Household Saving in East Asia and Latin America: Inequality, Demographics, and All That

Orazio P. Attanasio and Miguel Székely

This article compares household savings behavior in East Asia and Latin America by analyzing data on four economies—Mexico, Peru, Taiwan (China), and Thailand. It provides the first comparisons of saving in these two regions at the microeconomic level using synthetic cohort techniques. The article does not focus only on total household saving, as is common in the literature. Instead, the population is disaggregated into education groups to determine whether differences in savings behavior are related to the distribution of income. Forecasts of future aggregate household savings rates are constructed based on demographic projections. The article provides evidence that allows testing of the relevance of the life-cycle model for explaining differences in savings behavior.

Lack of efficient credit and insurance markets makes household savings a crucial determinant of welfare in developing countries. Without savings, households have few mechanisms to smooth unexpected variations in their income. Shocks may thus leave permanent scars (by interrupting human capital accumulation at early ages, for example). Since saving is one of the few means of accumulating assets in the absence of credit and insurance markets, the capacity to save becomes critical for increasing social mobility and enhancing future income-earning possibilities. Additionally, although there is controversy over the relationship between savings and economic growth, it is generally agreed that once savings start to rise—perhaps as a result of increases in income—they enhance the potential to finance investment and lead to the creation of more opportunities in the economy.

Few studies have examined savings behavior at the microeconomic level in developing countries; with very few exceptions, comparisons have focused mainly on

Orazio P. Attanasio is professor at the Institute of Fiscal Studies, University College London. Miguel Székely is research economist in the Research Department of the Inter-American Development Bank. The authors would like to thank Ricardo Fuentes for excellent research assistance, Christina Paxson for providing the Thai data and some of the programs for reading them, Martín Valdivia and Jaime Saavedra for providing the Peruvian data, and an anonymous referee for comments.

Annual World Bank Conference on Development Economics 2000

aggregate savings data. In this article we use data from household surveys to compare and try to understand savings behavior at the microeconomic level in two Latin American economies (Mexico and Peru) and two East Asian economies (Taiwan [China] and Thailand). These economies were chosen because they are among the few in each region for which household data on income and consumption are available for a series of years. The experiences of these regions have differed remarkably in many dimensions over recent decades. An analysis based on household behavior may shed some light on why they have diverged.

We contribute to the literature on household savings behavior in four ways. First, to our knowledge this is the first time that savings behavior has been compared at the microeconomic level for these regions. Analyzing household saving rather than aggregate saving is useful because without microeconomic data it is very difficult, if not impossible, to interpret aggregate savings trends and to discriminate among alternative models.

Second, rather than focusing only on total household saving, as is common in the literature, we disaggregate the population into groups by education level to determine whether differences in savings behavior are related to differences in income. By focusing on education levels rather than income levels, we focus on permanent income effects rather than cyclical fluctuations. This analysis is crucial for determining whether different sectors of the population have different motives for saving and different capacities to smooth variations in income in the face of shocks and to build up assets.

Third, we use our analysis of age profiles of household saving to forecast aggregate household savings rates based on demographic projections. We are less interested in providing an efficient forecast of future savings rates than in assessing the extent to which projected demographic changes have the potential to bridge (or intensify) the differences in household saving between these two regions. The hypothesis that changes in demographic structure can affect the process of growth and saving in fundamental ways has received renewed attention. It has been argued that recent demographic shifts in East Asia—with the age groups that produce and save more growing relative to other age groups—are among the main reasons why economic performance and savings in that region improved (see Bloom and Williamson 1999 and Behrman, Duryea, and Székely 1999a). Latin America is on the verge of experiencing similar—though somewhat smaller—demographic changes, so it is of interest to verify the extent to which these shifts will contribute to increased savings in the region.

Fourth, microeconomic data allow different theories of household saving to be tested, in particular, the relevance of the life-cycle model in its different incarnations. We provide evidence that helps assess the usefulness of the theory for explaining the facts.

The evidence we present is of interest because the two Latin American countries we study experienced economic shocks during the period under analysis that are similar to those experienced recently in East Asia. Thus our analysis may provide insights into the response of saving to shocks and the capacity of different groups of the population to smooth the effects of those shocks. These insights may help policymakers protect vulnerable groups in East Asia.

We complement our analysis of savings behavior by examining a number of household decisions that are closely related to the intertemporal allocation of resources. For example, savings are likely to be affected by demographic variables through household composition effects, so heterogeneity in this dimension is likely to be important. It is also very likely that savings behavior is linked to labor supply decisions, particularly labor force participation. Additionally, if different groups of the population, such as those with different endowments of human capital, face different life-cycle earnings profiles, they will have different incentives to save. It is therefore important to characterize the behavior of different groups, as well as changes in household composition and labor force participation over time.

An analysis of this type is not without limitations. Both micro- and macroeconomic data are affected by severe measurement problems. Not only are they measured with error, but they often measure different concepts. Differences in the definition of consumption (particularly for such items as housing and health expenditure), in the population of reference (which is typically much smaller for survey data), and in the treatment of income sources (especially for pension benefits, interest income, capital gains, and imputed rents) all prevent a direct comparison between aggregate measures of savings rates and those derived from microeconomic data sources. Moreover, in the national accounts of many developing countries national saving is not disaggregated into private and public saving, and private saving is not broken down into household and corporate saving. In the few microeconomic data sources available, data on asset ownership, pension entitlements, and the like are of very limited scope and quality. Thus matching aggregate private saving with microeconomic data is not easy. Even if we think of households as the ultimate owners of corporations and assume that they are able to "pierce the corporate veil," aggregate private saving and microeconomic data may differ if foreign investors own some firms.

In addition to these measurement problems, there are important conceptual problems. The main one is that some expenditure items, such as durables, housing, education, and health, have important savings components, but it is difficult to establish how large those components are. We cut through these issues by making some strong assumptions and trying alternative definitions of consumption and saving. In the end, however, given the data, some problems cannot be solved and the results will thus have to be interpreted with caution.

Another issue that needs to be borne in mind is that while saving is an intrinsically dynamic phenomenon, microeconomic data typically do not follow the same individuals over time because they lack a genuine longitudinal dimension. To obviate this problem, we make extensive use of the synthetic cohort techniques pioneered by Browning, Deaton, and Irish (1985). The basic idea is to follow the average behavior of groups whose membership is assumed to be fixed over time. This procedure allows us to study the dynamic behavior of the average of the variable of interest in different years. Even this technique, however, is not free from problems. The most important are the endogeneity of household formation and dissolution and the differences in mortality and migration rates across socioeconomic

groups (see Behrman, Duryea, and Székely 1999b and Attanasio and Hoynes 2000). We discuss some of these issues below.

Aggregate Savings Trends in East Asia and Latin America

One of the main differences between the development experiences of East Asia and Latin America is that East Asia has had much higher savings rates for some time. National savings rates were already slightly higher in East Asia by 1970, and rates in the two regions have diverged dramatically since then (figure 1). (Rather than presenting means by region, which are quite "noisy," we present smoothed profiles obtained by regressing the savings rate on a dummy variable for each year and controlling for all country characteristics.)[1] National savings rates increased in Latin America in the early 1970s, collapsed in 1982 with the onset of the debt crisis, and recovered slightly in the second half of the 1980s. In contrast, in East Asia saving increased almost continuously throughout the period, so that by 1997 the gap between the two regions' national savings rates was about 20 percentage points.

The story is similar when we look at domestic savings as a share of GNP—and the differences even more apparent (figure 2). The average domestic savings rates for the two regions were the same in 1965, but they soon diverged and the gap has widened continuously since then. By 1997 the average domestic savings rate in Latin America was about 17 percent, while the rate in East Asia was almost 40 percent.

Figure 1. National Saving as a Share of GNP in East Asia and Latin America, 1972–97

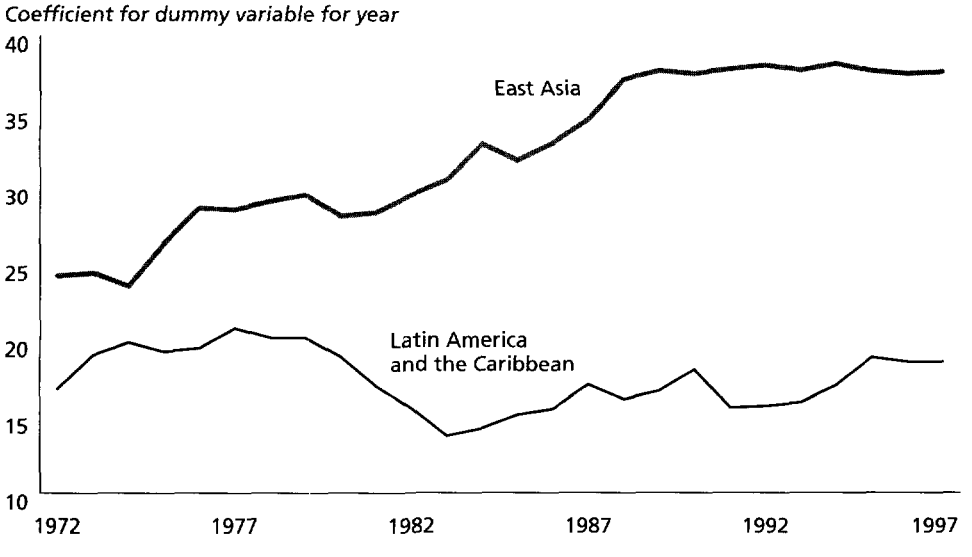

Note: The East Asian aggregate includes only Hong Kong (China), the Republic of Korea, Singapore, and Thailand, some of the fastest growing economies with high savings rates. Taiwan (China) is not included because of lack of data. The Latin American aggregate includes all countries in the region for which data were available.
Source: World Bank 1999.

Thus most of the difference between the regions' overall savings rates can be attributed to the different patterns in domestic saving.

Several interesting features emerge from an examination of the domestic savings rates in the East Asian and Latin American economies for which data are shown in figure 2. For our purposes, the most important is that the 4 East Asian economies started out with lower domestic savings rates than the 12 Latin American countries in 1960, but they experienced a huge increase during the following 28 years. Except for Chile, no Latin American country among the 12 studied enjoyed a savings rate that was significantly higher in 1997 than in 1960. (In Ecuador and Mexico the savings rate increased, but it did so by less than 10 percentage points—a much smaller increase than in East Asia, where the average savings rate rose 30 percentage points.)

Also significant is that the savings rate in Latin America is much more volatile than that in East Asia. The four East Asian economies (Hong Kong, the Republic of Korea, Singapore, and Thailand) show a much smoother pattern and one of continuous increases. In contrast, in Latin America the savings rate changed dramatically from year to year in some countries (such as El Salvador, Nicaragua, and Peru), and in no country was there a positive trend throughout the period. By 1997 all four East Asian economies had much higher savings rates than any of the Latin American countries.

These differences are especially important in light of our microeconomic analysis, because they indicate that we are comparing two economies that have gone through intensive crisis and volatility in their savings rates with two in which domes-

Figure 2. Domestic Saving as a Share of GNP in East Asia and Latin America, 1965–97

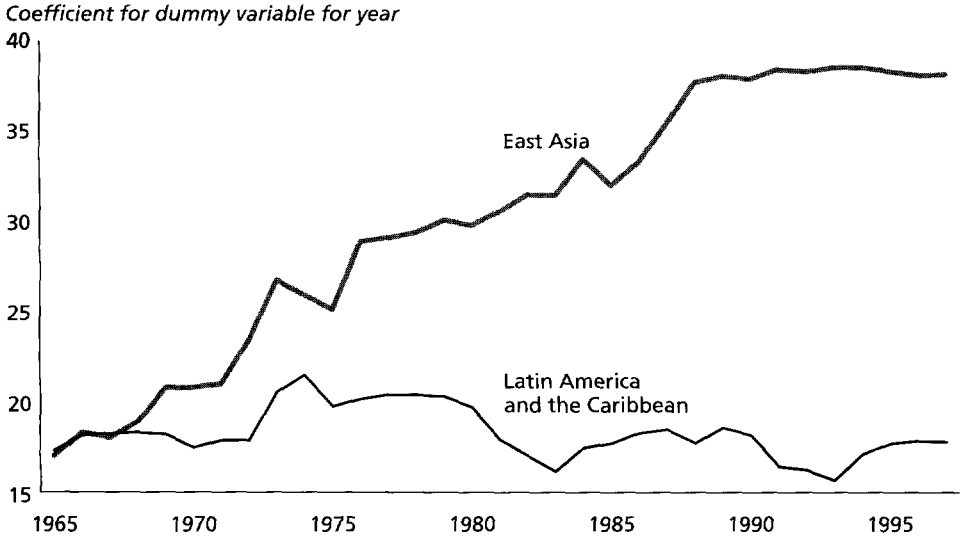

Note: The East Asian aggregate includes only Hong Kong (China), the Republic of Korea, Singapore, and Thailand, some of the fastest growing economies with high savings rates. Taiwan (China) is not included because of lack of data. The Latin American aggregate includes all countries in the region for which data were available.
Source: World Bank 1999.

tic saving has increased continuously and smoothly. In interpreting our results, we will have to keep this in mind.

At the same time that domestic savings patterns in East Asia and Latin America were diverging, there were also significant differences in other dimensions. Total fertility rates declined much faster in East Asia than in Latin America. In 1965 the demographic structures in the two regions were almost the same, but since then there has been a growing gap. Fertility in both regions started declining in the 1950s, but the reduction was much faster and steeper in East Asia, and by 1965 the cohorts entering working age were much larger than the newborn ones (United Nations 1998). Thus since 1965 a much larger share of the East Asian than the Latin American population has been entering the ages in which savings rates rise. This composition effect may be an important force behind the differences in saving.

The reductions in fertility in the two regions are highly correlated with sharp increases in female labor force participation. Fertility and female labor force participation are usually jointly determined, and they have a double effect on savings behavior: lower fertility rates mean that households have fewer children on average, while higher labor force participation means that more household members are in the workforce (and thus income is higher). The result of both effects is higher household per capita income and thus higher savings capacity. Female labor force participation was lower in Latin America than in East Asia in 1960 (World Bank 1999). Although it has increased slightly faster there than in East Asia, the gap between the regions remains large.

Another important transformation taking place in the two regions at the time their savings patterns were diverging was the transition toward higher levels of schooling. Here too East Asia enjoyed much faster progress than Latin America. In 1960 average schooling in East Asia was 0.7 years longer than in Latin America; by 1990 the difference had grown to 2.5 years (Behrman, Duryea, and Székely 1999b). More educated individuals usually have higher incomes and thus higher savings capacity, so this is another potentially important factor behind the differences in domestic saving.

Yet another difference between the two regions has been the rate of growth. It is not clear whether economic growth precedes higher savings rates or higher savings rates precede growth. But since the mid-1970s East Asia has had higher GDP per capita (adjusted for purchasing power parity). Differences in growth rates have been apparent since the mid-1960s, precisely when domestic savings rates surged in East Asia.

Methodology and Data

The main purpose of the analysis presented in the following sections is to characterize the patterns of household saving over the life cycle. The conceptual framework is the life-cycle model, although we do not take a stance on the version of the model (with or without precautionary saving, liquidity constraints, bequest motives, or habit formation) that might be most appropriate to describe the data. In addition to the life-cycle profiles of saving for the population at large, we focus on the differences in savings behavior among different groups of households. We also relate

the observed savings behavior to other variables, particularly demographic ones, which are likely to be important determinants of saving.

This type of analysis is useful for two reasons. First, the focus on different groups of the population—characterized by different earnings profiles, demographics, and shocks over the sample period—could shed light on the determinants of saving. We also emphasize important differences across groups that remain hidden in aggregate analysis. This is particularly important for Latin America, which is characterized by substantial inequality. Second, giving age profiles of saving a semistructural interpretation allows extrapolation of the relationship that links savings rates to demographic variables—and thus forecasts of future household savings rates.

The type of microeconomic data available dictated the technique we use. Unfortunately, most Latin American countries with household surveys that provide information on both income and consumption include only one or two data points. Because we believe that cohort effects are particularly important, we focus on the two Latin American economies (Mexico and Peru) and two East Asian ones (Taiwan [China] and Thailand) for which data are available for several points in time. Because we study a dynamic phenomenon and use time series of cross-sections, we are forced to use synthetic cohort techniques. These techniques allow us to follow the same groups of individuals over time, even in the absence of a genuine longitudinal dimension in the data.

Methodology

In the presence of strong cohort effects, the age profile of savings rates, consumption, or any other variable in a single cross-section may not correspond to the age profile of any individual. To circumvent this problem, we group the observations in each of several repeated cross-sections according to one or more variables chosen so that the group membership from which the observations are drawn is likely to be fixed. In this way, rather than following the behavior of individuals over time, we follow the average (or any other moment) of the variables of interest for the groups. In the context of a life-cycle model an obvious way to form groups is on the basis of the year of birth of the household head, so that we can follow the behavior of different cohorts as they go through different phases of their life cycle. We also consider education groups, under the assumption that the accumulation of human capital can be done only in the early phases of the life cycle.

While synthetic cohorts are extremely useful, the technique is not free from problems, since group membership might change over time and household formation and dissolution could be endogenous to the phenomenon under study. Different mortality and migration rates can also induce changes in composition. Endogeneity of household formation and dissolution is relevant if the propensity to form a household at the beginning of the life cycle is different in different groups of the population and household dissolution results in elderly individuals going to live with their offspring. Extended families and family arrangements in old age are particularly relevant for our analysis because they are directly related to life-cycle saving and the incentives to save.

Most of the analysis we present here is graphical. To identify the life-cycle profile of several variables of interest, we plot the average data for each cohort against age. Because different cohorts are observed over different parts of their life cycle, we are able to track the age profiles. Moreover, if the sample period covered by the time series of the cross-section is longer than the interval used to define a cohort, we observe different cohorts at the same age (although obviously at different points in time).

An important identification caveat should be kept in mind. While it is true that with a sufficiently long sample period one observes different cohorts at the same age, one should resist the temptation to interpret the resulting differences as always due purely to cohort effects. The obvious reason is the possible presence of year effects. We use smoothing techniques to present age and cohort effects, but strictly speaking, age and cohort effects can never be disentangled without additional information or restrictions from time effects because of the exact linear relationship linking age, time, and year of birth. While in some cases, such as demographic variables, it is natural to impose the absence of year effects, in other cases this assumption is a strong one and the results should be interpreted with caution. One should always remember that any combination of cohort and age effects can be obtained as a combination of age and time or time and cohort effects. We discuss these issues below.

