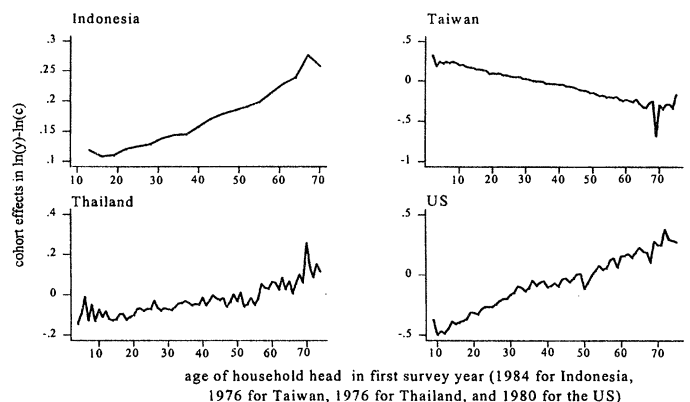
United States, Indonesia, and Thailand, and saving less in Taiwan. (The x axis of the figure is age in the first year of the survey, so older cohorts are farther to the right.) The presence of cohort effects contradicts the bequests as a fixed fraction of resources assumption and, if taken literally, implies that Americans, Indonesians, and Thais 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, especially in Taiwan and the United States. They imply that Taiwanese born in 1960 will bequeath 26 percentage points more of their lifetime wealth than their parents born in 1930, and that Americans born in 1960 plan to bequeath 20 percentage points less than those born in 1930.

The key to interpreting these findings is the fact that the graphs in figures 1 and 2 are approximately linear and that, within countries, the slopes of the age and cohort profiles are approximately equal and opposite. Note that the slopes would have the same sign if the cohort effects were graphed as a function of birth year. Write equation (14) in the form

$$s_{ab}/y_{ab} = \alpha_y(a) - \alpha_c(a) + \gamma_y(b) - \gamma_c(b), \qquad (16)$$

FIGURE 2.—COHORT EFFECTS IN LN (Y)-LN (C), AGE EFFECTS INCLUDED (GRAPHED IN FIGURE 1)

age of household head in first survey year (1984 for Indonesia, 1976 for Taiwan, 1976 for Thailand, and 1980 for the US)

where a is age and b is birth year. When the age and cohort profiles are linear and have slope β, say, equation (16) becomes

$$s_{ab}/y_{ab} = \beta^0 + \beta(a + b) = \beta^0 + \beta t \tag{17}$$

(because age plus year of birth is the calendar date t). The combination of a linearly decreasing (increasing) age profile and increasing (decreasing) cohort profile of the same slope is a time trend in which the saving rate of all age groups is falling (rising) through time. The offsetting age and cohort effects in figures 1 and 2 come from the large declines and increases in household saving rates in the U.S. and Taiwan, respectively, and the more modest downward trends in Indonesia and Thailand. (In contrast to the national aggregates, the household data in Indonesia and Thailand show decreasing saving rates over the relevant periods. For example, in Indonesia, the aggregate saving rate in the survey data is 11.7 in 1984, 9.8 in 1987, 8.6 in 1990, and 9.7 in 1993.) It is important to note that these offsetting age and cohort profiles should not exist if the LCH is giving a complete account of the data. According to the model, trends are supposed to be an aggregation effect associated with slower growth in the United States and faster growth in Taiwan. Instead, these results show that each American (Taiwanese) cohort is saving less (more) at each age than did his or her predecessors at the same age. If growth is causing saving, the mechanism is not macroeconomic, working by aggregation with 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 United States and Taiwan (or Indonesia or Thailand) 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 estimating not equation (14), with both age and cohort effects, but equation (15), in which the cohort effects are eliminated in spite of their statistical significance. (Year effects can also be included; although significant, they make little difference to the estimates of the age profile.) 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 resulting age profiles of saving shown in figure 3 are now a good deal more sensible, and in the United States and Taiwan have the general shape predicted by the theory, with saving rates among the elderly up to twenty percentage points lower than those at younger ages. For Indonesia and Thailand, by contrast, there is little evidence of any well-defined age profile of saving rates. But even the profiles estimated for Taiwan and the United States vary too little 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% to 4% a year is predicted to increase the U.S. household saving rate from 6.2% to 6.5%,

FIGURE 3.—AGE EFFECTS ON LN (Y)-LN (C), NO COHORT EFFECTS INCLUDED
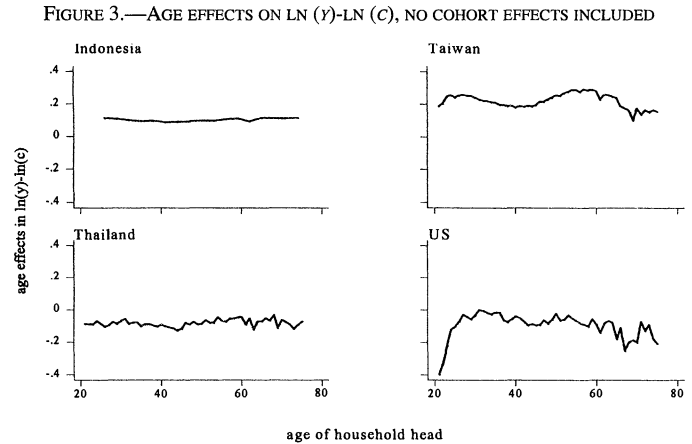
the Taiwanese rate from 15.9% to 16.2%, and to decrease the Thai saving rates from 17.6% to 17.5%. (See Paxson (1996).) The variation of saving rates with age is simply not large enough—nor always in the right direction—to allow growth to have much effect on savings through the aggregation effects postulated by the LCH.

