

**NBER WORKING PAPER SERIES**

**SAVING AND GROWTH: EVIDENCE  
FROM MICRO DATA**

**Christina Paxson**

**Working Paper 5301**

**NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
October 1995**

Material from the Family Expenditure Survey is Crown Copyright, has been made available by the Central Statistical Office through the ESRC Data Archive, and has been used by permission. Neither the CSO nor the ESRC Data Archive bear any responsibility for the analysis or interpretation of the data reported here. Thanks are due to the Directorate General of Budget, Accounting and Statistics for permission to use the Taiwanese Personal Income Distribution Surveys, and the Thai Bureau of Statistics for permission to use the Thai Socioeconomic Surveys. I benefitted from comments from Orazio Attanasio, Angus Deaton, Ross Levine, seminar participants at the XIII Annual Macroeconomics Conference in Frankfurt and seminar participants at Princeton University. Thanks for funding are due to the National Institute of Aging through grants P01-AG05842 and R01-AG11957. This paper is part of NBER's research program in Aging. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

© 1995 by Christina Paxson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

SAVING AND GROWTH: EVIDENCE  
FROM MICRO DATA

ABSTRACT

This paper examines whether the observed cross-country correlation between aggregate saving rates and economic growth can be explained by models in which higher growth increases saving rates, rather than the other way around. The paper focusses on two explanations of why growth might increase saving. First, standard life-cycle theory implies that higher growth will increase the life-time wealth of younger savers relative to older dissavers, thereby increasing the aggregate saving rate. Second, models of consumption with habit formation imply that consumption responds slowly to unexpected income growth, and so unanticipated growth can produce a higher saving rate at least in the short run. I assess the validity of these explanations using time-series of cross-sections of household income and consumption surveys from four countries: the US, Britain, Taiwan and Thailand. I find that although in three out of the four countries there is evidence that saving behavior is consistent with life-cycle theory, there is simply too little life-cycle saving for higher growth to have a large effect on the aggregate saving rate. The habit formation model also implies very small effects of growth on saving rates. A large portion of the observed cross-country correlation between saving and growth cannot be explained by these models.

Christina Paxson  
219 Bendheim Hall  
Princeton University  
Princeton, NJ 08544-1022  
and NBER

## 1. Introduction

There is ample cross-country evidence that aggregate rates of saving and rates of economic growth are positively correlated. However, the direction of causality -- from savings to growth, or from growth to saving -- is not particularly well-understood. There are several types of models that yield such a correlation, with very different implications for causality. First, many growth models imply that high-savings economies will grow faster, either in the short-run as one moves towards a steady state, as in the Solow model, or in the long-run, as in many endogenous growth models. Second, Modigliani's classic life-cycle model predicts a correlation between savings and growth, but with the direction of causality reversed. As economic growth increases, the life-time resources of the young increase relative to the old, and since (according to the theory) the young save a larger fraction of life-time wealth than the old, aggregate savings rates increase. Third, models of consumption with habit formation provide another avenue through which growth may affect saving rates. As in noted in Deaton (1992) and Carroll and Weil (1993), habits cause consumption to respond slowly to unanticipated growth in earnings, and the result is higher saving, at least in the short run. The first of these three types of models implies that higher saving rates results in higher growth. The latter two imply that higher growth results in higher savings rates.

In this paper I investigate how much of the link between saving and growth observed in the macro data can be accounted for by life-cycle and habit formation models. I use time-series of cross-sections of household-level data on income, consumption and savings from four countries -- the United States, the United Kingdom, Taiwan, and Thailand -- to estimate life-cycle and habit formation models of saving. The results are used to calculate the implied effects of increases in economic growth on aggregate saving rates.

The estimates of life-cycle models for each country indicate that increases in rates of economic growth will have only small effects on the aggregate saving rate. Specifically, the results show that an increase in economic growth of 10 percentage points would increase the aggregate saving rate by no more than 2.5 percentage points (for the US and Taiwan), and would have virtually no effect on the saving rate in Britain or Thailand. The reason why these effects are so

small is simple: although in all countries except Thailand there is some evidence that saving rates follow the hump-shaped pattern indicated by life-cycle theory, it is also the case that income and consumption move together closely over the life-cycle in each country. The age effects in saving rates are simply too small to imply a large effect of growth on aggregate saving rates. Estimates of the habit-formation models, which yield measures of the effects of past earnings growth on current saving rates within cohorts of individuals, also indicate small effects of growth on saving, although the results differ somewhat across countries.

These estimated effects of growth on saving are small relative to the size of the cross-country relationship between rates of saving and growth. Figure 1 uses data from the Penn World Table 5.6 to plot decade averages (1951-1960, 1961-1970, 1981-1990) of each country's ratio of investment to GDP against growth in real GDP. All countries from the data series are included, although decades for which any of the variables in any year were missing were excluded. The figure also shows the regression line through these points, with an intercept of .141, and a highly significant slope of 1.38. Thus, the cross-country evidence indicates that a 10-percentage point increase in growth is associated with a more than 13-percentage point increase in the aggregate saving rate. It should be noted that the points for the four countries used in this analysis (which are highlighted in the table) are not outliers, and that for Taiwan and Thailand it appears that there is a positive within-country correlation of growth and the investment share.

Of course, one could argue that the relationship graphed in Figure 1 implies nothing about the link between saving rates and growth because, in the presence of international capital mobility, investment need not be financed from domestic saving. However, there is evidence that the correlation between investment rates and saving rates are high (see, for example, Feldstein and Bacchetta, 1991) Furthermore, research that uses data on saving rates yields similar results. For example, Carroll and Weil (1993) use data from the Penn World Table to construct measures of savings rates, and find a strong cross-country correlation between average savings rates between 1950 and 1987, and average growth rates over the same period. The regression of the average savings rate on the average growth rate yields significant slope coefficients of 1.84 for the subset

of OECD countries, and 1.19 for the full set of countries. They also work with average rates of saving and growth over 5-year periods, and find that a significant correlation persists even when country-fixed effects are included, although the slope coefficients are somewhat smaller. Other studies with similar findings include Modigliani (1970), Maddison (1992), and Bosworth (1993).

Thus, the life-cycle and habit formation models account for only a small portion of the observed cross-country relationship between growth and saving. It should be noted, however, that the saving measure used in the analysis is quite narrow: it includes only household saving, excluding accumulation in employer-sponsored pension plans (for which information is not available at the household level). The exclusion of saving in pensions means that the life-cycle effects of growth on saving rates may be higher than I have calculated, especially in the US and Britain where pensions are an important source of income during retirement.

The paper proceeds as follows. Section 2 presents the analysis of life-cycle savings patterns in the four countries, and also discusses the relevant institutional features of and data from the four countries used in the analysis. Section 3 presents evidence on within-cohort earnings growth and savings, and Section 4 concludes.

## **2. Life-cycle saving in four countries**

### *2.1 Introduction*

Just as growth theory provides support for the proposition that saving drives growth, Modigliani's life-cycle model provides support for the opposite, that growth drives saving. The simplest version of the model, with complete certainty, assumes that consumption is a constant equal to the annuity value of life-time wealth. Consumers save during working years, and dissave during retirement. There may also be dissaving at young ages when incomes are low, although the assumption that consumers cannot borrow against future earnings would rule this out. Given these assumptions, economic growth that increases the life-time wealth of younger savers relative to older dissavers (holding the shape of age-earning profiles fixed) will increase the aggregate saving rate. Furthermore, if the age profiles of earnings and consumption are the same across

countries, then the positive relationship between growth and saving that is observed within countries should also appear in cross-country data.

The prediction of the simplest life-cycle model -- that higher growth will result in a higher saving rate -- does not necessarily hold up when the assumptions of the model are modified. For example, it is unlikely that individuals desire consumption to be constant over their lives: consumption “needs” may vary systematically over the life-cycle, driven by factors such as changes in health status and declines in work-related expenses experienced upon retirement. The rate of interest relative to the rate of time preference may also affect the desired age-consumption profile. Uncertainty (about incomes, life-spans, or future consumption needs) combined with precautionary saving motives may result in upward-sloping age-consumption profiles, see Deaton (1992, Ch. 6). If age-consumption profiles can in theory take on any shape, then so can age-saving profiles, and in this case the relationship between growth and saving implied by life-cycle theory is *not necessarily* positive. However, as long as the average age at consumption exceeds the average age at earning, higher growth will result in a higher aggregate saving rate.

There is an extremely large empirical literature that examines the validity of the life-cycle model. Overall, the evidence is quite mixed. First, there is strong evidence that consumption “tracks” income over the life-cycle, see Carroll and Summers (1991) and Attanasio and Davis (1994) for particularly compelling evidence using US data. This coincidence of movements in income and consumption over the life-cycle could be explained by borrowing constraints, although borrowing constraints should not prevent middle-aged people from saving for retirement. But perhaps social security and pension benefits in the US are generous enough so as to make retirement saving (over and above pension savings) unnecessary for many people. The simulations of Hubbard, Skinner and Zeldes (1995) indicate that a life-cycle model that includes precautionary motives for saving, social insurance programs, and borrowing constraints, can yield the result that it is optimal for individuals to accumulate little or no wealth over their working lives. Although the model they simulate is a “life-cycle” model, in the sense that individuals choose consumption optimally given expected life-time wealth (subject to a constraint that assets

be non-negative), the results do *not* imply that growth will have large effects on the aggregate saving rate.

The evidence on “tracking” is balanced by evidence that there are “humped-shaped” patterns in savings rates predicted by life-cycle theory. Poterba (1994a, Table 3), in the introduction to a recent volume on international comparisons of household saving, shows age-specific household savings rates calculated from household survey data (where age is defined as the age of the household head) from OECD countries. For each of the six countries, savings rates generally increase up to the age of retirement, a feature consistent with life-cycle theory. What the results do *not* show is dissaving during retirement years. However, the absence of dissaving among the old can perhaps be explained by selection bias. This could happen if elderly individuals move in with children or into institutions when their assets are drawn down to low levels, so that only higher-saving individuals continue to be household heads at the oldest ages. Excess mortality among the poor elderly could yield similar results. Although the question of whether or not the elderly actually dissave is far from settled, it should be noted that *negative* saving rates at old ages are not required to generate a positive relationship between growth and saving. As discussed above, all that is required is that the average age of consumption exceed the average age of earnings, which is consistent with no dissaving at older ages provided that there are positive bequests.

### 3.2 Empirical specification

The starting point of this work is a simple version of the life-cycle model with no uncertainty. Consumption for a consumer  $i$  aged  $a$  and born in year  $b$  is a function of life-time wealth  $W_{ib}$  and preferences  $f_i(a)$  is specified as:

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

The function  $f_i(a)$  determines how fixed life-time resources  $W_{ib}$  are allocated to consumption over the life-cycle, and in the “textbook” life-cycle model would be a constant. There is no reason, however, to believe that tastes for consumption do not vary with age, especially when (1)

is applied to *household* rather than *individual* consumption. Consumption “needs” would depend on life-cycle variation in the ages and numbers of household members, variables which display systematic changes over the life-cycle, and which could be also be included in the function  $f_i(a)$  as additional independent variables. Furthermore, interest rates may affect the allocation of consumption over the life-cycle, and these effects should be captured by  $f_i(a)$ .

The advantage of specifying consumption as the product of life-time wealth and the function  $f_i(a)$  is that the logarithm of consumption can be expressed as the sum of an age-specific component and a fixed life-time wealth component:

$$\ln(c_{iab}) = f_i(a) + \ln(W_{ib}). \quad (2)$$

With panel data on individuals, equation (2) could be easily estimated by regressing the logarithm of consumption on a set of dummy variables for each age, and an individual-specific fixed-effect to account for each person’s life-time wealth. Unfortunately, suitable panel data are practically nonexistent. However, many countries collect household expenditure surveys annually or at fairly regular intervals, and these data can be used to estimate a variant of equation (2) using information on the consumption of birth-year cohorts. Specifically, take averages of (2) across all individuals in of age  $a$  born in year  $b$ , to obtain:

$$\overline{\ln(c_{ab})} = \overline{f(a)} + \overline{\ln(W_b)}, \quad (3)$$

which can be estimated by regressing the average logarithm of consumption for each cohort-age cell on a set of age and birth-year dummy variables. Given that the life-cycle framework used here is correct, the estimated age effects measure the effects of life-cycle tastes and the rate of interest on consumption, and the cohort effects reflect the effects of economic growth of the life-time wealth of each cohort. Equation (3) can be further modified to include cohort-age averages of other variables, such as the numbers of children and adults in the household, that may affect consumption regardless of age. The inclusion of these demographic variables is likely to be of importance in Taiwan and Thailand, whose fairly recent demographic transitions mean that current households have sizes and age structures that are quite different from households of 15 years ago.



The average logarithm of income is also decomposed into age and cohort effects, although there is no strong theoretical justification for doing so. The assumption that underlies this additive decomposition is that growth works by shifting up age-income profiles, without changing their shapes. This assumption is somewhat problematic, because income includes income from assets in addition to earnings. If higher growth shifts up age-*earnings* profiles (up to the age of retirement) without changing their shapes, then it will also alter the age-profile of asset holdings. This will have the result of changing the age profile of *income*, which includes asset income, as long as the rate of interest is positive. A theoretically “cleaner” way to proceed would be to estimate age and cohort effects in earnings rather than income. There are two reasons I do not do this. First, I am interested in the how saving rates vary with age, and the saving rate is conventionally measured as the fraction of income (not earnings) that is not consumed. Second, estimating cohort averages of the logarithm of earnings is complicated by selection out of the labor force around the age of retirement. The difference between the logarithm of income and consumption is approximately equal to the saving rate, and this difference is also decomposed into age and cohort effects. The major point of the exercise is to determine the shape of the age-saving profile, since for the life-cycle mechanism from growth to saving to be valid, it must show higher savings rates among working-aged individuals than among those who are retired.<sup>1</sup>

Equation (3) is based on a life-cycle model with the assumption of certainty. The addition of uncertainty complicates things considerably. A key problem is that there may be common macro shocks that shift the life-time wealth of all individuals. Life-cycle theory indicates that the effect on consumption of common shocks to life-time wealth will vary across consumers of different ages, but these age/time interactions can not be identified with the data available. The problem can be at least partially dealt with by including year effects in equation (3) and the corresponding income and saving rate equation. However, because of the co-linearity of age, birth-year, and time, additional restrictions on the year effects are required. As is shown in Deaton and Paxson (1994), if year effects are entered in an unrestricted form, then any *trends* in the data can be arbitrarily attributed to year effects, or a combination of age and cohort effects. To see

this, consider an extreme example in which all individuals have identical consumption that increases by 5% per year. This 5% growth could imply a set of year effects that increase 5% in each subsequent year, or alternatively it could imply a set of age effects that increase by 5% with each year of age and a set of cohort effects that decline by 5% for each subsequent cohort. In most of the following analysis, I include year effects in all equations, but normalize them so that they sum to zero and are orthogonal to a time trend. Essentially, all trends in the data are forced to appear as a combination of age and cohort effects. It will be seen below that this method of normalization is not always satisfactory. In some versions of the savings rate equations, I restrict the cohort effects to be zero and allow for unrestricted year effects. This normalization makes sense given that standard life-cycle theory implies that there should be no cohort effects in savings rates. The existence of (non-zero) cohort effects mean that some cohorts have higher savings rates in *all* years of their lives than do other cohorts, and this does not make sense unless there are fixed differences across cohorts in bequest motives.

