Aging, saving and the welfare of the elderly in developing Asia

Angus Deaton
Christina Paxson

Research Program in Development Studies
Princeton University

First Rough Draft, June 5 1998

PLEASE DO NOT QUOTE WITHOUT PERMISSION

1. Introduction

As the demographic transition works its way through Asia, the fraction of the population that is elderly is rising. In China, 5.8 percent of the population was aged 65 or over in 1990; by 2010, the fraction will be 8.3 percent, and by 2025 13.3 percent. For India and Indonesia, the corresponding figures are 3.7 and 3.0 percent in 1990, 5.3 and 5.9 percent in 2010, and 7.8 and 9.8 percent in 2025. For many countries, this slow change in demographic structure is accented by the baby booms that took place immediately after the Second World War so that, in addition to trend aging, there is a large cohort of people whose oldest members will reach the ages of 60–65 in the next ten to fifteen years. By 2010, Singapore and Korea will have 10.4 and 9.0 percent of the population aged 65 percent or over, numbers that will rise to 20.6 and 15.2 percent by 2025. During their working years, the size of the baby boom cohort and its relatively low fertility ensured that unusually large fractions of these populations were economically productive, and this and their life-cycle savings have been seen as one of the causes of the rapid rates of economic growth in much of the region. As the cohort ages, there is a corresponding concern that these processes will reverse themselves, and that the future may not be as bright as the recent past.

In this paper, we take up two of the many possible economic issues associated with the aging of Asia's baby boom generation. We look first at the relationship between savings rates and demographic structure, and address the old question of the link between population structure and the aggregate rate of saving. We ask whether the empirical evidence is consistent with the story that credits demography with Asia's high saving rates, and whether future demographic changes are likely to lower them. Our second issue is specific to the baby boom generation itself, and concerns whether being the member of an unusually large cohort lowers wages, earnings, and lifetime welfare relative to that of members of smaller cohorts.

The theoretical link between demography and savings comes from the life-cycle hypothesis, according to which the main motive for saving is accumulation for retirement. Young people save and old people dissave, so that changes in the age structure of the population alter aggregate saving rates. More realistically, young couples may save little or dissave when they are bringing up children, only saving in middle-age for their retirement, so that aggregate saving depends not only on the balance between young and old, but on the fractions of children, of working-age adults, and of the elderly. Whether these theoretical patterns are important in practice depends on the importance of life-cycle relative to other forms of saving, for example bequests or precautionary motives, and on whether the family or financial markets are the usual providers of support in old-age. It is also important to note that, even if the evidence were to provide a clear link between aging and saving, and even if we could confidently expect the aging of Asia to drive down its saving rate, there would not necessarily be any cause for concern. We do not worry when individuals draw down their assets to support their retirement, and there is no need to worry when nations with a large fraction of elderly do likewise. In closed economies, decumulation of capital can be expected to cause a (temporary) decline in the rate of economic growth, though not even that would happen if international capital markets were fully-integrated, thus decoupling investment and saving on a country-by-country basis. Even so, there are concerns that, when many countries in the world are aging at the same time, there could be a global capital shortage, and that international funding will not be available to supplant domestic shortfalls. There are also many governments (and economists) who see falling rates of growth as a concern in their own right, whatever their cause. In this paper, we do not attempt to address these broader questions of the link between saving and growth, or the desirability of growth in its own right. Instead, we are concerned with the factual question: treating demography as given, can we attribute saving rates to changes in demographic structure?

The theoretical link between cohort size and its relative welfare comes from the supposition that workers of different ages are not perfect substitutes in the labor market. As a result, being a member of a large cohort can mean relatively low wages throughout life, and relatively low lifetime welfare compared with the smaller preceding or succeeding cohorts. As with the link with saving, the existence of a cohort-size effect is not necessarily a matter for concern for policy. In Asia, members of the baby boom generation are much better-off than were their parents or grandparents, and they have benefited from the economic growth of the region in general. Nevertheless, if this generation turns out to be relatively poorer than its immediate neighbors, and given that the within-group inequality in consumption, income, and earnings is larger among the elderly than among younger people, Deaton and Paxson (1994), there is an unusually large group of people that are vulnerable to unexpectedly poor economic outcomes.

The paper is laid out as follows. In Section 2, which is concerned with saving and demography, we look at a group of fourteen Asian countries, and document developments in their saving and demography since the mid-1960s. We also review the literature linking demography to saving in these countries, using cross-country comparisons, time-series within countries, and household survey data. Cross-country and time-series work using aggregate data has typically found (often large) negative effects on saving rates of both youth and elderly dependency rates. Yet the microeconomic data that would appear to be the ultimate source for documenting savings patterns typically provide no support for the aggregate findings. We review some of these studies, including our own work for Indonesia, Taiwan, and Thailand, and discuss some of the reasons why they may give misleading results. We argue that the household data are indeed consistent with the existence of a link from demographic structure to saving, but that the effects are too small to be consistent with most of the macroeconomic evidence. We then turn to an analysis of annual time-series data for the fourteen Asian countries for which we have adequate data, Bangladesh (73–94), China(65–94), Hong Kong (65–94), Indonesia (66–94), India (65–94), Korea (65–94), Sri Lanka (65–94), Malaysia (65–94), Nepal (75–94), Pakistan (68–94), Philippines (65–94), Singapore (65–94), Thailand (65–94) and Taiwan (65– 95). We readily reproduce the link between demography and saving that appears in much of the literature, but show that it does not survive a more detailed and appropriate analysis. Although the analysis of cross-country data, or of pooled time-series cross-section data show the conventional effects, we argue that such analysis does not provide consistent estimates of the quantity of interest, which is the effect of dependency rates on the aggregate saving rates for the countries taken together. The problem is that the estimates differ from country to country, so that the estimates from pooled regressions are neither readily interpretable nor interesting. A more appropriate procedure is to estimate on a country by country basis, and then enhance precision by averaging the results across countries. Once this is done, the link between dependency and savings vanishes, or becomes perverse.

Section 3 turns to the question of whether cohort size is a determinant of economic well-being, with an application to Taiwan. The analysis of this question requires micro-level data on the wages and incomes of individuals (or households) over a large number of years, together with information on the size of different birth cohorts. These data requirements are fulfilled in Taiwan. In addition, Taiwan experienced a very large post-war baby boom, followed by a rapid fertility decline. Because shifts in the age structure over the past half century have been so dramatic, Taiwan is an excellent country in which to look for cohort-size effects. Contrary to the research that has been done on the US baby boom, we find essentially no effect of cohort size on a large number of economic measures, including wage rates, earnings, household income and consumption, and individual income and consumption. The baby boom generation in Taiwan has not suffered relative to those in smaller cohorts.

2. Demographic structure and saving in Asia

2.1 Saving rates and dependency rates: data

Figures 2.1 and 2.2 show data on saving rates and demographic trends in the fourteen south and east Asian countries for which we have useful data. (The obvious geographical exclusion is Japan which, as an advanced industrialized country, is outside our current concerns.) Figure 2.1 shows *national* saving rates, defined as the difference between gross national disposable income and the sum of private and public consumption as a percentage of gross national disposable income. These data are taken from the World Bank's recently constructed database on saving, Loayza, Lopez, Schmidt-Hebbel, and Serven (1998). It is important to note that these saving rates include public saving, including central and local government, and enterprise saving, as well as saving by households. Although for many policy purposes, it is national saving that is the ultimate concern, the life-cycle theory, which generates the predictions about demography and saving, is a theory of individual or household saving. Data on aggregate private saving and on its disaggregation between households and enterprises are too sparse to permit a replication of our analysis, but it should not be assumed that household saving rates move similarly to national saving rates. Central governments run large surpluses in many Asian countries, and there are cases (such as

Thailand) where it is known that a substantial fraction of private saving is done by enterprises, not households.

The aggregate data show the familiar strong upward trends in saving rates in all but one or two cases. These are matched by *downward* trends in dependency rates in Figure 2.2. The bottom lines in this figure show the youth dependency rate—the fraction of the population aged less than 15—while the upper line shows the total dependency rate, which is the sum of the fraction young and the fraction older than 65. If we start out from the premise that children and the elderly are "burdens" that reduce the ability to save, then demographic change in Asia has been favorable towards the increases in saving rates that have taken place. This is what the Asian Development Bank (1998) calls Asia's "demographic gift," to which they attribute about a third of its economic growth in recent decades.