These results are consistent with other studies from developed economies. For the United States, Bosworth, Burtless, and Sabelhaus (1991) have 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 United States. The results for the United States and United Kingdom are consistent with those reported above, while those for Canada show both the matching cohort and age effects identified by Paxson for the United States 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 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.

III. Problems with the Use of Household Data

While it is obviously true that households are not individuals, and that there is no reason to suppose that household behavior varies with the age of the household head in the same way that individual behavior varies with the age of the individual, these assumptions are largely forced upon us by the data, and it is not immediately clear what sort of biases they might introduce. In this section, we first document the differences between the head's age and individual age and then discuss how these differences can bias the estimated

FIGURE 4.—AVERAGE AGE OF HOUSEHOLD HEAD BY AGE OF INDIVIDUAL

age profiles of consumption, saving, and income, as well as the effects of growth on aggregate saving.

The first point to note is that, unlike individual age, head's age depends on the way headship is defined. There is no agreed standard. For example, in the Taiwanese survey, headship is associated with being the main earner; in Thailand, the household head is defined as "the person recognized as such by other members whether he was responsible for financial support or welfare of the household members or not" (National Statistical Office (1981)). In such circumstances, behavioral changes indexed on the age of the head are not only internationally noncomparable, but are hard to interpret within countries, because the change in behavior can be driven as much by changing composition of households as by changes in behavior by individuals as they age.

That there is no simple relationship between head's and individual age is apparent in the data. For example, approximately 53% of 40-year-olds in Taiwan are household heads, but the fraction declines with age thereafter to only about 23% at age 70. In Thailand, headship rates rise with age until much later, exceeding 60% at ages 60 to 69, and declining rapidly thereafter. Figure 4 presents the data from Taiwan and Thailand in a different way. These graphs plot, for each survey year, the average age of each individual's household head against the age of the individual. If everyone were a household head, head's age would equal own age, and the plot would be the 45-degree line. As it is, the plot lies above the 45-degree line for young people (young adults living with their parents), then runs more or less along the line until around age 40 in Taiwan and for longer in Thailand, and then falls below the line and flattens out (as the elderly increasingly live with their children). Even though the lines are higher in the later surveys, particularly in Taiwan, where more of the elderly are living independently in recent years, head's age remains very different from individual age.

The disconnect between head's and individual age raises two separate issues. The first can be thought of as bias in the age profiles; if the processes determining headship are correlated with wealth and saving, saving to age profiles will be biased by working with head's age instead of individual age. The second, which is independent of the first, is the question of whether, when growth rates change, the age profiles estimated from household data can be expected to be invariant to changes in growth rates.

We have good reasons to expect selection bias in the estimates of age profiles of saving, as there is some evidence that the elderly prefer to live by themselves when they can, so that high savers may be more likely to survive as independent households. If low savers move in with their children and high savers continue as heads, the age profile of saving will not decline with age in a way that reflects the average behavior of the elderly. It is also likely that other processes of family formation—such as migration, setting up as an independent household, marriage, and divorce—are conditioned on wealth and saving. Similar arguments are often put about mortality, where there is good evidence that the wealthy have lower probability of death. (See, for example, Jianakoplos, Menchik, and Irvine (1989) and Attanasio and Hoynes (1996) for evidence from the United States.) However, to assess the effects of growth on saving, we are interested only in those who are alive, and the mortality selected profile is the relevant one.

Even without any link from saving to headship, there remains the question of whether the household age profiles can be expected to remain constant as growth rates change. That this will not be the case is most obvious for changes in the structure of the population. Suppose, for example, that there is a decline in the rate of population growth so that, in the new demographic equilibrium, there are more elderly relative to others, and fewer children. The new population structure must somehow be reflected in new structures of households, which will in general imply that, conditional on head's age, the composition of households will be different. When the typical 40-year-old has two elderly people living in his or her household instead of one, the household saving rate will be lower, even if there is no change in the individual age profiles of saving. The same is true when demography is held constant but there is a change in the rate of economic growth that redistributes lifetime resources among individuals of different ages within the household.