Once we estimate the age effects for savings rates, we extrapolate them to forecast future aggregate savings rates. In particular, we use the following relationship. If we indicate the aggregate savings and income at time t by S_t^{ag} and Y^{ag} and the savings, income, and size of group c (cohort) at time t by S^c, Y^c, and N^c , the aggregate savings rate will be given by the following expression:

(1)
$$sr_t^{ag} = \frac{S_t^{ag}}{Y_t^{ag}} = \frac{\sum_c S_t^c N_t^c}{\sum_c Y_t^c N_t^c} = \sum_c \frac{S_t^c}{Y_t^c} w_t^c$$

where
$$w_t^c = \frac{Y_t^c N_t^c}{\sum_c Y_t^c N_t^c} .$$

If we assume that group savings rates are a function of age and cohort effects, we can extrapolate from the group savings rates

$$\frac{S_t^c}{Y_t^c}$$

estimated in the microeconomic data using equation 1, relative income profiles, and demographic projections to forecast future savings rates. These forecasts should be treated with extreme caution, however, because they are based on the behavior in a given economic environment (that is, by households facing given earnings profiles and demographic profiles and a given set of institutions, including arrangements for

old age). Moreover, the data problems mentioned above make it very difficult to match the micro- and macroeconomic measures of saving. Nonetheless, these forecasts are informative about the potential effects of demographic trends and changes in the composition of the population on aggregate saving.

If the groups are defined not only by the year of birth of the household head but also by educational achievement, forecasting using the aggregate savings rate in equation 1 becomes even harder. This is because it involves forecasting not only the age structure of the population (for which demographic projections can be used) but also the accumulation of human capital. General equilibrium effects—in particular, the effect that the relative size of different skill groups might have on the returns to human capital—complicate this type of exercise even further. In principle, in the absence of cohort effects one can use the age profile from a single cross-section to perform the same exercise. Caution is needed in interpreting the results of such an exercise, however.

Even if one does not want to disentangle age and cohort effects and consider equation 1 at two points in time, one can use equation 1 to decompose the changes in aggregate savings rates into changes due to shifts in the cross-sectional age profile and changes in the weights:

$$sr_t^{ag} - sr_{t-1}^{ag} = \sum_c S_t^c (w_t^c - w_{t-1}^c) - \sum_c w_{t-1}^c (s_t^c - s_{t-1}^c) \quad (2)$$

Changes in weights can in turn be decomposed into changes in the relative sizes of different age groups and changes in their disposable income (an exercise we perform later).

Data and Definitions

Few high-quality data sources include information on both income and consumption; many of the sources that do so include data far apart in time. The problem affects both industrial and developing economies. For Mexico, Peru, Taiwan (China), and Thailand cross-sectional data observed at several points in time are available. For Mexico we have data from the National Survey on Income and Expenditure of Households, conducted by the National Institute of Statistics, Geography, and Informatics, for 1984, 1989, 1992, 1994, and 1996. The last year is of particular interest because it is the year after the peso crisis. The data for Peru are from the National Household Survey for the Measurement of Living Standards and cover 1985, 1991, 1994, and 1997. They include data from two surveys that surround the 1990 crisis. The data for Taiwan (China) are annual data for 1976–96 from the Family Income and Expenditure Survey, conducted by the National Statistics Office of Taiwan (China). The data for Thailand—for 1975, 1981, 1986, 1988, 1990, 1992, 1994, and 1996—are from the Socio-Economic Survey (SES) conducted by the National Statistical Office.[2]

Using these data, we construct measures of income, consumption, household composition, educational attainment, and labor supply. Some of the definitions relevant for the analysis of saving include the following:

- Income is defined as disposable household income. It includes earnings, transfers, capital income, and nonmonetary income.
- We use four definitions of consumption to calculate savings rates. For the first, s_1, we include all household expenditure; this is closest to the definition typically used in national accounts data. The second savings rate, s_2, excludes from consumption all expenditures on durable goods as well as expenditures on health and education. This measures tries to take into account the fact that some expenditure items have an important savings component. While this definition is far from perfect—it does not, for example, include services that accrue from durable goods—the analysis of consumption and saving based on this definition deserves attention. The third definition, s_3, excludes expenditures on durable goods but considers health and education current consumption rather than savings, as in s_2. Finally, a measure we label s_4 includes durable and nondurable expenditures in the definition of consumption but excludes health and education, which are considered savings.
- All surveys include some definition of human capital. We divide the population into three groups: primary education or less, some secondary schooling, and higher education. This classification takes into account institutional differences across economies.
- Household arrangements differ across economies. While we present some evidence on this and document the extent of possible problems with endogenous household formation and dissolution, in the end we use the standard definition of declared household head across economies.

Throughout the analysis we divide the samples into year-of-birth cohorts. To work with cells of reasonable size, we use a five-year definition. The cohort definition is homogeneous across economies.

Cell sizes range from 86 to 2,528 households. Interestingly, one of the smallest samples is observed in Mexico, the country with the largest population: Mexico has about 70 million more people than Taiwan (China), yet we have about 50 percent more observations for Taiwan (China). The differences are larger at older ages. Since Mexico and Peru have the smallest samples, estimates derived for these economies will have less precision, especially for the older cohorts, than those for the two East Asian ones. Except for Thailand (where weights are not available), we use population weights to compute our results.

Static Analysis of the Most Recent Data

The data on aggregate private savings rates presented in figure 2 are far from perfect, for many reasons. One of the main problems is that in the national accounts savings are calculated as residuals of other aggregates that are also measured with some error. In this section we first complement the aggregate evidence by presenting household savings rates calculated from the microeconomic data for the most recent year available for each of the four economies. We also characterize

saving for different population groups and investigate whether the differences in savings rates between East Asia and Latin America are due to demographic differences.

To link the aggregates presented earlier and the results from the microeconomic data, we concentrate on the most standard definition of savings—total disposable income minus total expenditures, divided by total disposable income (s_1). The figures we present are computed as ratios of average saving and average income, rather than the average of the ratio (table 1). This procedure reduces the effect of outliers.

Based on the aggregate data, Thailand appears to have the highest domestic savings rate (almost 36 percent), while Mexico, Peru, and Taiwan (China) show roughly similar savings rates. When we turn to the microeconomic data, a different picture emerges. Mexico and Peru register an s1 of about 9.5 percent, while Thailand and Taiwan (China) have rates of 29.7 percent and 49.1 percent. Thus Thailand has a much higher rate of domestic saving than Mexico and Peru, and much of the difference seems to originate at the household level; the difference between Taiwan (China) and the Latin American countries is even greater. So, if we set aside the issue of the comparability of micro- and macroeconomic sources, the small difference at the aggregate level must reflect lower public and corporate saving in Taiwan (China).

As expected, dependency ratios are much higher in Mexico and Peru (by about 20–30 percentage points), reflecting the fact that these countries are at an earlier stage of the demographic transition, with a smaller proportion of the population of working (and saving) age. Thailand has a much higher female labor force participation rate than Mexico and Peru; since higher participation is associated with higher savings capacity, the result is consistent with the huge difference in household savings rates. Thailand also registers fewer average years of schooling than the two Latin American countries and similar per capita GDP; these variables therefore cannot account for the difference. Institutional factors, such as the lack of compulsory retirement benefits in Thailand until 1999 and the tradition of public pension provision in Mexico and Peru, may partly explain the difference in household savings rates.

Table 1. Household Saving and Selected Other Indicators, by Economy, 1996
(percent, except where otherwise indicated)

Economy	*Aggregate domestic savings rate*	*Household savings rate*	*Youth dependency ratio*	*Female labor force participation rate*	*Average years of schooling*	*PPP GDP per capita (international dollars)*[a]
Mexico	25.4	9.5	0.59	44.4	7.1	5,757
Peru[b]	19.0	9.6	0.60	64.5	8.5	2,993
Taiwan	26.8	49.1	0.31	57.4	9.3	14,634
Thailand	35.9	29.7	0.41	79.7	5.8	5,080

Note: Data on aggregate domestic saving and PPP GDP per capita for Mexico and Peru are from World Bank (1999).
a. Gross domestic product converted to international dollars using purchasing power parity (PPP) rates.
b. Data for peru refer to 1997.
Source: Authors' calculations based on household survey data. Data for Taiwan (China) are from the National Statistics Office of Taiwan (China). For Taiwan (China) data on aggregate domestic saving refer to gross national savings, and GDP per capita is not adjusted for purchasing power parity.

For Taiwan (China), the lower dependency ratio, the greater average years of schooling, the larger per capita GDP, and the higher female labor force participation rate (relative to Mexico, although not to Peru) are consistent with the huge difference in household savings rates between Taiwan (China) and the two Latin American economies.

Thus one possible explanation for the differences in household savings rates is that the East Asian economies are at a later stage of the demographic transition, with larger shares of their populations at ages at which savings rates typically peak. For all ages the savings rate is higher in Taiwan (China) and Thailand than in Mexico and Peru. The weight of the middle-aged groups is also somewhat greater, as expected. As a gross approximation to assess the role of demographics in the differences in saving, we recompute the savings rates in Mexico and Peru using weights from Taiwan (China) and vice versa. Although the rates for Mexico and Peru increase and the rate for Taiwan (China) declines, the difference relative to the original savings rate in each economy is only marginal. This suggests that having a larger share of household heads in age groups that save more does not account for the major differences in saving.

This exercise only partially accounts for demographic differences, because it considers the distribution of households across working ages but ignores the fact that the major source of demographic differences is the larger share of the population in Latin America in the 0–15 age range. Both Latin American countries have smaller population shares of prime working age than the East Asian economies. A second accounting exercise attempts to take this into account, at least in a very general way. The calculations consist of weighting the age-specific savings rates in Mexico, Peru, and Thailand with the total population weight of Thailand and weighting the Taiwan (China) rates with the Peruvian weights. Doing so narrows the difference between the Latin American countries and Taiwan (China) from 37 percentage points to about 14 percentage points, which gives us a gross idea of the importance of differences in demographic structure.

In examining savings rates by education group, we observe that Taiwan (China) and Thailand have much higher rates than Mexico and Peru, with the largest differences between Thailand and the Latin American countries among the most educated households (Thai savings rates are 28 percentage points higher for the most educated households, 22 percentage points higher for the least educated). Clearly, the differences in household savings rates are given not by the size of the education groups but by the differences in the group-specific rates. The simple accounting exercise in which we recompute the Latin American averages using weights from Taiwan (China) and vice versa confirms this.

In comparisons of saving by income quintile, the most striking difference between Thailand and the Latin American countries is at the top of the distribution, where savings rates in Thailand are considerably higher (table 2). Differences appear beginning in the second quintile, but the difference is greatest in the top quintile. Surprisingly, the only group in which Mexico has higher (less negative) savings rates than Thailand is the poorest quintile. This may indicate the importance of transitory

income components in the two countries or larger measurement error for the lowest income quintile. Much of the difference between the aggregate savings rates in these two countries is explained by the extremely high savings rates among Thai households in the two richest quintiles.

The only economy in which differences along the income distribution are small is Taiwan (China), where there is only a 15 percent gap between the savings rates of the poorest and richest quintiles. This is not surprising, since inequality in Taiwan (China) is very low (as indicated by the Gini coefficient of just 0.30).

In comparisons of the quintile savings rates in Taiwan (China) with those in each of the other economies, the most interesting feature is that in all cases the largest difference is in the poorest quintile. The comparison between Peru and Taiwan (China) is especially illustrative. Almost all of the difference in the total household savings rate (about 40 percentage points) is explained by the disparities in the first four quintiles. In fact, the richest 20 percent of households in Peru have a savings rate only 15 percentage points lower than that of their counterparts in Taiwan (China). Although the differences between Mexico and Taiwan (China) are less extreme, the conclusion still holds that the largest part of the difference in household savings comes from poorer households.

Differences in the relative importance of education groups do not, then, account for the large disparities in household saving between East Asia and Latin America, while differences in demographic structure have the potential to explain part of the gap. But an important part of the story is the difference in savings rates across the income distribution. It is difficult to determine which way the causality runs (do all groups of society have similar earnings and savings potential when income is more equally distributed, or do income-earning assets increase when more individuals are able to save for reasons other than income?). But regardless of the direction of causality, the finding is important because it reveals that in Mexico, Peru, and Thailand the richest sectors of the population have much greater capacity to accumulate assets and smooth unexpected variations in income than the rest of the population, while in Taiwan (China) this is not the case. In Mexico and Peru nearly all household savings originate among the richest 20 percent of the population. After

Table 2. Household Savings Rate by Income Quintile, 1996

Economy	*Total*	*1 (poorest)*	*2*	*3*	*4*	*5 (richest)*	*Gini coefficient*
Mexico	9.5	–0.17	–0.05	–0.01	0.05	0.21	0.53
Peru[a]	9.6	–1.81	–0.52	–0.18	0.07	0.40	0.51
Taiwan (China)	49.1	0.39	0.43	0.46	0.48	0.54	0.30
Thailand	29.7	–1.32	0.00	0.15	0.29	0.49	0.53
Taiwan (China)–Mexico difference		0.55	0.48	0.46	0.43	0.33	
Taiwan (China)–Peru difference		2.20	0.95	0.63	0.41	0.15	
Taiwan (China)–Thailand difference		0.71	0.43	0.30	0.19	0.05	

a. Data are for 1997.
Source: Authors' calculations based on household survey data.

accounting for these differences along the distribution, it is much less surprising that Taiwan (China) has savings rates well above those registered in the two Latin American countries.

Aggregate Household Saving over Time

Savings rates changed significantly over the 28-year period studied. In Mexico s_1 increased between 1984 and 1989, a period characterized by stagnation and the partial recovery of growth after a substantial decline in 1986. The 1989–94 period was characterized by a consumption boom and a simultaneous increase of 3 percentage points in the savings rate. Changes in s_2, s_3, and s_4 were similar in 1984–89, but in 1989–94 the increase was much smaller that that in s1 and there was even a decline in some cases. This pattern suggests that the consumption of durable goods was increasing faster than expenditures on nondurables, education, and health. Between 1994 and 1996 average s_1 and s_3 declined sharply, while average s_2 and s_4 rose. The medians for all four definitions of savings rates declined during this period. They increased much less than the average or declined between 1984 and 1989 and dropped in 1992. The savings rates for the poorest 50 percent of the population were less responsive to the increase in income after 1989, while the 1995 tequila crisis had a larger negative effect on their savings, perhaps because of the limited income-smoothing mechanisms available to this group.[3]

Household saving in Peru has been much more erratic. Average rates increased between 1985 and 1991, declined in 1994, and rose in 1997. This trend is surprising, since Peru experienced sharp declines in per capita GDP during the 1980s and more stable growth during the first half of the 1990s. The only period in which savings rates behaved as expected was 1994–97, when per capita GDP grew at a much faster pace. In the case of median rates, the data reveal that the 1990 crisis had a larger negative effect on the poorest 50 percent of the population, while the 1994–97 recovery did not have the positive effect on this group that was observed at the upper part of the income distribution.

The picture for Taiwan (China) and Thailand is quite different. Comparing 1985–86 and 1996, we find that with few exceptions, household saving increased smoothly. The four savings rates peaked in 1996 at 49.1 percent, 84.0 percent, 74.8 percent, and 58.5 percent in Taiwan (China) and at 33 percent, 62.3 percent, 56.9 percent, and 38.8 percent in Thailand. The patterns followed by the average savings rates and the median rates show that the increases occurred across definitions and across the income distribution.

We use annualized household savings from the surveys to estimate the relative importance of household savings at the national level. These estimates are only gross approximations: household surveys normally suffer from misreporting or underreporting of income, so the data may not reflect total household savings with precision. Furthermore, the calculation does not account for the savings that households have in firms, the value of pensions, or other important items. The results are nevertheless useful for identifying differences across countries.

As expected, household savings account for the largest share of GDP in Taiwan (China). Interestingly, however, although household savings in Mexico are lower than those in Thailand and similar to those in Peru, they represent a much larger share of GDP. The results confirm that a very large part of total domestic savings in Taiwan (China) originates at the household level.

We use equation 2 to decompose the changes documented earlier into three effects (table 3). The first is an age profile effect, which reflects the increase in saving due to the fact that cohorts age through their life cycle and save more or less depending on their needs and future prospects. The second is a demographic effect, which measures the change in saving due to the fact that the population weight of different age groups shifts. The third is an income effect, which reflects the fact that as individuals age, their income tends to rise and they thus have greater savings capacity.

The contrast between East Asia and Latin America is stark. In Mexico and Peru all of the change in household saving is driven by the age profile effect of cohorts moving through the life cycle; in Taiwan (China) and Thailand the age profile effect is positive, but most of the change in saving is accounted for by the demographic and income effects. These results support the view that the East Asian economies have experienced much larger increases in saving because of the increase in their income and the shift in the population toward age groups that save more, not because of households moving through the life cycle.

It must be kept in mind that the time span in table 3 is longer for East Asia than for Latin America and, more important, that the period under analysis in Latin America is characterized by economic instability and low economic growth. A large part of the difference may thus reflect the fact that the economic environment in East Asia has been much more favorable for building up savings.

A Life-Cycle Analysis of Household Savings Behavior

We are now ready to analyze the life-cycle patterns of household saving in the four economies under study. To do so, we use the time series of repeated cross-sections to construct synthetic panel data. We divide each survey into 12 birth cohorts and 3 education groups. Much of the analysis in this section is graphical. In particular, we plot the life-cycle profiles of several variables of interest. It is therefore worthwhile to briefly describe the way in which we construct the graphs.

Table 3. Decomposition of the Change in Household Saving
(percentage points)

Economy	*Period*	*Total change*	*Age profile effect*	*Effect of change in weights*	*Effect of change in weights*	
					Demographic	*Income*
Mexico	1984–96	1.0	1.3	–0.2	–1.7	–1.5
Peru	1985–97	12.7	14.3	–1.7	–1.8	0.1
Taiwan (China)	1976–96	20.0	3.74	16.3	7.6	8.6
Thailand	1975–96	19.5	9.03	10.4	4.6	5.8

Source: Authors' calculations based on household survey data.

We plot the cohort averages at different points in time against the age of the household head and connect the points for each cohort. Visually, we follow the average behavior of each cohort as it ages. If the interval that defines a cohort is shorter than the period covered by our sample, we observe different cohorts at the same age. While it is tempting to interpret such a difference as a cohort effect, one should remember that these figures refer to different years and therefore could be equally explained by a time effect.

Changes in Household Formation and Composition

We start the analysis by looking at household formation. For this reason the first graph we present is based on individual data. In each panel of figure 3 we plot the average age of the head of the household in which an individual lives against the age of the individual. To produce this graph we use all individuals in the sample, regardless of their position in the household. For each economy we display four panels: the top-left panel graphs the data for the entire sample, while the other three graph the three education groups.

If all individuals in a given cohort were household heads (or living in a household with a head of the same age), the 45-degree line would coincide with the cohort profile of the average household age. In figure 3 the cohort profiles diverge from the 45-degree line at the beginning and the end of the life cycle. Naturally, headship rates are low at the beginning of the life cycle, so that the cohort profile lies above the 45-degree line, indicating that some young adults still live with their parents. The speed with which the profile approaches the 45-degree line is an indication of how early new families are formed. Toward the last stages of the life cycle headship rates decline and fall below the 45-degree line, because the elderly merge into households in which the head is younger.