The old-age and child dependency ratios have evolved differently in the past, and will evolve differently in the future. The ratio of children in these countries has typically been falling with fertility decline, and is predicted to continue to do so at least for another 10 or 15 years, see for example Heller and Symansky (1997, Chart 1). In the countries where the transition is most advanced, and where the baby booms were most pronounced—China, Korea, Hong Kong, Singapore, and Taiwan—the child-dependency rates will begin to rise after around 2010–15. This will eventually happen in the other countries too, but not until much later in the century; for example, in Indonesia, Malaysia, and Thailand, the child dependency ratio is projected still to be falling in 2050. By contrast, the old age dependency ratio is already rising in nearly all of these countries, as can be seen from the widening gaps between the two lines. In Singapore and Hong Kong, this effect is already large enough for the total dependency rate to have begun rising, which will eventually also be the case in the other countries. By 2015 or shortly thereafter, the total dependency rate will also be rising in China, Indonesia, Korea, Malaysia, Taiwan, and Thailand. The demographic "gift" will have become a demographic "levy."

In the analysis below, we shall attempt to assess the extent to which the evidence supports a link between the demographic gift and high and growing saving rates. This is an obvious first step in assessing what might lie in the future. Before doing so, and before looking at the some of the previous literature, it is important to note some obvious caveats. First, our task involves out-of-sample prediction, and suffers from all the associated dangers. In particular, political and other institutions may change in response to shifts in demographic trends, and these institutional changes are likely to change behavior—indeed, that is likely to be their purpose. Social security systems are an obvious example. In none of these countries does a large pay-as-you-go social security scheme present the same challenge to savings and to public policy as is the case in many industrialized countries, but social security arrangements are likely to become ore important over time, especially in the richer countries. Second, the out-of-sample exercise is made more difficult by the slowly changing nature of the demographic variables, and the general econometric difficulty of identifying the effects of trend-like changes. Third, our analysis, like the rest of the literature, treats demographic changes as exogenous, in spite of the fact, obvious even from Figure 2.2, that the richer and more rapidly growing countries are more advanced in the demographic transition. We do not believe that there is any econometric way of dealing with this endogeneity, but we would nevertheless argue that it is important to examine whether it is possible to tell a story of saving, conditional on demographics. Much analysis and policy prescription is based on the "demographics causes economics" view, and it is important to test whether such views are internally consistent on their own terms.

2.2 Saving rates and dependency rates: literature

There is a large literature on the effects of demographic structure on savings rates. Coale and Hoover (1958) took it for granted that saving was hindered by the high dependency rates associated with high fertility and high population growth. Much subsequent research has been motivated by the life-cycle hypothesis. Leff (1969) was one of the first to find empirical support for dependency effects based on cross-section international evidence, though subsequent work has cast considerable doubt on the robustness of his results. Since then, there have been numerous studies based on cross-section analysis, on time-series analysis within individual countries, and on pooled time-series cross-section data. Recent reviews of this material are given in Meredith (1995), Heller and Symansky (1997), and Turner, Giorno, De Serres, Vourc'h, and Richardson (1998). For example, Turner et al, Table 2, provide a range of estimates from cross-section and time-series studies on industrial countries in which a one percent point increase in the youth dependency ratio decreases the private saving rate by between 1.10 and 0.13 percentage points, whereas the effects of the elderly dependency rates are typically larger, with a one percent increase decreasing private saving rates by from 1.61 to 0.14 percentage points. Typical estimates over all these studies might be –0.75 for the youth ratio and –1.00 for the elderly ratio.

Heller and Symansky (1997) and Higgins and Williamson (1997) discuss similar studies that use data for developing countries. While there are fewer such studies, the results appear to be roughly consistent with those from the industrialized countries. Higgins and Williamson's own results are based on the same countries used here, plus Myanmar and Japan, but less Hong Kong, Nepal, and Taiwan. Although the results are not presented in a form that allow ready comparison with other studies, Higgins and Williamson find large and (apparently) significant effects of the age structure, both at young and old ages, and their results are the basis for the Asian Development Bank's attribution of one-third of growth to the demographic "gift." Rather than critique these previous studies, we shall offer in the next subsection our own analysis of the Asian data, where we also try to address at least some of the econometric issues that have compromised the earlier work.

Another important issue in the literature is the extent to which aggregate analyses are at odds with analyses that use household survey data. Turner et al (1998), citing the studies of Britain, the U.S., Japan, and Germany in Poterba (1994), emphasize that household survey data show no clear traces of dissaving associated either with children or with old age. In our own work, Paxson (1996) and Deaton and Paxson (1997, 1998a), we have consistently been unable to find evidence of the predicted age pattern in household saving rates, not only in the U.S. and Britain, but also in Taiwan, Thailand, and Indonesia. Most notable is the singular lack of evidence that households dissave in old age; indeed several studies, particularly those for Germany and Japan, find precisely the opposite of the life-cycle prediction, that elderly households have high saving rates.

If the microeconomic results are correct, there is something wrong with the aggregate regressions, a possibility that we take up in the next subsection. But there are also a number of other reasons why the macro and micro results might be different.

First, the data are different. In the developing countries, aggregate data typically do not distinguish household from enterprise saving, so that if the latter is important in the aggregate, there is no automatic presumption that the results should be the same. This argument clearly has even greater force when we are forced to work with national saving, so that government and private activities are lumped together. Many social expenditures, for health, education, and pensions, are tied to the demographic structure of the population, so that we might expect to find effects of demographics on government saving, even in the absence of any such effect by households.

Second, even in industrialized countries, where data typically exist on household saving in aggregate, there are important items of income that are handled differently in the micro and macro data. Contributions to pension funds by employers on behalf of individuals are (correctly) counted as household saving in the aggregate data—and in the U.S. they are a large fraction of household saving—but typically are not collected in the household surveys. At the other end of life, the payouts from these same pensions are essentially annuities, a fraction of which is derived from the decumulation of assets. As a result, some dissaving among the elderly is counted as income, biasing upward their observed saving rates.

Third, both sources of data are subject to large margins of uncertainty. In both aggregate and micro data, saving is a residual between two large magnitudes, income and consumption, so that even if the large numbers are relatively precise, their difference may be out by orders of magnitude. Savings data are among the worst measured in both national accounts and survey data. Even in the industrialized countries, there is often little relationship between household saving in the national accounts and household saving as estimated from survey data. This may even be an area where the high saving countries of Asia are at an advantage, because their saving rates are so high, and therefore are easier to trace in the survey data.

Fourth, there are a number of more subtle failures in the survey approach. Weil (1994) has argued that the contradiction between micro and macro data might be resolved through the effects of anticipated bequests: the dissaving associated with the elderly is real enough, but it is being done not by the elderly themselves, but by their younger heirs in anticipation of inheritances. There is also a long-standing concern about selection effects among older households. The survey data on income, consumption, and saving come from *households*, not individuals, and the process of household formation and household headship is itself likely to be conditioned by savings and wealth. In particular, it is possible that high saving households, or perhaps those that have been fortunate in their investments, are likely to survive longest as independent households. There is a well-established inverse relationship between mortality and wealth, but more important is the way in which elderly people can give up being householders, and move in with their children or relatives. It is quite possible that the lack of dissaving with age is an artefact of the increasingly severe selection with age in favor of high savers.

In Deaton and Paxson (1998a), we address this question in the context of Taiwan by estimating a life-cycle model in which it is individuals who satisfy the life-cycle hypothesis, and not households. Under this approach, we find a much more pronounced life-cycle profile of saving, and we find much clearer evidence of dissaving associated with children and with the elderly. Our calculations show that, for Taiwan, the survey data are consistent with a model in which demography is capable of exerting important effects on saving though, as first suggested by Fry and Mason (1982), the effects interact with the rate of growth of income. What happens when the child dependency ratio is increased depends on the ratio of lifetime incomes of the young relative to those of the old, which in turn is determined by the rate of economic growth. At high rates of growth, children are relatively well-funded, so that increases in their relative numbers should have a much larger negative effect on the saving rate than when economic growth is slow, and they are relatively poor. Similarly, increases in the elderly dependency ratio have larger negative effects on saving when economic growth is slow or negative. Even so, when we use the model to calibrate the contribution of Taiwan's demography to Taiwan's national saving rate in Deaton and Paxson (1998b), we find that, at the actual historical growth rates, demographic structure contributed very little. So these new results, although consistent with an important potential role for demographics, do not resolve the difference between the household survey results, which show little effect on saving, and the aggregate econometric evidence, which seems to show large effects.