### *3.3 The countries and data*

The countries used in this study have several important features that should be kept in mind when interpreting the results. First, they have had quite different growth and saving experiences over the past several decades. The US and Britain have grown slowly, with average annual growth in per capita output of less than 2% (in the US) and less than 3% (in the UK) during the 1970's and 1980's.<sup>2</sup> In contrast, the average annual growth rate in Thailand during this period exceeded 4%, and for Taiwan was around 7%. In addition to having higher growth rates, the two Asian economies have also recently had higher gross saving rates and personal saving rates (see Table 1). The experience of these countries is consistent with the observed positive cross-country correlation between saving and growth

Second, the four countries have different institutional features that may affect private saving behavior. The US and the UK both have well-developed public pension programs, and private pensions are also common. In the US in 1988, public pensions from the Social Security

system accounted for 38% of income for households with a head aged 65 or more. Private pensions, government pensions, and annuities accounted for an additional 17% (Poterba, 1994b). In the UK in 1989, retired households derived 20% of their incomes from annuities and pensions, and 41.5% from social security (Banks and Blundell, 1994b). In contrast, neither of the Asian countries have adopted “Western-style” old-age social security systems. Private pensions are not common, in part because a high fraction of workers in both countries are self-employed. Taiwan’s Labor Insurance Act and Labor Standards Law mandate that employers pay a lump-sum retirement benefit equal to 3 months of pay for each year of service, but not to exceed 95 months of pay (see Levenson (1994)). However, the Survey of Personal Income Distribution does not distinguish these benefits from earnings, so it is difficult to assess their importance. Pensions are quite rare in Thailand. Data from the 1992 Socio-Economic Survey indicate that for households headed by someone aged 60 or older, about 2.0% of income is derived from pensions. One might expect that individuals in the two Asian countries will have much stronger incentives to save for retirement, since public and private pension systems are not well developed. However, in both of these countries responsibility for the care of the elderly has traditionally fallen on their children, and it may be that informal family-based institutions fill much the same role as the formal social security and pension systems of the West.<sup>3</sup>

Each of the four countries used in this analysis collected household income and expenditure surveys throughout the 1980’s, and for several of the countries the surveys began even earlier. Although the details of the surveys differ across countries, they have several common characteristics. First, they are all large, surveying between 7,000 to 15,000 households per year. The size of the surveys makes it possible to construct information on birth-year cohorts. Second, they all collect fairly detailed information on consumption expenditures and total household earnings and income, and basic demographic information (such as household size and numbers of children).

For all of the countries, I measured consumption as total expenditure on all goods and services including durable goods, but excluding housing purchases, and income as after-tax

income.<sup>4</sup> The life-cycle theory discussed above strictly applies only to nondurable consumption, and the “best” procedure in theory would be to measure consumption as nondurable expenditure plus the flow of consumption services from durable goods. There are three reasons why I used total consumption rather than nondurable consumption. First, information on the flow of services from nondurables is not available from any survey. Second, the Taiwanese survey uses fairly coarse expenditure categories, making it difficult to distinguish durable from nondurable expenditure. Third, I wanted the saving measure (defined as the difference between income and consumption) to correspond as closely as possible to the definition of personal saving used in National Income and Product Accounts. The saving measure therefore does not include changes in wealth caused by capital gains or losses. It should also be noted that the saving measure does not count contributions to employer-provided pension plans as saving, and this omission may bias down the estimated effects of growth on saving, at least in the US and UK where private pensions are important.

The general procedure followed in all countries was to use the household-level data to construct “cohort” data sets, containing averages of necessary variables (logarithms of income and consumption, the numbers of adults and children in households, etc) for each cohort-year cell. In all countries I defined a cohort to consist of all households with heads (or, as in the US, “reference members”) born in the same year. The samples were restricted to households with heads between the ages of 25 and 75, since selection into and out of household headship is likely to be most severe at the youngest and oldest ages. However, as is described below, selection is still major issue in Taiwan and Thailand. The following paragraphs describe basic features of each of the four surveys.<sup>5</sup>

### *3.2.1 The United States*

The US data are drawn from the *Consumer Expenditure Surveys* (CEX) from 1980 through 1992. I exclude rural households, since these were not surveyed in 1983 and 1984. The CEX is a “rolling panel” that interviews approximately 7,000 households per year. Each

household is interviewed 5 times at quarterly intervals, with different households beginning their series of interviews in different quarters. The results of the 2nd through 5th interviews are made available. Not all households complete the full series of surveys, however, and the number of quarterly observations ranges from 15,613 to 20,716 per year. The information on each household in each quarter consists of expenditure in the 4 months preceding the interview. Information on income was collected only at the 2nd and 5th survey (unless a job change was experienced between the 2nd and 4th survey), and refers to income in the previous year. The income figures for households in their 3rd and 4th interview refer to the 2nd-interview income figure.

I made use of the data at the annual level and ignored their panel feature. The average of the logarithm of after-tax income for cohort of a given age in a given year is simply measured as the (weighted) average over all households with reference members of the given age in all quarters of the survey year. The expenditure and earnings variables are treated similarly. This procedure is not without its problems. For households whose second interview falls early in the calendar year, a large portion of their reported income covers income earned in the year *before* the survey, and this does not precisely "match" the expenditure measures which are for the quarter preceding the survey.<sup>6</sup>

Top coding is another problem in the CEX, particularly for income. The upper limit on specific income categories (e.g. self-employment income, salary income, etc.) was \$100,000 by 1992, but was more severe in the early 1980's. Values greater than the top coded values are set to the upper limits. Total after-tax income is the sum of the specific income components (some of which may be top-coded) , and so it can exceed the upper limit although its components cannot. Expenditure is less subject to top coding, although some items, such as various types of medical expenditure, face upper limits. However, these upper limits are generous. For example, in 1992 there was a \$200,000 limit for specific medical care items, such as hospital services. I made no attempt to impute the true values of top-coded figures, and all values of income and expenditure were used regardless of top coding. The more severe top coding problems for income will tend to bias down savings rates, conceivably by quite a bit if the richest households do the bulk of the

saving. A less serious problem is caused by my use of logarithms of income and expenditure: households with negative values are excluded from the calculations. In the results that follow, I discuss several things that were done to ensure that the results are reliable despite the numerous problems with the CEX. First, I compare the basic results with results that use cohort medians of the dependent variables. (This is done for all countries.) Cohort medians are less sensitive to top coding and outliers, although they have the obvious drawback that they do not sum up to national aggregates. Second, I estimate the US income equations using data from both the CEX *and* the Current Population Survey (CPS). The CPS covers more households and is generally thought to provide more accurate income figures than the CEX.

The major variables are defined as follows. Income is the CEX's measure of after-tax income (called "fincatax" in CEX documentation). It is measured as gross household minus state and local taxes. A worker's social security contributions are not counted as taxes. Earnings is before tax, and is measured as the sum of wage and salary income, plus profits from farm and non-farm self-employment. The expenditure measure is the "total expenditure" measure constructed by the Bureau of Labor Statistics, and consists of expenditure on both durable and nondurable goods. The final "cohort" data set contained 663 cohort-year observations. The cell sizes (i.e. numbers of households used to compute cohort-year averages) ranged from 110 to 488, with a median value of 241.

### 3.2.2 *Great Britain*

The data from Great Britain are drawn from the 1970 through 1992 rounds of the *Family Expenditure Survey* (FES). The sample is representative of Great Britain, which is the United Kingdom excluding Northern Ireland. Approximately 7000 households are surveyed per year. Like the CEX, the FES is done on a quarterly basis, but households are not interviewed more than once. This means that cell size for the British "cohort" data are smaller than in the US. The smallest number of observations used to construct a cohort/year cell is 63 (for 75-year-Olds in 1972). The median cell size is 197.

Unlike the CEX, the FES does not have the problem that income and expenditure may refer to different time periods. The earnings measure (which is used when constructing income) is based on a question about “normal weekly earnings.” Consumption includes durables, and the nondurable components of total consumption are reported on a *weekly* basis. Income is total disposable income, again with a weekly reference period. It is defined in the FES documentation as “normal gross income, excluding tax and National Insurance contributions, but including income in-kind.” The final cohort data set contained 1,173 cohort-year observations.

### *3.2.3 Taiwan*

The data from Taiwan are drawn from the *Survey of Personal Income Distribution* (PID). These surveys have been conducted annually since 1976. In 1976 and 1977 the sample size was approximately 9,000 households. This was increased in 1978, to over 14,000 households per year. The survey design is described in Republic of China (1989). Each household is surveyed once, and asked questions about income and consumption in the past year. Consumption is defined as total expenditure on all goods (both durable and nondurable). Income is after-tax. Earnings includes wages, salaries, and self-employment income. Although the number of households surveyed is large, some of the cell sizes are quite small, especially for the oldest age groups in the earlier years of the survey. It is common for old people in Taiwan to live with their adult children. Since the household head is defined to be the primary earner in the family, households headed by people in their 70's are rare. The smallest cell contains information for 5 households (with heads aged 74 in 1978). The median cell size is 291.

### *3.2.4 Thailand*

The Thai data are drawn from the *Socio-Economic Surveys* conducted in 1976, 1981, 1986, 1988, 1990, and 1992. Like the other surveys used for this paper, these are large cross-sectional income and expenditure surveys, with between 11,000 and 13,500 households surveyed per year. The 1975/76 survey was started in November of 1975, and ran for 12 months with

approximately 1/12 of the households surveyed in each month. The other surveys ran from January through December of the survey year.

The Thai surveys are in some ways more detailed than those from the other countries. A large fraction of households is engaged in farming, and for these a large fraction of food is home-produced. The survey asked details about farm revenues and costs, where revenues include the value of all production regardless of whether it was sold or consumed, and this information is used to measure farm profits. Total after-tax income includes farm profits, earnings, rents received (both in cash and in-kind), and income from financial assets. Consumption includes both home-produced and purchased items. The reference period for income is a year. The reference period for consumption items differed across items. Food consumption had a reference period of a week, other small expenditure items (such as clothing, personal supplies, and entertainment) had a reference period of a month, and major expenditure items (such as bicycles and major appliances) had an annual reference period. The cohort data set based on the Thai surveys has 306 observations. The cell sizes vary from 45 to 453, with a median of 226.

### *3.2.5 Accuracy of the survey data*

Studies of household saving based on household survey data invariably ask whether savings rates calculated from the survey data are comparable to those in the national accounts. Unfortunately, it is typical to find that they are not. Table 1 presents information, for each country, on the household savings rates derived from the survey data, together with corresponding information from national accounts. Of the four countries, the Taiwanese survey results most closely match the national accounts figures in all years except 1976. Savings rates from the CEX show the same general decline over the 1980's that appears in the national accounts figures, but in many years the two rates are quite different. Bosworth et. al. (1991) has a careful discussion of differences between the national accounts and CEX savings measures. The survey-based savings rates from Thailand are the least closely matched to the national accounts figures. Under-reporting of income relative to expenditure is thought to be a common feature surveys



from developing countries, see Visaria (1980), and the Thai surveys may suffer from this problem: the survey actually shows *negative* savings rates in 1981, 1986 and 1988.

The British data show a marked divergence between the survey and national accounts information after 1986: from 1986 to 1991 the survey data indicate a steady rise in household savings rates not matched in the national accounts. This appears to be due to greater income growth in the FES than in the national accounts (see Attansio and Weber (1992)). One possible explanation is that the sample design was changed in 1987. In the early years of the survey, households were selected from voter registration lists, and in the later years households were selected from address lists. However, this would not explain the continuing rise in FES savings rates from 1987 to 1991, and the sharp decline in 1992. In what follows I estimate equations for Britain using the full 1970-1992 sample, as well as a sample from 1970-1986, to make sure that apparently spurious increase in savings in the later years of the survey does not affect the results.

Another issue is whether selection into and out of headship is likely to substantially bias the results. In principle, the sample of observations used to construct information on a specific cohort should be drawn from a population of individuals that remains fixed over time. This will not be the case in practice, because the cohorts are formed on the basis of the age of the household head. Movements into and out of headship caused by household formation, marriage, divorce, and death will alter the underlying population from which the observations for a specific cohort are drawn. This is likely to be a special problem in Taiwan and Thailand, where it is common for parents to reside with adult children. Furthermore, the definition of "headship" differs across the surveys of these two countries. In Taiwan, the household head is defined to be the primary earner, so that retired parents who live with children are not typically household heads. In Thailand there is no clear-cut rule for defining headship, but it appears from the data that adult children who live with their parents are often not designated to be heads. The different patterns of headship with age can be seen from Figure 2, which shows the fraction of people within each age group who are household heads. The upper panels, for the US and Britain, indicate fairly stable headship rates after the age of 27 or so. The rates increase at very old ages, reflecting the greater

percentages of people who are widows and widowers living alone. In Taiwan headship rates increase through age 40, and then decline steadily to age 80. Only about 25% of individuals aged 70 years old are heads. Thailand, in contrast, shows an almost steady increase in headship rates with age. The results from Thailand and Taiwan must be treated cautiously, since it is plausible that selection into and out of headship is correlated with income, consumption, and saving behavior.