Two patterns are apparent. First, in all four economies there are strong differences across education groups, especially in the extent to which the profile falls below the 45-degree line. Less educated individuals seem to be more likely to move in with younger household heads. In contrast, the most educated have much higher headship rates and continue to be heads of households even at older ages. These findings suggest that the family plays an important role in smoothing consumption for the elderly, especially among those with lower income-earning capacity.

Second, there are strong differences across economies at both ends of the life cycle. In particular, the phenomenon of elderly individuals living in households headed by younger ones is much more prevalent in the two Asian economies. In contrast, at the beginning of the life cycle the differences cut across regions. Household formation seems to occur much later in Peru than elsewhere. At the opposite extreme is Taiwan (China), where by age 30 the profile already coincides with the 45-degree line.

The results shown in figure 3 are also interesting from a methodological standpoint, because they show that household composition changes in important ways through the life cycle, especially among the least educated. This means that even though we are tracing the same type of household in the repeated cross-sections, the

Figure 3. Age of Household Head for Household in Which Each Individual Lives

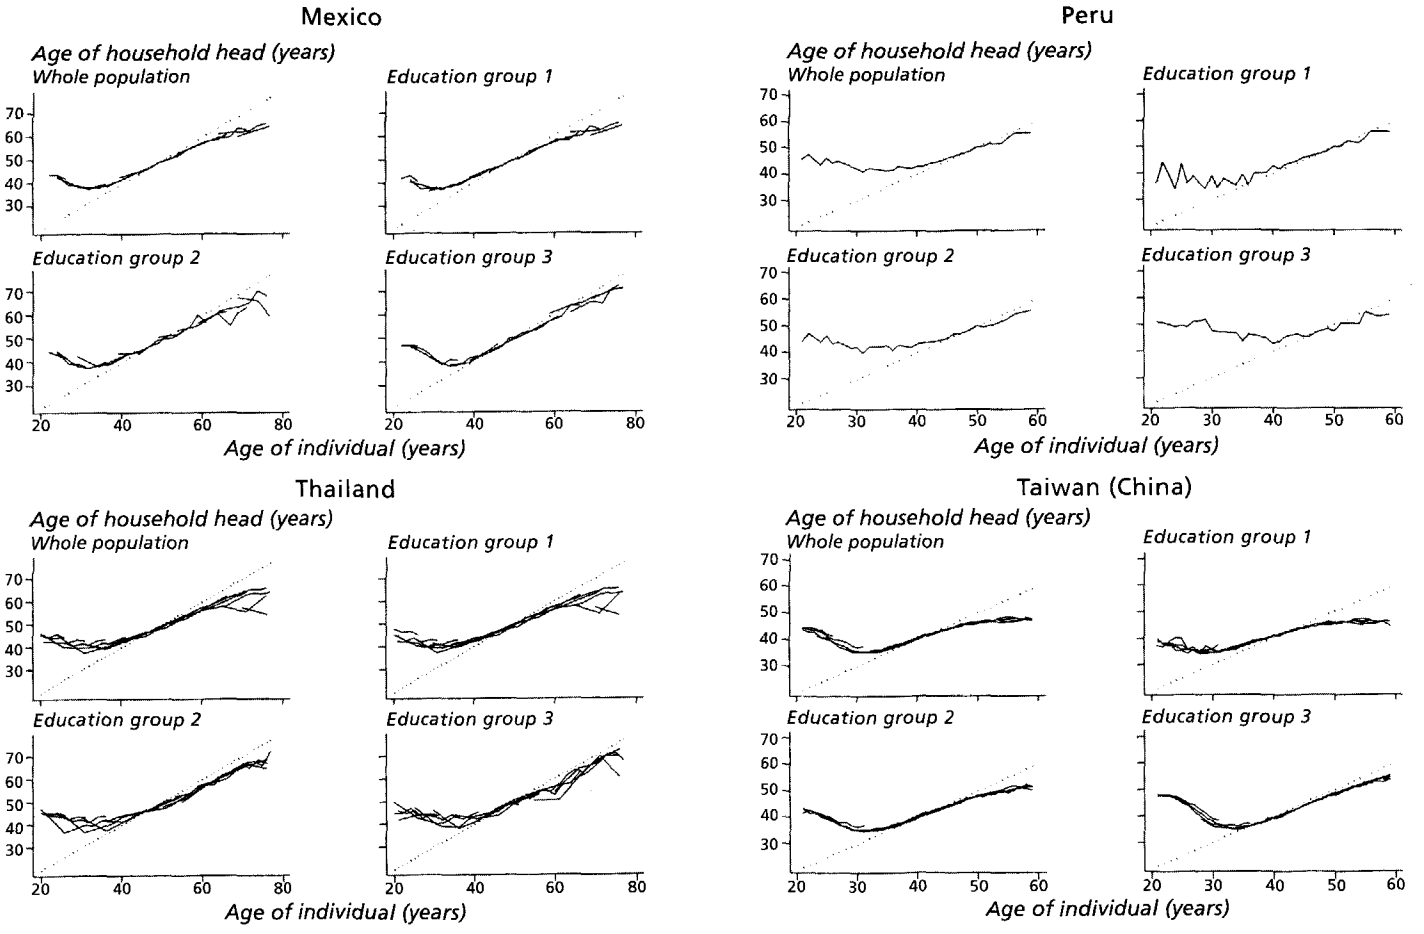

Source: Authors' calculations based on household survey data.

composition of the group is changing, blurring our inferences about the behavior of cohorts as they age.

Figure 4 plots the average years of schooling of household heads, as well as the proportion of household heads with secondary and higher education. Since after age 26 very few individuals continue to acquire formal education, we plot the cohorts from age 26 on. If there were no composition effects in the cross-sections under analysis, there would be differences in levels of education across cohorts, but the age pattern of each cohort would be close to a horizontal line.

Two patterns emerge from figure 4. First, the cohort effects for average years of schooling and for the proportion of individuals with secondary and higher education seem to be much larger in East Asia. Second, there are some compositional changes in the surveys. While some of these changes could be attributed to sampling error, systematic positive trends in years of education or the proportion of well-educated individuals could be a symptom of differential mortality effects. While these effects exist for Mexico, Taiwan (China), and Thailand, they are not particularly strong or significant. There are significant shifts in Peru, however, suggesting that the results for this country should be interpreted with more caution.

In figure 5 we analyze household data on household size. In the four panels for each economy we plot the log of household size against the age of the household head (for the entire population and the three education groups). Three features deserve emphasis. First, there are large differences in household size among education groups, with the least educated having the largest families in all cases. Second, there are large cohort effects, especially for Taiwan (China) and Thailand, with the youngest cohorts having much smaller families. Third, families are much smaller in the East Asian economies. In Peru and Thailand household size declines more slowly in the last part of the life cycle than in Mexico and Taiwan (China), perhaps because children leave home much later or because older adults join what become extended families.

The patterns in figure 5 are mirrored in figure 6, which plots the average number of children against the age of the household head. It is not implausible to interpret the differences between cohorts shown in figures 5 and 6 as pure cohort effects, since it is plausible to rule out the existence of systematic year effects.[4] Cohort effects are stronger in the two East Asian economies, and the number of children is smaller. The differences across economies tend to be larger in the groups with primary and secondary schooling.

Income, Consumption, and Savings Profiles

In figure 7 we plot the average of the log of disposable household income and the log of total consumption expenditure (note that the scale is different for each economy). In all four economies, as in other countries, consumption and income track each other closely. Moreover, differences in the shape of the income profile across education groups are mirrored in differences in the consumption profile.[5] Once again there are important differences across both education groups and economies. First, the most educated not only have higher incomes but their income profile is

Figure 4. Average Years of Schooling of Household Heads and Share with Secondary and Higher Education, by Age

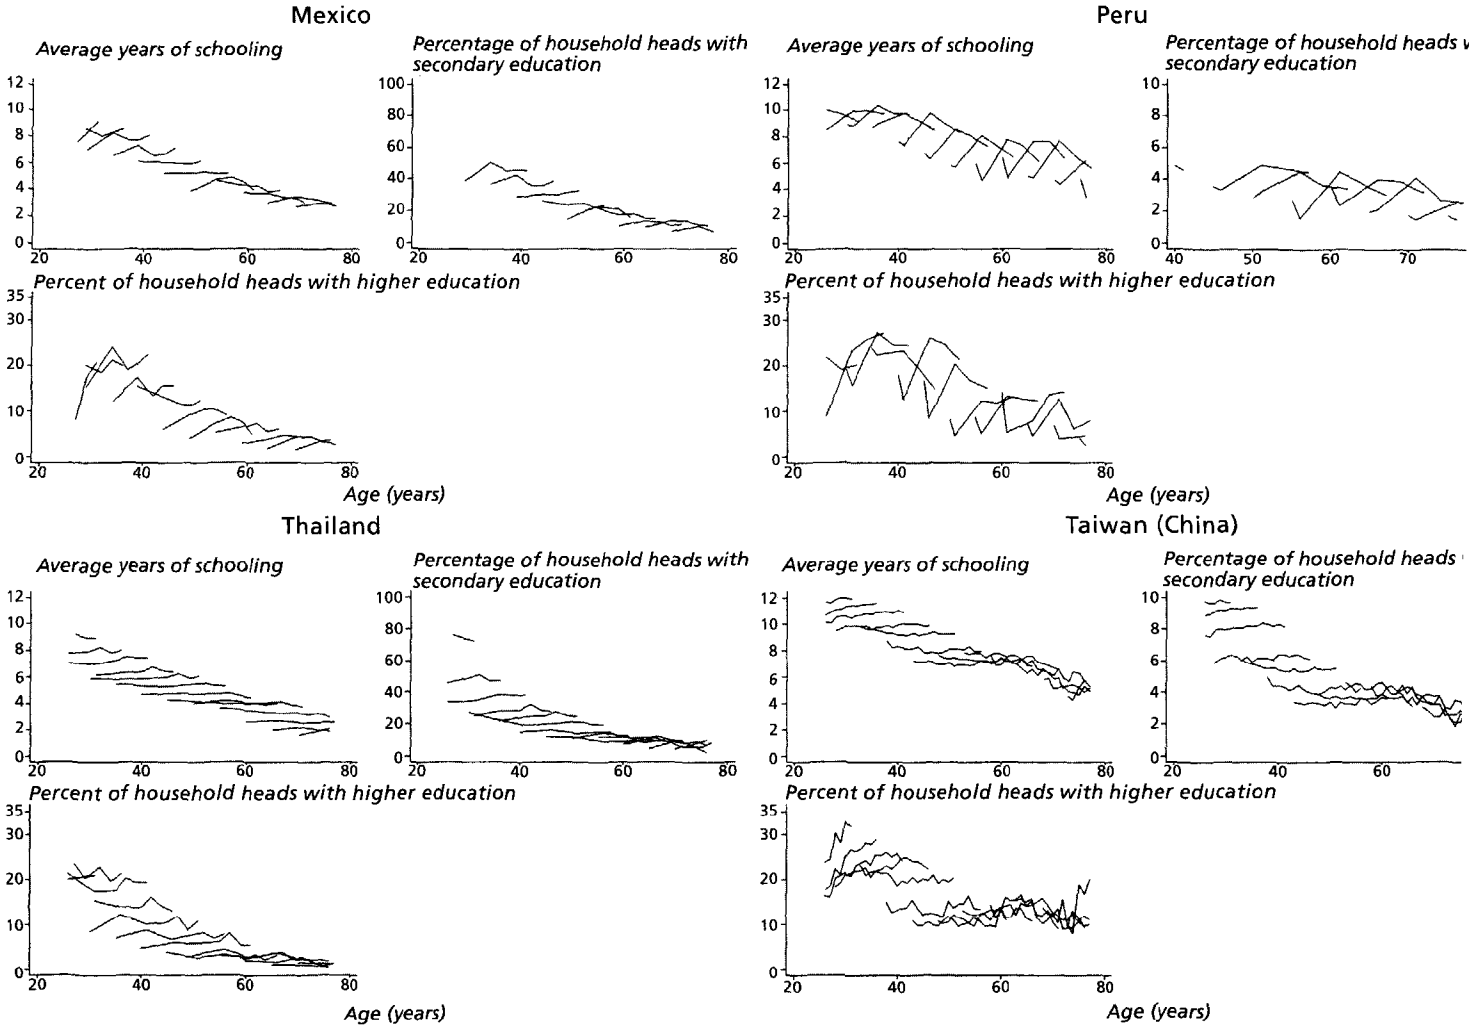

Source: Authors' calculations based on household survey data.

Figure 5. Size of Household, by Household Head's Age, Cohort, and Education Level

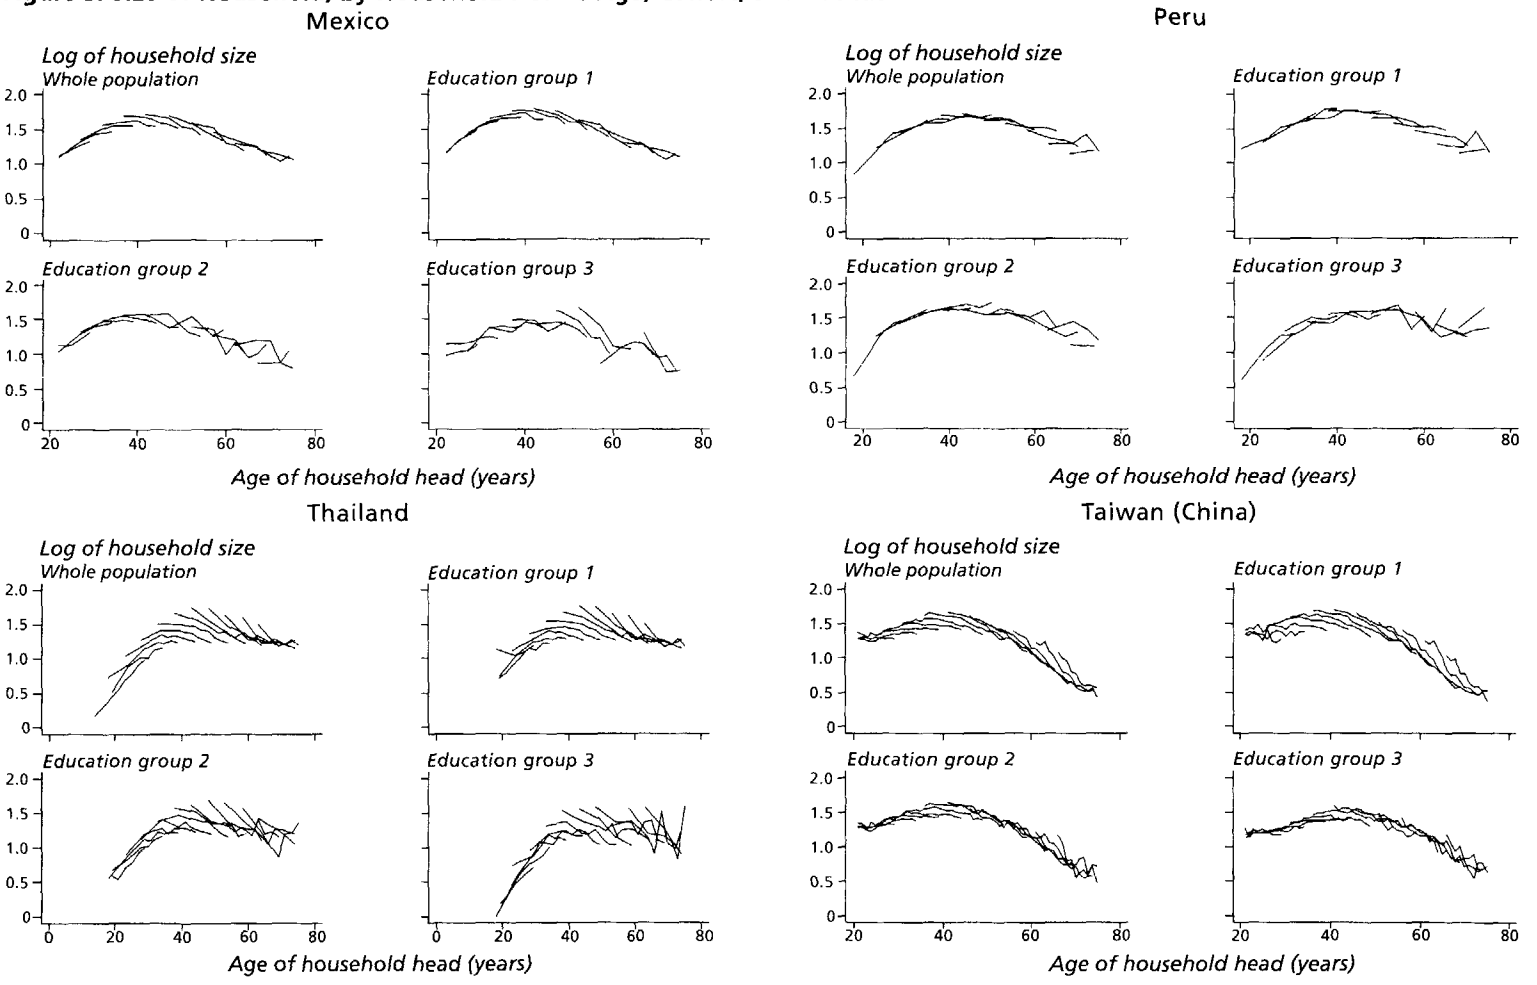

Source: Authors' calculations based on household survey data.