The following model was constructed by Pierre-Olivier Gourinchas, and illustrates clearly how the household specification can cause difficulties. Suppose that each individual lives for 80 years, is born in year 0, goes to work at 15 until age 54, and retires from year 55 to year 79. Assume that there is no population growth, so that there is a uniform distribution of ages. Consumption is constant throughout life, and the interest rate is zero, so that consumption in each year of life is half of earnings in each year of life, $c_b = 0.5y_b$ for labor income y_b indexed on year of birth b. Lifetime resources W_b is $40y_b$, so that we can also write the consumption and age profiles in terms of age a as

$$c_a = \frac{1}{80} W_b = \gamma_a^c W_b, \qquad a = 0, \ldots, 79$$

$$y_a = \frac{1}{40} W_b = \gamma_a^y W_b, \qquad a = 15, \ldots, 54$$

$$= 0, \qquad \text{otherwise.}$$

(18)

FIGURE 5.—COMPARISONS OF THE RELATIONSHIP BETWEEN SAVING AND GROWTH, USING THE HOUSEHOLD AND INDIVIDUAL METHODS

Suppose that earnings grows at rate g, so that $y_b = y_0 (1 + g)^b$, and lifetime resources also grows at rate g across cohorts. If we normalize initial lifetime resources to be unity, and write G for $(1 + g)$, we have $W_b = G^b$. Aggregate consumption over all individuals in the economy is given by summing over all individuals currently alive, so that, at time t,

$$C_t = \frac{1}{80} \sum_{k=0}^{79} G^{t-k} = \frac{G^t}{80} \frac{1 - G^{-80}}{1 - G}.$$ (19)

Total income (which with a zero interest rate is total labor income) is given by

$$Y_t = \frac{1}{40} \sum_{k=15}^{54} G^{t-k} = \frac{G^t}{40} \frac{G^{-15} - G^{-40}}{1 - G}.$$ (20)

The aggregate saving rate is therefore given as a function of the growth rate by

$$\frac{S}{Y} = \frac{(1 - G^{-40}) - 0.5\,(G^{15} - G^{-65})}{1 - G^{-40}},$$ (21)

the graph of which is the inverted U in figure 5. At low rates of growth, higher growth redistributes lifetime resources from old dissavers to younger savers, and increases national saving. But at high enough growth rates, the elderly are sufficiently (relatively) poor not to determine the result, and higher growth redistributes resources away from middle-aged savers towards dissaving children. Note that in this example, we assume that children consume as much as adults, which is why the saving to growth relationship turns down at such a low growth rate.

Suppose now that we rearrange these people into households. We assume that each person moves out of the parental home at age 25, immediately bears a single child who also remains at home until age 25, earning from age 15. At age 65, people move in with their 40-year-old child and remain there until death. Households therefore live for 40 years,

during which the head ages from 25 to 64. The household lifecycle has four regimes. In regime 1, the head is aged 25 to 39, and there is a working adult and a nonworking child. In regime 2, the head is 40 to 49, there is a working head, a working child, and a retired elderly person. In regime 3, the head is 50 to 54, there is a working head and a retired elderly person. Finally, in regime 4, the head is aged 55 to 64, lives alone, and has no earnings. We assume that households act only as veils so that each person behaves exactly as they do in the individual model, and that household behavior is simply the aggregate of their individual behavior. However, when we come to do the analysis, we measure age by head's age, and resources by household resources. On this basis, household total resources are $35y_b + 5y_{b+25}$, from the 35 years of work by the head and the five years of work by the child. In each regime, we can then calculate total household consumption and income (earnings) as a ratio of total household lifetime resources, which gives us the age profiles, or γ's, for household consumption and earnings. This is not worth laying out for each regime, and we illustrate only for regime 2, in which there is an adult, an elderly person, and a child. Second regime household consumption and income are given by

$$c_2^h = \frac{1}{80} (40y_0 G^b)(1 + G^{25} + G^{-25})$$

$$y_2^h = y_0 G^b(1 + G^{25}).$$ (22)

In the first expression, household consumption is the consumption of the head (who is born in b) which is given by the product of the first two terms, plus the consumption of the child and the parent, who are 25 years younger and older, respectively. Similarly, income is the sum of the adult's income and the child's income. If we express these as fractions of household lifetime resources, we obtain the corresponding components of the age profiles:

$$c_2^h = \frac{1}{80} \frac{(40y_0 G^b)(1 + G^{25} + G^{-25})}{y_0 G^b(35 + 5G^{25})} W^h = \tilde{\gamma}_2^c W^h$$

$$y_2^h = \frac{y_0 G^b(1 + G^{25})}{y_0 G^b(35 + 5G^{25})} W^h = \tilde{\gamma}_2^y W^h.$$ (23)

If we do this for the other three regimes, we get a compete household age profile, which can then be aggregated over all households to give the relationship between growth and the aggregate saving rate for the economy as a whole.