### *3.3 Results*

Life-cycle patterns in income and consumption can be seen by looking at raw data for the various cohorts. Figures 3 and 4 plot the logarithm of income and consumption against age for every 5th cohort in each country. Two features of the data are important. First, the income and consumption profiles appear to track each other, consistent with the other evidence of tracking discussed above. In all countries, consumption and income increase up to middle age, and bend down at old ages, although the “peaks” in income and consumption occur later in Taiwan and Thailand than in the US and Britain. Second, the differences in growth between the two Asian countries and the two Western countries are evident in the data. In Taiwan, the age profiles of income and consumption shift up for each successive cohort, consistent with the rapid growth (over 5% per year on average) that occurred during the sample period. The cohort diagram for Thailand also indicates some growth, although there is also a clear decline in income and consumption in 1986, the third survey year, for all cohorts. The British and US results show little growth over the sample periods.

The patterns of consumption and income across ages and cohorts can be seen more clearly by regressing the cohort averages of the logarithm of each variable on sets of age, cohort and restricted year dummies.<sup>7</sup> Table 2 provides F-statistics for the joint significance of each set of dummies, and indicates that each is jointly significant in all countries. The age effects are graphed in Figure 5. To maintain the same scale across countries, the age effects have been normalized to zero at the lowest age of 25. Therefore, the vertical difference between the two curves is

approximately the saving rate at each age plus or minus a cohort-specific constant. Note that there are two sets of age effects in income for the US. One is based on income data from the CEX, and the other is based on before-tax income data from the March Current Population Surveys (CPS).<sup>8</sup> Measures of before-tax income from the CPS are, on average, higher than in the CEX. However, the patterns of income with age are virtually identical across the two surveys, providing evidence that the results for the CEX are not too unreliable. The cohort effects are graphed in Figure 6.

The major features of the raw data -- that consumption tracks income over the life-cycle, and that younger cohorts have higher consumption controlling for age -- are also evident in Figures 5 and 6. Within each country, the shape of the age-consumption and age-income profiles are similar. Thailand and Britain, in particular, show a very close correspondence between movements in income and consumption with age. Furthermore, the age-consumption profiles differ across countries. Taiwan, in particular, has an age-consumption profile that is much steeper than that of the other countries. These results are not particularly supportive of standard life-cycle theory. Under a life-cycle interpretation of the data, the age effects in consumption reflect only age-specific tastes and the rate of interest. For the results to be consistent with the life-cycle model, one must believe that these factors differ across countries (producing different age-consumption profiles), and that the similarities in age-consumption and age-income profiles within each country are due solely to common effects of interest rates on both -- for example, if higher real interest rates result in both steeper age-income profiles and age-consumption profiles. This explanation for tracking seems implausible given the high degree of co-movement of the income and consumption profiles shown in Figure 5

One possible explanation for the differences in age-consumption profiles across countries is that life-cycle patterns in household size and age structure vary across countries, producing differences in desired consumption at each age. The data indicate that the countries do have pronounced differences in household structure. Figures 7 and 8 show the average numbers of adults and children for every 5th cohort in each country. Two features of these figures stand out. First, the graphs for Taiwan and Thailand show the rapid demographic transition that has taken

place over the last several decades. The number of children in the Asian households declines noticeably across cohorts, given the age of the household head. Second, the Asian households are much larger than the Western households, especially at older ages. In Thailand, the average number of adults per household is around 3 for households headed by 60 to 70 year Olds. Households headed by older people in the Asian countries are also more likely to contain children. The results for Taiwan in Figure 7 indicate that the numbers of adults first decreases as the age of the household head goes from 20 to 40, and then rises. This pattern reflects the fact that many young adults in Taiwan live with their parents, and if the income of a younger adult exceeds that of his or her parents then he or she will be designated the household head.

To see if these differences across countries in household size and composition are responsible for differences in age-consumption profiles, I estimated variants of the consumption and income equations which included the average numbers of adults and children as additional regressors. F-statistics for the age, cohort, and year effects, as well as the coefficients for numbers of adults and children, are in Table 2. When the demographic variables are included, the sets of age, cohort and year effects are still jointly significant in all countries. As might be expected, extra adults increase both income and consumption. Extra children sometimes increase and sometimes decrease both income and consumption, but in any case the coefficients for children are much smaller than those for adults.

Figure 9 graphs the age effects in consumption and income when controls for the numbers of adults and children are included. Including the demographic controls makes little difference to the age effects in all countries except Taiwan. In Taiwan, the inclusion of controls eliminates the feature that consumption declines with age for the oldest households, and results in an even more steeply-sloped age-consumption profile. Thus, instead of reconciling differences in age-consumption profiles across countries, the inclusion of demographic controls *increases* the cross-country differences. These differences can be seen most clearly by plotting the profiles for the four countries on one graph. This is done in Figure 10, which shows that, controlling for numbers of adults and children, consumption declines in old age in Britain, is fairly flat in the US and

Thailand, and increases with age in Taiwan. It is possible that the differences in age-consumption profiles across countries are accounted for by differences in interest rates across countries, or differences in tastes not captured by the demographic variables used here. However, a more plausible story is that the simple version of the life-cycle model is incorrect, and that the age-consumption profile is affected by the timing of income receipts over the life-cycle.<sup>9</sup>

The main goal of decomposing consumption and income into age, cohort and (restricted) year effects is to identify age effects in saving rates. I estimated regressions of the “saving rate”, measured as the difference between the logarithm and income and consumption, on age, cohort and year effects, with and without controls for numbers of adults and children.<sup>10</sup> The age effects in savings rates are graphed in Figure 11; the heavier lines show the age effects when numbers of adults and children are not included as additional regressors. Note that the age effects are normalized to be 0 at age 25, and that the scales of the graphs differ across countries.

The age-saving profiles do not display the “hump-shaped” pattern predicted by basic life-cycle theory. The Taiwanese results with demographic controls included show saving rates that increase up to age 40 but do not decline at old age. In Britain, there is a small decline in savings rates around the age of retirement, but this is followed by a large increase after about age 65. The age effects in savings rates in Thailand are small and show no clear pattern with age. The results for the US are the most at odds with the theory: the figure shows huge declines in savings rates with age. Taken at face value, the results imply that the average saving rate of 25-year-Olds exceeds that of 65-year-Olds by 0.5. Since we know that the cross-sectional data does not show large declines in saving rates by age in the US, it must be that these large age effects are offset by cohort effects that indicate lower rates of saving among younger cohorts. Figure 12 indicates that this is the case: the cohort effect in the savings rate for those aged 25 in 1980 is approximately 0.5 less than the cohort effect for those aged 65 in 1980.

Overall, these results are not encouraging for the life-cycle model. However, it may be that the age and cohort effects are actually picking up general trends in savings rates. These trends may be real, in the sense that they match secular trends in aggregate savings rates that appear in

the national accounts data, or (especially in the case of Britain) they may reflect trends that are likely due to measurement error in the household survey data. I did two things to examine whether trends in saving rates are influencing the age-saving profiles. First, I re-estimated the British equations using only data from 1970 to 1986, thereby eliminating the years over which the British survey data show unrealistically large increases in saving rates. These results are graphed in Figure 13. The left-hand panel reproduces the age effects shown in Figures 9 and 11 (which include demographic controls.) The right-hand panel shows results when the years 1987 through 1992 are excluded. When the later years are excluded, the age-saving rate profile is “tilted” down, showing a smaller increase in the saving rate up to age 60, and a larger decline immediately after the age of 60. Basically, some of the increase in saving rates between 1986 and 1992 appeared as increases in saving rates with age, and eliminating these years from the analysis flattens the age-saving rate profile. However, the large increase in the saving rate after age 65 or so appears regardless of the sample used. The result that the very oldest have the highest saving rates is perhaps explained by selection out of headship by individuals who save little, but without panel data on individuals over time the validity of this explanation cannot be tested.

The second thing I did was to re-estimate the saving rate equations for all countries (using 1970-1986 for Britain), including only age effects, unrestricted year effects, and numbers of adults and children, but excluding cohort effects. The life-cycle model does not provide a strong rationale for including cohort effects in saving rates, and in the case of the US, the combination of increasing age effects and even larger decreasing cohort effects could simply reflect the fact the saving rates in the US survey data have declined over time. The age effects from these saving rate equations are graphed in Figure 14 (again, the heavier lines show results when demographic variables are not included in the regressions.) Removing the cohort effects substantially changes the results. The US numbers are now much more consistent with life-cycle theory. Saving rates are fairly flat until retirement, and then show a subsequent decline. The Taiwanese and British results also show a decline in saving rates around the age of retirement, although in Britain saving rates still increase with age for the oldest ages. The Thai age-saving profile is flat, and unlike the

other countries the age effects are not jointly significant.

Cohort averages of saving rates may be influenced by a few households with very high or low saving rates. To make sure that outliers are not affecting the estimates of age effects in saving rates, I re-estimated the saving equations using median rather than average saving rates. The age effects from these regressions (which include demographic controls) are shown in Figure 15, and are remarkably similar to the age effects based on average saving rates.

### 3.4 Calculating the effect of growth on the aggregate saving rate

A useful question to ask is whether the results discussed in the last section are supportive of a “growth-causes-saving” interpretation of the macro data. In what follows, I use the estimates of age effects in income and saving rate profiles to calculate the effects of growth in life-time wealth on the aggregate saving rate.

Suppose that growth has the effect of shifting up the profile of the logarithm of income for successive birth-year cohorts, without changing its shape. The income of a cohort aged  $a$  in year  $t$  can then be expressed as:

$$y_{at} = e^{h(a)} (1+g)^{t-a}, \quad (4)$$

where the function  $h(a)$  measures age effects in the logarithm of income, and the term  $(1+g)^{t-a}$  measures the effect of growth on income on cohorts born in year  $t-a$ . The fraction of aggregate income earned by those aged  $a$  can be expressed as:

$$y_{at}^s = \frac{e^{h(a)} (1+g)^{t-a}}{\sum_{\alpha=1}^A e^{h(\alpha)} (1+g)^{t-\alpha} \eta_{\alpha}} \quad (5)$$

where  $\eta_{\alpha}$  is the fraction of households aged  $\alpha$ . Notice that the assumption that growth does not change the shape of age-income profiles implies that the share of income earned by each age group is independent of time. Likewise, the aggregate saving rate ( $s/y$ ) is independent of time, and can be expressed as:

$$(s/y) = \frac{\sum_{a=1}^A s(a) e^{h(a)} (1+g)^{-a} \eta_a}{\sum_{a=1}^A e^{h(a)} (1+g)^{-a} \eta_a} \quad (6)$$

where  $s(a)$  is the age-specific saving rate. Equation (6) can be used to derive the partial effect of changes in the growth rate  $g$  on the aggregate saving rate:

$$\frac{\partial (s/y)}{\partial g} = \frac{1}{(1+g)} \frac{\sum_{a=1}^A a (s/y - s(a)) e^{h(a)} (1+g)^{-a} \eta_a}{\sum_{a=1}^A e^{h(a)} (1+g)^{-a} \eta_a} \quad (7)$$

Equations (6) and (7) can be calculated for different values of  $g$  given information on the age profile of the logarithm of income  $h(a)$ , the age profile of the saving rate  $s(a)$ , and the fraction of the population aged  $a$ . Values for  $h(a)$  and  $s(a)$  are taken from the estimates of age effects in income and saving rates discussed in the last section. Two basic methods are used.

First, I filled in values for  $h(a)$  and  $s(a)$  using regressions results that did *not* control for the effects of demographic variables. The age effects in income  $h(a)$  correspond to those in Figure 5, and the age effects in saving rates are from Figure 14, which excluded cohort effects from saving rates. To derive  $\eta_a$ , I used the actual fractions of households of each age in the latest survey year for each country. Because there are year and (in the income equations) cohort effects, intercept terms must be included for both  $s(a)$  and  $h(a)$ . I specified these intercept terms so that the income level and saving rate matched the actual average income level and actual saving rate for 35-years olds averaged over all survey years.

The second method of computing (6) and (7) was similar, but used income and saving-rate equations that included controls for demographic variables. In this case, the age effects  $h(a)$  and  $s(a)$  included the effects of adults and children on income and consumption, and the numbers of adults and children at each age were taken from actual sample averages in the latest survey year. Thus,  $h(a)$  is measured as:



$$h(a) = \beta_a + \delta_1 ADULTS_{at} + \delta_2 CHILDREN_{at}.$$

where  $t$  denotes the last survey year, and  $\delta_1$  and  $\delta_2$  are the estimated coefficients on adults and children shown in Table 2, and  $\beta_a$  are the age effects in income shown in Figure 9. The age-specific saving-rate terms  $s(a)$  are similarly defined.

The results of tabulations of (6) and (7) for different growth rates are shown in Table 3. These results indicate that for all countries except Thailand, the effect of an increase in growth that does not change the shape of the age-earning profile will be to increase the aggregate saving rate. However, these effects are very small. Consider, for example, the results that include controls for demographics. In the US, a 10 percentage point increase in growth increases the aggregate saving rate by .1 to 1.9 percentage points, an effect that is much smaller than the increase of 14 percentage points or so implied by cross-country regressions. The Taiwanese results show slightly larger effects of growth on the saving rate, and those for Britain show very small effects. The Thai results indicate negative and extremely small effects of changes in growth rates on aggregate saving rates. This is not surprising, given the absence of age effects in Thai saving rates shown in Figure 14.

Overall, in none of the countries is there evidence that higher growth will result in large increases in aggregate saving rates. Of course, these tabulations rest on the strong assumption that growth results in higher age-earnings profiles for younger cohorts, but does not change their shapes. However, if growth increases the slope of age earnings profiles, then it is likely that the effects of growth and saving predicted by the life-cycle model will be even smaller. In this case the youngest (and wealthiest, in a life-time sense) will have incentives to save little and possibly to borrow, in anticipation of much higher earnings in middle age.