Figure 6. Average Number of Children, by Household Head's Age, Cohort, and Education Level

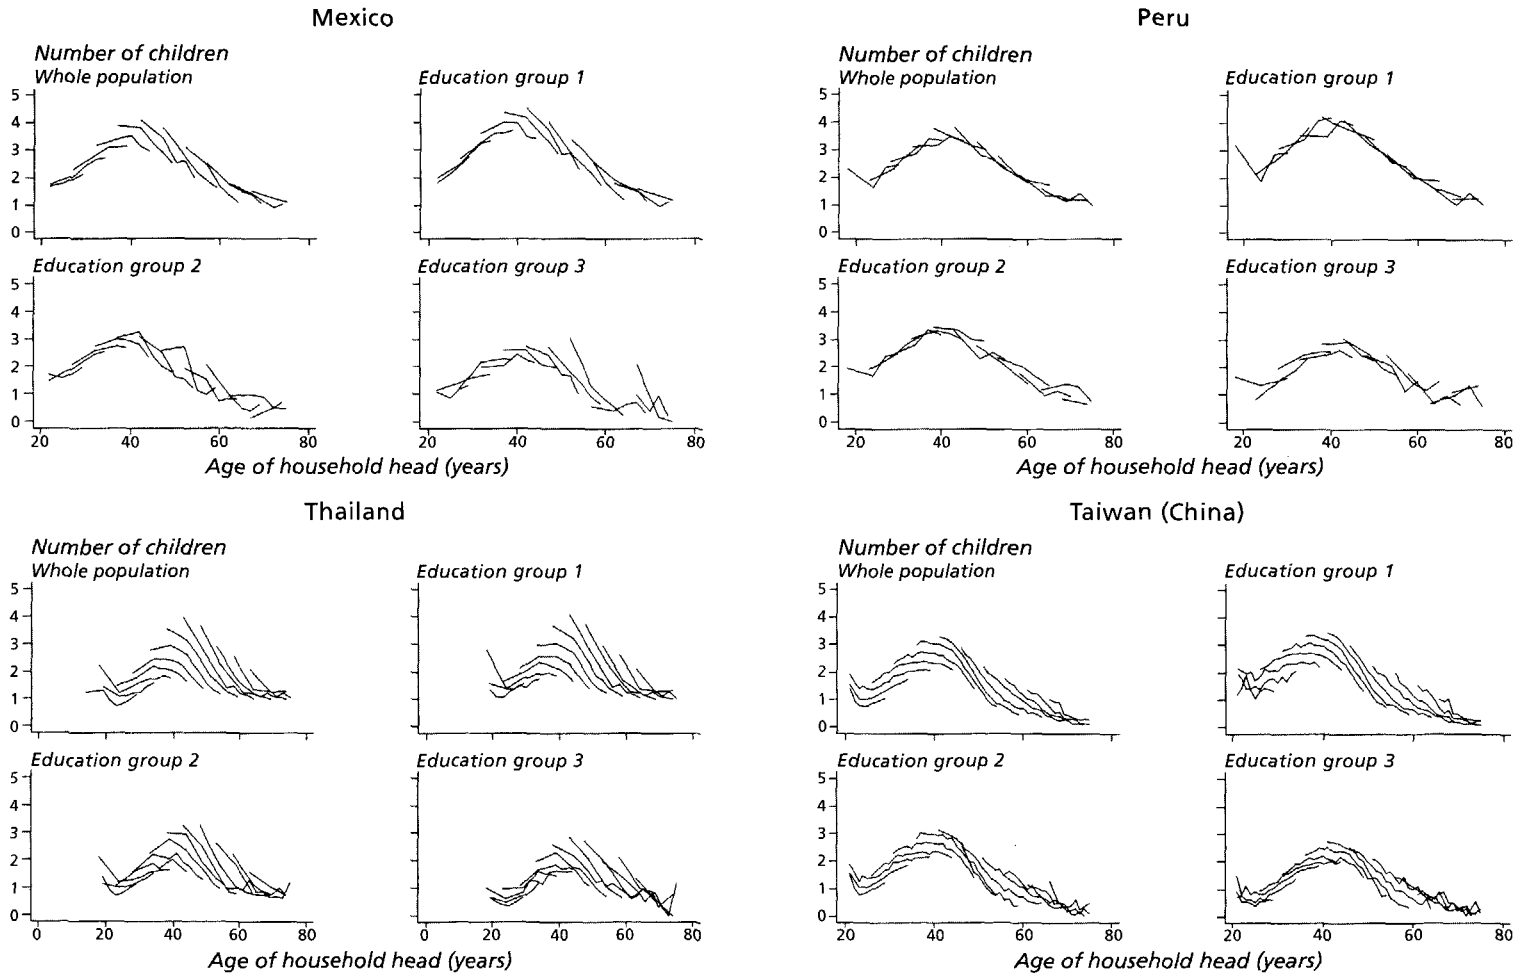

Source: Authors' calculations based on household survey data.

Figure 7. Disposable Household Income and Consumption, by Household Head's Age, Cohort, and Education Level

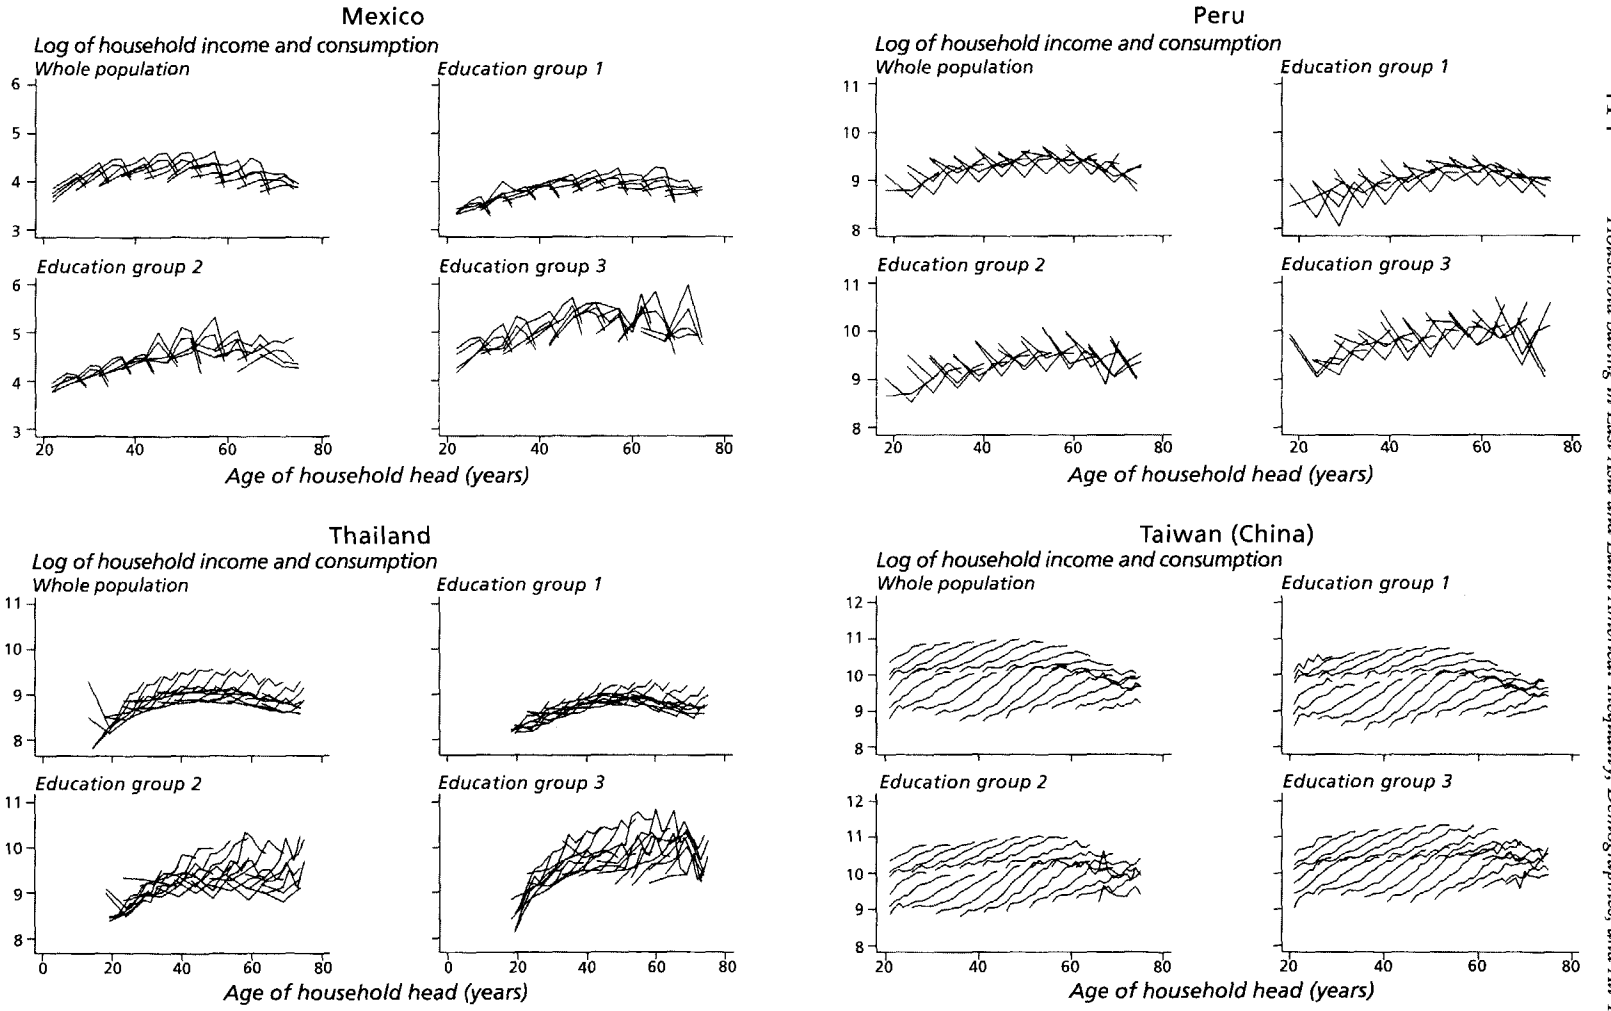

Source: Authors' calculations based on household survey data.

steeper over the life cycle. This difference is particularly apparent in Mexico and Thailand. Second, the experience of the two Asian economies, particularly Taiwan (China), is marked by the impressive growth of all cohorts. The entire life-cycle profile seems to shift up year on year starting in the early 1980s. Third, the experience of Peru, for which we have only four data points per group, is marked by the crisis of 1990, which is reflected in the 1991 survey. The decline in disposable income seems to have affected all education groups, and it seems more pronounced for the youngest cohorts. In comparison with this drop in income, even the declines observed in Mexico after the 1995 crisis (reflected in the 1996 survey) look small. In Mexico, unlike in Peru, the decline is more apparent for the most educated group.

In figure 8 we plot the log of per capita total and nondurable consumption. Overall, the per capita profiles are flatter than the corresponding household graphs in figure 7, although strong aggregate growth masks this fact in Taiwan (China). This result is consistent with the evidence reported by Attanasio and Browning (1995) for the United Kingdom and by Attanasio (1994) for the United States. An important difference between the regions emerges from figure 8: while the 1995 crisis in Mexico and the 1990 crisis in Peru are evident in the data, the data for East Asia are much smoother.

In figure 9 we plot two definitions of savings rates. The first (s_1) includes all expenditure items, while the second (s_3) excludes durable goods from the definition of consumption. The shape of the profiles is roughly similar, regardless of the definition used. In Mexico and Peru more educated households do most of the saving. This feature is not inconsistent with the life-cycle model, as the more educated face a steeper income profile. There seems to be no strong tendency for life-cycle profiles of saving to decline (and become negative) in the last part of the life cycle. The differences between the two definitions are greatest in Taiwan (China), where the hump shape is also more apparent.

As Deaton and Paxon (2000) recently emphasized, one reason that a clear hump shape consistent with the life-cycle hypothesis is not observed may be the changes in household composition documented earlier. Individuals may behave as the theory predicts, but aggregation into households that change in size, composition, and needs may mask these shifts in behavior. Deaton and Paxon suggest a method for identifying individual savings profiles from household data under some assumptions, and they conclude that a clearer hump is observed for individuals than for households in Taiwan (China) and Thailand.

Labor Supply and Wages

Households' capacity to save is determined largely by their income. Households' labor income reflects the wages paid in the market and labor force participation rates. In figure 10 we plot male wages for Mexico, Peru, and Taiwan (China) (Thailand is not presented because of data limitations). The results are unsurprising. First, the profile for more educated individuals is not only higher but also much steeper, especially in Mexico and Peru. Second, for all education groups the effects

Figure 8. Per Capita Total and Nondurable Household Consumption, by Household Head's Age, Cohort, and Education Level

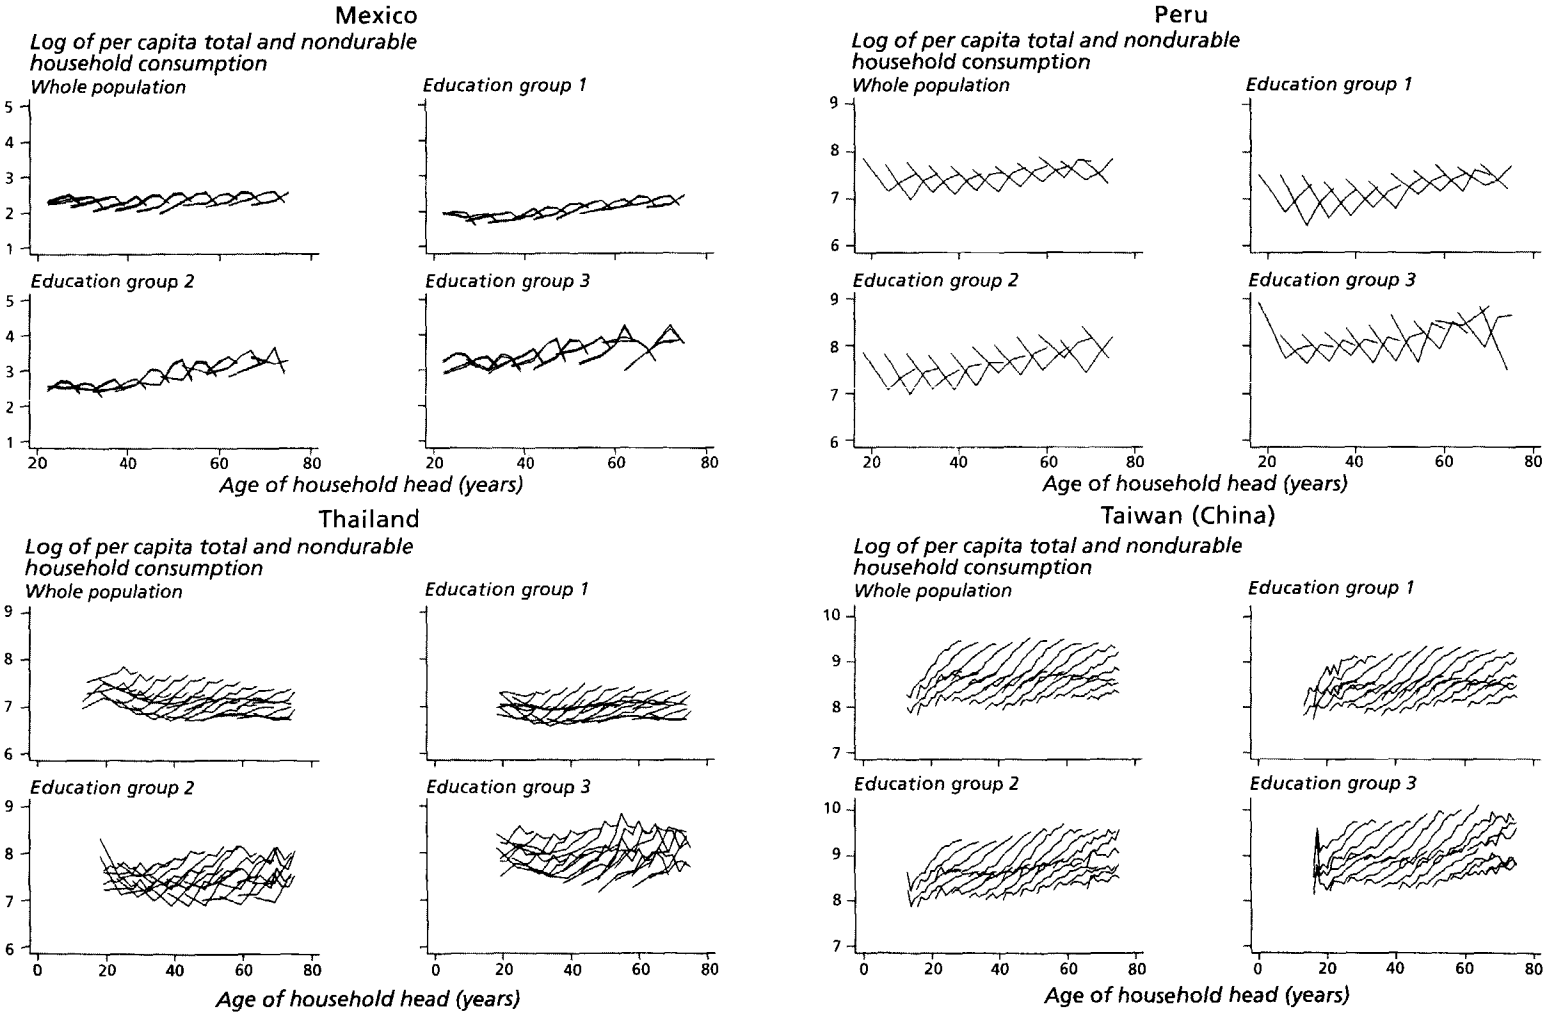

Source: Authors' calculations based on household survey data.