2.3 Saving rates and dependency rates: an econometric analysis

In this section we examine the link between demography and saving, using annual time-series data for fourteen Asian countries: Bangladesh, China, Hong Kong, Indonesia, India, Korea, Sri Lanka, Malaysia, Nepal, Pakistan, Philippines, Singapore, Thailand and Taiwan. We use these data to estimate, in a variety of ways, the effects of the youth and old-age dependency rate on the saving rate, controlling for income growth. Our first specification follows that of Modigliani (1990), who works with decadal averages rather than individual years. For each country, we have up to four observations for the decades of the 60s, 70s, 80s, and 90s. For most of the countries (see the footnote to Table 2.1), the decade of the 60s contains only 5 years and, for all countries, the decade of the 90s contains only five years. Although the decadal approach discards some potentially relevant information, it has some compensating advantages. The hypothesis linking demographics to saving is not something that we expect to explain year to year fluctuations, and the demographic ratios change only slowly, so the approach asks whether the broad patterns of trends in the saving ratios in Figure 2.1 can be matched to the decade by decade changes in the age structure of the population. The same argument holds for the other variable included in the regressions, the rate of growth of real national income. According to the life-cycle model, income growth increases saving by redistributing resources from old dissavers to younger dissavers and, once again, this is a process that does not work on a year by year basis. Modigliani also avoids using the growth of income in the current decade, replacing it with the growth in the previous decade in an attempt to rule out reverse causality, from saving rates to growth.

Table 2.1 presents the results of a number of specifications in which the decadal data are pooled over the 14 countries. In column 1, the pooled OLS regression of saving rates on growth and the dependency rates gives results that are close to those reviewed in the previous subsection. An increase in the youth dependency rate by one percentage point decreases the saving rate by just over one percentage point, and the parameter estimate is well-determined, with an absolute *t*-value of 4.3. Again in line with the literature, the elderly dependency ratio has a larger coefficient—a one percent increase in the ratio decreases the saving ratio by 1.8 points—but the *t*-value is only 1.5. Results like these will appear in many of our specifications. In line with life-cycle theory, income growth has a marked effect on the saving rate, but note that this is the concurrent decadal growth rate so that there is the usual concern with simultaneity. There is also the further possibility (better, certainty) of spurious correlation through measurement error. Over- or underestimation of income will simultaneously affect both income growth and the saving ratio, inducing a spurious positive correlation between them. Taken at face value, the coefficient is consistent with other literature; for example, Modigliani (1990) obtains a coefficient of 1.30 on the lagged growth rate for a "sample" of 85 less-developed countries.

The second column adds interaction terms between growth rates and the dependency rates; neither attract significant coefficients. Subject to nothing else being wrong, this finding is a strike against the life-cycle model. If the model is correct, so that consumption is determined by life-time resources, with periods of dissaving associated with children and with the elderly, the size of the dependency effects *must* vary with the rate of growth of real income.

The third and fourth columns repeat the first and second with the addition of country specific fixed effects, so that the intercept is allowed to be different for each country. These country dummies are jointly significant (as we shall see again below), but admitting them has only minor effects on the estimated parameters. The growth effect is somewhat smaller, and the (negative) effect of the elderly dependency rate is somewhat larger, sufficiently so to move the *t*-value into the significant range. As was the case with OLS, the interaction terms add nothing to the fit, either singly or jointly. Random effects models were also estimated but are not shown here; not surprisingly, the results are close to both OLS and fixed-effect estimates.

In our view, it is close to impossible to devise a satisfactory econometric treatment for the potential endogeneity of the growth rate, either because saving causes growth, or because both contain a common measurement error. We repeated Modigliani's procedure of dropping the current growth rate in favor of its lag (from the previous decade). In none of the regressions does the estimated coefficient approach statistical significance and there is little effect on the coefficients of the dependency rate. The last column reports a (failed) attempt at another strategy. We estimate a fixed-effects-IV regression, in which the lagged values of growth, the ratio of gross domestic investment to GDP, and the ratio of exports to GDP are used as instruments for the current growth rate. But none of these, singly or jointly, has predictive power for the growth rate, so neither the IV estimates nor their t-values in the last column can be taken seriously. More fundamentally, we see no grounds on which these or other choices of instruments could be defended. Our theory posits that the growth rate and the dependency rates should be in the regression, but rules almost nothing out. As a result, any potential instrument could equally well be argued to belong in the substantive regression. With panel data as here, there is the further difficulty that, in the presence of fixed effects, exogeneity cannot rely on temporal priority. In consequence, and as with the demographics themselves, we have little choice but to condition on the current growth rate. There is some comfort in the fact that, conditional on country fixed effects, growth is not correlated with the demographics—the partial correlation coefficient between growth and the youth (elderly) dependency ratio is 0.07(–0.08)—whose coefficients are thereby not much affected by how it is treated.

Table 2.2 moves from the decadal averages to the annual time-series data themselves. Growth rates are now entered in levels as well as in their first and second (annual) lags; otherwise the specification is the same as in the decadal averages in Table 2.1. Because the demographic terms are so highly autocorrelated, we do not attempt to sort out their contemporaneous from lagged effects. These results are close to those obtained from the Modigliani-type procedure in Table 2.1. The sum of the growth rate effects are somewhat smaller than before, but the estimated effects of the two dependency rates are very much the same. The youth dependency rate is about half the size of the elderly dependency rate, but has a much smaller standard error. The random effects and fixed effects estimates shown in the second and third columns do not suggest any bias in the OLS results in the first. In all three columns, we can reject the hypothesis that the youth and elderly dependency rates have the same effect on the rate of saving.

The right-hand part of the table shows the corresponding estimates when we move from national to private saving rates; see the notes to the table for exactly how these were computed. At the same time, the growth rate is of real private income, defined as gross national income less public consumption and public saving. Making this switch forces us to drop five countries, Bangladesh, China, Hong Kong, Sri Lanka, and Nepal, for which we either have no data, or too little data. The changes in the results are mostly confined to the growth rates, which show little effect on saving except in the fixed-effect regressions. But the coefficients on the demographic terms are much the same as before, so that, at least as far as this limited evidence goes, it does not appear as if the effects of demographic structure are working through the reactions of the public sector.

So far, the results are quite supportive of the hypothesis that demographic structure is a major determinant of national (and private) saving rates, that the demographic structure of Asia has been a major contributor to its saving rate, and that demographic events in the next century are likely to bring saving rates down. Nevertheless, a closer analysis brings out a rather different message. In particular, all of the analysis so far has assumed that the relationship between savings and demography is the same in each country, at least up to an intercept term. But there are no good reasons to suppose that this will be the case. Institutions differ from one country to another, as do many other variables that we have made no attempt to include, and even according to a literal version of the theory, there should be nonlinearities and interaction terms that are likely to show up in cross-country differences. In fact, the homogeneity assumption can readily be tested;

the F-test for the fixed effect model versus OLS is 32.71, while that for the fixed effect model as a special case of the fully restricted model in which all coefficients are allowed to vary from country to country, is 14.98. Both of these statistics give overwhelming rejections of the homogeneity hypothesis. The effects of growth and demographics on saving rates are different from one country to another.

The heterogeneity of the relationship can readily be dealt with provided we abandon the pooled approach. While it is true that country by country regressions are likely to be quite imprecise, in at least some cases, there is no need to accept this imprecision, nor to discard the role of the cross-country information in reducing it. In such a case, the appropriate procedure, both from a statistical and substantive point-of-view, is to estimate the relationships country by country, and then to average the results. The averaging, if done correctly, results in an economically meaningful estimate, and restores much of the statistical precision that is lost in the country by country approach. To see how this works, consider the pooled OLS estimator, which can be written in the form

$$\hat{\beta}_{(1)} = [\sum_c X_c' X_c]^{-1} [\sum_c X_c' y_c] \tag{2.1}$$

where c refers to a country, X_c is the matrix of independent variables for country c, and y_c is the time-series vector containing its saving rates. The country specific coefficients, which the F-tests tell us we need, are given by

$$\hat{\beta}_c = (X_c' X_c)^{-1} X_c' y_c. \tag{2.2}$$

The two measures are linked by the formula

$$\hat{\beta}_{(1)} = [\sum_c X_c' X_c]^{-1} [\sum_c X_c' X_c \hat{\beta}_c] \tag{2.3}$$

so that the pooled estimate is a *matrix-weighted average* of the individual estimates. In the presence of country specific intercepts (fixed effects), we are effectively weighting by *variances*; the country whose growth rates and demographic structure varies most around its mean gets the most weight in the pooled estimate. Such an estimator has little to recommend it.

Consider instead the estimator that weights the individual estimators, written as

$$\hat{\beta}_{(2)} = \sum_c w_c \hat{\beta}_c \tag{2.4}$$

where the weights w_c are proportional to the gross national income of country c measured in U.S. dollars. (We use the World Bank Atlas conversion factors to make the conversion.) What (2.4) does is to estimate the effect of the demographic ratios on the share of saving in gross domestic income for the 14 countries as a whole. The coefficient for (say) the youth dependency rate in Indonesia estimates how much Indonesia's saving rate will change for a one point change in the dependency rate. Multiplying by gross national income converts the effect to rupiah, and by the Atlas conversion factor to US dollars, so that adding up over countries gives the effect of the dependency rate on regional saving. Dividing by total gross national income in US dollars brings us back to the effect on the total saving ratio. Note also that the average coefficient (2.4) is likely to be more precisely estimated than any of its individual country components, at least if the country by country residuals are not perfectly correlated.