#### **4. Growth and saving within cohorts**

Recent literature has focussed on the role of habit-formation in consumption decisions. As is noted in Carroll and Weil (1993), and Deaton (1992), habit formation can produce a positive correlation between growth and saving even if consumers are forward-looking. The simplest

model of consumption under habit formation, presented in Deaton (1992), is one where utility today is a positive function of consumption today minus some fraction of last period's consumption:

$$u_t = u(c_t - \delta c_{t-1}). \quad (9)$$

Assuming additive utility across periods and quadratic sub-utility functions, utility maximization results in a consumption equation of the form:

$$\Delta c_t = \delta \Delta c_{t-1} + \epsilon_t, \quad (10)$$

where  $\epsilon$  is the effect on consumption of an innovation in earnings. If  $\delta=0$ , equation (10) produces the familiar result of the permanent income model, that consumption follows a martingale. With  $\delta$  greater than 0, innovations in earnings produce changes in consumption that persist for more than one time period. The result of an unexpected increase in earnings is an increase in consumption that takes many years to complete, and the result is higher saving in the short run.

The habit model might potentially explain why saving rates in fast-growing Taiwan are higher than those in the US and Britain: Taiwanese consumers may not have adjusted their habit stocks to be consistent with the higher wealth levels produced by rapid growth, and US and UK households may not have adjusted their habit stocks down in response to the productivity slowdown of the 1970's. For this story to be believable, however, it must be that habits adjust very slowly, or that consumers in these countries are continually surprised that income growth (or, in the US and UK, lack of growth) has persisted over time.

The link between saving and growth that is implied by the habit model works quite differently than that of the life-cycle model. Specifically, the simple life-cycle posits that the link from growth to saving hinges on changes in the distribution of wealth across age groups: growth does not increase the savings rates of individuals at specific ages, but merely redistributes wealth towards higher-saving age-groups. The habit model, in contrast, implies that unanticipated earnings growth within cohorts over time results in higher savings rates. It should also be kept in mind that there are other possible ways that earnings growth affects saving. For example, it could be that during higher growth periods consumers face more uncertainty in future income. This

could produce higher saving rates if consumers have precautionary saving motives. The work that follows does not distinguish between the habit model and other mechanisms.

I start by simply examining whether total within-cohort earnings growth is correlated with higher subsequent saving rates. Specifically, for each country I estimate equations of the form:

$$(s/y)_{at} = \beta_0 + \sum_{j=1}^4 \beta_j \Delta \ln(e_{a,t-j}) + \epsilon_{at}, \quad (11)$$

where  $(s/y)$  is the “saving rate”, measured as usual as the difference between the cohort/year average of the logarithm of income and consumption. The term  $\Delta \ln(e_{a,t-j})$  is the change in earnings between  $t-j$  and  $t-j-1$ , where earnings is measured to include wage and salary income, and self-employment income. Because earnings drops off rapidly for all cohorts around the age of retirement, I estimate equation (11) for cohorts aged 25 to 60 in year  $t$ . The current earnings change between  $t$  and  $t-1$  is not included, since measurement errors in current earnings will be correlated with measurement errors in the saving rate. Controls for the numbers of adults and children are included in all equations. The British estimates exclude the years 1987-1992. The Thai estimates must be done somewhat differently, since the surveys are unevenly spaced in time. This is discussed in more detail below.

The estimates of the effects of earnings growth on saving rates for the US, Britain and Taiwan are shown in Table 4. For each country, the top row of results shows the effects of increased earnings growth on saving rates when no controls for age, cohort, or time are included. In all countries, earnings growth is correlated with higher subsequent saving rates. The sum of the coefficient  $\beta_1$  through  $\beta_4$  ranges from .04 (for the US) to .52 (for Britain). The coefficients for the US are not statistically significant, and there is no evidence that past earnings growth produces higher current saving rates. Although the coefficients for Britain and Taiwan are significant, their sizes are quite small: recall that the cross-country evidence implies a slope of over 1.

This first set of results does not distinguish between anticipated and unanticipated earnings growth, and it may be that the results merely reflect the fact that middle-aged households, with the highest savings rates, are those that have experienced the largest (anticipated) earnings

increases. To account for the standard age effects in saving rates implied by life-cycle theory, I estimate a variant of equation (11) that includes a set of age effects. The results of these are shown in the second row for each country in Table 4. The inclusion of age effects yields slightly smaller effects of earnings growth on savings in Britain. In Taiwan, the effect of growth on saving is cut in half. The addition of cohort effects (row 3) has little additional effect on the results. The final row shows results that include year effects. It is not clear that year effects should be included, since these are likely to pick up the effects of unanticipated earnings growth on saving rates that I am trying to measure. As might be expected, when the year effects are included there is no longer a significant correlation in any country between past earnings growth and saving rates in any country.

Results for Thailand are shown in Table 5. The Thai data are treated differently than those of the other countries, for two reasons. First, in Thailand it is virtually impossible to disentangle “earnings” from “income.” Few households report income from financial assets, and a large fraction of household income comes from farms. Income from these sources includes both earnings and returns from investments in farms. I use income growth rather than earnings growth as the right-hand-side variables when estimating (11). This is problematic if rapid income growth reflects increases in farm assets that result from past saving behavior, and so the results for Thailand should be treated cautiously.

The second problem is that the Thai surveys were not conducted every year, and were not evenly spaced: since 1986 the surveys have been done every other year, and before then every 5 years, so measuring income growth over a single time span is not possible. The general procedure I followed was to regress the saving rate on average annual income growth over two past surveys. In all cases past income growth in 1992 is measured as average annual growth between 1986 and 1990, and past income growth in 1988 is measured as average annual growth between 1981 and 1986. Past income growth for 1990 is measured in one of two ways: in the top panel of Table 5, it is average annual growth between 1981 and 1988, and in the lower panel it is average annual growth from 1986 to 1988.

These two definitions of income growth yield fairly similar results. In Thailand, as in the other countries, past growth is correlated with higher current saving rates. The coefficient on income growth ranges from .10 to .19, slightly smaller than the sums of coefficients for Britain and Taiwan. Adding controls for age effects and both age and year effects does not reduce the coefficients, although they do become insignificant when the year effects are included. Cohort effects cannot be included because of the small sample size.

Overall, the results provide some support that past earnings growth increases the current saving rate in all countries except the US, but by only a small amount.

## **5. Summary**

This paper has investigated how much of the observed correlation between saving rates and growth can be explained by the effects of growth on household saving behavior. The answer is that although growth may increase saving rates, it will not increase them by much. First, even if the life-cycle model is valid, the estimates of age-specific saving rates imply only small effects of growth on the aggregate saving rate. The reason is simple: although there is some evidence that saving rates follow the hump-shaped pattern indicated life-cycle theory, it is also the case that income and consumption move together closely over the life-cycle. The age effects in saving rates are simply too small to imply a large effect of growth on saving. Second, even if a habit-formation model is correct, so that (unanticipated) earnings growth produces higher saving rates, the effects of past earnings growth on saving rates are again quite small. One must look elsewhere to explain a large chunk of the cross-country correlation between saving and growth.

## Endnotes

1. Throughout the paper, labor supply is assumed to be exogenous, and I do not consider how the large increases in labor force participation among women have affected income and consumption. The cohort effects will absorb the effects of these changes insofar as they result in cohort-specific shifts in age-income or consumption profiles. However, changes in age-income and/or age consumption profiles caused by changes in participation patterns will be not identified.
2. These numbers are based on data from the Penn World Table.
3. Declines in fertility and increases in longevity in Taiwan might be partly responsible for the high saving rates among the currently young, since they must plan for a longer old age with fewer children for support, see Deaton and Paxson (1994).
4. There are discrepancies across surveys in how social security contributions (in the US) and National Insurance contributions (in the UK) are treated. These are discussed below.
5. The methods used here are similar to those employed in Deaton and Paxson (1994), for Taiwan, and in the papers on individual countries in Poterba (1994a). Two of the countries I examine - the US and Great Britain - are the subject of two chapters of that book, by Attanasio (1994) for the US, and Banks and Blundell (1994a) for Great Britain, and these papers are extremely useful sources of information about detailed features of the US and the British data. Poterba (1994b) and Banks and Blundell (1994b) provide information on public policies that affect saving in the US and Britain, respectively.
6. An alternative procedure, used by Bosworth et. al. (1991), is to construct annual expenditure for each household by summing up each household's quarterly expenditure records to arrive a full reporting year value of expenditure, and then using the income figures from the 5th interview only (since these pertain to the quarters over which expenditure is measured.) Although this procedure results in income and expenditure figures that always refer to the same time period, the large number of households that skip interviews or drop out of the survey results in the loss of large numbers of observations. Of the 75,726 households whose surveys began between 1980 and 1991, only 37,586 completed all surveys.
7. These and the regression results that follow are estimated using weighted-least-squares, where the weights are inversely proportional to the standard deviation of the cohort means of each variable. Weighted-least-squares and ordinary least-squares results are quite similar.
8. The March CPS's from 1981 through 1993 (which collected information on income from 1980 to 1992) were used to construct cohort data. No information on taxes is collected, so the numbers are based on before-tax income. The CPS covers a large number of households (roughly 55,000 to 60,000 per year).

9. The finding of different age-consumption profiles bears on the evidence presented in Carroll and Summers (1991). Carroll and Summers have argued that if the life-cycle model is valid, and *if* age-specific tastes for consumption are identical across countries, then cross-sectional age-consumption profiles in more rapidly growing countries should show higher consumption among the young relative to the old, since younger cohorts have higher life-time wealth. They find that cross-sectional age-consumption profiles are not tipped toward the young in more rapidly growing economies, and take this as evidence against the life-cycle model. However, the results in Figure 10, which indicate that age-consumption profiles differ across countries even after cohort effects have been removed, makes it difficult to assess Carroll and Summer's results. One interpretation is that age-specific tastes differ across countries, in which case the life-cycle model does not necessarily imply that age-consumption profiles should be tipped toward the young in more rapidly growing economies. Another very plausible interpretation is that the simple life-cycle model is incorrect, and age-consumption profiles are influenced by age-income profiles.

10. Using the actual average saving rate yielded very similar results.

## References

- Attanasio, Orazio, 1994, "Household saving behaviour in the US," in J. Poterba (ed.) *International Comparisons of Household Saving*, Chicago: University of Chicago Press.
- Attanasio, Orazio, and Steven Davis, 1994, "Relative wage movements and the distribution of consumption," NBER Working paper number 4771.
- Attanasio, Orazio, Luigi Guiso, Tullio Jappelli, and Guglielmo Weber, 1992, "The consumption boom in the UK and Italy in the late 1980s," Stanford University, Bank of Italy, Universita' di Napoli, and University College London (May), mimeo.
- Banks, James, and Richard Blundell, 1994a, "Household saving behaviour in the UK," in J. Poterba (ed.) *International Comparisons of Household Saving*, Chicago: University of Chicago Press.
- Banks, James, and Richard Blundell, 1994b, "Taxation and personal saving incentives in the UK," in J. Poterba (ed.) *Public Policies and Household Saving*, Chicago: University of Chicago Press.
- Bosworth, Barry P. 1993, *Saving and investment in a global economy*, Washington, D.C. The Brookings Institution.
- Gary Burtless, and John Sabelhaus, 1991, "The decline in saving: some microeconomic evidence," *Brookings Papers on Economic Activity*, 183–241.
- Carroll, Christopher and Lawrence H. Summers, 1991, "Consumption growth parallels income growth: some new evidence," in B. Douglas Bernheim and John B. Shoven (eds.), *National Saving and Economic Performance*, Chicago. Chicago University Press for NBER, 305–43.
- and David Weil, 1993, "Saving and growth: a reinterpretation," Carnegie-Rochester series, forthcoming.
- Deaton, Angus, 1992, *Understanding consumption*, Oxford. Oxford University Press.
- and Christina H. Paxson, 1994, "Saving, growth, and aging in Taiwan," in D. Wise (ed.) *Studies in the Economics of Aging*, Chicago: University of Chicago Press.
- Feldstein, Martin, and Philippe Bacchetta, 1991, 'National saving and international investment,' in B. Douglas Bernheim and John B. Shoven, eds., *National Saving and Economic Performance*, Chicago. Chicago University Press for NBER, 201–26.



Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1992, 'Precautionary saving and social insurance,' Columbia University, University of Virginia, and University of Pennsylvania, manuscript (forthcoming, *Journal of Political Economy*.)

Levenson, Alec R., 1994, "Do consumers respond to future income shocks? Evidence from social security reform in Taiwan," Milken Institute for Job and Capital Formation, manuscript.

Modigliani, Franco, 1970, "The life-cycle hypothesis of saving and intercountry differences in the saving ratio," in W. A. Eltis, M. FG. Scott and J.N. Wolfe, eds., *Induction, trade, and growth: Essays in honour of Sir Roy Harrod*, Oxford. Clarendon Press, 197-225.

Poterba, James, 1994a, Introduction to *International Comparisons of Household Saving*, Chicago: Chicago University Press.

Poterba, James, 1994b, *Public Policies and Household Saving*, Chicago: Chicago University Press.

Republic of China, 1989, *Report on the Survey of Personal Income Distribution in Taiwan Area of The Republic of China*, Taipei. Directorate General of Budget, Accounting and Statistics Executive Yuan.

Visaria, Pravin, 1980, *Poverty and living standards in Asia*, World Bank: Living Standards Measurement Study.

**Table 1: National income accounts and survey estimates of personal savings rates**

year	US		UK		Taiwan		Thailand	
	NIA	survey	NIA	survey	NIA	survey	NIA	survey
1970	.	.	6.1	3.9	.	.	.	.
1975*	.	.	8.1	6.0	20.0	7.8	15.4	6.5
1980	7.6	14.2	10.2	8.8	17.9	17.1	17.4	.
1981	8.2	6.0	8.4	8.6	20.2	19.0	14.5	-1.1
1982	10.4	15.3	6.7	7.2	20.1	16.7	17.3	.
1983	8.6	14.2	5.0	6.3	21.4	17.3	13.7	.
1984	9.7	3.0	6.3	4.6	22.9	16.3	14.2	.
1985	8.2	13.5	5.7	7.3	23.9	17.1	15.7	.
1986	7.9	2.9	3.9	7.3	28.5	18.6	16.0	-6.1
1987	6.2	5.7	2.0	10.7	28.3	19.2	16.9	.
1988	6.2	3.7	0.2	12.9	23.6	20.5	19.0	-3.2
1989	5.9	4.3	2.0	12.8	19.5	22.2	19.5	.
1990	6.0	4.5	3.2	13.7	20.0	23.3	13.2	2.5
1991	6.5	6.3	5.0	16.5	.	.	13.1	.
1992	6.9	4.8	7.8	5.2	.	.	.	8.0

\* For Taiwan and the UK survey data, the numbers in the "1975" row are from the surveys conducted in 1976.