Figure 9. Household Savings Rates, with and without Durables, by Household Head's Age and Education Level

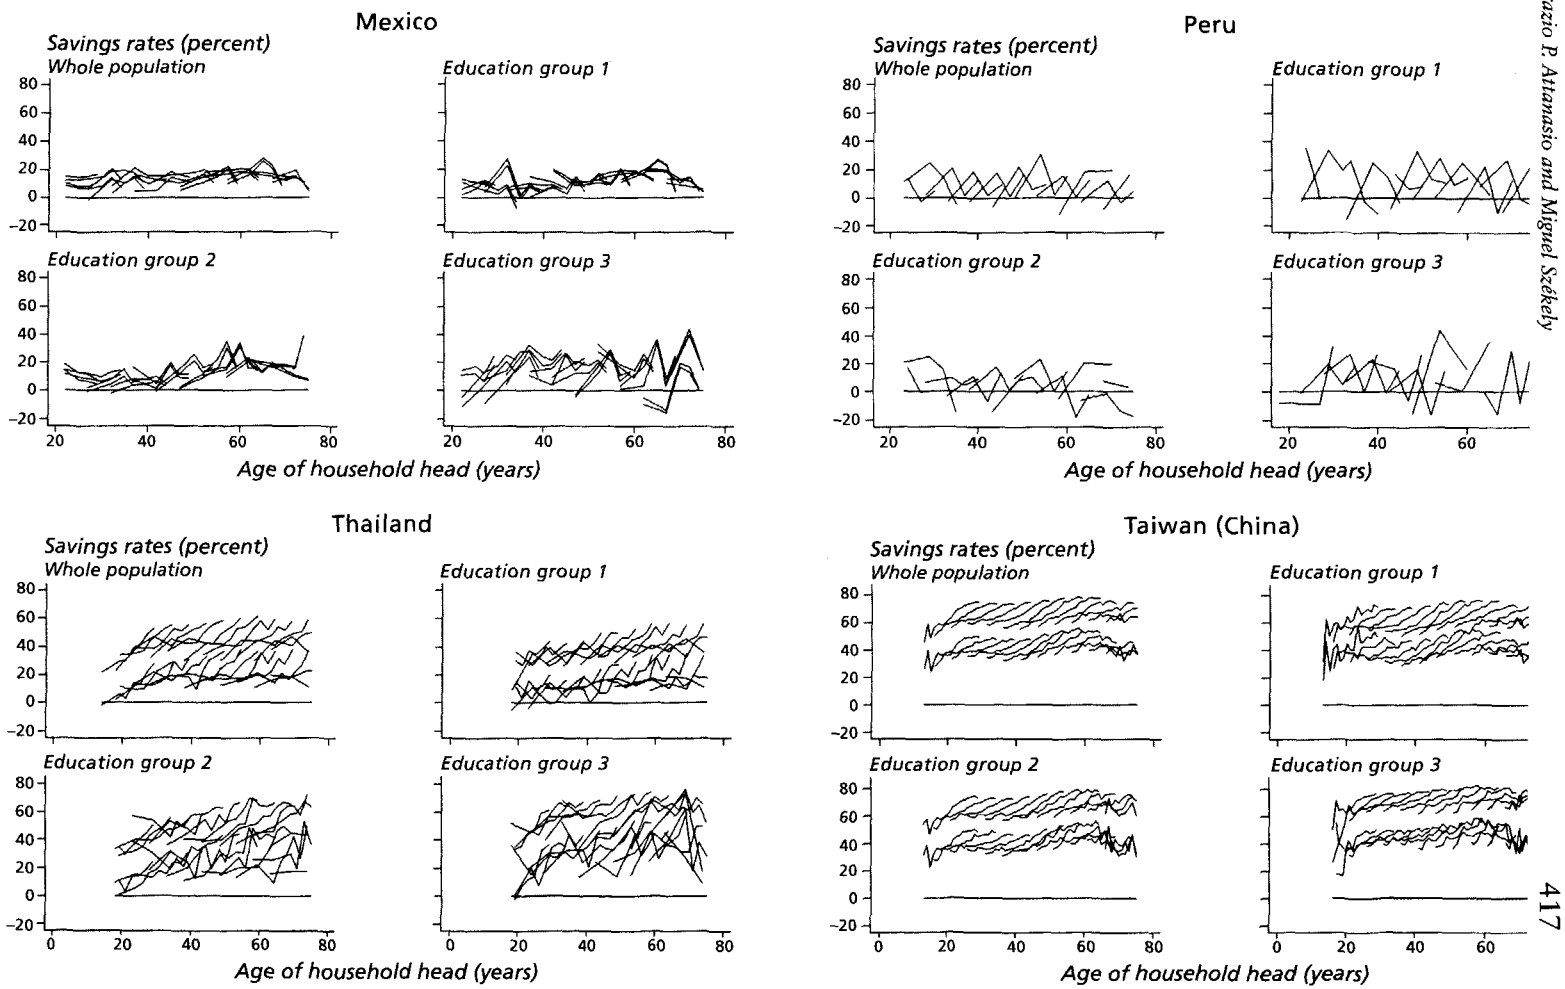

of the 1995 crisis are apparent in Mexico, and a strong negative effect from the 1990 crisis is apparent in Peru in 1991. The patterns for Taiwan (China) are much smoother and show a continuous increase.

Perhaps most interesting is that these income profiles for individuals are much more hump shaped than the household profiles in figure 7. This is consistent with the argument by Deaton and Paxon (2000) that even though individuals may behave according to the life-cycle model, their behavior may be blurred by aggregation into households.

In figure 11 we plot male and female labor force participation rates by education level. The results indicate strong differences between men and women, especially in groups with less education. Relatively large cohort effects are visible in the labor force participation of women in the least educated households. For men, retirement is much more gradual among the least educated men, especially in Mexico, Peru, and Thailand, than among the more educated who tend to retire around the same age. This phenomenon may be related to the fact that more educated individuals are more likely to work in the formal sector and therefore to be covered by pensions. Higher labor force participation rates enhance savings capacity, so these results are compatible with the apparent cohort effects in saving presented earlier.

Female labor force participation has several interesting features. First, it is much higher in Thailand than anywhere else. Among educated Thai women, labor force participation is essentially equal to that of men. Second, female labor force participation seems high in Peru, at least by Latin American standards. Third, there are strong cohort effects in Mexico, especially for educated women. Finally, Taiwan (China) is the only economy in which labor force participation seems to decline in the most fertile years of women's life and then to recover. This dip seems to be absent elsewhere and is much less pronounced for educated Taiwanese women.

The evidence on female labor force participation should be interpreted with care, as cross-country differences might reflect differences in the way survey questions are asked. But it is important to consider the effect that female labor force participation might have on measured saving, as it is typically linked to a substitution of market goods for home production. Moreover, female labor supply could also have the effect of diversifying risk and therefore reducing the incentive for precautionary saving. But there are also reasons that households in which the wife works may save more. If female labor force participation is temporary, for example, it makes sense to save more to smooth income over time.

Another key determinant of saving is access to pension benefits. Individuals who believe that they will receive a pension after retirement have less incentive to save. Because household surveys contain only limited information on this issue, we are unable to document differences among the economies in detail. But broadly speaking, the evidence is consistent with some general differences between East Asia and Latin America. While most Latin American countries (including Mexico and Peru) have long traditions of providing pension benefits in the formal sector, such benefits have not been standard in Asia (Lora and Pagés 2000). In Thailand compulsory retirement benefits were introduced only in 1999. Thus in addition to the lower

Figure 10. Male Wages, by Age, Cohort, and Education Level

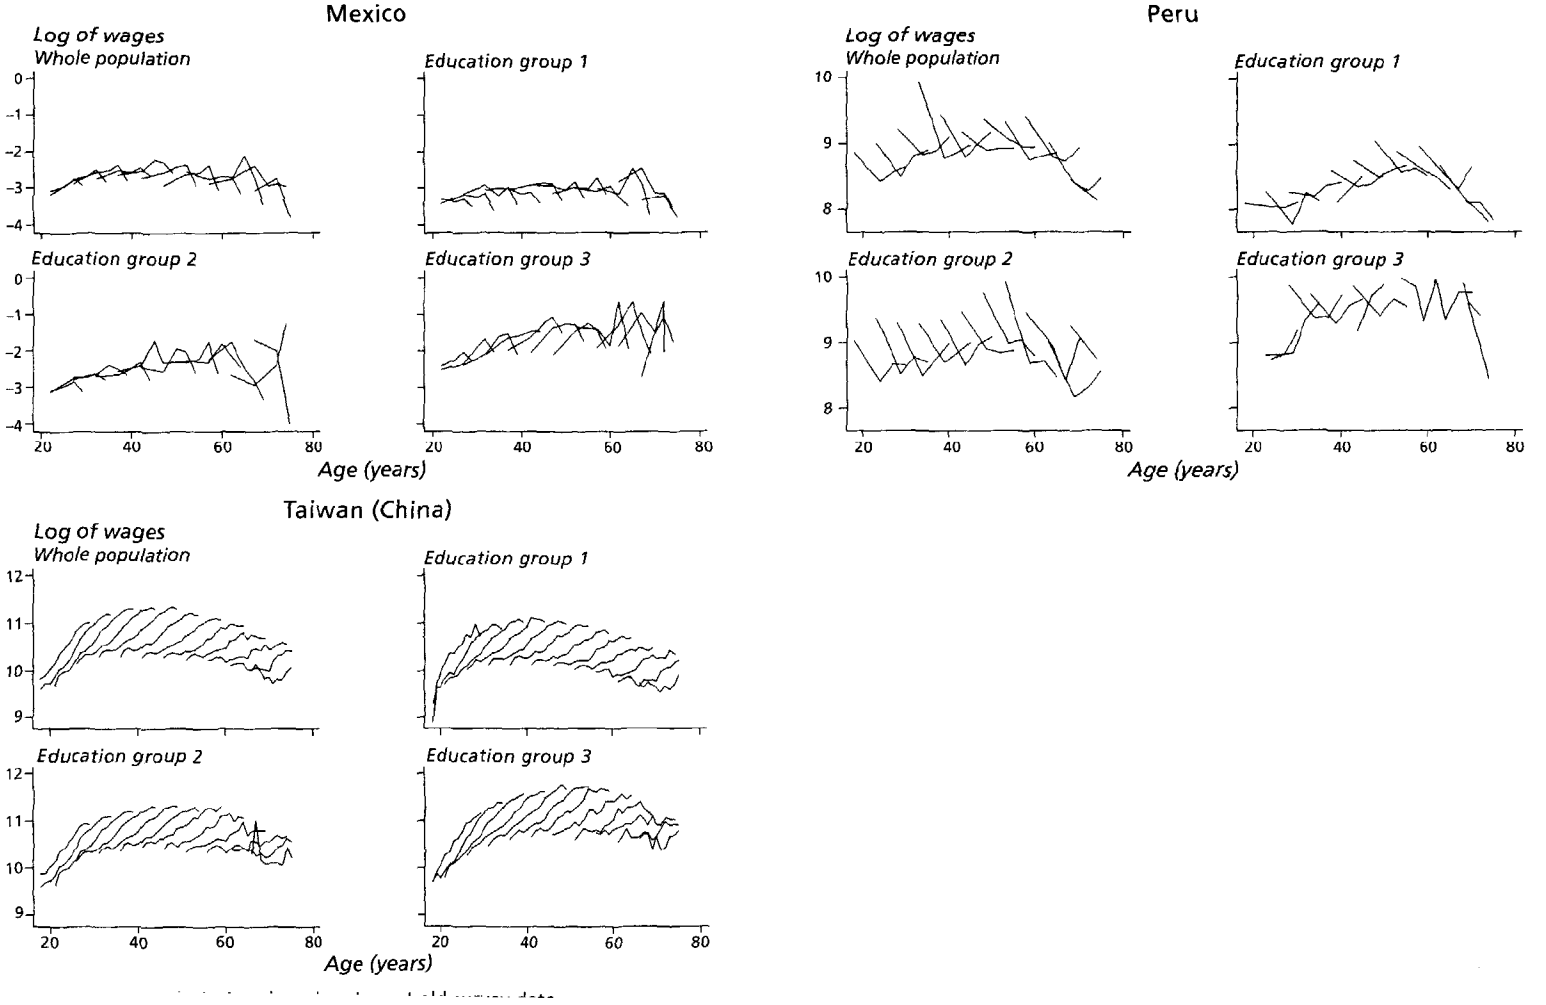

Figure 11. Male and Female Labor Force Participaton Rates, by Age, Cohort, and Education Level

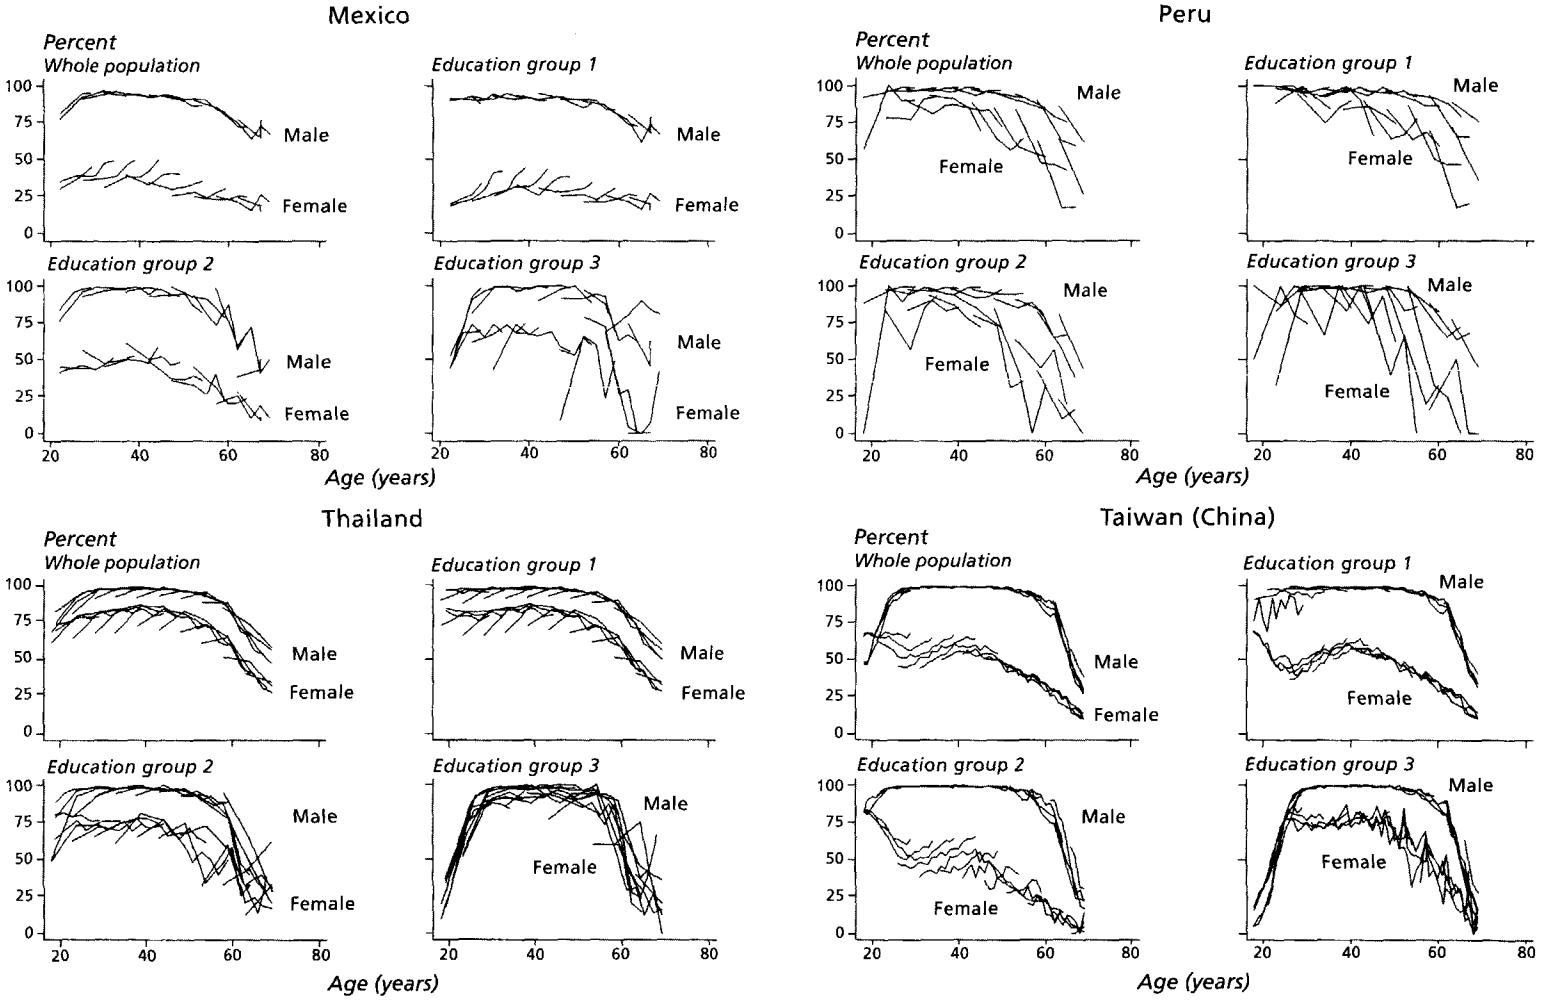

Source: Authors' calculations based on household survey data.

wage levels and female labor force participation rates in Mexico and Peru, the fact that formal jobs in these economies have traditionally provided pensions may be another explanation for the lower household savings rates in the two Latin American economies than in the two East Asian ones.

Smoothing Savings Profiles

To identify the age profile of savings rates, we smooth the savings profiles presented in figure 9. However, as discussed earlier, one cannot separately identify time, age, and cohort effects without additional restrictions. We therefore regress the data points plotted in figure 9 on an age polynomial, a set of dummy variables for cohorts, and a set of dummy variables for years constrained to have zero mean and to be orthogonal to a time trend. The restriction we impose on the data is equivalent to assuming that all the deterministic trends in the savings rate data originate from a combination of cohort and age effects (see Heckman and Robb 1987, MaCurdy and Mroz 1995, Attanasio 1998, Deaton and Paxon 1994a, and Paxson 1996). We also assume that the age profile of saving is the same across cohorts, except for an intercept shift. While this assumption is not necessary for identification, it is forced on us by the fact that at least for Mexico and Peru, each cohort is observed for only a few years.

In figure 12 we plot the polynomial for an arbitrary cohort. Given our assumptions and restrictions, in none of the economies or education groups considered does the life-cycle profile have a marked hump shape. The profile for least educated individuals is very flat in Mexico, Peru, and Thailand, but not in Taiwan (China). For the most educated, the profile increases monotonically with age in Mexico, Taiwan (China), and Thailand and increases toward the end of the life cycle in Peru.[6] Only in Taiwan (China) and in the middle group in Mexico and Peru do the smoothed profiles resemble the hump-shaped profile implied by some versions of the life-cycle model.[7] Even for these groups, however, the decline starts only after age 65 and is very gentle—surprising in light of the reductions in labor force participation documented earlier. Regardless of whether the evidence is consistent with the life-cycle model, the lack of a hump in the middle of the life cycle has important implications for the effect of projected demographic trends on aggregate savings rates, as we show below.

In figure 13 we plot the dummy variables for the cohorts obtained from the same regression. The pattern indicates substantial differences across cohorts. In particular, the youngest cohorts (those with the lowest cohort numbers) seem to have much higher savings rates than older cohorts, although in Mexico this effect is not observed for the population at large.