Table 2.3 takes up this empirical story and lists the regression results of saving rates on growth and the dependency rates for each country, with the weighted average at the bottom. We can see immediately why the F-tests rejected the hypothesis of homogeneity. Ten of the youth coefficients remain negative, eight of them significantly so, but only six of the elderly coefficients are negative, and only three are significantly different from zero. Particularly notable are the very large positive coefficients for Indonesia, and to a lesser extent India; all of these coefficients have large standard errors, but all are significantly different from zero even without pooling. The averages in the last row, which are the numbers in which we are most interested, show a small, but significantly positive effect of income, and *positive* average effects of both dependency rates on the average saving rate.

These results illustrate very well the effects of treating heterogeneity seriously, but they are not without their own problems. In particular, the dynamics of saving are clearly inadequately captured, as shown by the low Durbin Watson statistics in the final column. A simple, if atheoretical, remedy is to exploit the putative cointegration of consumption and income, and to estimate an error-correction model in which the saving rate adjusts to a desired value that depends on the demographic ratios. In order to make this as straightforward as possible, it is convenient to make the dependent variable (minus 100 times) the logarithm of the ratio of consumption to income (which is approximately the saving rate in percent), and to regress it on the growth rate, the dependency rates, and its own lag. The coefficients on the dependency terms are now interpreted as the short-run effects, and must be divided by one minus the coefficient on the lagged dependent variable to give the (larger) long-run effect once the saving rate has adjusted to the new demographic structure. From a statistical perspective, the exercise also has the advantage of eliminating the overestimation of t-values in Table 2.3 that comes from the (untreated) serial correlation.

As in Table 2.3, the results in Table 2.4 show a great deal of country heterogeneity and much of the difference in cross-country effects is preserved from the static model. But the standard errors are much larger than before, so that only three of the nine negative youth dependency effects are significantly different from zero, and only one of the seven negative elderly dependency effects. When we take the weighted averages, the overall effects are still *positive*, but they are only marginally significantly different from zero. The growth effect is positive and significant, and if not as large as required by standard aggregate life-cycle models, is well within the range that is consistent with estimated habit formation or life-cycle models estimated from the micro data, see e.g. Paxson (1996).

One of the countries for which the standard result continues to hold is Taiwan, where the dependency rates have large and significant negative effects on the saving rate. Although these results do not carry through to the average, we have done a great deal of work on the micro data in Taiwan, and have not been able to find evidence that would support this finding on the aggregate data. One possibility is the difference between household and national saving, something that can be investigated for Taiwan. When we rerun the regression using household saving and growth rates of household real disposable income, the coefficients on growth and the lagged saving ratio remain virtually unchanged, but the coefficients on the demographic ratios are reduced to –0.51 ($t = 2.5$) and –0.84 ($t = 0.8$), figures that are much more closely compatible with our microeconomic results. Perhaps some of the difficulty here is in the differences between household, private, and national savings, and at least some of the demographic effects may have nothing to do with household behavior or the life-cycle hypothesis.

To return to the overall findings, the results in Table 2.4, which show little or perverse effects of dependency rates on aggregate saving rates, are in our view the appropriate ones to use, absent further research and evidence. The pooled regressions, even with fixed effects, are rejected against the data, and the statistically and economically appropriate estimates do not find that dependency rates reduce aggregate savings for these countries as a whole. While there are some countries where the traditional effects appear to exist, there are others that show the opposite, and there is no net effect. Nor can the pooled results be taken seriously, even after allowing for fixed effects in the estimation. To at least this extent, the aggregate data are consistent with the micro data. Neither gives much credit to demography for Asia's high savings rate, and by the same token, neither suggests that future demographic change will reduce those saving rates.

2. The Welfare of the Asian Baby Boom: Does Cohort Size Affect Welfare?

2.1 Background

A feature common to many of the Asian countries discussed above is the existence of a "baby boom" in the 1950's, followed by a subsequent fertility decline. As discussed in Section 2.1, the aging of these large cohorts of baby boomers over the next two to three decades will result in sharp increases in the share of elderly in the population. Although we have reliable predictions about *how many* old people there will be in Asia 10, 25, or 50 years from now, we have less information about what their economic status will be. In particular, we do not know whether the large cohort of baby boomers will be at any particular disadvantage as they age. There are several reasons to think that the living standards of the elderly relative to those in younger age groups may decline in the next few decades as the baby boom ages:

First, the average earnings of the baby boom generation, that is working now but will age and retire over the next 30 years, may be depressed by its size, which in turn will make it more difficult for baby boomers to accumulate assets for their old age. In addition, to the extent that pensions (both public and private) exist and pay benefits that are positively related to previous earnings, lower earnings for baby boomers will result in smaller pension receipts upon retirement.

Second, many Asian countries have experienced rapid technological change and large increases in educational attainment over the past several decades, and the gains of these changes may accrue largely to younger cohorts. In the presence of rapid growth, it is likely that younger cohorts will have much higher life-time wealth than older cohorts. Of course, rapid growth is unlikely to be harmful to any cohort, but it may decrease the living standards of the elderly *relative* to the non-elderly, and also has the potential to increase inequality. The effects of rapid growth on the relative well-being of different cohorts is independent of the size of the different cohorts. However, these growth effects on relative living standards may accentuate the effects of cohort size on the earnings and income of baby boomers, discussed above.

Third, evidence from a variety of countries indicates that inequality in living standards increases with age within cohorts (see Deaton and Paxson, 1994.) Given average living standards, greater dispersion will increase the fraction of people who are below a poverty threshold. This effect will be even larger if elderly baby boomers have *lower* average living standards due to their large size and lower average skill level.

Fourth, because of fertility declines that followed the baby boom, the Asian baby boomers have relatively few children to rely on for support in their old age. There is evidence that, in at least some countries, traditional family support networks are weakening. This may not pose a threat to the welfare of older Asians, especially if independent living is the result of higher wealth among the elderly. However, with fewer children and weaker family support networks, it may be that the elderly will become more vulnerable to adverse shocks unless public support systems replace family support systems.

Not all of these issues can be easily investigated. It is extremely difficult to predict how much support elderly baby boomers will receive from family members thirty years from now, or the extent to which public old-age social security programs will have developed to replace traditional family support systems. Furthermore, poverty among the Asian elderly (and all groups) in the future hinges on whether Asian rates of economic growth return to the levels of the past two decades, or whether the current crisis marks the beginning of a long-run decline in growth rates. However, several of the issues discussed above are amenable to empirical investigation. Since the members of the baby boom are currently of working age, data on the earnings, income, and consumption of current workers can be used to investigate whether the currently-working baby boom generation is earning relatively less than other generations, and this information can be used to draw inferences about the effects of the baby boom on the living standards of the future elderly.

The question of whether the size of baby boom cohorts depresses the life-time earnings of their members has been investigated in the US, by Welch (1979), Freeman (1979), Murphy, Plant, and Welch (1988). The following simple model is based on Freeman (1979) and captures cohort-size effects on wages in the simplest possible way. Suppose that output at time t is a CES function of capital and an aggregate of labor, so that

$$Q_t = [(1-\alpha)L_t^{\theta} + \alpha K_t^{\theta}]^{1/\theta} \tag{3.1}$$

where $1/(1-\theta)$ is the elasticity of substitution between capital and labor. Total labor is defined as a homogeneous of degree one aggregate of the labor at each age, so that

$$L_t = (\sum_a \beta_a L_{at}^{\rho})^{1/\rho} \tag{3.2}$$

and $1/(1-\rho)$ is the elasticity of substitution between the different types (ages) of labor. The β's represent the efficiency of labor at different ages and play an important part in shaping the age profile of wages. If we suppose that labor is paid its marginal product, and the price of output is normalized to unity, the real wage for someone aged a at time t is easily shown to satisfy

$$\ln w_{at} = \ln\beta_a + (1-\theta)\ln(Q_t/L_t) - (1-\rho)\ln(L_{at}/L_t) \tag{3.3}$$

which shows that the log wage at any given age depends positively on aggregate output per head, and will grow with it at a rate that depends on the elasticity of substitution between labor and capital, and negatively on the size of its cohort relative to the total labor force. Especially if labor at different ages is not very substitutable, a baby boom cohort can expect to have relatively low wages throughout life.

Models of this type have been estimated for the US by Freeman (1979) and by Murphy, Plant, and Welch (1988), and these authors (and others) have found the effects predicted by the theory, with larger cohorts having wages tipped against them. The Murphy, Plant, and Welch model is more sophisticated than the simple version sketched above; wages are modeled as the return to a small number of unobservable factors, factors that different age groups are endowed with in different proportions. Such a formulation allows workers that are closer in age to be better substitutes than workers who are much older or younger than one another, and indeed Murphy, Plant and Welch find that most of the effect of being in a large cohort comes in the first years in the labor market, and is dissipated relatively quickly with age.