Notes: The NIA personal savings rates for the US, the UK, and Thailand are from United Nations, National Accounts Statistics: Main Aggregates and Detailed Tables, various issues. The NIA savings rates for Taiwan are from DGBAS, National Income in Taiwan Area, 1991. For all countries the NIA savings rates are defined as the fraction of after-tax income saved by "households and non-profit institutions." The household survey for the UK actually covers only Great Britain (i.e. it excludes Northern Ireland.)

**Table 2: Cohort, year, and age effects in  $\ln(y)$ ,  $\ln(c)$  and  $\ln(y)-\ln(c)$**

Dependent variable		US (obs=663)	Great Britain (obs=1173)	Taiwan (obs=765)	Thailand (obs=306)
$\ln(y)$	cohort effects	3.04	27.50	298.91	5.24
	age effects	12.24	151.55	301.92	16.83
	year effects	3.70	24.91	85.69	151.64
$\ln(y)$	cohort effects	3.12	28.88	44.34	2.26
	age effects	6.81	50.22	300.07	7.89
	year effects	3.13	36.48	94.61	187.72
	adults	.263 (5.22)	.360 (16.35)	.185 (7.77)	.199 (6.26)
	children	-.109 (2.39)	.0346 (2.17)	-.048 (2.23)	-.021 (.077)
$\ln(c)$	cohort effects	15.45	14.59	558.47	9.55
	age effects	33.84	139.32	632.45	24.62
	year effects	41.29	12.32	99.80	86.85
$\ln(c)$	cohort effects	16.88	17.23	73.78	3.30
	age effects	25.92	37.10	677.88	12.49
	year effects	46.29	18.05	123.77	110.98
	adults	.270 (8.78)	.345 (14.89)	.141 (9.32)	.172 (5.81)
	children	-.051 (1.78)	.052 (3.14)	.048 (3.60)	-.050 (1.93)
$s/y$	cohort effects	4.70	5.04	28.34	1.86
	age effects	6.20	7.63	24.38	2.03
	year effects	20.55	11.81	45.05	66.10
$s/y$	cohort effects	4.65	2.91	4.58	1.35
	age effects	6.27	6.02	12.70	2.01
	year effects	20.89	12.05	48.88	55.60
	adults	-.013 (0.31)	.024 (1.32)	.0412 (2.32)	.026 (1.17)
	children	-.070 (1.81)	-.025 (1.75)	-.093 (5.58)	.029 (1.44)

Notes: The dependent variable "s/y" denotes the difference between  $\ln(y)$  and  $\ln(c)$ , approximately equal to the saving ratio. There are two regressions for each of the three dependent variables. The first is a regression on a set of cohort, age, and year effects, where the year effects are constrained to sum to zero and be orthogonal to a time trend. The second adds as regressors the average numbers of adults and children for households in the cohort/age cell. The table reports F-statistics and associated p-values for the cohort, age, and year effects. The regression coefficients for adults and children are reported, along with t-statistics (in parentheses). For all countries the number of age effects is 50. The number of cohort (year) effects is 60 (9) for the US, 72 (21) for the UK, 64 (13) for Taiwan, and 66 (4) for Thailand. The sample consists only of age groups between 25 and 75, inclusive. The regressions are estimated using weighted least squares, as described in the text.

**Table 3: Estimated effects of growth rate changes on aggregate household savings rates**

	s/y no demographics	$\partial(s/y)/\partial g$ no demographics	s/y demographics	$\partial(s/y)/\partial g$ demographics
United States				
growth rate = 0.02	.048	0.27	0.062	0.19
0.04	.053	0.18	0.065	0.11
0.06	.056	0.12	0.067	0.05
0.08	.057	0.07	0.067	0.01
Great Britain				
growth rate = 0.02	0.131	0.02	0.132	0.03
0.04	0.131	0.01	0.132	0.03
0.06	0.131	0.01	0.133	0.01
0.08	0.131	0.01	0.133	0.01
Taiwan				
growth rate = 0.02	0.150	0.12	0.159	0.15
0.04	0.153	0.16	0.162	0.18
0.06	0.156	0.19	0.166	0.22
0.08	0.160	0.23	0.171	0.25
Thailand				
growth rate = 0.02	0.178	-0.08	0.176	-0.05
0.04	0.176	-0.06	0.175	-0.03
0.06	0.175	-0.03	0.175	-0.01
0.08	0.175	-0.01	0.175	-0.00

Notes: For the US, Taiwan, and Thailand, these estimates are based on the age effects in  $\ln(y)$  and  $s/y$  graphed in Figures 4, 8 and 9. For Great Britain, I used the estimates from the 1970-1986 sample period shown in Figure 9. The calculations are described in the text.

**Table 4: Savings and growth within cohorts, US, Great Britain, and Taiwan**

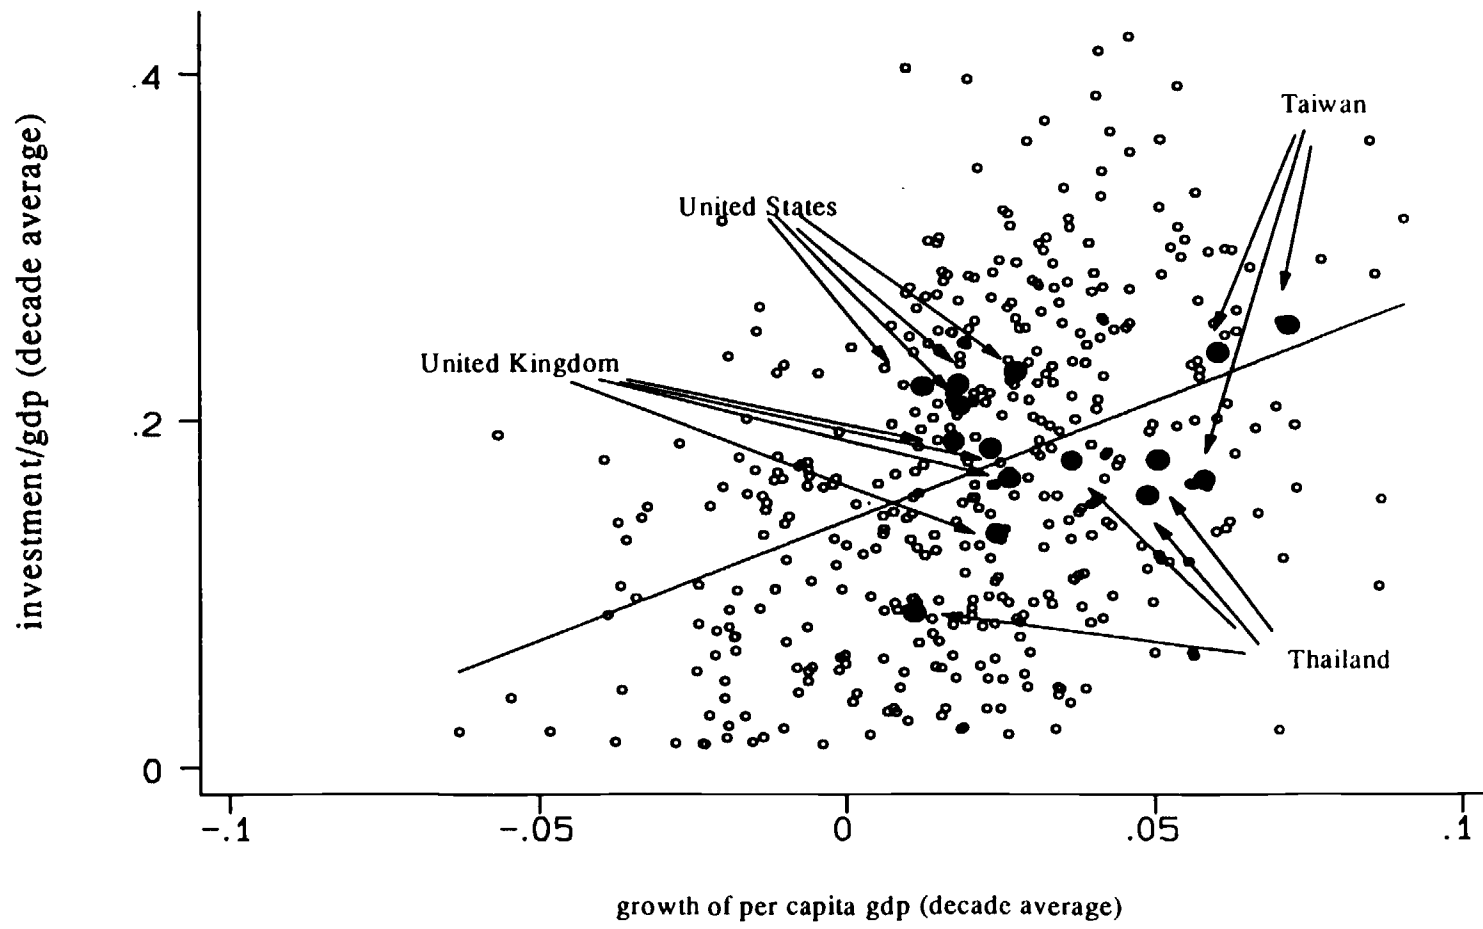
dependent variable: s/y	$\Delta \ln(e)_{t,1}$	$\Delta \ln(e)_{t,2}$	$\Delta \ln(e)_{t,3}$	$\Delta \ln(e)_{t,4}$	sum of coefficients	Age, cohort and year effects?
United States (obs=328)	0.0443 (1.4)	0.0483 (1.4)	-0.0089 (0.3)	-0.0440 (1.3)	0.040 (0.5)	None
	0.0193 (0.5)	0.0349 (0.7)	-0.0150 (0.3)	-0.0554 (1.3)	-0.016 (0.1)	Age
	-0.0527 (1.3)	-0.0528 (1.0)	-0.0974 (1.8)	-0.1237 (2.7)	-0.327 (2.2)	Age+cohort
	-0.0179 (0.5)	-0.0085 (0.2)	-0.0175 (1.4)	-0.0957 (2.2)	-0.194 (1.3)	Age+cohort+year
Great Britain (obs=492)	0.1382 (5.6)	0.1811 (7.1)	0.1485 (6.1)	0.0532 (2.3)	0.521 (8.8)	None
	0.0866 (3.6)	0.1542 (5.9)	0.1267 (4.8)	0.0434 (1.9)	0.411 (5.8)	Age
	0.0810 (3.4)	0.1556 (5.9)	0.1372 (5.3)	0.0467 (2.0)	0.421 (5.8)	Age+cohort
	-0.0014 (0.1)	0.0216 (0.9)	0.0268 (1.1)	0.0103 (0.5)	0.057 (0.8)	Age+cohort+year
Taiwan (obs=410)	0.0969 (9.1)	0.1033 (9.5)	0.0899 (8.4)	0.0497 (5.1)	0.340 (13.1)	None
	0.0270 (2.8)	0.0502 (4.8)	0.0534 (5.4)	0.0250 (3.0)	0.156 (5.7)	Age
	0.0257 (2.7)	0.0619 (5.8)	0.0714 (6.8)	0.0355 (4.2)	0.195 (6.6)	Age+cohort
	-0.0153 (1.6)	-0.0037 (0.3)	0.0125 (1.0)	0.0169 (1.9)	0.0104 (0.3)	Age+cohort+year

Notes: The dependent variable for each regression is "s/y", measured as  $\ln(y) - \ln(c)$ . The dependent variables are lags of earnings growth. The variable  $\Delta \ln(e)_{t,j}$  is defined as  $\ln(e)_{t,j}$  minus  $\ln(e)_{t,j-1}$ . The contemporaneous earnings growth term  $\Delta \ln(e)_t$  is not included since, by construction, measurement errors in  $\ln(e)_t$  will be correlated with those in s/y. The first three columns shows parameter estimates and t-statistics in parentheses. The last column shows the sum of the coefficients reported in the first four, and the number in parenthesis in that column is a t-statistic to test that the sum is significantly different from 0. The sample consists of all age/cohort groups between 25 and 65 in year t. Controls for the numbers of adults and children are included in all regressions. The results for Britain exclude 1987-1992.