There are also important differences across education groups. In Mexico and Peru positive cohort effects are observed only for the middle group and the most educated. In Thailand there are positive cohort effects in all groups, but they are stronger for the middle group and the most educated.[8] Taiwan (China) is the only economy in which differences across groups are not apparent.[9]

Figure 12. Age Effects on Household Saving Rates

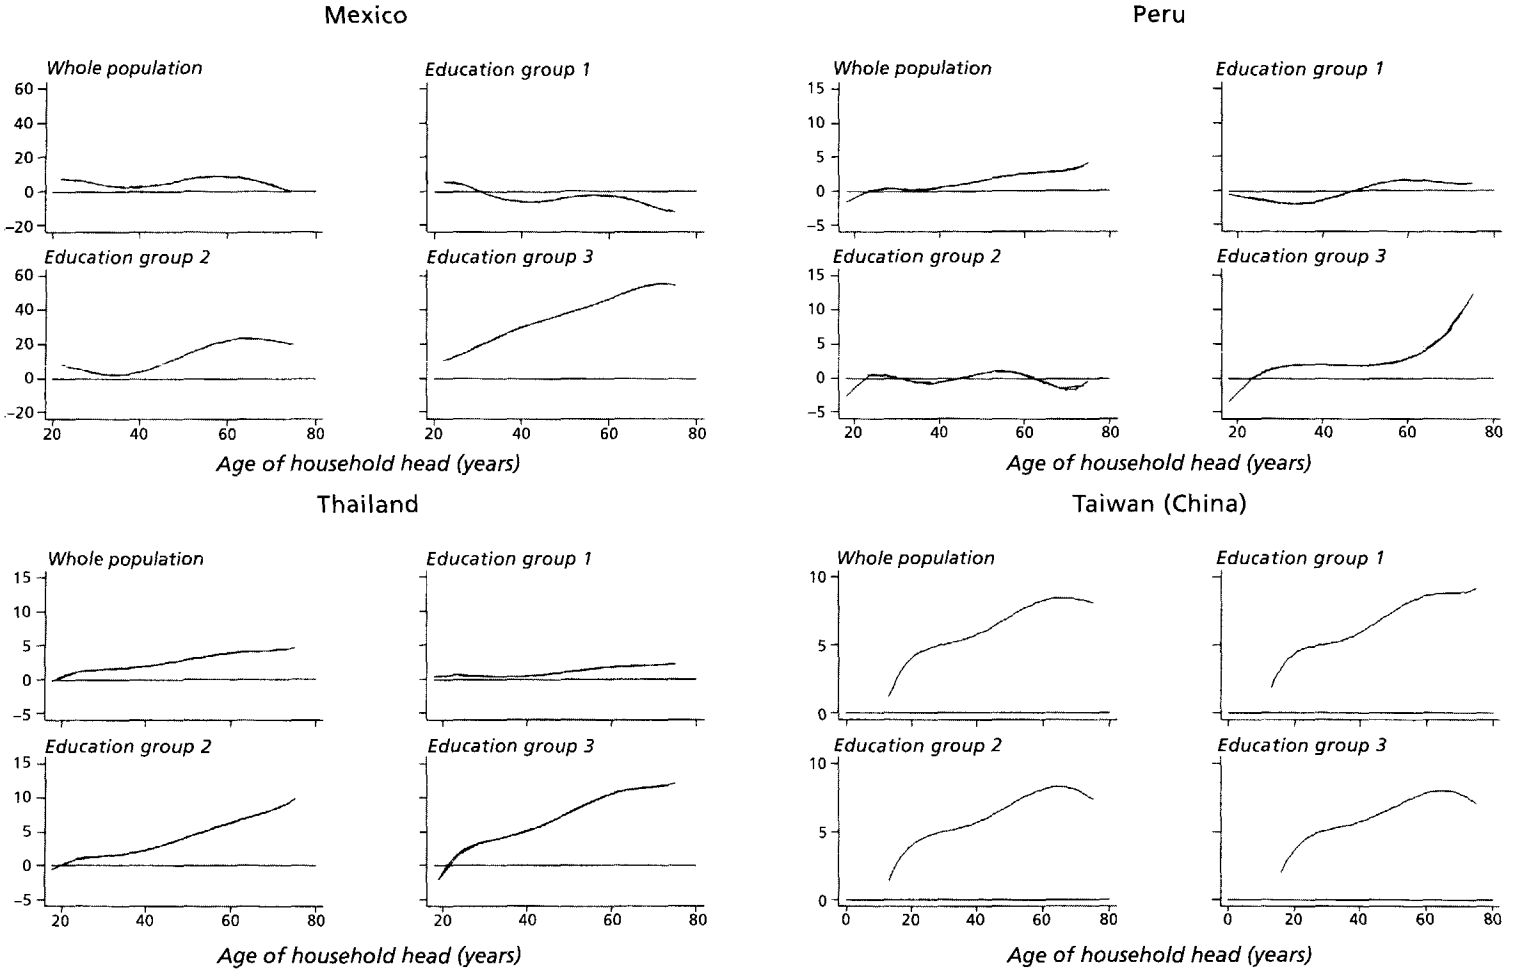

Source: Authors' calculations based on household survey data.

Figure 13. Cohort Effects on Household Saving Rates

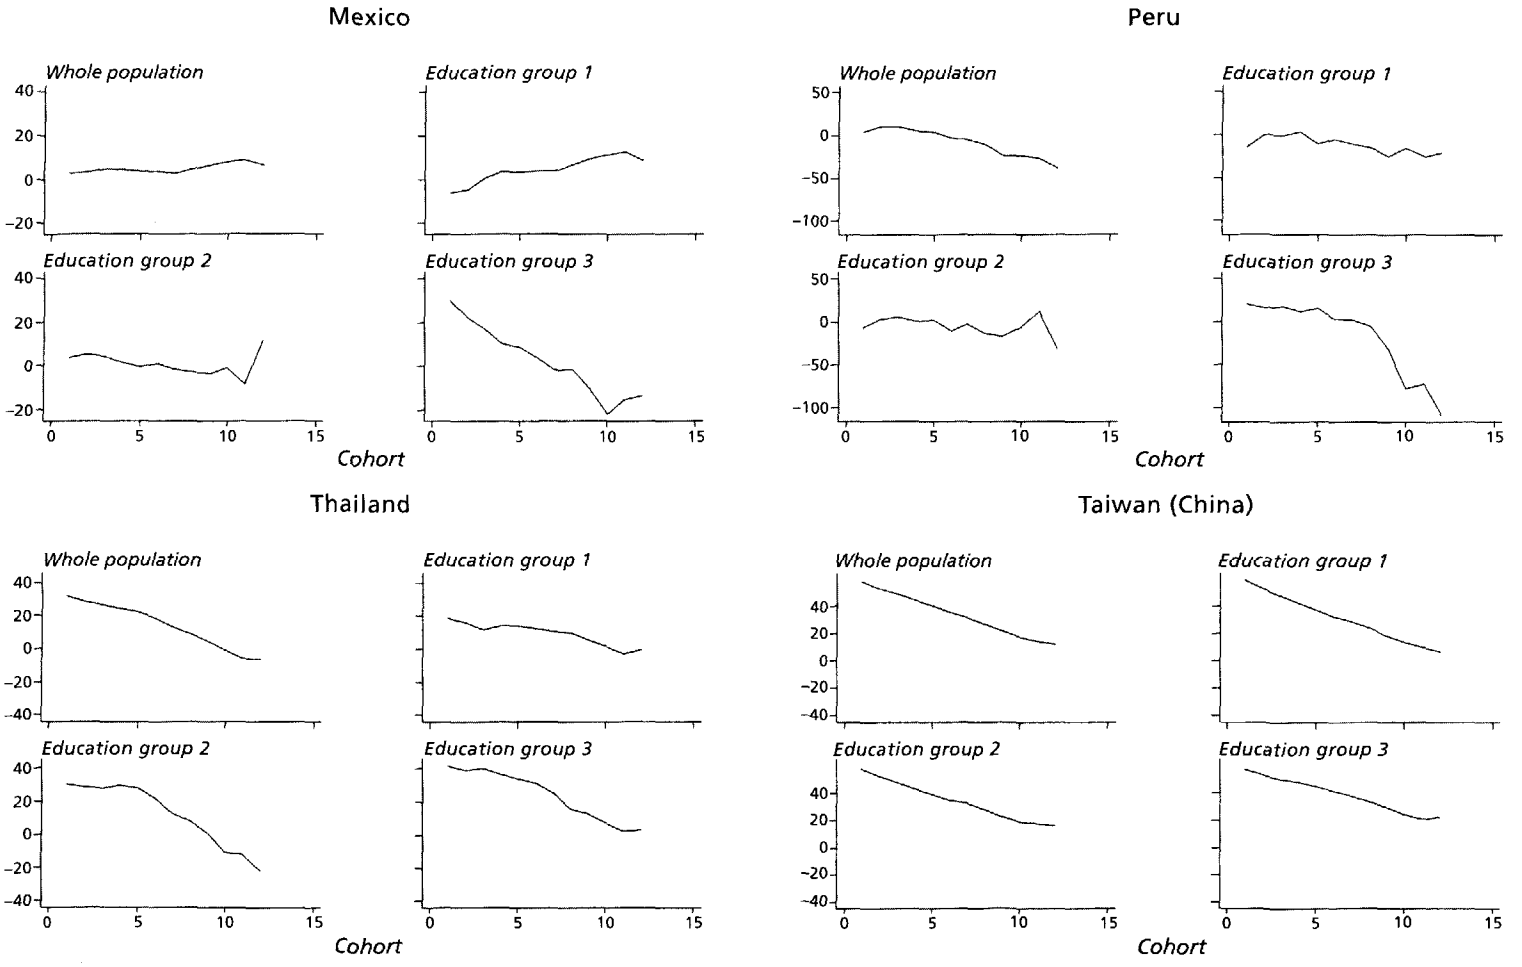

Note: Cohorts at the low end of the range are younger; those at the high end older.
Source: Authors' calculations based on household survey data.

Figure 14 plots the restricted year effects estimated together with the age profile plotted in figures 12 and 13. A large negative shock is evident for all education groups in Mexico in 1996 and in Peru in 1994. In Mexico the effect of the 1996 shock on the most educated is very large and negative, while the effects in previous years are all mildly positive (or very small and negative). For the least educated the aggregate shock is negative in 1992, and the positive time effects in the years before 1996 are smaller than for the rest of the population. The 1996 effect is similar for all groups.

The results for Peru are surprising in that in 1991, the year after the main crisis, the estimated residual is positive. This is a direct consequence of the fact that consumption falls more than income in the survey data. All groups experience a strong negative shock in 1994; only the most educated have a positive time effect in 1997. In Thailand time effects are smaller and stronger for the least educated. In Taiwan (China) the effects are very small, and time effects vary little across education groups.

While Mexico and Peru experienced severe economic crisis during the period under analysis, the economic environment in Taiwan (China) and Thailand was much more stable. Households in Latin America have been exposed to an economic environment characterized by the use rather than the accumulation of savings, which seems to be one reason that they have lower savings rates and that savings have increased much less than in East Asia.

The age and cohort profiles estimated so far depend on the arbitrary normalization that year effects sum to zero and are orthogonal to a linear trend. An alternative restriction that allows the identification of age profiles is the assumption that there are no cohort effects in savings rates. This would be the case if cohort effects in income and consumption exactly canceled each other out. Deaton and Paxon (1994a) suggest that some versions of the life-cycle model do imply such a restriction. By imposing it, we can identify unrestricted year effects.

We plot these age profiles in figure 15, using the first year of the sample as the intercept. For the entire sample the assumption of no cohort effects increases the size of the hump in all economies, with savings rates peaking just after age 60. This pattern is roughly consistent with the life-cycle model. For the least educated the effect is roughly similar: a profile that looked basically flat now reveals a modest hump, with a peak just past age 60. The profile for the middle education group is almost unaffected. The largest effect is observed for the most educated group. The assumption of no cohort effects implies a hump-shaped profile with a dramatic decrease after age 65 in Mexico, Taiwan (China), and Thailand, while the previous profile increased monotonically with age. In Peru there is a clear hump shape, with an increasing trend after age 60.

The unrestricted time effects in figure 16 are also different. These dummy variables for time now capture all the trends in the data that were previously interpreted as cohort effects. In Mexico the strongest negative effect in 1996 is observed among the least educated household heads. Time effects are also nega-

Figure 14. Year Effects on Household Saving Rates

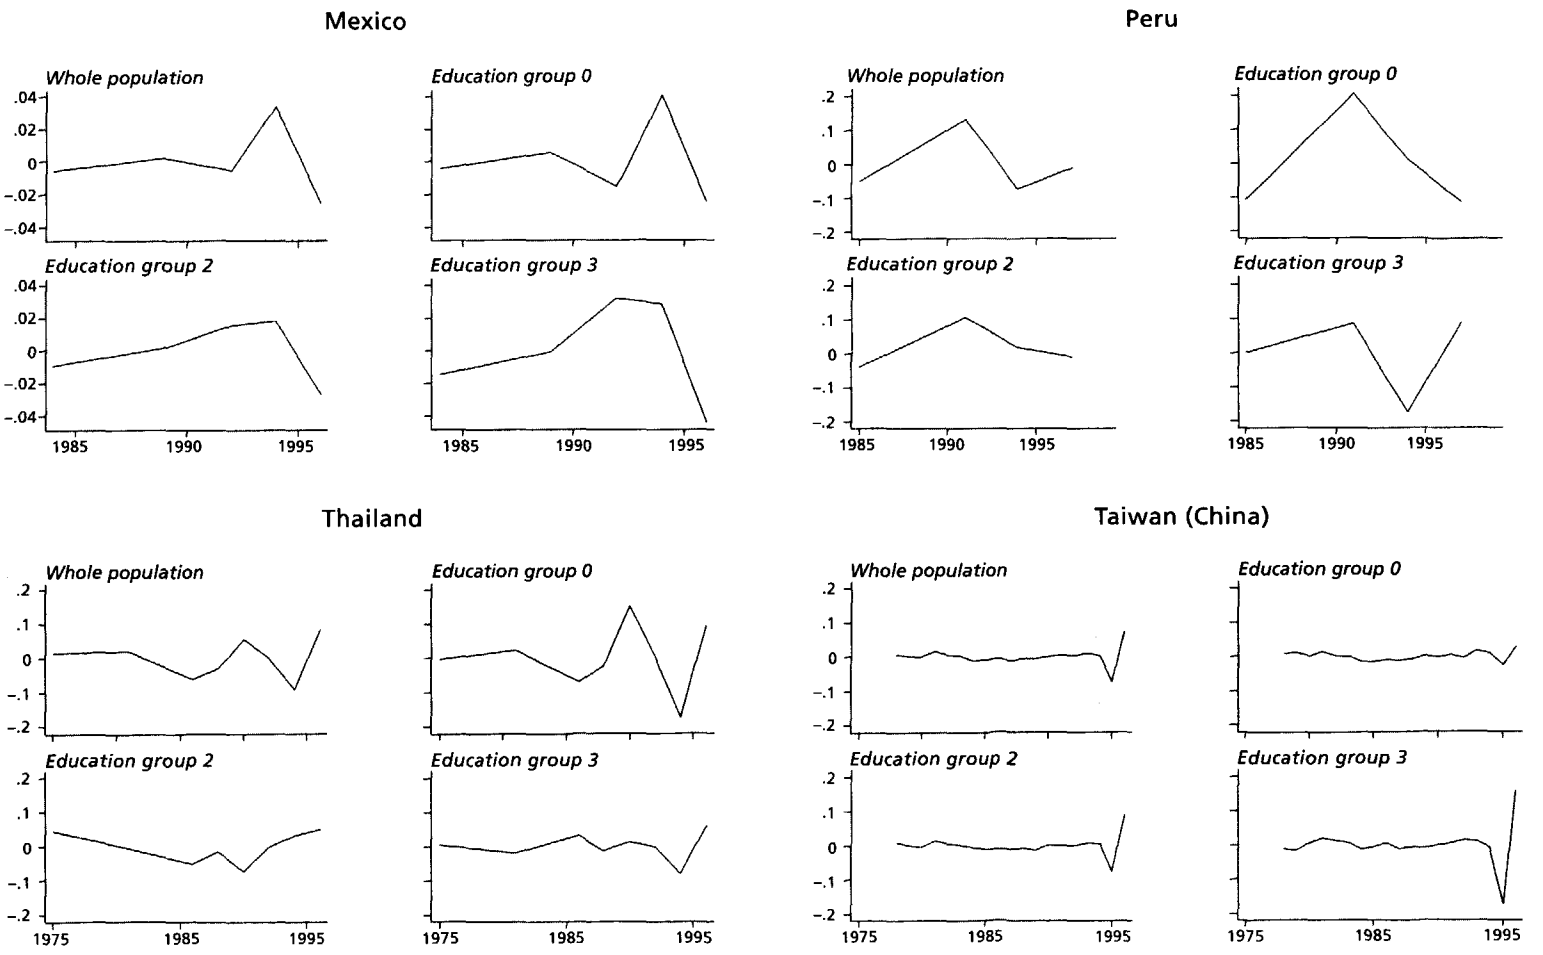

Source: Authors' calculations based on household survey data.

Figure 15. Age Profile of Household Saving Rates, No Cohort Effect Assumption

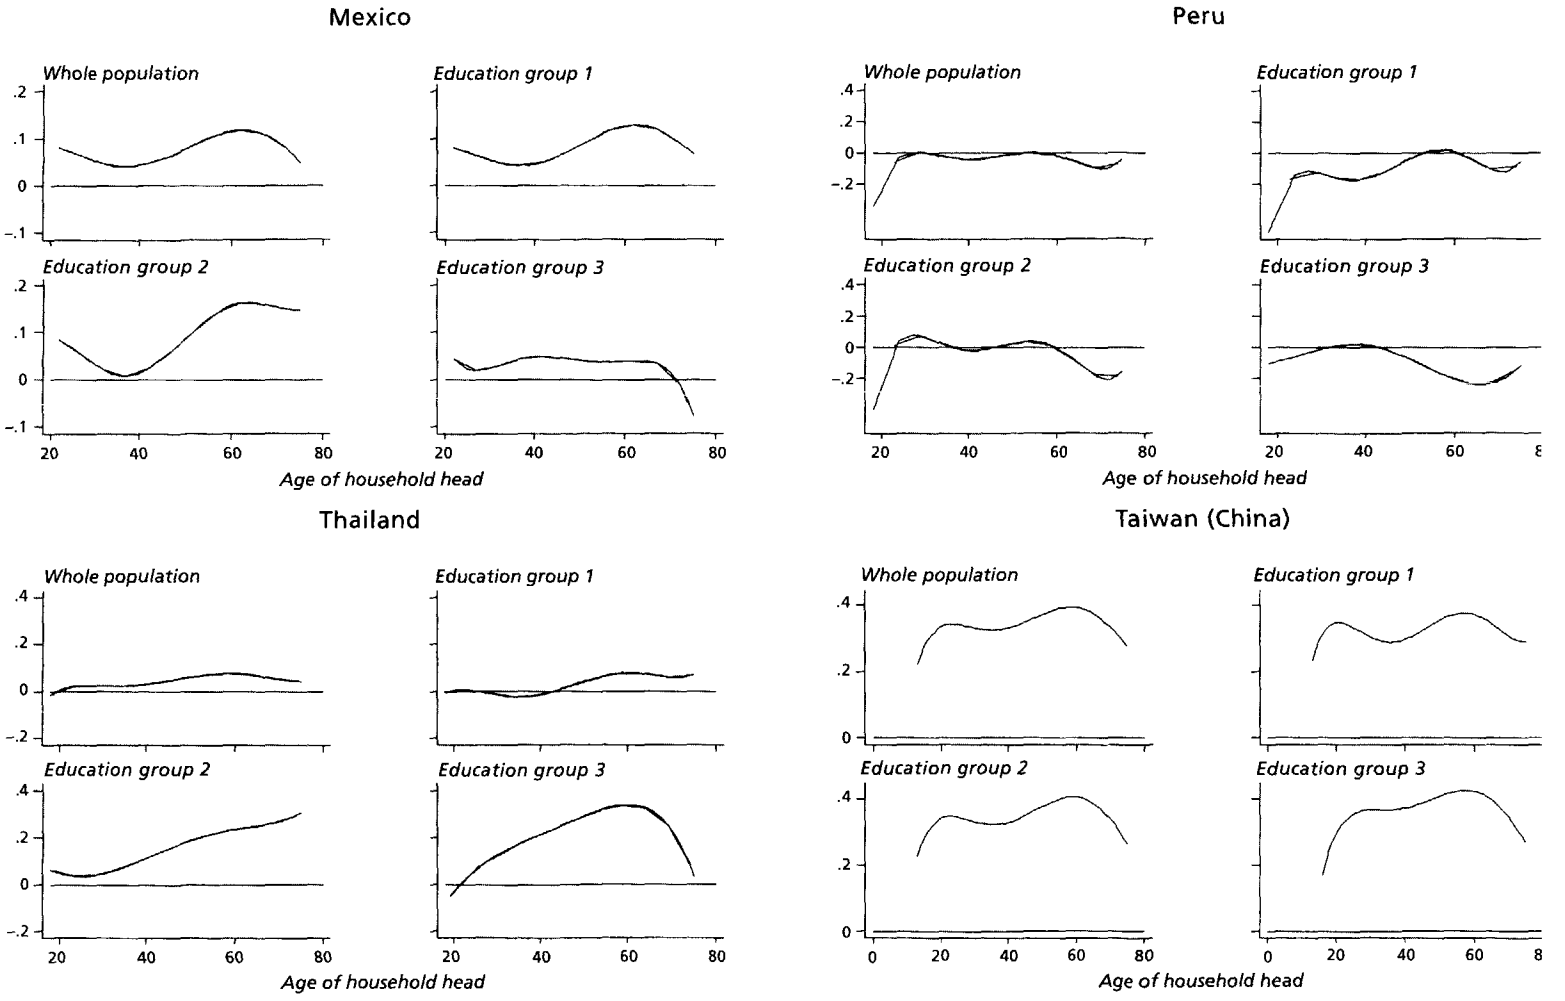

Source: Authors' calculations based on household survey data.