Because the households are simply a veil for the individuals within them and because in our model the existence of households has no effect on anyone's consumption, the "individual" and "household" economies are identical, as is the relationship between saving and growth in each. Indeed, if the aggregation procedure outlined in the previous paragraph is followed through, we will eventually get back to

equation (21). But that is not what will happen if an econometrician analyzes the household data to recover the age profile by age of head. Such an exercise will recover the age profiles of consumption and income—the $\tilde{\gamma}^c(a)$ and $\tilde{\gamma}^y(a)$ parameters in equation (23)—and, then, in a false analogy with individual profiles, will treat them as constants when calculating the relationship between saving and growth. But, in fact, as equation (23) illustrates, the γ's are a function of the growth rate, because growth changes the relative economic positions of people in the household according to their age. Exactly what happens when we follow through the household method depends on the rate of growth in the economy for the period originally analyzed, and three possibilities are shown in figure 5, corresponding to growth rates of 1%, 1.5%, and 2%. These three curves are quite different from the true curve, equation (21), and quite different from one another, showing a positive relationship with growth at 1%, no relationship at 1.5%, and a negative relationship at 2%. Note that each of these curves intersects the true curve at the growth rate used in its construction. But it is clear from the figure that, if households truly act as a veil, the relationship between growth and saving constructed from the household data may not be at all informative about the relationship between growth and saving in the economy.

IV. Specification and Estimation of an Individual Lifecycle Model

We now take up the challenge of proposing an individual version of the lifecycle model that is empirically tractable and that we can apply to the data for Thailand and Taiwan. As in section III, we suppose that each individual follows his or her own lifecycle trajectory, in which both consumption and income at each age is the product of an age effect and a lifecycle 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 lifecycle 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, nor for other interactions between household

members in consumption. It is not difficult to think of other, richer specifications, although all of those with which we have experimented present serious econometric challenges that remain to be solved. We think of the version of the model presented here as a first step away from the household version of the model which we hope will serve as a base from which to make further exploration.

Our econometric procedures are inspired by the work of Chesher (1997, 1998), who is also concerned with recovering individual behavior from household data, although 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 equation (2) in the form

$$c_{iab} = c_{ab} + \epsilon_{iab} = f(a)\,W_b + \epsilon_{iab}, \tag{24}$$

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 it 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

$$c_{ht} = \sum_{a=1}^{A} n_{aht} f(a)\,W_{t-a} + \sum_{i \in h} \epsilon_{iat-a} \tag{25}$$

where n_{aht} is the number of people aged a in household h at time t,

A is the maximum age in the population, and

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 nonparametrically using dummy variables, but we must now proceed in two steps. At the first stage, define

$$\beta_{at} = f(a)\,W_{t-a} \tag{26}$$

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 equation (25),

$$c_{ht} = \sum_{a=0}^{N} n_{aht}\beta_{at} + v_{ht}, \qquad t = 1 \ldots T, \tag{27}$$

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 0 to 99. These β's are our estimates of "individual" consumption. We also estimate

equation (27) using household income as the dependent variable to obtain measures of "individual" income; in this case, we restrict the coefficients to be zero for ages 80 and above. The restriction cannot be rejected for all but a few years of the data, and attempts to estimate the income coefficients beyond gave unacceptably noisy estimates. At the second stage, these measures of individual consumption and income are decomposed according to equation (26) into age effects and cohort (wealth) effects. Provided that the β's are positive, this decomposition can be performed by pooling the β's from all of the years, taking logarithms, 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. However, we do not restrict the β's to be positive when we estimate equation (27), and in practice they are sometimes negative for those at very young and very old ages. We therefore need an estimation technique that works in general.

The method works as follows. First, suppose that we have estimates of consumption for C birth cohorts at each of the 100 ages, running from 0 to 99. The β's from equation (27) can then be stacked into a $C \times 100$ matrix (denoted B), the rows of which represent different birth cohorts and the columns of which correspond to each of the 100 possible ages. Equation (26) can be rewritten in matrix form as

$$B = c\alpha', \tag{28}$$

where c is a $C \times 1$ vector of cohort effects, and α is a 100×1 vector of age effects. The vectors c and α can be estimated by minimizing the sum over all cohorts and ages of the squared residuals between the β's and the product of the age and cohort effects or, equivalently, by minimizing the trace of

$$(B - c\alpha')'(B - c\alpha'), \tag{29}$$

subject to a suitable normalization of c or α. With the normalization that $c'c = 1$, it is straightforward to show that c is estimated as the eigenvector of $B'B$ corresponding to the largest eigenvalue. The estimate of α is $B'c$.

In practice, this general method cannot be implemented exactly as described above because we do not observe all cohorts at all ages. For example, in the Thai data, we observe the cohort of those born in 1950 only between the ages of 26 (in 1976, the first year of the survey), and 42 (in 1992, the last year of the survey). Furthermore, because the survey was conducted only six times over a 22-year period, we observe the cohort only at four ages between 26 and 42. Many elements of the matrix B are therefore missing. The problem is solved using an iterative estimation procedure, in which the missing values in the matrix B are filled in using the values predicted by estimates of c and α from the previous iteration. More specifically, let β_{ab} represent the individual consumption (or income) for people aged a who were born in year b, α_a denote the age effect for people aged

a, and c_b denote the cohort effect for people born in year b. The estimation procedure described above is equivalent to minimizing:

$$\sum_{b=1}^{C} \sum_{a} (\beta_{ab} - c_b\alpha_a)^2 \tag{30}$$

where the ages are summed only over the appropriate range for each birth cohort. Define β_{ab}^* to be equal to β_{ab} if β_{ab} is observed, and to $c_b\alpha_a$ if it is not observed. Then equation (30) can be written as a sum over all possible cohorts and ages:

$$\sum_{b=1}^{C} \sum_{a} (\beta_{ab}^* - c_b\alpha_a)^2 \tag{31}$$

with no change in the resulting solution. Minimizing equation (31) is equivalent to minimizing the trace of equation (29), provided that the matrix B contains the elements β_{ab}^* rather than β_{ab}.