The results discussed above are based on US data. Although there is no general reason to think the theory should *not* apply to the Asian countries in which we are interested, there are several reasons why differences might exist. First, the growth experience of the US over the past several decades has been extremely different from that of Asian countries. The US baby boomers entered the labor market in the mid-1970's, at the beginning of a productivity slowdown that followed the first oil shock. Although Asian economies also suffered from the 1974 oil shock, for many of these countries growth rates rebounded quickly. For example, during the 1970's, growth in real per capita income exceeded 7% in Taiwan, compared to about 2% for the US. For the US, it may be difficult to disentangle the effect of the productivity slowdown on the baby boom's wages from the effect of the size of the baby boom cohort. Second, there are large differences in the sectoral composition of output between the US and Asian countries (as well as large differences between different groups of Asian countries.) It is possible that the substitutability of labor of different ages differs greatly across sectors, so that different countries effectively have different values for ρ.

In what follows we present estimates of the model discussed above using data from Taiwan. The Taiwanese case is of interest because the baby boom was quite large, and the subsequent fertility decline was very rapid. The fraction of the population that was aged 5 or less is graphed in Figure 3.1 for the years 1947 through 1995. The percentage of young children in the population peaked at 23% in 1956, and then began a sharp decline. By 1970, only 15.6% of the population was age 5 or younger. If large cohorts actually depress their own wages, these effects should be discernable in Taiwan.

Another advantage of studying Taiwan is the availability of several excellent data sources. We use the 1976-1995 Surveys of Personal Income Distribution and the 1979-1994 May Manpower Surveys to construct measures of the wages, earnings, income, and consumption of different cohorts at different ages, and we examine whether these different measures of well-being are affected by cohort size as indicated by (3.3). These are large surveys that collect information on over 50,000 individuals per year. Because both of these surveys are repeated cross-sections, we can track the wages, earnings, income and consumption of different cohorts through time as they age, by constructing the averages of the wage and income measures (or, their logarithms) for each cohort (defined by year of birth) in each survey year. In addition, Taiwan has excellent demographic data. Annual information on population by age is published in the *Taiwan-Fukien Demographic Factbook*, and is also available in electronic format from the Directorate General of Budget, Accounting and Statistics. Taiwan has a household registration system that keeps track of the numbers of Taiwanese residents of different ages in each year, and these numbers provide information on cohort size that may be more reliable than the census-based estimates available for many countries.

2.2 Methods and Data

Our starting point is equation (3.3), which with some algebraic manipulation and the addition of an error term can be rewritten as:

$$\ln w_{abt} = \ln \beta_a + [(1-\theta)\ln(Q_t) + (\theta-\rho)\ln(L_t)] - (1-\rho)\ln(L_{abt}) + \epsilon_{abt} \tag{3.4}$$

where b represents year of birth. We make two modifications to (3.4). First, because estimating the parameter θ is not a concern, we substitute a set of year dummies for the middle term in 3.4, i.e. $[(1-\theta)\ln(Q_t) + (\theta-\rho)\ln(L_t)]$. Second, we decompose cohort size $\ln(L_{abt})$ into a fixed component that measures the size of the cohort, either at birth or some other fixed age, plus age and time components that pick up changes in cohort size over the life-time of the cohort due to death and migration:

$$\ln(L_{abt}) = \ln(L_b) + \ln(\beta'_a) + \ln(\gamma_t) + v_t . \tag{3.5}$$

This can be substituted back into (3.4), to obtain:

$$\ln(w_{abt}) = B_a + \Gamma_t - (1-\rho)\ln(L_b) + e_{abt} \tag{3.6}$$

where the term B_a equals the sum of $\ln(\beta_a)$ from (3.4) and $\ln(\beta'_a)$ from (3.5). Likewise, the year effects and the error term in (3.6) are the sums of the year effects and error terms from the two preceding equations. Equation (3.6) indicates that the log wage is composed three additive terms (excluding the error term): the first term, B_a, is the same for all cohorts in all time periods, and determines the shape of the age-wage profile; the second term is the same for all cohorts in the same year, and measures the component of wages determined by levels of capital and labor in the economy (as well as any short-term macro shocks); the third term reflects the effects of fixed cohort size on the wage. Our main interest is in estimating the coefficient on $\ln(L_b)$; this can be done by regressing the log wage on a set of age dummies, year dummies, and a measure of cohort size.

We start by estimating wage (and earnings, income, and consumption) equations that are more general than (3.6) and do not explicitly control for cohort size, but simply decompose the wage into age, year, and cohort effects. This is done completely non-parametrically, by regressing the average of the log wage of each cohort in each year on a set of dummy variables that indicate the age, birth year, and calendar time for the observation:

$$\ln w_{abt} = intercept + D^{a}\alpha_{a} + D^{t}\Theta_{t} + D^{b}\delta_{b} + e_{abt} \tag{3.7}$$

where D^{a}, D^{t}, and D^{b} are matrices of age, birth year and calendar year dummies, and the terms α_{a}, Θ_{t}, and δ_{b} are the age, calendar year, and birth year effects to be estimated. The variables in (3.7) are co-linear and several restrictions are necessary for estimation. As usual, one age dummy, one birth year dummy, and one calendar year dummy must be omitted. One addition restriction is also necessary, since calendar year is equal to birth year plus age. Many restrictions are possible, and which one we choose does not alter the overall fit of the model. We impose the restriction that the year effects Θ_{t} sum to zero and are orthogonal to a time trend (see Deaton and Paxson 1994 for a fuller explanation.) This normalization implies that the calendar year effects reflect only short-term macro shocks, whereas the growth in real wages across cohorts is measured by the cohort effects δ_{b}. This normalization is sensible if one wants to measure how much higher on average a later-born cohort's wages are than an earlier-born cohort's, controlling for age.

After estimating (3.7), our next step is to impose the restrictions that are implicit in (3.6), and re-estimate. The simplest method is to regress the log wage on a set of age dummies, unrestricted year dummies, and the log of cohort size. However, it is useful to know the average growth in wages, controlling for age and population size, so the model we actually estimate includes the restricted set of year dummies and a time trend:

$$\ln(w_{abt}) = B_a + \Gamma_t + \eta t + \alpha \ln(L_b) + e_{abt} \tag{3.8}$$

where the year effects Γ_t sum to zero and are orthogonal to a time trend, and η measures the growth rate. The parameter α measures the effect that cohort size has on the logarithm of the wage, which according to (3.6) equals $-(1-\rho)$.

Equations (3.7) and (3.8) are estimated using a variety of dependent variables. First, we use the Manpower Surveys to construct a measure of the hourly wage of each person who works for money. This survey asks workers about their monthly earnings from their main job, and also collects information on the hours per week worked on the main job. The wage is computed as monthly earnings divided by 4.2 times weekly hours, and is denoted as w_i, where the "i" subscript stands for "individual" (and in what follows "h" subscripts will stand for "household.") The logarithm of the wage is then averaged over all workers with positive wages within cohort-age cells. An alternative dependent variable is obtained by repeating this process using monthly earnings (denoted e_i) rather than the wage.

The theory discussed above implies that cohort size should affect the wage that individuals in the cohort command on the labor market. However, because individuals often live in households with people from other cohorts, it could be that cohort size has only small effects on household income and consumption; these variables may be more reliable indicators of welfare than individual earnings or wages. To examine the effect of cohort size on *household* welfare, we use the Personal Income Distribution Surveys to obtain measures of household total income (y_h) and consumption expenditure (c_h). Total income is after tax, and includes labor income, self-employment income, property income, and transfer receipts, aggregated over all members of the household. Expenditure is total expenditure on all goods and services, both durable and nondurable. As before, logarithms of these variables are averaged within cohort-age cells, but now cohort membership is determined by the birth year of the household head rather than the birth year of the individual. These dependent variables can be used to examine whether households headed by individuals from larger cohorts fare worse relative to other households.