Figure 16. Year Effects on Household Saving Rates, No Cohort Effect Assumption

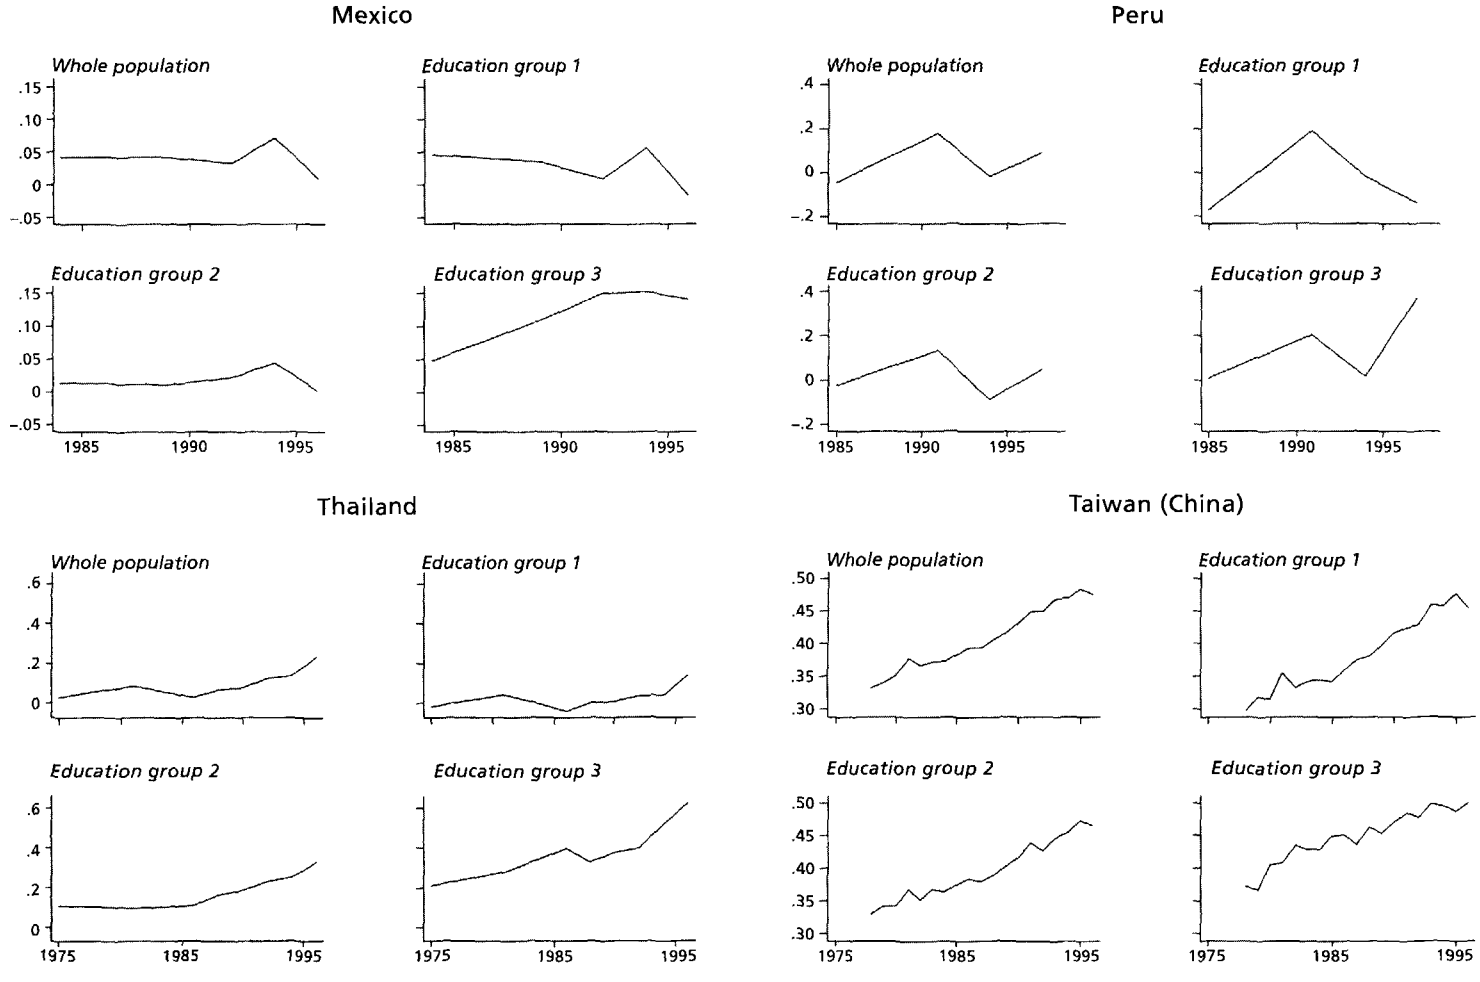

Source: Authors' calculations based on household survey data.

tive in 1996 for the intermediate group; surprisingly, the most educated show a positive time effect in 1996. This result may indicate their greater capacity to smooth variations in income. Similar results are obtained for Peru in 1997. In Taiwan (China) and Thailand the strong cohort effects are mirrored in strongly increasing time effects.

Savings Profiles in East Asia and Latin America

Given the exercise we propose in the next section, it is worth comparing the demographic factors and smoothed savings profiles of East Asia and Latin America. The demographic factors confirm the results of the aggregate data: the demographic transition is much more advanced in East Asia, where fertility rates are much lower and families smaller. The size of the cohort effects in household size and number of children, however, suggests that Mexico and Peru will be in a similar position in a few years.

The picture emerging from the analysis of savings behavior is more complex and more difficult to interpret, if for no other reason than that the estimated effects are conditional on strong and untestable identification assumptions. If we assume that all the trends in the data are explained by age and cohort effects and compare the age profiles of aggregate saving, in Mexico we find a mild hump shape with a peak around age 60, while in Peru and Thailand the profiles rise steadily with age. This picture hides strong differences among education groups, however. In these three countries more educated individuals seem to show no tendency to reduce savings in the last part of the life cycle. In Thailand similar patterns emerge for the other two groups, although the profile for the least educated is basically flat rather than rising. In Mexico and Peru the two lower groups show a mild hump in the last part of the life cycle. This surprising result presumably occurs because although the most educated do most of the saving, their population weight is much lower.

If we move on to cohort effects, the differences between the two regions are even more apparent. In Taiwan (China) and Thailand the two groups with more education show strong positive cohort effects; in Mexico and Peru a similar pattern is observed only for the most educated. At the other extreme, the least educated have negative cohort effects in Mexico and no cohort effects in Peru, while in Taiwan (China) and Thailand they have mildly positive cohort effects. Time effects, as expected, also differ markedly. Strong negative shocks were experienced in Mexico in 1996 and in Peru in 1994; in Taiwan (China) and Thailand time effects were mostly positive. The effects are reinforced when cohort effects are assumed to be equal to zero. Interestingly, time effects in both Latin American countries are stronger (more positive) for the most educated.

The picture changes, however, once we use a different identification assumption—the absence of cohort effects. Under this assumption the estimated profiles do exhibit a hump in the middle of the life cycle, with a peak just before retirement. While this identification assumption is questionable, the results, which

indicate a slight hump in the middle of the life cycle, are not inconsistent with the life-cycle model.

Projections

Is the gap in aggregate saving between East Asia and Latin America likely to narrow in the future, or will it continue to expand? For believers in a standard version of the life-cycle model, there are reasons to believe that the gap will narrow. Latin America is on the verge of a rapid demographic transition that will result in a shift in the population toward age groups expected to save more. Although less pronounced, the predicted demographic changes are similar to those that preceded the boost in private savings rates in East Asia.[10] Whether this expectation will be fulfilled depends largely on the shape of the age profile of saving and how it evolves in the future.

The basic idea we investigate in this section is as follows. If one believes that life-cycle saving is important and that saving is concentrated in a certain age group, then an economy or region in which the share of the population in that age group increases because of the demographic transition (neglecting possible general equilibrium effects on factor prices) might experience an increase in aggregate saving during that transition. Because it is plausible that most life-cycle saving occurs in the middle of the life cycle, economies that experience an increase in the share of middle-aged individuals (and a decline in both youth and old age dependency ratios) might also be expected to experience a temporary increase in saving. This argument is relevant for Latin America, since that region is about to enter such a demographic phase. We attempt to quantify these effects, using the evidence on savings behavior discussed in the previous section.

The projections we present should be interpreted with care. The simulations are based on simple reduced-form relationships, and their use in forecasting future aggregate savings rates assumes that they are stable over time. Specifically, we need to assume that the savings profiles identified in the previous section (and the income profiles used to weight them) do not change in the future. Changes in age profiles of saving could be induced by changes in the shape of the life-cycle profiles of earnings, by changes in factor prices (wages and interest rates), and, in models with habit formation, by the process of growth. Moreover, in some cases we use the last year of our sample as a benchmark, even though it may not be representative.[11] These caveats notwithstanding, the exercises proposed should serve as benchmark calculations for quantifying the potential effects of demographic trends on aggregate saving given the evidence on life-cycle saving.

One important difference between our analysis and earlier studies that have used either macroeconomic aggregate data (such as Behrman, Duryea, and Székely 1999b) or simulations of general equilibrium models (see Attanasio and Violante 2000) is that while we focus only on household saving, those studies focus on domestic saving, which includes public and firm saving. Moreover, those studies are more likely to take into account pension assets and liabilities. In some sense, then,

the exercise based on microeconomic data is more limited in scope. As noted, however, even if the analysis with microeconomic data excludes some important elements of household saving, it also brings considerable gain in that it allows identification of some of the mechanisms driving the dynamics of saving.

Forecasting Household Saving with Microeconomic Data

To forecast aggregate household savings rates, we use the accounting identity given by equation 1, the smoothed profiles we estimated earlier, and demographic projections. Specifically, for any year beginning in the late 1990s, we compute:

(1)
$$\hat{s}_t^{ag} = \sum_c \hat{s}_t^c \frac{\hat{Y}_t^c N_t^c}{\sum_c \hat{Y}_t^c N_t^c}$$

where $\hat{s}_t^c$ is the savings rate of group c at time t predicted by the smoothing procedure used to produce figure 12. In particular, we use the estimated age profile (if c is the year of birth of a cohort, its age will be $t - c$ at time t) and the relevant cohort effects. We use an analogous procedure to compute $\hat{Y}_t^c$. That is, we estimate age and cohort effects using the same procedure used to identify the age and cohort effects for savings rates reported in figures 12 and 13. The N_t^c are from United Nations (1998).

We define the aggregate savings rate as the savings rate of households whose head is 23–75 years old. Because we forecast aggregate savings rates far into the future, new cohorts will join the sample and some cohorts will leave it. For the new cohorts we use the same age profile as for the other cohorts and the cohort effect of the youngest cohort in the sample. This exercise can then be extended to consider different education groups, by simply repeating the exercise for each education group and then aggregating across education groups given some projections about each group's relative size. Because we do not have forecasts of future generations' educational attainment, for future cohorts we use the proportions observed in the youngest cohort. This procedure ignores the fact that future generations are likely to be more educated. But ignoring education groups is equivalent to assuming that changes in the composition of future households will leave the shape of the life-cycle profile of savings rates unchanged, which is obviously unrealistic.

Our aim is not to reproduce the level of aggregate savings rates or efficiently forecast their evolution. As we discussed earlier, there are many reasons why microeconomic data do not match up exactly with aggregate statistics. These reasons are compounded by the fact that the shape of the life-cycle profile is likely to change as a result of changes in its determinants. Our more modest aim is to understand the implications of our estimated age profiles and the predicted demographic trends for the evolution of aggregate savings rates. The results of even this limited exercise must be interpreted with caution, however, for several reasons. First, as emphasized, we identify age and cohort effects under the arbitrary normalization that year effects have zero mean and no trend. Second, we assume that the age profile of savings rates

and income is the same across cohorts, except for an intercept shift—an assumption forced on us by the lack of a long time series of cross-sections. Third, even if existing cohorts (within an education group) have the same age profile of savings rates and income, it is likely that changes in wage profiles, household size and composition, labor force participation, and institutional factors will have an effect on the age profiles of saving. Fourth, changes in the stock of human and physical capital are likely to change wage and interest rates, inducing further changes in age profiles.

Evidence Based on the Microeconomic Simulations

We start with the forecast for Mexico, performing two different exercises. In the first we use the age profile and cohort effects for the overall population (the top-left panel in figures 12 and 13). In the second we use education-specific profiles and cohort effects (the other three panels in figures 12 and 13). To aggregate across education groups, we use the proportions in the sample for cohorts currently alive and the proportions in the last cohort for future cohorts. We plot both sets of forecasts in the first panel of figure 17. Both forecasts show a marked increase in aggregate savings rates that starts leveling off only around 2040. The forecast that does not take into account the education split shows a larger increase, although the forecast that uses the education-specific profiles starts to catch up.

The exercise for Peru shows an increase similar to that for Mexico, although the forecast using education-specific profiles and cohort effects catches up more rapidly. The most interesting result, however, is that the forecasts for the two East Asian economies reveal even greater increases in saving than those for Latin America.

While at first sight these forecasts seem to support the hypothesis that demographic trends will lead to an increase in aggregate savings rates in Latin America, a more careful consideration of the mechanics behind the forecasts shows that this is not necessarily the case. Demographic shifts actually play a small role. The main reason for this is the lack of a hump in the age profiles of saving. Therefore, even when the share of the population between ages 40 and 60 increases—as it is projected to do in the next 40 years in Latin America—the change has little effect on aggregate savings rates; most of the increase in aggregate savings rates is driven by the cohort effect. In figure 12 the cohort effect is strongest for the youngest cohort. Because we give the intercept for that cohort to future cohorts that enter our computation, as older cohorts (with lower intercepts) disappear, the aggregate savings rate increases. The increase is smaller for the education-specific profiles because the cohort effect for the youngest cohort averaged across education groups is smaller than that estimated for the entire population. The catching up occurs because, as the population becomes more and more similar to the first cohort, it has both the intercept and the proportions of educational attainment of that cohort. Because the youngest cohort will be more educated, the savings rate increases slightly more quickly after the first few years.