Our procedure is to obtain starting values of the age and cohort effects using the standard log-linear decomposition, after setting any negative values of individual consumption and income to a low value of 1 (so that we can take logs). These estimates are converted from logs to levels to obtain initial estimates of c and α, which are used to calculate the elements of B for pairs of cohorts and ages that are not observed. We then estimate c and α as described above, and again use these estimates to recalculate and replace the elements of B that are not observed. This is repeated until the estimates converge.

In our empirical work, we impose the restriction that the cohort effects in income and consumption are identical. Estimating the cohort and age effects for consumption and income subject to the restriction is done by defining B to be a $C \times 180$ matrix: The first 100 columns contain information on consumption of individuals from ages 0 to 99, and the last 80 columns contain information on income of individuals from ages 0 to 79. The vector α has 180 elements, the first 100 of which contain age effects in consumption and the last 80 of which contain age effects in income. Otherwise, the estimation procedure is as described above.

V.   Results from Thailand and Taiwan

The first step is to estimate equation (27) and the corresponding equation for income for each cross section for the two countries. There are twenty cross sections for Taiwan (from 1976 to 1995) and six for Thailand (1976, 1981, 1986, 1988, 1990, and 1992). Figure 6 plots the estimated parameters: for Taiwan in the top two panels, for Thailand in the bottom panels, on the left for consumption (showing only the coefficients up to age 79) and on the right for income.

FIGURE 6.—ESTIMATES OF AGE-SPECIFIC INCOME AND CONSUMPTION (B_{AT}'S) BY SURVEY YEAR
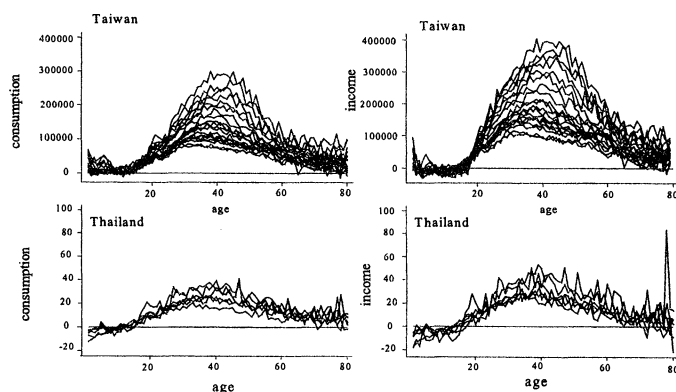
Because these estimates conflate age and cohort effects, the graphs should not be interpreted as longitudinal age profiles. The consumption profiles are low for young children, rise to approximately age 45, and then fall with age. Among the elderly, there are increasingly few observations, and the estimates become relatively 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 negative values. For income, and especially for Thailand, the coefficients are frequently negative for young children. In this individual version of the lifecycle model, the reduction in parental earnings associated with young children can be captured only by assigning them negative income. Even if we do not want to treat this literally, these estimates will ensure that, when we calculate saving rates with different growth rates or with different demographic structures, the effects that work through reductions in parental earnings will be accounted for. Lower fertility will raise household earnings, and, the richer are children in lifetime resources, the more their care reduces household earnings.

The difference between the Taiwanese and Thai profiles comes from the much more rapid rate of growth of household incomes and consumption in Taiwan. Because incomes and consumption are so much higher in the later than in the earlier surveys, and because the consumption of children and of the elderly remains close to zero, the Taiwanese curves are stretched further and further up as time passes.

Figure 7 shows the estimated age profiles for savings (or, more precisely, the differences in the logarithms of the α's of income and consumption estimated from equation (31) for the individual method, and the differences of the estimates themselves in the household case) for the two countries. These individual age effects are shown together with the household age effects estimated from the original household methodology, and shown previously in figure 3. Because the individual profiles are estimated from the occasionally noisy estimates of the β's, they are themselves somewhat noisy,

especially in Thailand where there are many fewer surveys. However, we have resisted the urge to smooth them; we do not need smooth estimates to calculate the growth effects, and it is easy enough to imagine smoothed versions of the profiles in the figure. The major result here is that the age profiles computed on the individual basis are a good deal more in accord with the lifecycle story than are the age profiles computed on the original household basis. In particular, saving rates (or, more precisely, the difference between the logarithms of income and consumption) decline sharply and become negative after about age 60 and show evidence of rising throughout the work years, gradually in Thailand and more sharply between 20 and 25 in Taiwan. Adding in the negative saving effects for those younger than 20 (not shown in the figures because they would dominate the scale) would further emphasize the hump-shaped saving profiles from the individual approach. By contrast, the household saving-to-age profiles are essentially flat throughout the lifecycle.