Finally, we use the Personal Income Distribution Survey data to estimate *individual* income from *household* data, using methods described in Deaton and Paxson (1998a). Ideally, we would like to have a measure of individual income that includes income from all sources, not just money earnings on the main job. This information is not available from the Manpower Surveys, not only because the survey simply does not collect information on non-labor income, but also because many Taiwanese workers work "for free" in family enterprises; since they report no earnings they are not included in the calculations of w_i and e_i described above. Although information on self-employment income, property income, and transfer income is available from the Personal Income Distribution Surveys, this survey also has the problem that it is not clear to whom different types and amounts of income should be attributed. For example, income from a family-run business is usually all attributed to the household head, even if other family members work for the business. Our method of estimating individual income (denoted y_i) is to regress, for each year of the survey, household income on a set of variables that measure the numbers of people in the household of each age. Specifically, we estimate:

$$y_{ht} = \sum_{a=1}^{A} \lambda_{at} n_{ht}^{a} + v_{ht} \qquad (3.9)$$

where y_{ht} is the income of household h in year t, and n_{ht}^{a} is the number of people aged a in household h. The term λ_{at} measures the expected value of the contribution of a person aged a in year t to household income, and can therefore serve as a measure of average income of people aged a in year t (and born in $b=t-a$.) A complication is that the estimates of λ_{at} are imprecisely estimated for very old ages (where there are few observations) and are very small and sometimes negative for both very old and very young ages. We counter these problems by imposing the restriction that λ_{at} is zero for ages of less than or equal to 16, and older than 79; we also smooth the estimates of λ_{at} using a kernel smoothing technique; again, see Deaton and Paxson (1998a) for more detail. Our final measure of the cohort average of the logarithm of individual income, denoted $\ln(y_i)$, is calculated as the logarithm of the smoothed values of λ_{at}. This procedure is repeated to obtain measures of individual *consumption*, denoted c_i, from the household data. For consumption, none of the parameters in (3.9) are restricted to be zero, so that people are assumed to consume at all ages. As before, there is no guarantee that the estimates of λ_{at} for consumption are positive, and in fact they are sometimes negative for very young ages (where consumption is low) and very old ages (where consumption is estimated imprecisely.) However, after smoothing, only a handful of the estimates remain negative; these are dropped from our analysis.

Equation (3.7) and (3.8) are estimated using all six dependent variables: logarithms of the wage (w_i), earnings (e_i), household income (y_h), household consumption (c_h), individual income

(y_i) and individual consumption (c_i.) Estimation of (3.8) requires information on cohort size, L_b. Population counts by age are available for Taiwan from 1947 to the present, and so using cohort size in the year of birth to measure L_b would eliminate all cohorts born earlier than 1947; this is not desirable, since individuals from these older cohorts make up a large fraction of the labor force in the late 1970's and 1980's. Instead, we measure cohort size L_b as the number of people born in year b in the year in which they were age 25 (so, L_{1960} is measured as the number of people aged 25 in 1985.) This allows us to use information on cohorts born as early as 1922. The age range is restricted to those aged 25 to 70, inclusive, to eliminate ages when large numbers of people are retired, in school, or in the military. When we use the Manpower Survey, which runs from 1979 to 1994, we have 645 cohort-age cells. The cohorts represent people born from 1922 (aged 25 in 1947, and 70 in 1992, after which they were excluded from the sample) to 1969 (aged 25 in the 1994, the last survey year available.) The Personal Income Distribution Survey runs from 1976 to 1995, and from this data set we constructed 784 cohort-age cells on people born from 1922 to 1970.

3.3 Results

The first question we address is whether older people have lower living standards than younger people. The estimates of equation (3.7) provide information that is relevant to this question. The estimates of cohort effects tell us how much poorer earlier-born cohorts are than later-born cohorts, at any age. The estimates of age effects tell us whether, controlling for year of birth, individuals experience declines in income or consumption at old ages. We are most interested in the age profile of consumption; consumption is a better indicator of living standards than income, especially at old ages when consumption may be financed by asset declines rather than income.

Before looking at the regression results, it is useful to examine the "raw" cohort-age averages of the data. Figure 3.4 graphs the averages of household and "individual" log consumption against age; consumption is chosen since it is the best summary measure of welfare, although graphs for the income and earnings variables look similar. Each line connects the points for a single *year* of data, with the lower lines representing earlier years. The graph on the left, for household consumption by age of the household head, indicates that, within any year, the elderly have much lower consumption than the young. Consumption rises slowly with age until about age 55, after which shows a large decline. However, there are several reasons why this graph may be misleading. First, it is for *household* consumption by age of the *household* head. The sample of household heads is increasingly selected at older ages, as older individuals either move in with children or give up headship to their co-resident children. (The Taiwanese surveys designate the "primary earner" to be the head of household.) Although in theory this selection problem could cause the living standards of the elderly to be overstated or understated, it is probably the former if wealthier people are more likely to retain headship at old ages. Second, there are differences in household size with age which could account for the age patterns in household consumption. Households headed by older people are on average smaller, and the decline in consumption with age might be less pronounced if consumption is put on an individual basis.

Individual consumption, graphed on the right side of Figure 3.4, is not subject to these selection and household size problems. The graph indicates that, within years, individual consumption is highest for those in middle age, and lowest for the very young and very old. If consumption is a measure of welfare, than in any year older people are on average worse off than those in middle age. (The lower consumption level of children may not reflect lower welfare, if children have lower needs.) However, because the graph traces out age-consumption profiles for each year, it does *not* necessarily indicate that individuals experience consumption declines as they grow older. The relatively low consumption of the elderly could reflect lower life-time wealth of people who are in earlier birth cohorts, which depresses consumption at all ages; alternatively, it could be that consumption does genuinely fall with age.

The estimates of (3.7) allow us to sort out these cohort (wealth) and age effects. Figure 3.3 graphs the cohort effects from estimates of equation (3.7) for each of the dependent variables. The graph shows the large and relatively constant growth in wages, earnings, income, and consumption that has occurred across birth cohorts. Later-born cohorts are much richer than earlier-born cohorts: controlling for age, those born in 1922 are more than 250% poorer than those born in 1970. In a life-time sense, the currently old in Taiwan are poor relative to the currently young. Figure 3.4 shows the age effects from (3.7). These indicate that most of the variables grow with age, at least up until 60 or 65, at which point they flatten out and decline slightly for some measures. However, we do not see a large decline in living standards with age, holding cohort fixed, which is reassuring. The (relatively) low consumption of the elderly seen in Figure 3.2 is largely the product of the high rate of economic growth in Taiwan, which has increased the wealth and consumption of younger relative to older cohorts. Given this result, the lower relative living standards of the elderly are not a cause for concern; lower growth that might improve the *relative* living standards of the elderly would make everyone worse off, and is surely not desirable.

The second question we address is whether cohort size affects wages, earnings, income, and consumption levels. Taiwan is a good laboratory in which to examine this issue, since it has experienced very dramatic shifts in cohort size through the sample period. Figure 3.5 provides information on cohort sizes by graphing $\ln(L_b)$, the log of cohort size at age 25, against year of birth. The figure clearly shows the "baby bust" during the war years, the following baby boom, and then the rapid fertility decline in the late 1950's and 1960's.

Estimates of (3.8) are in Table 3.1. The first column shows the estimate of η, the trend in log wages (holding age and cohort size fixed.) Consistent with the results discussed above, there is a large time trend in all of the measures, on the order of 5.7% to 6.8% per year. The second column shows the coefficient on cohort size, which is expected to be negative. Contrary to the theory, there is no evidence that cohort size affects wages, earnings, income or consumption, at either the household or individual level. For none of the variables is the estimate of the effect of cohort size statistically significant. The model is most directly applicable to the first two dependent variables, the log wage and log earnings, and for these two variables the estimated effect of cohort size is *positive*. This can be seen even more clearly in Figure 3.6, which shows detrended values of the cohort effects from estimates of (3.7) and detrended cohort size against birth year, for each of the wage, income and consumption measures. In all cases, we see no clear association between cohort size and the variable in question, and the detrended wages, earnings, and income of the baby boom cohort are extremely similar to those of their neighboring cohorts.

Overall, our empirical results do *not* indicate that special concern about the welfare of the Taiwanese baby boomers is warranted. The earnings and income of the baby boomers have not been depressed by their size. More generally, there is no reason to think that living standards decline in old age in Taiwan. Within birth cohorts, our analysis indicates that consumption and other welfare measures *increase* with age up until the age of retirement age and then flatten out, but never exhibit large declines. There is no reason to think this pattern will not hold for the baby boomers when they retire, although some attention should be paid to shifts in sources of old-age support for the elderly as the baby boom ages. For example, we might expect to see declines in family-based support (due to fertility declines), which may or may not be offset by increases in pensions and income from assets.

List of works cited:

Asian Development Bank, 1997, *Emerging Asia: changes and challenges*, Manila. Asian Development Bank.

Coale, Ansley J. and Edgar J. Hoover, 1958, *Population growth and economic development in low-income countries*, Princeton. Princeton University Press.

Deaton, Angus, and Christina H. Paxson, 1994, 'Intertemporal choice and inequality,' *Journal of Political Economy*, 102, 000–00.

Deaton, Angus, and Christina H. Paxson, 1997, "The effects of economic and population growth on national saving and inequality," *Demography*, **34** (1), 97–114.

Deaton, Angus and Christina H. Paxson, 1998a, "Saving and growth: another look at the cohort evidence," Princeton University, Research Program in Development studies, processed.