Similar considerations hold for Taiwan (China) and Thailand, with two notable differences. First, the increase is more marked than in Mexico. And second, the forecast using education-specific profiles catches up faster. We should not read much

Figure 17. Simulated Saving Rates

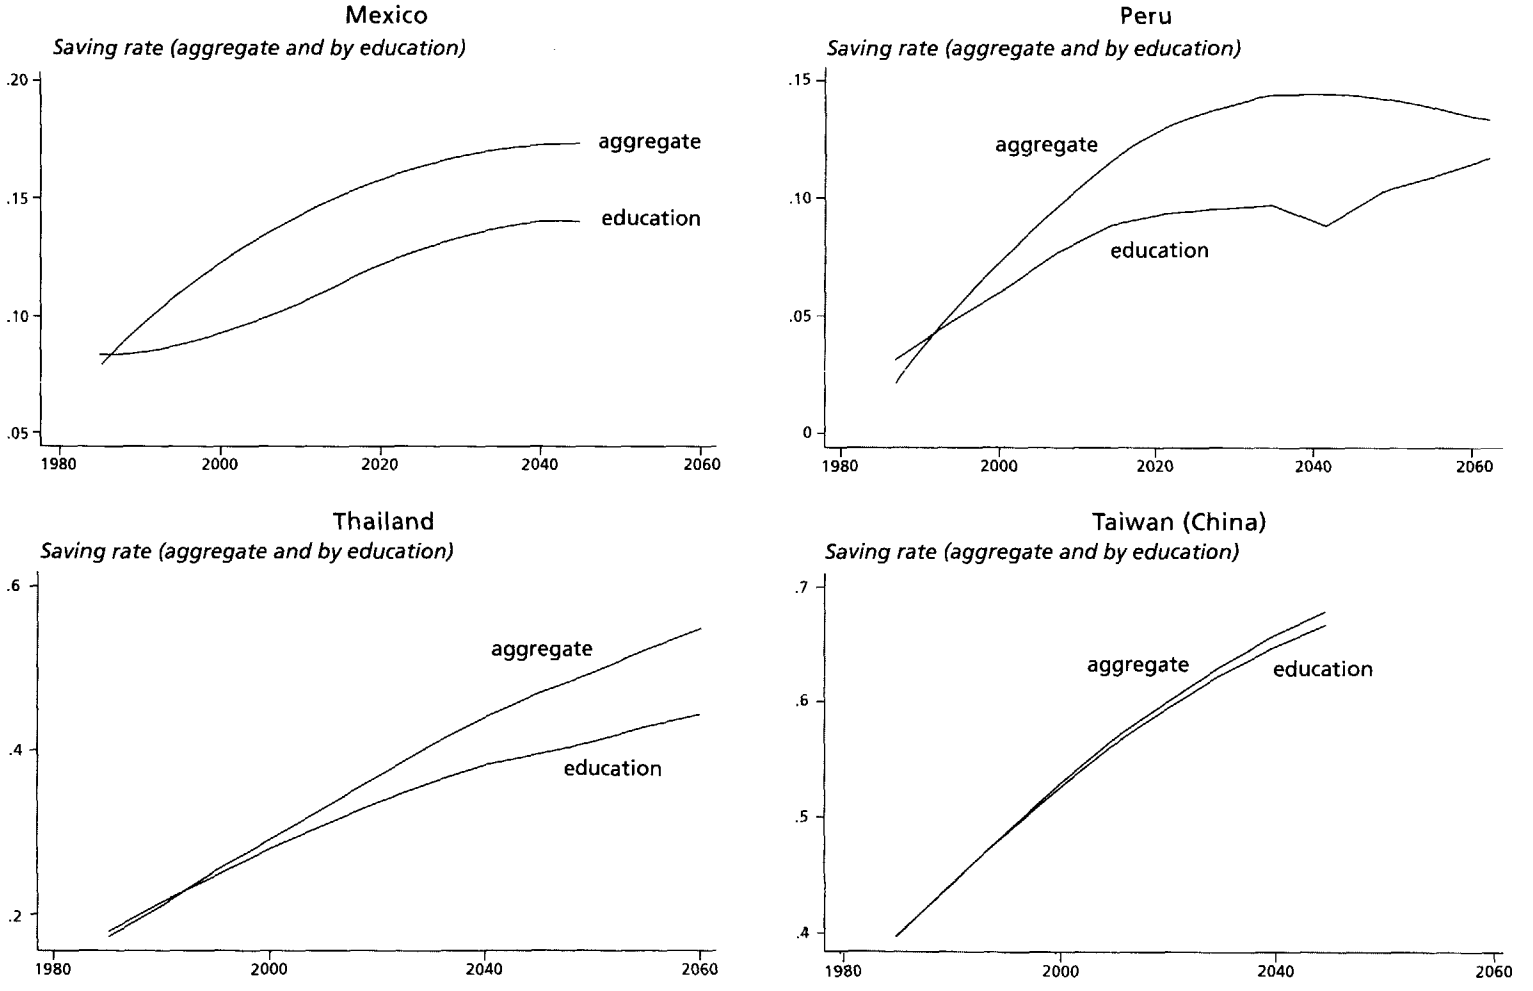

Source: Authors' calculations based on household survey data.

into the first effect because the increase is driven mainly by the estimated intercept for the youngest cohort, as it was for Mexico. Moreover, the absolute level of these profiles does not have a straightforward interpretation because of the definitional and measurement issues raised earlier.

Because the projected increases in aggregate savings rates are driven by the estimate of a single parameter, these results should be interpreted with extreme care. Nevertheless, given our identification assumptions, younger cohorts do seem to be saving more today than their predecessors. Whether this pattern can be maintained in the future will depend on the evolution of the determinants of savings.

We showed earlier that the shape of the estimated profiles depends strongly on the assumption made to identify the age profile of savings rates. If we assume that all the trends observed in the data come from either cohort or age effects and that year effects have zero mean and are orthogonal to such trends, there is no strong evidence of a hump-shaped savings profile in Mexico, Peru, or Thailand. But if we assume that there are no cohort effects, so that year effects can be estimated in an unrestricted fashion, the Mexican and Peruvian age profiles show a marked hump, as do those for the most educated households in Thailand. Both sets of profiles can be used to forecast future household savings rates to check the extent to which the demographic transition is likely to affect aggregate household savings rates.

We perform the exercise described above for Mexico, using the numbers plotted in figure 15 as the age profile of savings rates. For the population as a whole and for two of the three education groups, the profiles show a marked hump. This approach, however, assumes that there are no cohort effects, shutting down the main source of increase in the aggregate savings rate in figure 17.[12] The result we obtain is that savings rates start to decline around 1995, continue to decline for about 20 years, and then begin to increase around 2020 as the population share corresponding to the hump in the savings rate profile starts to increase. But the increase in savings rates is minuscule (about 0.2 percentage point). There are two reasons for the small size of this effect. First, the hump in figure 15 is not very pronounced. Second, although the demographic change we project is relatively large, it results only in a change in weighting, which cannot have a very large effect given the size of the hump. Although the scope of this exercise is limited, we can conclude that, given the estimated shapes of the age profiles of saving (and the projected demographic trends), these forces by themselves are unlikely to result in a large increase in aggregate savings rates if the current economic environment prevails.

Conclusion

In this article we compare East Asia and Latin America in an important dimension in which they have diverged markedly during the past three decades—household savings behavior. In addition to the life-cycle profile of household saving and other variables of interest, we characterize differences in savings behavior across education

groups in two economies in each region. The evidence indicates large differences, both across economies and, within each economy, across education groups.

We document the huge disparity in the level and growth of household saving between Mexico and Peru on the one hand and Taiwan (China) and Thailand on the other. Normalizing time effects to have zero mean and no trend, we identify cohort effects in the data and confirm that relative to older generations, younger generations in East Asia are saving much more than their counterparts in Latin America. Our analysis suggests three main reasons why cohort effects are stronger and total household savings are much higher in East Asia.

First, East Asian households, especially those in younger generations, have had greater capacity to save because of higher income growth, lower fertility rates, different household structure (with more elderly individuals living in extended households, which prevents us from observing a decline in saving toward the end of the life cycle of individuals), and a much more advanced demographic transition (with larger shares of the population in the ages at which productivity and saving peak).

Second, the macroeconomic environment in Latin America was highly volatile, and the two countries we analyze were subject to severe shocks during the period under study. In contrast, East Asia was much more stable economically. Thus in Latin America the context was one in which household savings were typically used rather than created, while in East Asia savings were built up smoothly to accumulate resources.

Third, in Latin America almost all household savings are generated by the richest 20 percent of the population, while in East Asia saving is much more widespread. As a result, in Latin America the savings rate of households in the richest quintile does not differ significantly from that of their East Asian counterparts, but there are huge disparities between the richest households and the rest of the population. For the poorest 50 percent of households in Latin America, savings have been much less responsive to increases in income during periods of economic growth and more sensitive to declines during downturns. Thus differences in the capacity to save across the income distribution account for an important part of the difference in total household saving between the two regions.

We do not find strong evidence of negative saving or even declining saving in the last part of the life cycle in any of the economies studied. While this finding contradicts a simple version of the life-cycle model, a conclusive judgment can be obtained only if one explicitly takes into account the variation in needs induced by changes in household size and composition over the life cycle as well as changes in labor supply behavior. In this study we document differences in life-cycle profiles in these variables, but we do not explicitly consider their effect on savings rates.

Another issue we have ignored is the effect of different institutional settings, particularly pension arrangements, on savings behavior over the life cycle. Using evidence on household saving to test alternative models of consumption and savings decisions does not obviate the need to consider these factors. However, some of the evidence—particularly that from Taiwan (China) and Thailand—suggests a need for a more complex model than the simple version of the life-cycle model. Paxson

(1996) and Deaton and Paxon (1994a) suggest considering models with habit formation, in which growth in itself induces an increase in savings rates, at least until the stock of habits "catches up." This issue is extremely important for understanding the relationship between growth and saving, and more detailed study is needed.

In all four economies we find that households headed by the most educated save considerably more than households headed by those with less education (the exception could be Taiwan [China]). This finding is consistent with the fact that more educated individuals experience more variation in lifetime income.

We also identify age, cohort, and time effects in the data for each of the education groups. Regardless of the identification assumptions, we find that time effects are small and mostly homogeneous across groups in East Asia. The effects are much stronger in Latin America, however, where there are large differences across groups, with larger negative effects in downturns and smaller positive effects in upturns for the least educated households. This is especially true under the assumption of no cohort effects, suggesting that less educated individuals in Latin America have more limited capacity to smooth unexpected variations in income and to build up income-earning assets in good times.

Cohort effects, identified under the assumptions on time effects mentioned above, are strong and positive in Taiwan (China) and Thailand for all education groups and strong and positive only for the most educated in Mexico and Peru. Cohort effects among the most educated in Latin America appear to be even stronger than those experienced by their counterparts in East Asia, but overall they are almost flat because of negative or no cohort effects for the other groups.

As with the total population, age effects are estimated with two alternative specifications. When dummy variables for the years are constrained to have zero mean and to be orthogonal to a time trend, the main difference is that in Mexico, Peru, and Thailand age profiles for the least educated and those with secondary education are mostly flat, while they increase monotonically for the most educated. In Taiwan (China) the monotonic increase is observed across all groups. When the age profiles are estimated under the assumption of no cohort effects, the profile for the most educated in Mexico, Peru, and Thailand appears to be more hump shaped than the profiles for those with less schooling. In Taiwan (China) the age profile also becomes much more hump shaped, but there are no differences across education groups. Under some identification assumptions, we thus conclude that the savings behavior of the most educated households in Mexico, Peru, and Thailand is more in line with the life-cycle hypothesis than is the savings behavior of households with less education.

It has been argued that because Latin America is on the verge of a demographic transition similar to that already experienced by East Asia, future demographic trends might bring about an increase in aggregate saving that will reduce the savings gap between the regions. We present simulations indicating that although our life-cycle profiles and cohort effects predict an increase in aggregate savings rates, the changes cannot be attributed to current demographic trends. The increase is driven mainly by strong cohort effects, identified under

the assumption that all the trends in the data can be interpreted as either age or cohort effects. Moreover, although the shape of the age profile is much more in line with the life-cycle model when we use the alternative identifying assumption that there are no cohort effects in savings rates, the estimated hump in the middle of the life cycle is not large enough to generate increases in aggregate savings rates in Latin America large enough to catch up with savings rates in East Asia in the next 20 years. As we emphasize, the simulations should be interpreted with great caution because they are based on very strong identification assumptions and assume that the economic environment in East Asia and Latin America will remain unchanged in the future. This is a rather pessimistic scenario for Mexico and Peru—and perhaps an overly optimistic one for Taiwan (China) and Thailand.

Our results suggest that the projected demographic trends by themselves are unlikely to generate large increases in savings rates under current circumstances. This does not mean that these trends are unimportant. They will play an important role in determining the ability of developing regions to attract capital flows from industrial countries, where capital-labor ratios are projected to be much higher in the future than in developing countries (an issue discussed by Attanasio and Violante 2000).

It is also possible that the estimated age profile of saving, whose shape is responsible for our results, will change as a consequence of the structural changes Latin America is experiencing. Two of these changes are particularly important. The first is changes in labor supply behavior, in particular, an increase in female labor force participation. The second is the shift from public to privately funded pension schemes that has occurred in many Latin American countries in recent years.

Notes

1. The graphs plot the coefficients for the dummy variables for the years that result from putting together two panels with a different mix of countries depending on the region and then estimating fixed effects regressions on each panel, where the dependent variable is the savings rate.

2. The Mexican, Thai, and Taiwanese surveys considered here are strictly comparable in terms of questionnaires, objectives, and sampling techniques. The data for Peru are largely comparable. The 1991 survey excludes some rural areas from the sample, but restricting the comparison to exactly the same geographic areas in all four years does not change any of our conclusions.

3. For Peru we are able to compute only s1 because of the lack of data by expenditure item. The estimates of s_1 and s_3 for Mexico for 1984, 1989, and 1992 do not coincide exactly with the estimates reported in Székely (1998), although they are produced with the same data. The difference is that Székely measures saving as the difference between disposable income and nondurable consumption but adjusts consumption to include interest payments on debt. We have not made the same adjustment here. As a result, the savings rates are 2.6, 1.8, and 2.3 percentage points higher for s_1, and 2.4, 1.7, and 2.6 percentage points greater for s_3, than those reported in Székely (1998). All estimates for Mexico are compatible with those in Attanasio and Székely (1998).

4. Sampling error, induced by small cell sizes, could be interpreted as a time effect. However, it is plausible to assume that sampling error has zero mean and does not exhibit any time trend.

5. Carroll and Summers (1991) interpret such results as a failure of the life-cycle model. Attanasio and Browning (1995) and Attanasio and others (1998) offer reasons, in the context of the United Kingdom and the United States, why the life-cycle model is not necessarily at variance with the data in figure 7.

6. The increase in the last part of the life cycle in Peru is implausibly strong. We tried different specifications for the age polynomials, obtaining similar results. The result is driven by the marked increase in savings rates for the two oldest cohorts.

7. There are two issues. On the one hand, the life-cycle model implies that with a hump-shaped income profile, savings should also be hump shaped, if needs are constant over the life cycle. But as we saw earlier, household size—and therefore needs—change considerably over the life cycle. On the other hand, simple versions of the model imply that saving should decline (and become negative) after retirement. In principle, several data issues might explain why we typically do not observe negative savings rates in the last part of the life cycle. However, the evidence from the four economies seems to indicate that at least for the most educated, saving increases in the last part of the life cycle.

8. Our results for age and cohort effects in Thailand differ from those reported by Paxson (1996). Paxson, who considers the entire population and does not look at education groups, finds an age profile of savings rates that is very flat over the life cycle. While we use slightly different selection criteria (and unlike Paxson do not use expansion factors because of difficulties computing them), the main reason for the difference seems to be that Paxson does not use the data from 1994 and 1996. When we drop those years, we obtain results much closer to hers. The reason is that in the last two years savings rates increase considerably for all cohorts. Our smoothing procedure forces us to interpret these trends as either age or cohort effects, as time effects are assumed to be orthogonal to a time trend. This gives us the rising age profiles for the entire population and for each of the education groups and should remind us of the caveats discussed earlier. Interestingly, our results for Thailand are similar to those that Paxson reports for Taiwan (China).

9. Our results for age and cohort effects in Taiwan (China) are very similar to those reported by Deaton and Paxon (1994b), even though they use data only up to 1990. Our results are also consistent with those reported in Deaton and Paxon (2000).

10. Bloom and Williamson (1999) and Behrman, Duryea, and Székely (1999a) have argued along these lines. Attanasio and Violante (2000) simulate the effects of demographics on private savings and predict a large increase in Latin America as a result of future reductions in the old age dependency ratio. That study, however, focuses on the general equilibrium effects (that is, changes in wage and interest rates).

11. This is a particularly important issue for Mexico, for which 1996, the year we use as a benchmark, is by no means a standard year for comparison. As shown in the previous section, there are strong negative effects from the shock the economy faced in 1995. Assuming that the same conditions will prevail is thus a rather pessimistic scenario. For Taiwan (China) and Thailand the benchmark is 1996, two years before the recent financial crisis. In those economies economic conditions may have changed for the worse, and projections based on the conditions in 1996 can be regarded as rather optimistic.

12. In this exercise the level of savings rates is particularly difficult to pin down because year effects, by definition, are unpredictable.

References

Attanasio, Orazio P. 1994. "A Cohort Analysis of U.S. Household Saving Behavior." NBER Working Paper 4454. National Bureau of Economic Research, Cambridge, Mass.

———. 1998. "Cohort Analysis of Saving Behavior by U.S. Households." *Journal of Human Resources* 33 (3): 575–609.

Attanasio, Orazio P., and Martin Browning. 1995. "Consumption of the Life Cycle and over the Business Cycle." *American Economic Review* 85 (5): 1118–37.

Attanasio, Orazio P., and Henry Hoynes. 2000. "Wealth Accumulation and Differential Mortality." *Journal of Human Resources* 35 (1): 29.

Attanasio, Orazio P., and Miguel Székely. 1998. "Household Saving and Income Distribution in Mexico." OCE Working Paper Series 390. Inter-American Development Bank, Office of the Chief Economist, Washington, D.C.

Attanasio, Orazio P., and Giau Lucca Violante. 2000. "The Demographic Transition in Closed and Open Economies: A Tale of Two Regions." RES Working Paper Series 412. Inter-American Development Bank, Research Department, Washington, D.C.

Attanasio, Orazio P., James Banks, Costas Meghir, and Guillermo Weber. 1998. "Humps and Bumps in Lifetime Consumption." *Journal of Business and Economic Statistics* 17 (1): 22–35.

Barro, Robert J., and Jong-Wha Lee. 1993. "International Comparison of Educational Attainment." *Journal of Monetary Economics* 32 (3): 363–94.

Behrman, Jere, Suzanne Duryea, and Miguel Székely. 1999a. "Aging and Economic Opportunities" OCE Working Paper Series 405. Inter-American Development Bank, Research Department, Washington, D.C.

———. 1999b. "Schooling Investments and Macroeconomic Conditions: A Micro-Macro Investigation for Latin America and the Caribbean." OCE Working Paper Series 407. Inter-American Development Bank, Research Department, Washington, D.C.

Bloom, David E., and Jeffrey G. Williamson. 1999. "Demographic Transitions and Economic Miracles in Emerging Asia." *World Bank Economic Review* 12 (3): 419–56.

Browning, Martin, Angus Deaton, and M. Irish. 1985. "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle." *Econometrica* 53 (3): 503–43.

Carroll, Chris, and Lawrence H. Summers. 1991. "Consumption Growth Parallels Income Growth: Some New Evidence." In B. D. Bernheim and J. B. Shoven, eds., *National Saving and Economic Performance.* Chicago: University of Chicago Press.

Deaton, Angus, and Christina Paxon. 1994a. "Intertemporal Choice." *Journal of Political Economy* 102 (3): 437–67.

———. 1994b. "Saving, Growth, and Aging in Taiwan." In D. Wise, ed., *Studies in the Economics of Aging.* Chicago: University of Chicago Press.

———. 2000. "Growth and Saving among Individuals and Households." *Review of Economics and Statistics* 82 (2): 213–25.

Heckman, James, and Richard Robb. 1987. "Using Longitudinal Data to Estimate Age, Period, and Cohorts Effects in Earnings Equations." University of Chicago, Department of Economics.

Lora, Eduardo, and Carmen Pagés. 2000. "Hacia un envejecimiento responsible: Las reformas de los sistemas de pensiones en America Latina." Inter-American Development Bank, Research Department, Washington, D.C.

MaCurdy, Tom, and Thomas Mroz. 1995. "Measuring Macroeconomic Shifts in Wages from Cohort Specifications." Stanford University, Department of Economics, Stanford, Calif.

Paxson, Christina. 1996. "Savings and Growth: Evidence from Micro Data." *European Economic Review* 40 (2): 255–88.

Székely, Miguel 1998. "Nivel y distribución del ahorro de los hogares en México." *El Trimestre Economico* 65 (2): 263–314.

United Nations. 1998. *World Population Prospects 1950–2050: The 1998 Revision* (Database). United Nations New York: Population Division.

World Bank. 1999. World Development Indicators 1999. Washington, D.C.