However, before concluding that switching from households to individuals validates the lifecycle approach, it must always be kept in mind that the individual profiles, like the household profiles, are computed under the restriction that there are no cohort effects in saving rates, a restriction that is strongly rejected in the data. As a result, the trends in actual saving rates in both countries are mechanically accounted for by time effects, and are left unexplained by either version of the lifecycle model. The individual method yields something like a lifecycle version of saving rates, but only once we have abandoned the attempt to explain the major features of the data.

VI. Saving and Growth

The age profiles of saving, using both the individual and household methods, can be used to examine how growth affects the aggregate saving rate. As before, 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,

$$W_b = W_0(1 + g)^b = W_0(1 + g)^{(t-a)}, \qquad (32)$$

so that, for the corresponding income expression, we have

$$y_{ab} = (1 + g)^{(t-a)} W_0 \gamma_{ay} \qquad (33)$$

FIGURE 7.—AGE PROFILES OF LN (Y)-LN (C), USING THE HOUSEHOLD AND INDIVIDUAL METHODS

where γ_{ay} denotes the age effect in the *level* of income for individuals (or, for the household model, for household heads) aged a. The γ_{ay} are as estimated in the individual model, although the income coefficients are set to zero for those older than 79; we do not complicate the notation by noting this exception again, but it is taken into account in all the calculations. When we use the household method, γ_{ay} is the exponent of the age effects estimated in logarithms. Consumption for cohort b at age a is

$$c_{ab} = (1 + g)^{(t-a)} W_0 \gamma_{ac}, \tag{34}$$

where the age effects for the different countries and models are defined as for income. The aggregate saving ratio at time t is therefore given by

$$\left(\frac{S}{Y}\right)_t = \frac{\sum_{a=1}^{A} \eta_{at}(1 + g)^{-a}[\gamma_{ay} - \gamma_{ac}]}{\sum_{a=1}^{A} \eta_{at}(1 + g)^{-a}\gamma_{ay}}, \tag{35}$$

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.

The fractions η_{at} are calculated under the assumption that the population is in demographic equilibrium, so that, instead of using the actual age distribution of the population from survey data, we use the age distribution that is implied by a given rate of population growth and a fixed set of probabilities of surviving to each age. We do this so we can examine how changes in the rate of population growth affect the relationship between growth and aggregate saving. In the individual case,

$$\eta_{at} = \frac{(1 + n)^{-a} p_a}{\sum_a (1 + n)^{-a} p_a} \tag{36}$$

where n is the rate of population growth and p_α is the probability of survival to age α. For Taiwan, the survival probabilities are from the Taiwanese life table in Keyfitz and Flieger (1990). For Thailand, the life table is from NSO (1997). In the household case,

$$\eta_{at} = \frac{(1 + n)^{-a} p_a^h}{\sum_a (1 + n)^{-a} p_a^h} \tag{37}$$

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 by Deaton and Paxson (1997).

FIGURE 8.—RELATIONSHIP BETWEEN AGGREGATE SAVING RATE AND THE RATE OF ECONOMIC GROWTH, FOR DIFFERENT VALUES OF THE RATE OF POPULATIONS GROWTH (N).

rate of economic growth
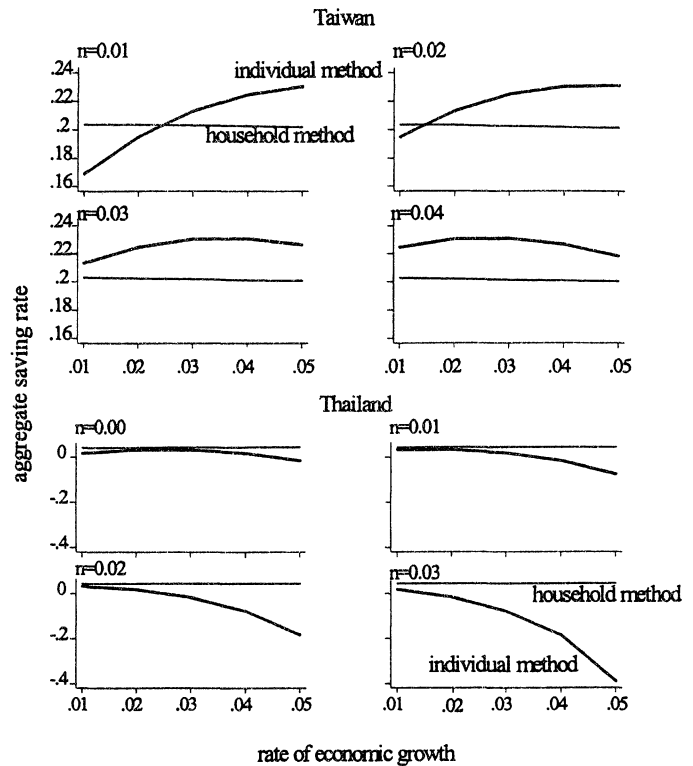
We are also interested in the derivative of the aggregate saving ratio with respect to the rate of growth g, since it is this number that must be compared with the slope of the line linking saving and growth in the international comparisons. From equation (35), we have

$$\frac{\partial \left(\frac{S}{Y}\right)_t}{\partial g} = \left(\frac{1}{1 + g}\right) \times \frac{\sum_a a\left[\left(\frac{S}{Y}\right)_t - (\gamma_{ay} - \gamma_{ac})\right](1 + g)^{-a}\eta_{at}}{\sum_a \eta_{at} \gamma_{ay}(1 + g)^{-a}}. \tag{38}$$