Deaton, Angus and Christina H. Paxson, 1998b, "Growth, demographic structure and national saving in Taiwan," Princeton University, Research Program in Development studies, processed.

Freeman, Richard B., 1979, "The effect of demographic factors on age-earnings profiles," *Journal of Human Resources*; 14(3), Summer, pps 289-318.

Fry, Maxwell J., and Andrew Mason, 1982, "The variable rate-of-growth effect in the life-cycle saving model," *Economic Inquiry*, **20**, 426–441.

Heller, Peter S., and Steve Symansky, 1997, "Implications for savings of aging in the Asian 'Tigers'", Washington, D.C. International Monetary Fund, Fiscal Affairs Dept., Working Paper 97/136, processed.

Higgins, Matthew and Jeffrey G. Williamson, 1997, "Age structure dynamics in Asia and dependence on foreign capital, " *Population and Development Review*, **23**, 261–93.

Leff, Nathaniel,1969, "Dependency rates and saving rates," *American Economic Review*, **59**, 886–96.

Loayza, Norman, Humberto Lopez, Klaus Schmidt-Hebbel, and Luis Serven, 1998, "The World Saving Data Base," Washington,D.C. The World Bank, processed.

Meredith, Guy, 1995, "Demographic change and household saving in Japan," in U. Baumgartner and G. Meredith, eds., *Saving behavior and the asset price "bubble" in Japan*, IMF Occasional Paper No. 124.

Modigliani, Franco, 1993, "Recent declines in the savings rate, a life-cycle perspective,"in Mario Baldassarri et al., eds., *World saving, prosperity, and growth*, New York, St. Martins Press, 249–86.

Murphy, Kevin M., Mark Plant and Finis R. Welch, 1988, "Cohort size and earnings in the USA," in Lee, Ronald D., Brian W.Arthur and Gerry Rodgers, eds. *Economics of changing age distributions in developed countries*. International Studies in Demography series. Oxford: Oxford University Press, Clarendon Press.

Paxson, Christina H., 1996, "Saving and growth: evidence from micro data," *European Economic Review*, **40,** 255–88.

Poterba, James, 1994, ed., *International comparisons of household saving*, Chicago, Il. Chicago University Press for NBER.

Turner, Dave, Claude Giorno, Alain De Serres, Ann Vourc'h and Pete Richardson, 1998, "The macroeconomic implications of ageing in a global context," OECD Economics Department Working Papers No 193, processed. (Mar.)

Weil, David N., 1994, "The saving of the elderly in micro and macro data sets," *Quarterly Journal of Economics*, **109**, 55–81.

Welch, Finis, 1979, "Effects of cohort size on earnings: The baby boom babies financial bust," *Journal of Political Economy*; 87(5), Part 2, pps. S65-97.

Table 2.1: Saving and demography regressions for 14 countries using decadal averages for the 60s, 70s, 80s, and 90s

	(1) OLS	(2) OLS	(3) Fixed Effects	(4) Fixed Effects	(5) Fixed Effects:IV
growth	1.27	–2.98	1.04	5.44	–0.54
	(3.6)	(0.6)	(2.8)	(1.0)	(0.4)
youth dep. rate	–1.05	–1.28	–1.33	–1.37	–4.40
	(4.3)	(2.4)	(5.3)	(2.4)	(2.2)
old dep. rate	–1.79	–3.98	–2.86	–0.04	–1.71
	(1.5)	(1.5)	(2.1)	(0.0)	(4.1)
growth* young		0.06		–0.03	
		(0.6)		(0.3)	
growth* old		0.52		–0.78	
		(0.9)		(1.5)	
p-value for equality	0.47	0.24	0.18	0.56	0.11

Dependent variable is the national saving rate as a percentage of gross national disposable income. Decades are 60s, 70s, 80s, and 90s up to 1994, with varying numbers of observations for the 60s. Countries and dates are Bangladesh (73–94), China(65–94), Hong Kong (65–94), Indonesia (66–94), India (65–94), Korea (65–94), Sri Lanka (65–94), Malaysia (65–94), Nepal (75–94), Pakistan (68–94), Philippines (65–94), Singapore (65–94), Thailand (65–94) and Taiwan (65– 95). Growth rates are per capita and, like the dependency rates, are measured as percentages. Instruments for the FE IV estimates are the current values of the ratios of exports to GDP, and gross domestic investment to GDP; the F(2,37) for the IVs in the first-stage regression is 3.06, which has a *p*-value of .0.06 Instrumental variables estimates using the lagged values (i.e. previous decade) of these variables as well as lagged growth for instruments in the FE regressions founder on the insignificance of the three variables jointly in the first-stage regressions. Random effect regressions look like FE regressions, or indeed OLS; none of the Hausman tests reject RE in favor of FE.

Table 2.2: Saving rates and demography, pooled cross-country annual time-series data

	National savings rate			Private saving rate		
	OLS	Random Effects	Fixed Effects	OLS	Random Effects	Fixed Effects
growth	0.33	0.24	0.24	0.08	0.08	0.15
	(4.6)	(4.3)	(4.2)	(1.4)	(1.4)	(2.9)
growth(−1)	0.23	0.18	0.18	−0.07	−0.07	−0.02
	(3.1)	(3.3)	(3.3)	(1.1)	(1.1)	(0.1)
growth(−2)	0.21	0.10	0.09	−0.09	−0.09	−0.02
	(3.0)	(1.9)	(1.7)	(1.6)	(1.6)	(0.3)
youth dep	−1.21	−1.36	−1.36	−1.14	−1.14	−0.97
	(13.8)	(15.4)	(15.1)	(8.2)	(8.3)	(5.9)
old dep	−2.48	−2.87	−2.89	−3.20	−3.20	−2.00
	(5.3)	(5.7)	(5.6)	(3.1)	(3.1)	(2.0)
p-value for equality	0.0014	0.0004	0.0006	0.02	0.24	0.40

Notes: National saving rates and growth in the first panel as defined in the notes to Table 1. Private saving rates are unavailable for adequate samples for Bangladesh, China, Hong Kong, Sri Lanka, and Nepal, so that the regressions in the second panel are for 9 countries, Indonesia (75–94), India (65–94), Korea (71–94), Malaysia (65–94), Pakistan (70–93), Philippines (65–94), Singapore (70–94), Thailand (67–94), and Taiwan (65–94). In the second panel, the growth variable is defines as gross national disposable income less the sum of public saving and public consumption (including non-central government), deflated by the implicit price deflator of GDP. Hausman tests do not reject the random effects specifications. Regressions were also run excluding once and twice lagged growth, but including interactions between growth and the youth and old dependency rates; in none of the six regressions were the coefficients on the interaction terms singly or jointly significant.

Table 2.3: Country by country regressions of savings and demographics, with weighted averages over all countries

	growth		youth dependency		old dependency		R^2	*d.w.*
Bangladesh	–0.04	(0.2)	–3.34	(2.9)	–5.99	(1.1)	0.59	1.42
China	0.12	(1.5)	–0.90	(1.6)	–0.60	(0.1)	0.69	0.66
Hong Kong	0.28	(2.7)	–0.61	(2.1)	–0.32	(0.3)	0.74	1.32
Indonesia	0.73	(2.4)	14.34	(6.6)	124.54	(7.4)	0.81	0.90
India	–0.03	(0.3)	4.84	(2.9)	28.52	(3.6)	0.62	0.83
Korea	0.48	(4.4)	–1.06	(4.4)	0.58	(0.3)	0.93	1.00
Malaysia	0.48	(3.1)	–1.01	(2.5)	5.19	(0.7)	0.75	1.57
Nepal	0.17	(1.2)	1.30	(0.6)	4.30	(0.4)	0.09	1.23
Pakistan	0.38	(2.0)	–5.45	(6.2)	5.97	(1.0)	0.82	1.34
Philippines	0.47	(3.8)	–1.07	(3.4)	–18.61	(5.1)	0.65	0.96
Singapore	0.09	(0.6)	–0.58	(1.7)	7.89	(3.4)	0.97	1.11
Sri Lanka	–0.42	(1.6)	–2.43	(3.5)	–10.52	(2.2)	0.71	1.26
Thailand	0.37	(2.2)	2.09	(2.1)	27.85	(2.7)	0.83	0.82
Taiwan	0.53	(4.1)	–1.89	(10.3)	–6.74	(8.5)	0.84	0.98
weighted average	0.28	(9.5)	1.02	(5.4)	14.04	(9.0)		

Notes: Each regression is run separately for each country, using data periods listed in the notes to Table 1. The weighted averages are computed using current price GDP of 1990 in US $ evaluated at World Bank Atlas conversion factors. The figures in brackets are t-values, computed to take into account correlations in residuals across countries in any given year (macro shocks.)