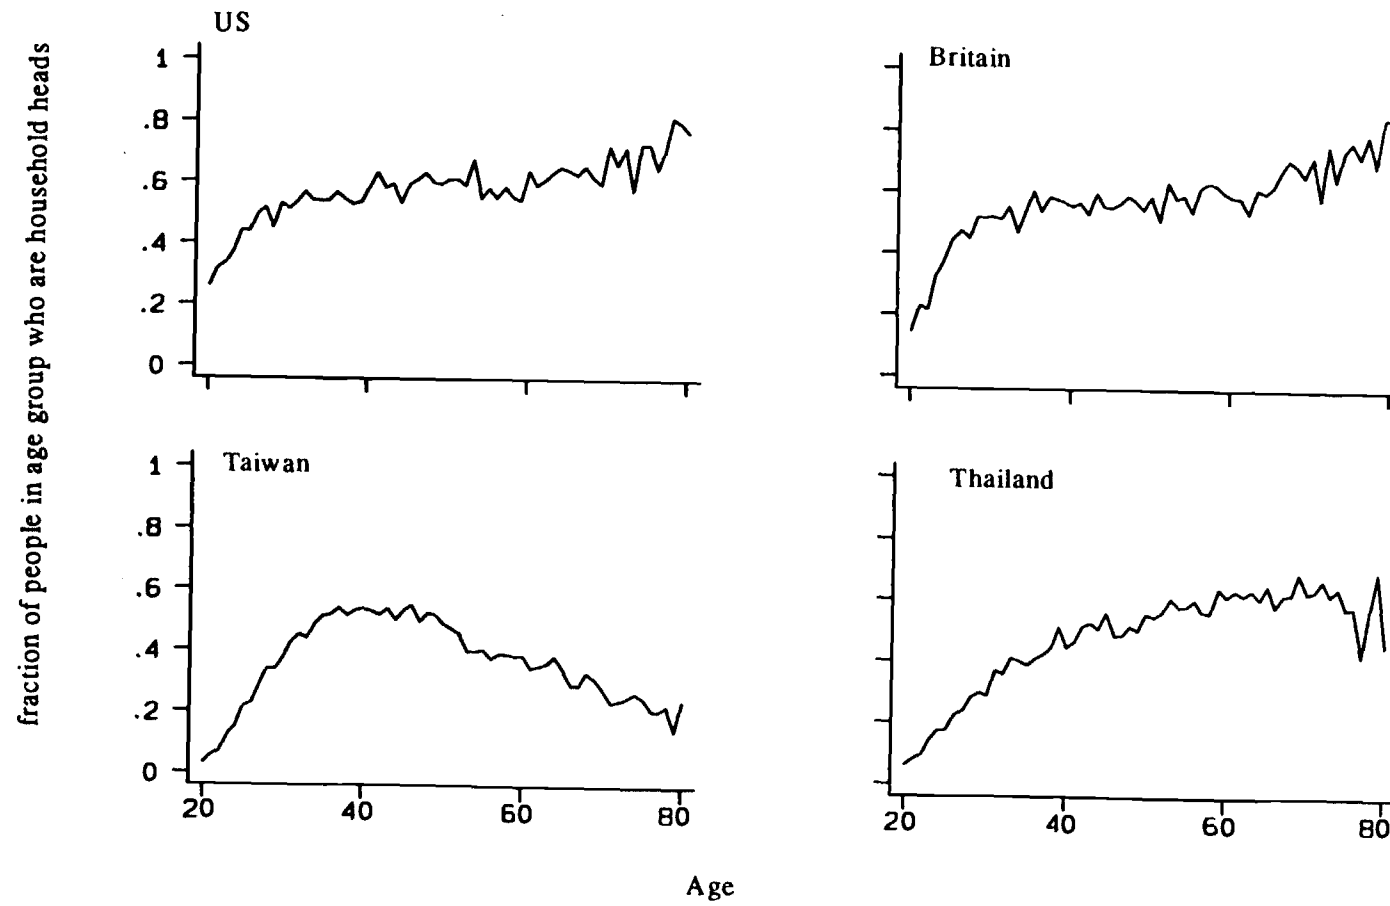
**Table 5: Savings and growth within cohorts, Thailand**

Variable definitions For 1992: $[\Delta \ln(y) = \ln(y)_{1990} - \ln(y)_{1986}]/4$ For 1990: $[\Delta \ln(y) = \ln(y)_{1988} - \ln(y)_{1981}]/7$ For 1988: $[\Delta \ln(y) = \ln(y)_{1986} - \ln(y)_{1981}]/5$		
dependent variable	coefficient (t-stat) for $\Delta \ln(y)$	Age/year effects?
$s/y = \ln(y) - \ln(c)$	0.185 (2.41) 0.249 (2.81) 0.245 (1.25)	None Age effects Age and year effects
Variable definitions For 1992: $[\Delta \ln(y) = \ln(y)_{1990} - \ln(y)_{1986}]/4$ For 1990: $[\Delta \ln(y) = \ln(y)_{1988} - \ln(y)_{1986}]/2$ For 1988: $[\Delta \ln(y) = \ln(y)_{1986} - \ln(y)_{1981}]/5$		
$s/y = \ln(y) - \ln(c)$	0.109 (1.81) 0.099 (1.64) 0.122 (1.07)	None Age effects Age and year effects

Notes: The sample included those aged 25 to 65. The regressions include only savings measures from 1988, 1990 and 1992, although the income growth terms are based on information from earlier years. Absolute values of t-statistics are in parentheses. The sample size is 123 observations. The estimates are calculated using weighted least squares.



**Figure 1:** Investment/GDP and growth of per capita GDP, by decade, 1950-1990. (Penn World Tables)



**Figure 2:** Fraction of people in age group who are household heads (or, in the US, “reference people”.)






Figure 3:  $\ln(y)$  and  $\ln(c)$  for every 5th cohort, UK and US

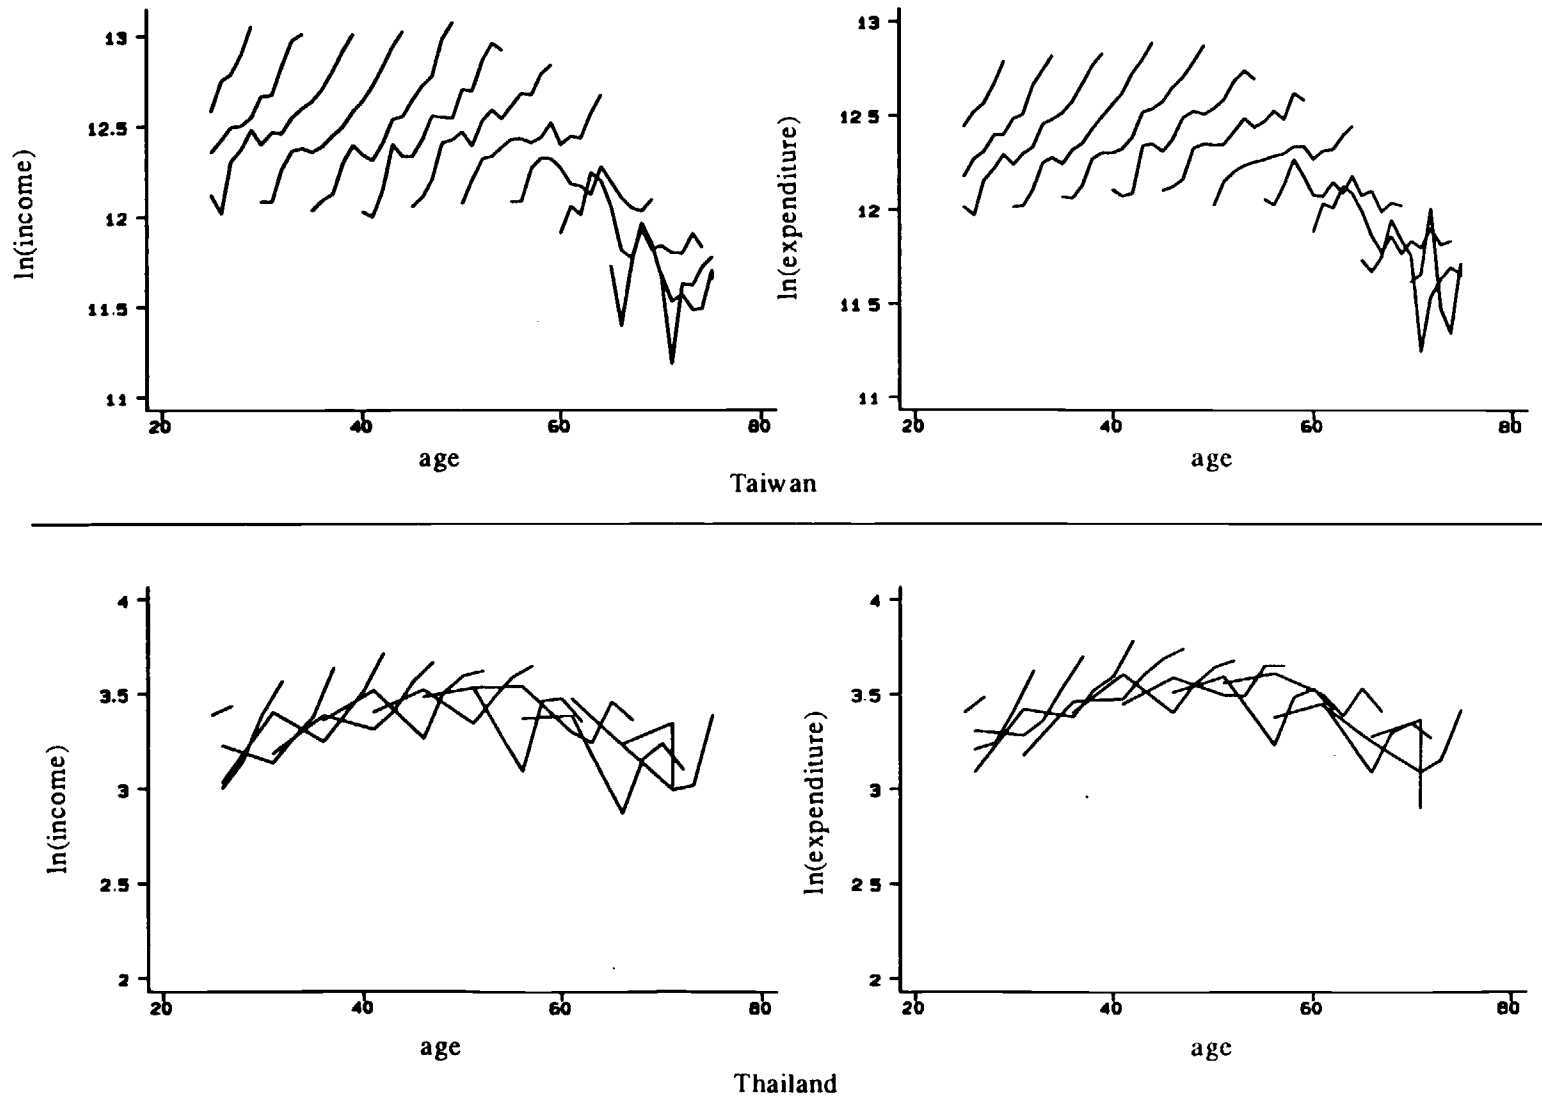


Figure 4:  $\ln(y)$  and  $\ln(c)$  for every 5th cohort, Taiwan and Thailand

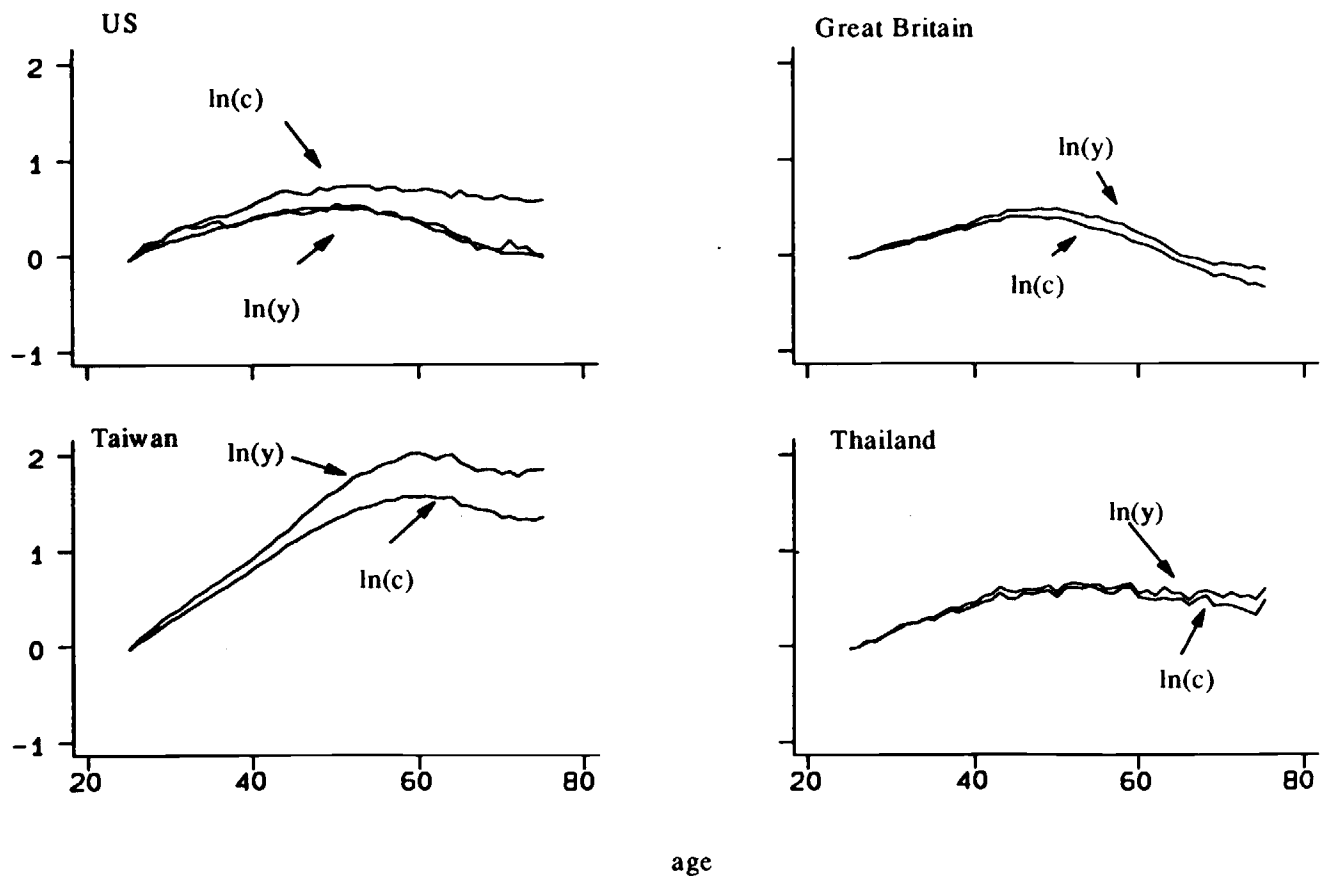
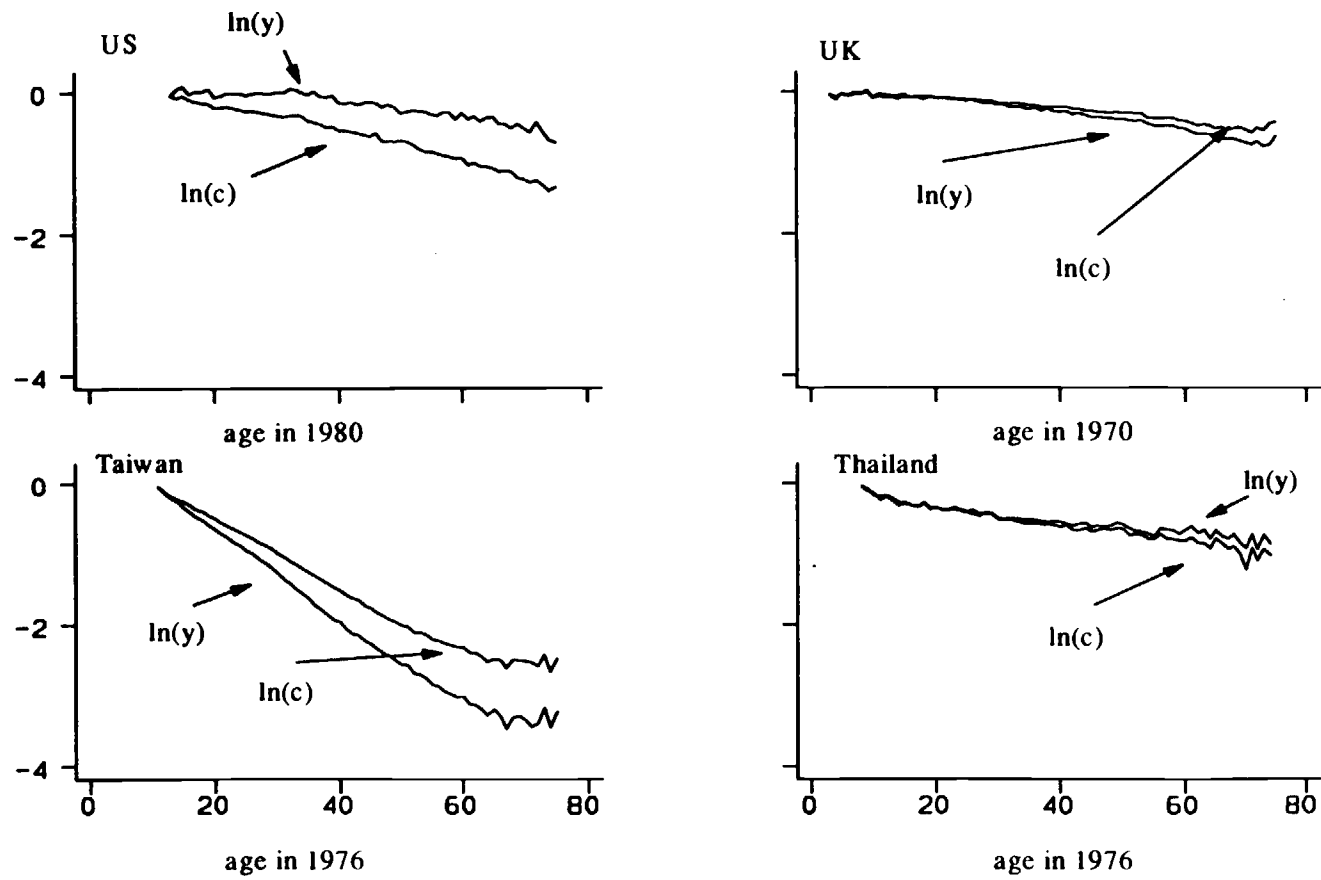