The results of the growth calculations for Taiwan and Thailand are shown in figure 8, which shows the relationship between saving and growth for four different rates of population growth, one in each panel. Tables 1 and 2 show the saving rates and derivatives of the saving rate with respect to growth for Taiwan and Thailand. There are three main points. First, there are strong interaction effects, so that the effects of economic growth depend on the rate of population growth; 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

TABLE 1.—SAVING RATES AT DIFFERENT RATES OF ECONOMIC GROWTH (g) AND POPULATION GROWTH (n), TAIWAN

Rate of Economic Growth (g):	S/Y		$\partial(S/Y)/\partial g$	
	Household	Individual	Household	Individual
	$n = 0.01$			
0.01	0.204	0.169	−0.014	3.031
0.02	0.203	0.195	−0.035	2.202
0.03	0.203	0.213	−0.050	1.476
0.04	0.202	0.225	−0.059	0.846
0.05	0.202	0.230	−0.062	0.300
	$n = 0.02$			
0.01	0.203	0.195	−0.036	2.224
0.02	0.203	0.213	−0.051	1.484
0.03	0.202	0.225	−0.060	0.842
0.04	0.202	0.230	−0.063	0.288
0.05	0.201	0.231	−0.061	−0.191
	$n = 0.03$			
0.01	0.203	0.213	−0.051	1.506
0.02	0.202	0.225	−0.060	0.851
0.03	0.202	0.230	−0.063	0.286
0.04	0.201	0.231	−0.061	−0.201
0.05	0.201	0.227	−0.055	−0.623
	$n = 0.04$			
0.01	0.202	0.225	−0.061	0.871
0.02	0.202	0.230	−0.064	0.294
0.03	0.201	0.231	−0.062	−0.203
0.04	0.201	0.227	−0.056	−0.633
0.05	0.200	0.218	−0.047	−1.009

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 is more likely to dominate, with higher growth of income redistributing income towards the high savers and causing saving rates to rise. Second, at low rates of population growth, the traditional growth-to-saving effect operates in Taiwan, and 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. With a population growth rate of 1% a year and a growth rate of per capita economic growth of 3%, an increase of one percentage point in the rate of economic growth (from 3% to 4%) would increase the aggregate saving rate by about 1.2 percentage points. The effect at initially lower rates of economic growth are larger, implying that an increase in the rate of economic growth from 2% to 3% would increase the aggregate saving rate by nearly two percentage points, which is close to the number that comes out of the simplest stripped-down lifecycle model. In Thailand, even at low rates of population growth, there is little or no positive effect of growth on saving, and, at high rates of population growth, higher economic growth drives down the aggregate saving rate.

The third point to note is that the effects of growth on saving are quite different when using the individual and household methods. The household method, which is based on an age-saving profile that does not capture dissaving at old and young ages, shows close to a zero effect of increased growth on saving. (The household results differ from those in Paxson (1996) for two reasons. First, for Taiwan, we have extended the number of surveys with which we work. Second, for both countries, we work here with stationary populations rather than actual populations.) The numerical example in figure 5 appears to be more than a theoretical curiosity.

While the results for Taiwan are consistent with a positive relation between growth and saving, the results from Thailand are not. Although the individual method yields a lifecycle saving pattern that has a pronounced lifecycle shape, it implies that higher economic growth will reduce the aggregate saving rate. This is true even at very low rates of population growth. The reason is that the negative effects of children on saving predominate. Our results indicate that children have a larger negative effect on saving in Thailand than in Taiwan. Furthermore, because survival probabilities to old age are lower in Thailand, the dissaving of the elderly is given less weight. Increases in the rate of economic growth that shift resources toward the young are therefore predicted to reduce the rate of aggregate saving. Although this result is consistent with lifecycle theory, it is not consistent with the observed (positive) relationship between growth and saving observed in the international aggregate data.

TABLE 2.—SAVING RATES AT DIFFERENT RATES OF ECONOMIC GROWTH (g) AND POPULATION GROWTH (n), THAILAND

Rate of Economic Growth (g):	S/Y		$\partial(S/Y)/\partial g$	
	Household	Individual	Household	Individual
	$n = 0.00$			
0.01	0.042	0.018	0.089	2.307
0.02	0.043	0.033	0.073	0.756
0.03	0.044	0.033	0.059	−0.713
0.04	0.044	0.018	0.046	−2.299
0.05	0.045	−0.014	0.035	−4.308
	$n = 0.01$			
0.01	0.043	0.033	0.074	0.748
0.02	0.044	0.033	0.059	−0.751
0.03	0.044	0.018	0.046	−2.375
0.04	0.045	−0.016	0.035	−4.447
0.05	0.045	−0.075	0.026	−7.603
	$n = 0.02$			
0.01	0.044	0.033	0.060	−0.758
0.02	0.044	0.017	0.046	−2.416
0.03	0.045	−0.017	0.035	−4.541
0.04	0.045	−0.077	0.026	−7.802
0.05	0.045	−0.182	0.019	−13.871
	$n = 0.03$			
0.01	0.044	0.018	0.047	−2.422
0.02	0.045	−0.017	0.035	−4.585
0.03	0.045	−0.078	0.026	−7.920
0.04	0.045	−0.185	0.019	−14.179
0.05	0.045	−0.389	0.013	−29.309