Figure 5: Age effects on  $\ln(y)$  and  $\ln(c)$



**Figure 6:** Cohort effects in  $\ln(y)$  and  $\ln(c)$


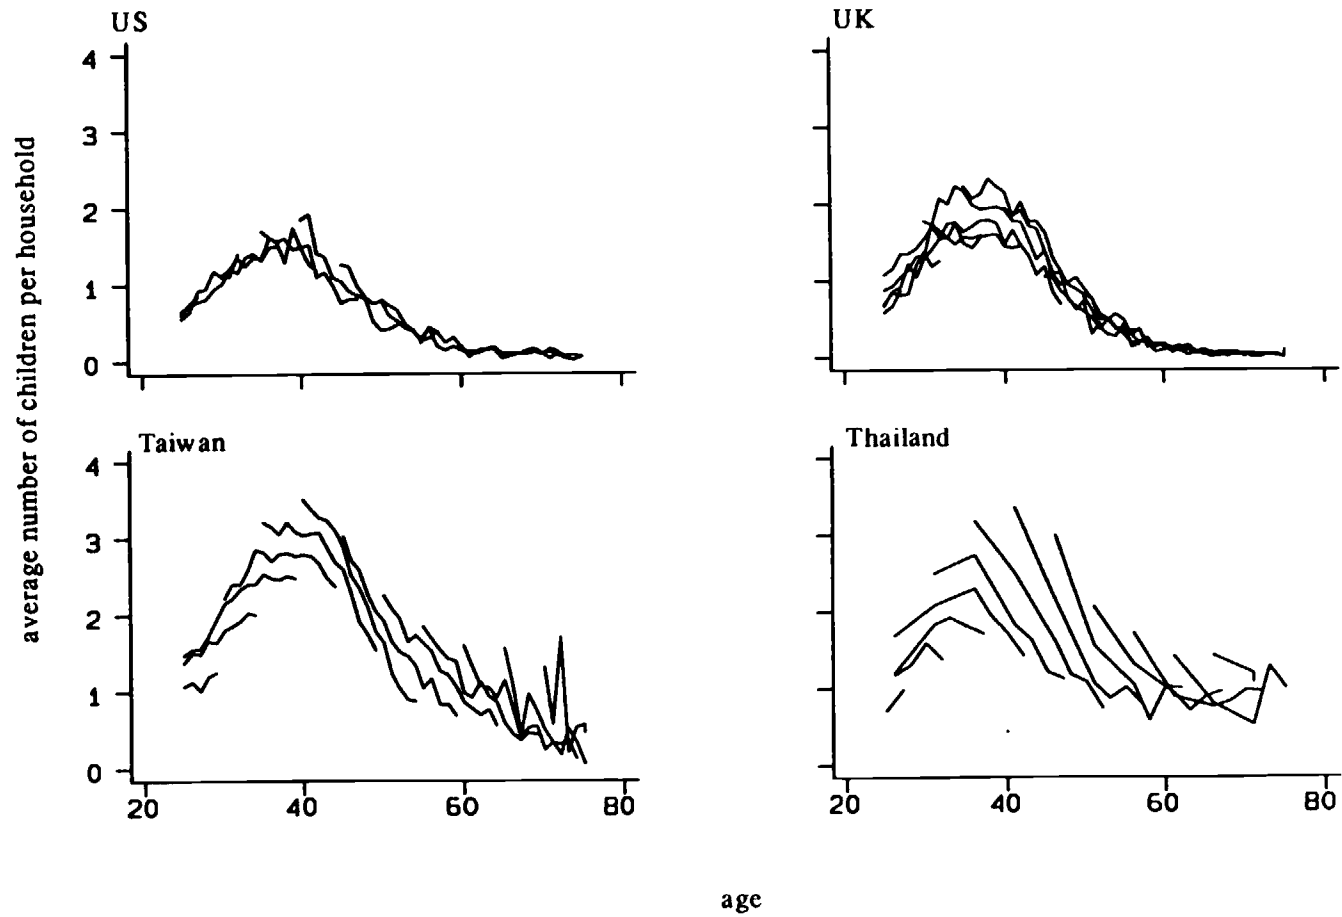
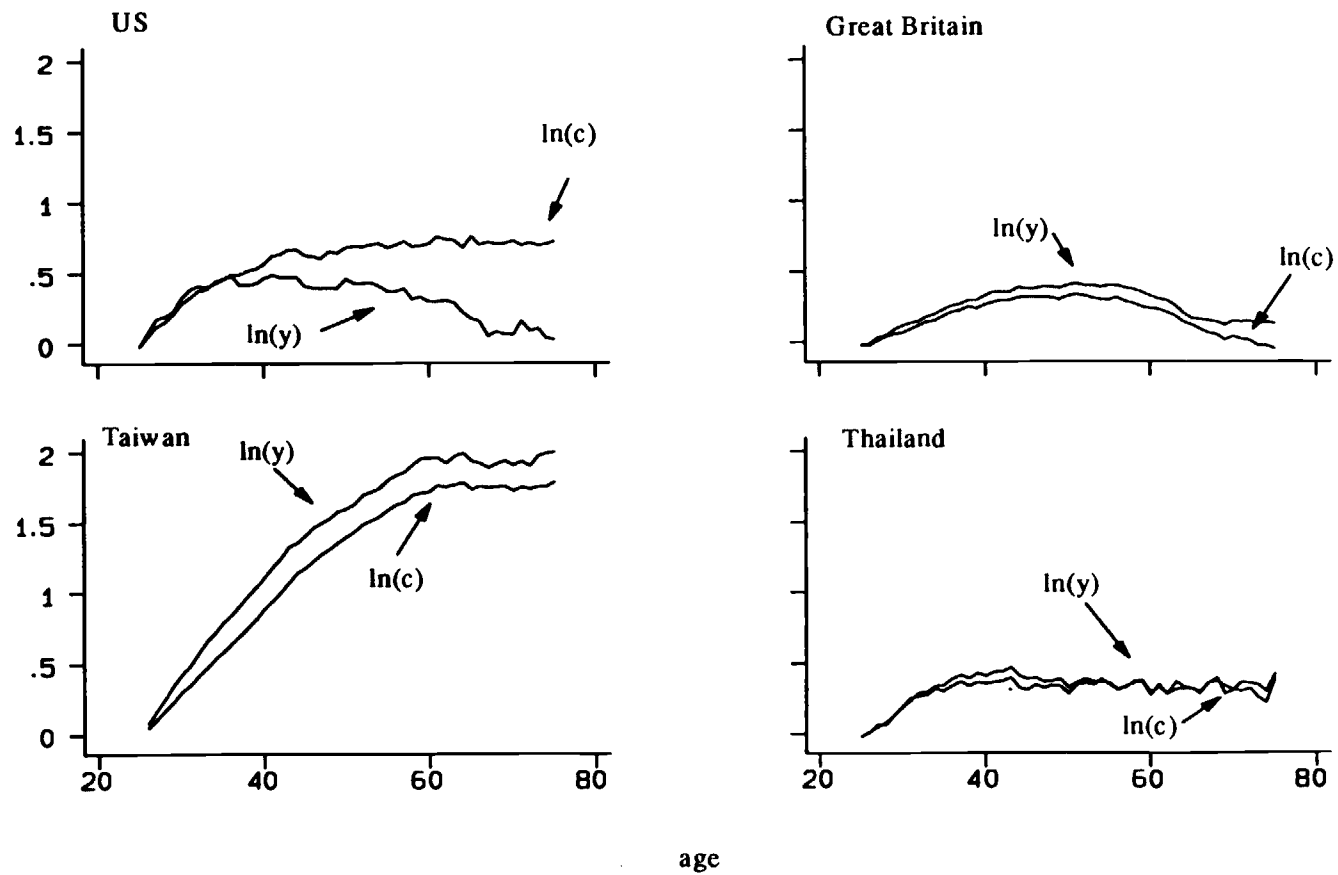


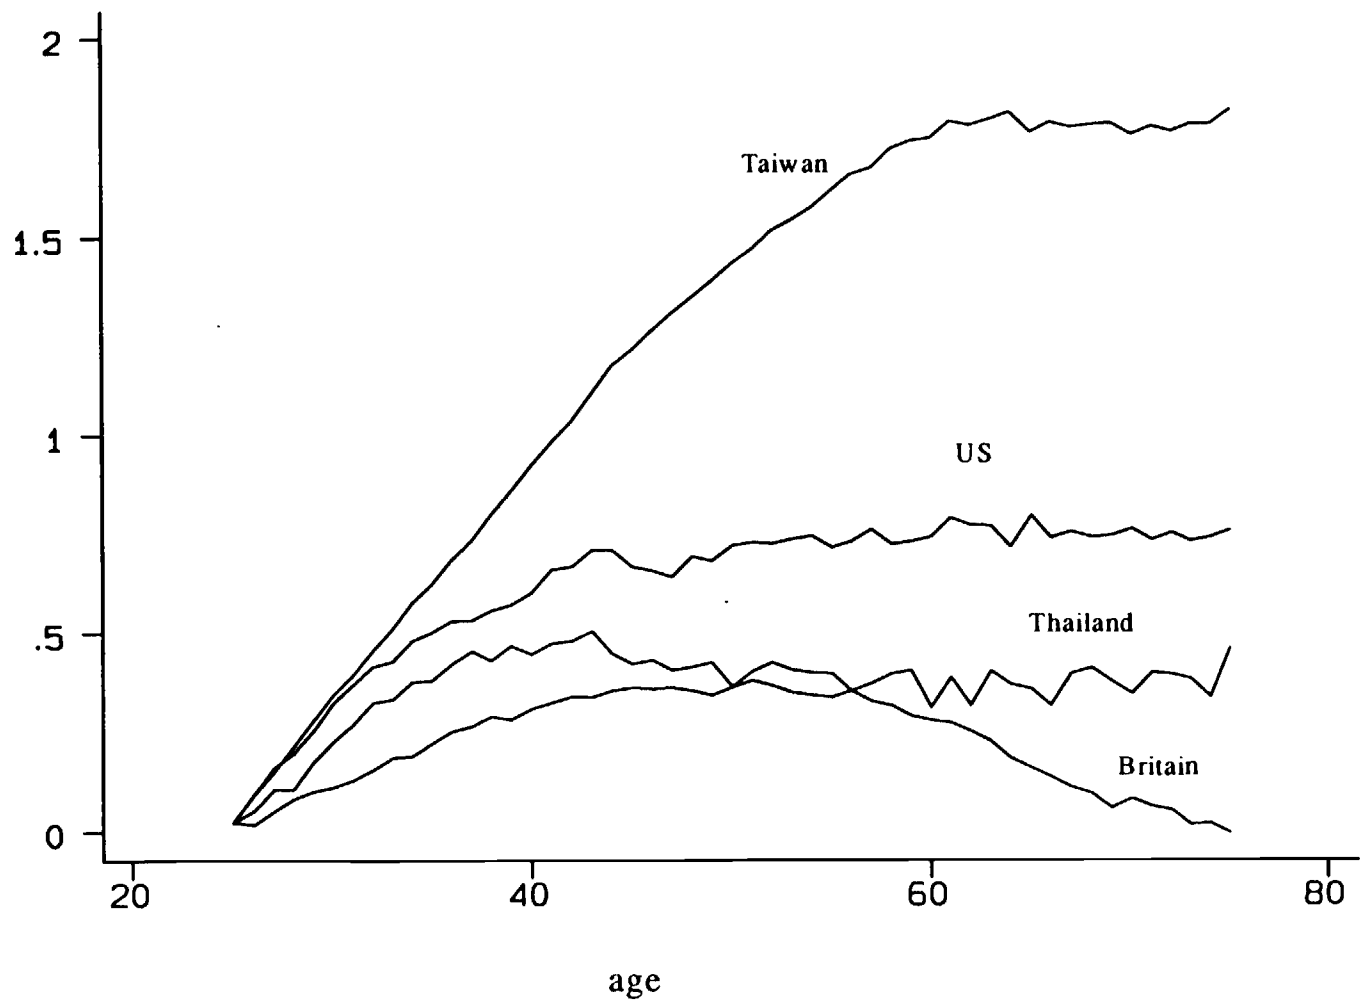
Figure 7: Average numbers of adults per household for every 5th cohort



**Figure 8:** Average numbers of children per household, every 5th cohort



**Figure 9:** Age effects in  $\ln(y)$  and  $\ln(c)$ , controlling for numbers of adults and children



**Figure 10:** Age effects in  $\ln(c)$ , with controls for demographics



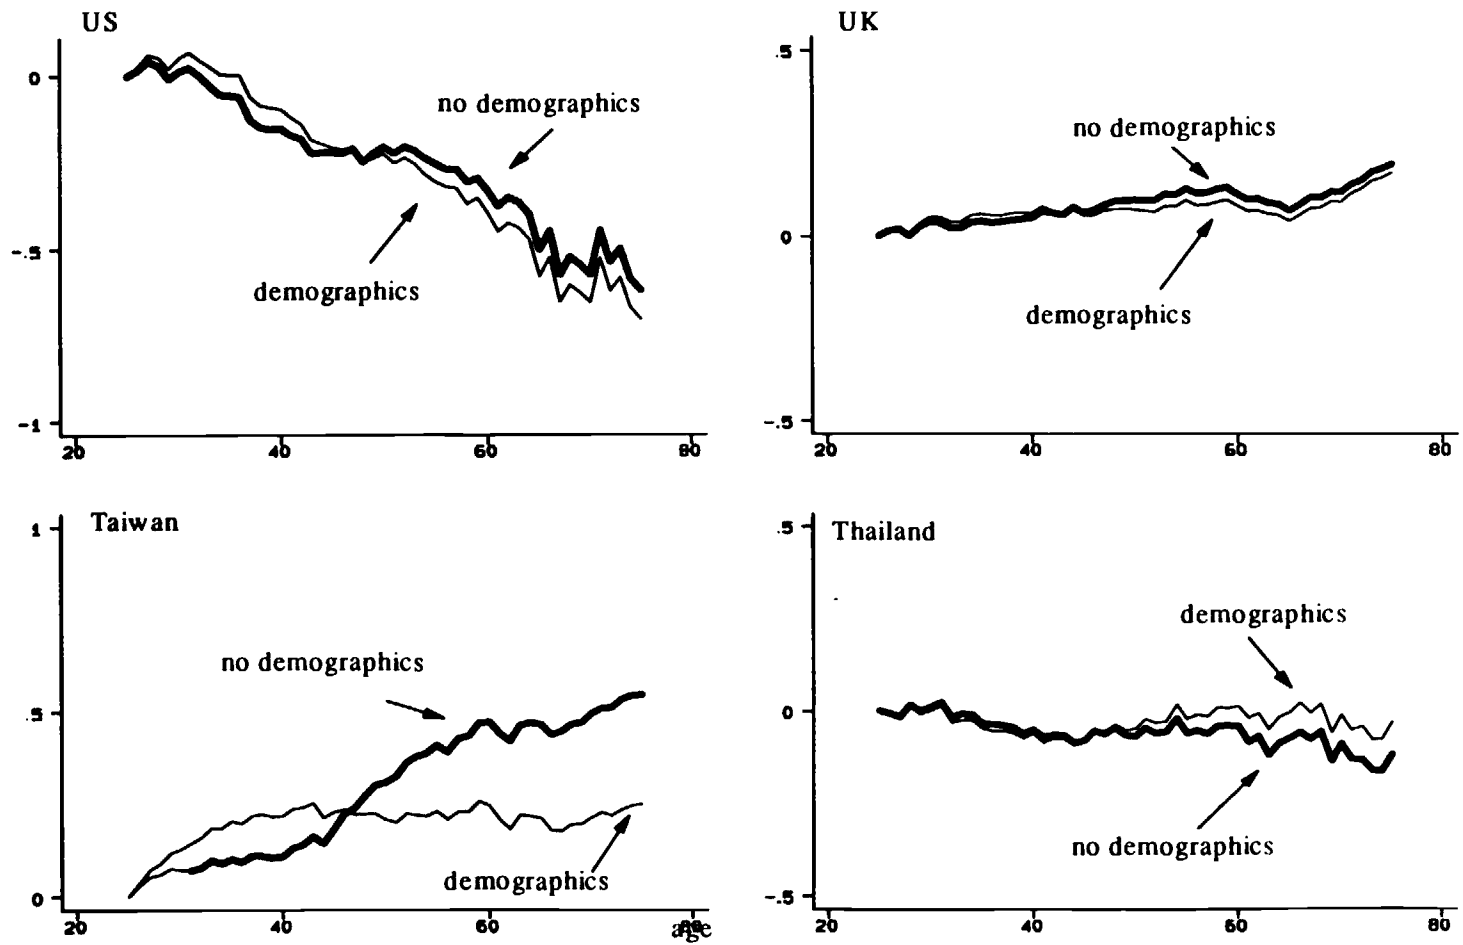


Figure 11: Age effects in savings rates - measured as  $\ln(y) - \ln(c)$  - with and without control for adults and children

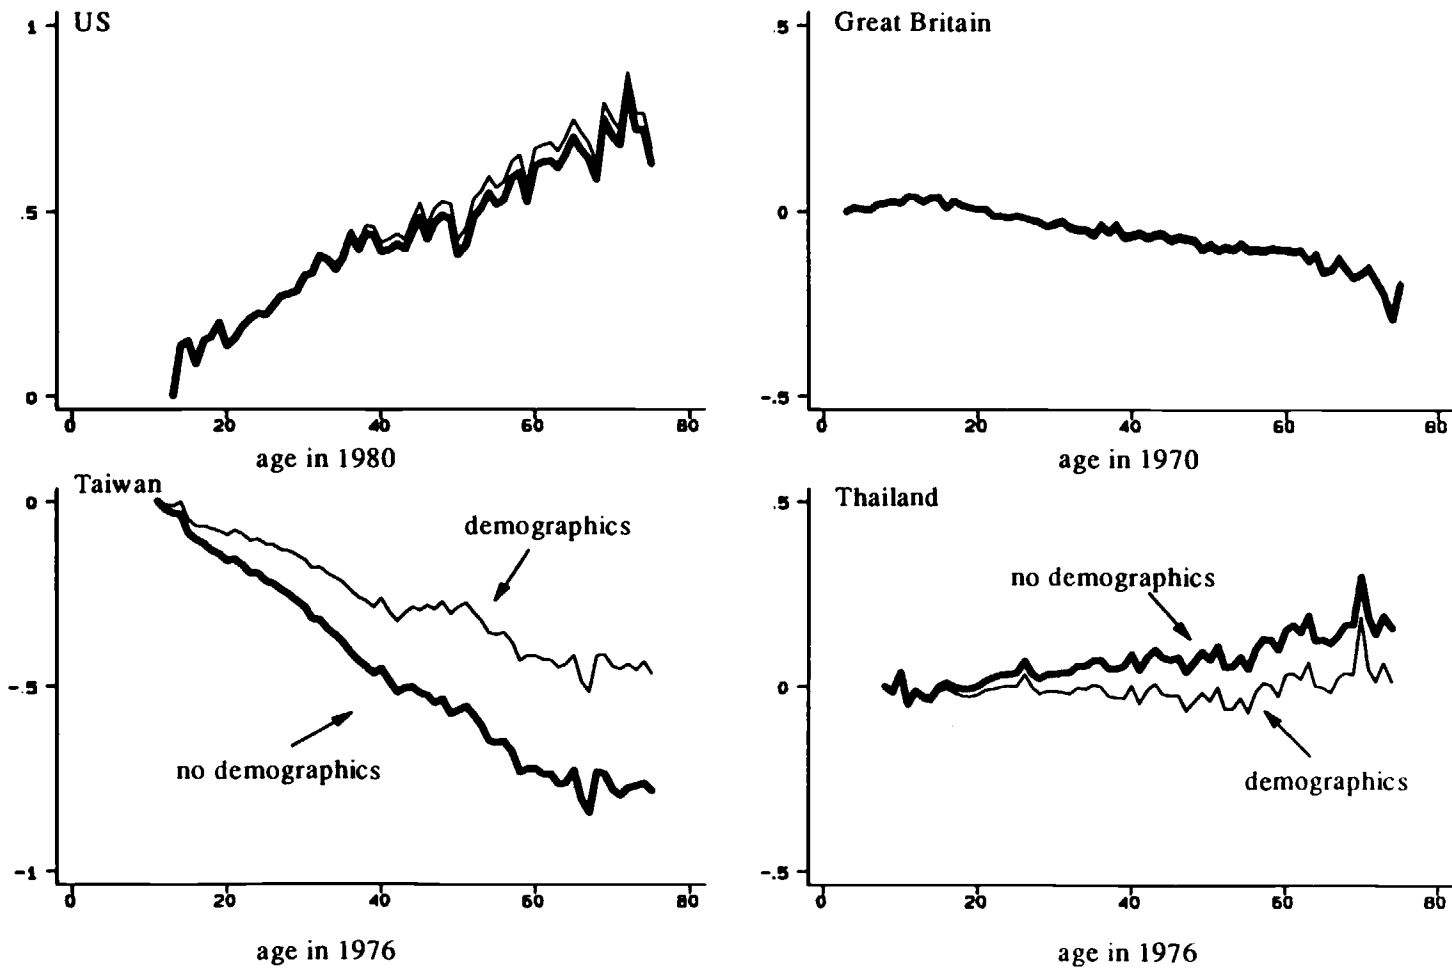


Figure 12: Cohort effects in savings rates


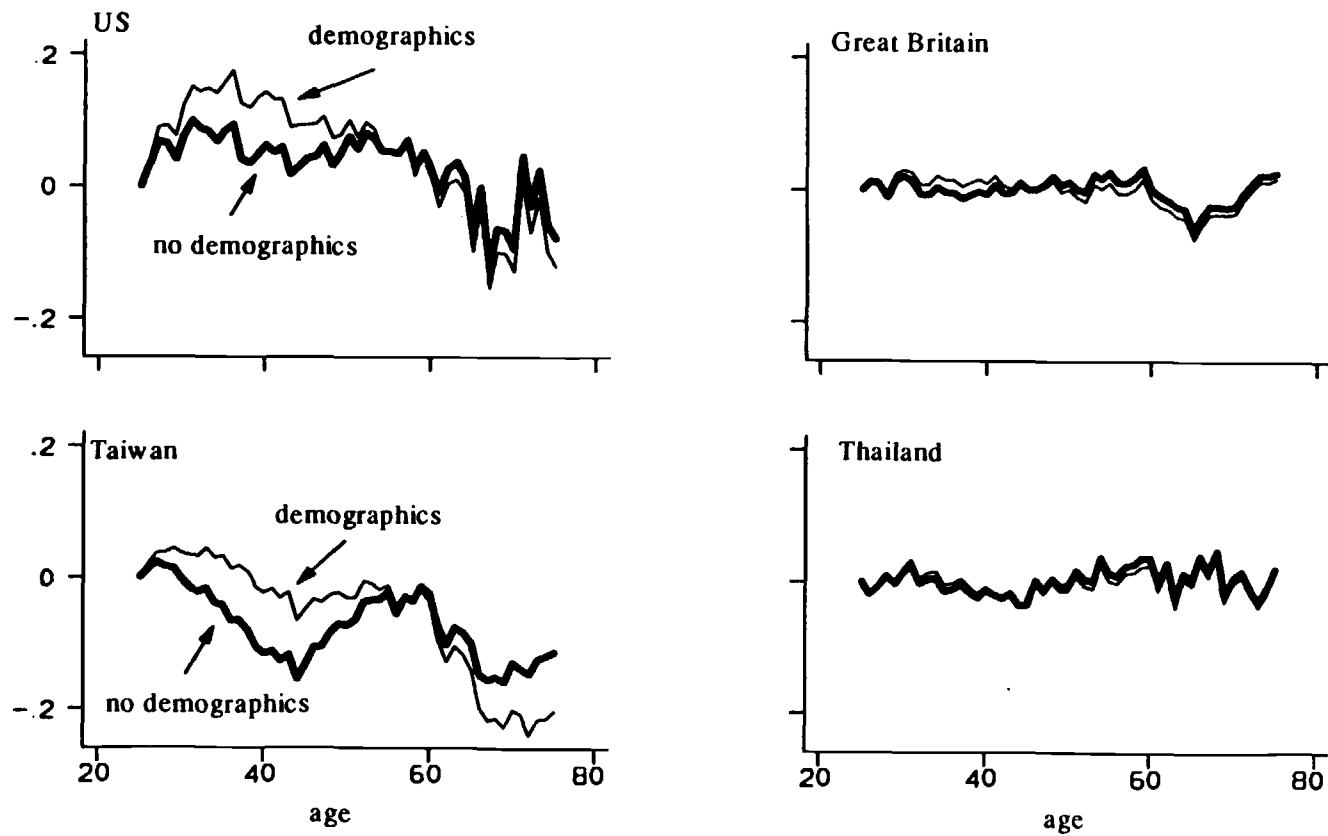


Figure 13: Age effects in  $\ln(y)$ ,  $\ln(c)$  and  $s/y$  in Britain, 1970-1992 and 1970-1986



**Figure 14:** Age effects in saving rates, no cohort effects and unrestricted year effects




Figure 15: Age effects in saving rates using cohort means and medians, demographic effects included