Note finally that all the results in this section are calculated by comparing situations in which the rates of economic and population growth are assumed to have been constant for a long time. In reality, it takes many years for demographic equilibrium to be reestablished after (for example) a drop in fertility, and transitional effects can be quite different from those obtained by comparing equilibria. In Deaton and Paxson (2000, forthcoming) we use the individual model estimated here to assess the contribution that demographic change in Taiwan has made to the increase in its saving ratio and find that, although the effects are in the right direction, they account for only a small fraction of the increase.

VII. Conclusions

We were surprised by the differences between the results of this paper and those from our earlier work. 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 the potential difficulties with the household approach, we had not anticipated 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 of income, consumption, and saving 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. Although our results indicate that growth may have large effects on aggregate saving, the size and sign of these effects differs across the two countries we study. We conclude that the positive relationship between growth and saving that is observed in international aggregate data is far from obviously consistent with the micro evidence.

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 lifecycle path from birth to death, and that the level of that path is set by the individual's lifetime resources. The assumption has some (perhaps) implausible consequences. For example, our results for both countries indicate that 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 lifecycle 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 nonsavers. It is this feature of the model that drives the negative relationship between growth and saving for Thailand. 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. 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, Thailand, or elsewhere.

It is also likely that our story of households as a veil for individual behavior is one that maximizes the differences between the household and the individual model, so that the truth may lie somewhere in between. People do not form households randomly, and people living together almost certainly do not consume, earn, and save in the same way that they would if living apart. Children do not really have negative incomes, and the existence of household economies of scale in consumption will change consumption patterns, as will the way in which resources are shared among household members. Intrahousehold sharing rules may themselves change as the rate of growth changes, so that the individual age profiles may change with growth rather than being fixed (as argued here). We see our model as only a starting point, albeit a tractable one, in the project of deconstructing the household for intertemporal analysis.

REFERENCES

Attanasio, Orazio P., and Hillary W. Hoynes, "Differential Mortality and Wealth Accumulation," NBER Working Paper No. 5126 (May 1996).

Bosworth, Barry, Gary Burtless, and John Sabelhaus, "The Decline in Saving: Evidence from Household Surveys," Brookings Papers on Economic Activity (1991), 183–241.

Chesher, Andrew, "Diet Revealed? Semiparametric Estimation of Nutrient Intake—Age Relationships," Journal of the Royal Statistical Society, Series A, 160:3 (1997), 389–420.

Chesher, Andrew, "Individual Demands from Household Aggregates: Time and Age Variation in the Composition of Diet," Journal of Applied Econometrics 14 (1998), 505–524.

Deaton, Angus, and Christina Paxson, "Saving, Growth, and Aging in Taiwan," in David A. Wise (Ed), Studies in the Economics of Aging (Chicago: Chicago University Press, 1994).

———, "The Effects of Economic and Population Growth on National Saving and Inequality," Demography 34:1 (1997), 97–114.

———, "Growth, Demographic Structure, and National Saving in Taiwan," Population and Development Review (2000, forthcoming).

Fry, Maxwell J., and Andrew Mason, "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 (1982), 426–442.

Gersovitz, Mark, "Saving and Development," Chapter 10 in Hollis Chenery and T. N. Srinivasan (Eds.), Handbook of Development Economics 1 (1988), 381–424.

Jappelli, Tullio, and Franco Modigliani, "The Age-Saving Profile and the Life-Cycle Hypothesis," manuscript (1998).

Jianakoplos, Nancy A., Paul L. Menchik, and F. Owen Irvine, "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, 1989).

Keyfitz, Nathan, and Wilhelm Flieger, World Population Growth and Aging: Demographic Trends in the Late Twentieth Century (Chicago: University of Chicago Press, 1990).

Mason, Andrew, "Saving, Economic Growth, and Demographic Change," Population and Development Review 14 (1988), 113–144.

National Statistical Office, *Handbook for Users of the 1981 Socio-Economic Survey Data Tape* (Bangkok: Office of the Prime Minister, 1981).

———, *Report on the 1995–1996 Survey of Population Change* (Bangkok: Office of the Prime Minister, 1997).

Paxson, Christina, "Saving and Growth: Evidence from Micro Data," *European Economic Review* 40 (1996), 255–288.

Poterba, James M. (Ed.), *International Comparisons of Household Saving*, National Bureau of Economic Research Project Report series (Chicago: University of Chicago Press, 1994).