Table 2.4: Country by country dynamic regressions of savings and demographics, with weighted averages over all countries

	growth		youth dependency		old dependency		lagged ratio		*F: zero*	*F: equal*
Bangladesh	−0.14	(0.6)	−2.97	(1.9)	−3.51	(0.5)	0.29	(0.9)	2.39	0.00
China	0.31	(3.5)	0.07	(0.1)	2.72	(0.6)	0.73	(5.2)	0.85	0.43
Hong Kong	0.42	(2.8)	−0.65	(1.5)	−0.34	(0.2)	0.28	(1.8)	7.93	0.10
Indonesia	0.52	(1.9)	4.99	(1.6)	44.03	(1.7)	0.58	(4.5)	1.68	2.87
India	0.01	(0.1)	2.74	(1.6)	15.29	(1.6)	0.59	(3.8)	1.58	2.69
Korea	0.66	(6.4)	−0.72	(2.9)	−0.82	(0.4)	0.59	(6.0)	8.39	0.00
Malaysia	0.69	(3.1)	−1.16	(1.8)	6.31	(0.6)	0.14	(0.8)	7.49	0.43
Nepal	0.14	(0.9)	3.00	(1.3)	14.38	(1.3)	0.41	(1.7)	0.90	1.81
Pakistan	0.22	(0.9)	−3.34	(1.9)	−4.14	(0.5)	0.38	(2.0)	5.18	0.01
Philippines	0.51	(4.0)	−0.61	(1.6)	−10.30	(2.0)	0.53	(4.2)	2.10	4.19
Singapore	0.36	(1.4)	−0.14	(0.3)	6.61	(1.3)	0.59	(3.1)	2.78	2.17
Sri Lanka	−0.44	(1.4)	−2.07	(2.2)	−9.16	(1.6)	0.31	(1.7)	5.26	2.05
Thailand	0.43	(2.6)	0.97	(0.9)	13.21	(1.2)	0.63	(5.2)	3.18	1.43
Taiwan	0.80	(6.4)	−1.12	(3.4)	−4.28	(3.5)	0.58	(6.1)	6.10	12.13
SR average	0.39	(9.5)	0.45	(1.2)	6.40	(2.6)				
LR average			1.32	(1.5)	16.64	(2.8)				

Notes: Each regression is run separately for each country, using data periods listed in the notes to Table 1. The dependent variable is minus the logarithm of the sum of public and private consumption plus the logarithm of gross national disposable income, which approximates the saving rate. The lagged ratio is the one-year lag of this variable. The averages are weighted as before, and are computed using current price GDP of 1990 in US $ evaluated at World Bank Atlas conversion factors. The short-run (SR) averages are the weighted averages of the coefficients shown for each country. The long-run averages are the weighted averages over countries of the coefficients divided by (1−coefficient on the lagged ratio), and estimate the long-run effect of changes in the dependency ratios on the ratio of aggregate saving to aggregate income for the 14 countries as a whole. The figures in brackets are t-values, computed to take into account correlations in residuals across countries in any given year (macro shocks.) The F-statistics test the hypotheses that the coefficients of the dependency ratios are jointly zero, in the penultimate column, and are equal to one another, in the last column.

Table 3.1: Estimates of growth and cohort size effects (t-statistics in parentheses)

dependent variable	data source	η across-cohort growth effect	α effect of cohort size
$\ln w_i$: ln of individual wage	Manpower Survey	.0584 (86.12)	.0265 (1.61)
$\ln e_i$: ln of individual earnings	Manpower Survey	.0570 (77.22)	.0211 (1.20)
$\ln y_h$: ln of household income	Personal Income Distribution Survey	.0679 (141.46)	-.0240 (2.12)
$\ln c_h$: ln of household consumption	Personal Income Distribution Survey	.0584 (129.09)	-.0015 (0.14)
$\ln y_i$: ln of individual income	Personal Income Distribution Survey	.0670 (76.32)	-.0194 (0.99)
$\ln c_i$: ln of individual consumption	Personal Income Distribution Survey	.0586 (53.40)	.0225 (0.89)

Note: The first two rows show results based on cohort-level data constructed from the Manpower Surveys from 1979 to 1994. The cohort-level data set has 645 observations. The last four rows are based on data from the 1976 to 1995 Personal Income Distribution Surveys. The cohort-level data set had 784 observations. The t-statistics are based on "Huberized" standard errors.

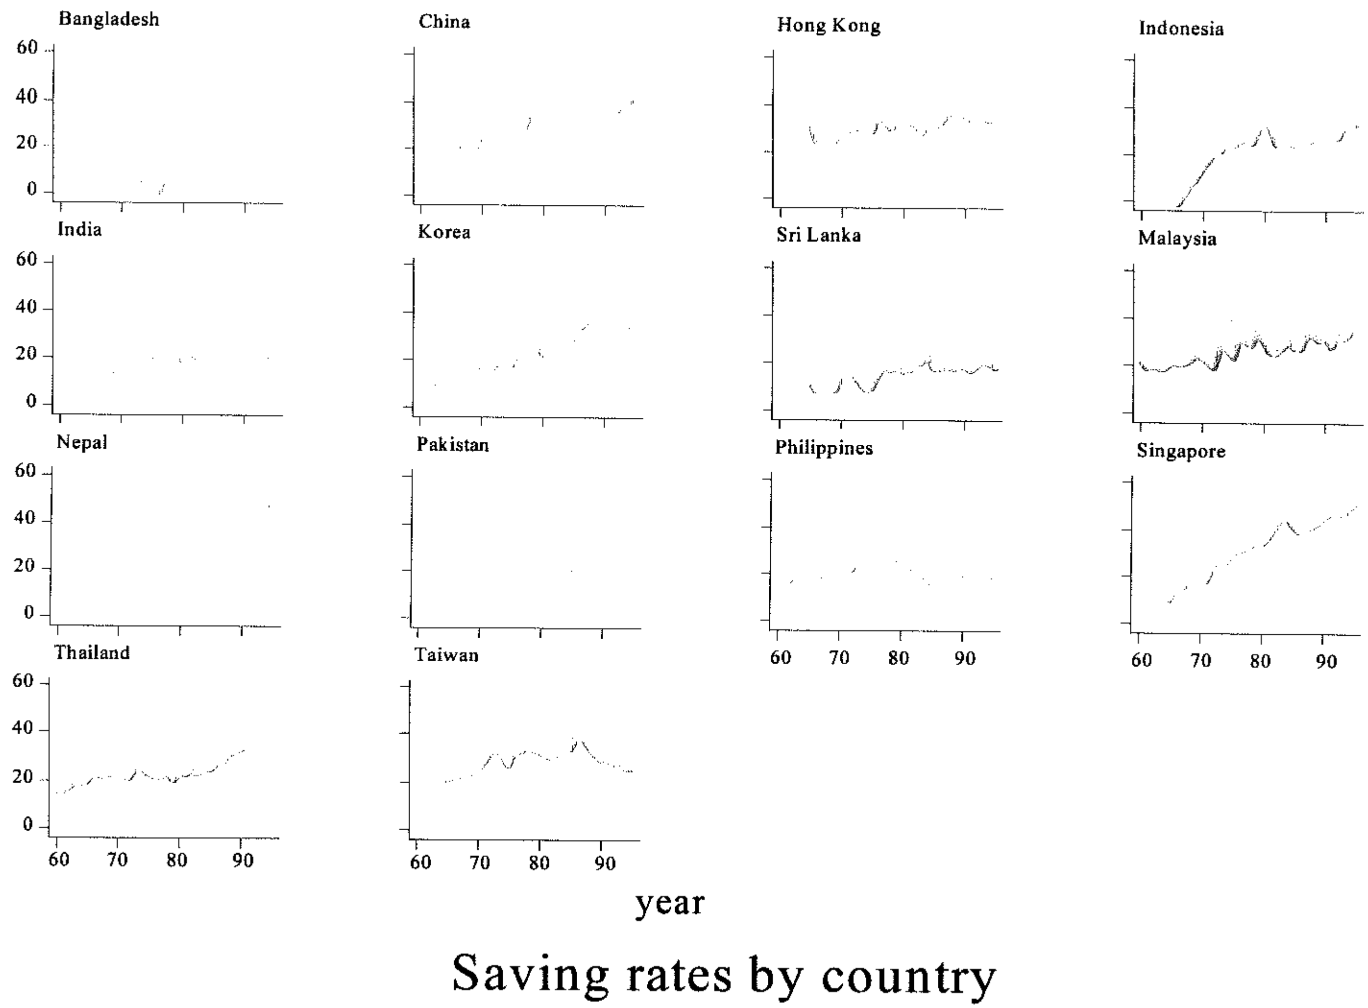

Figure 2.1: Aggregate national saving rates over time, 14 Asian countries

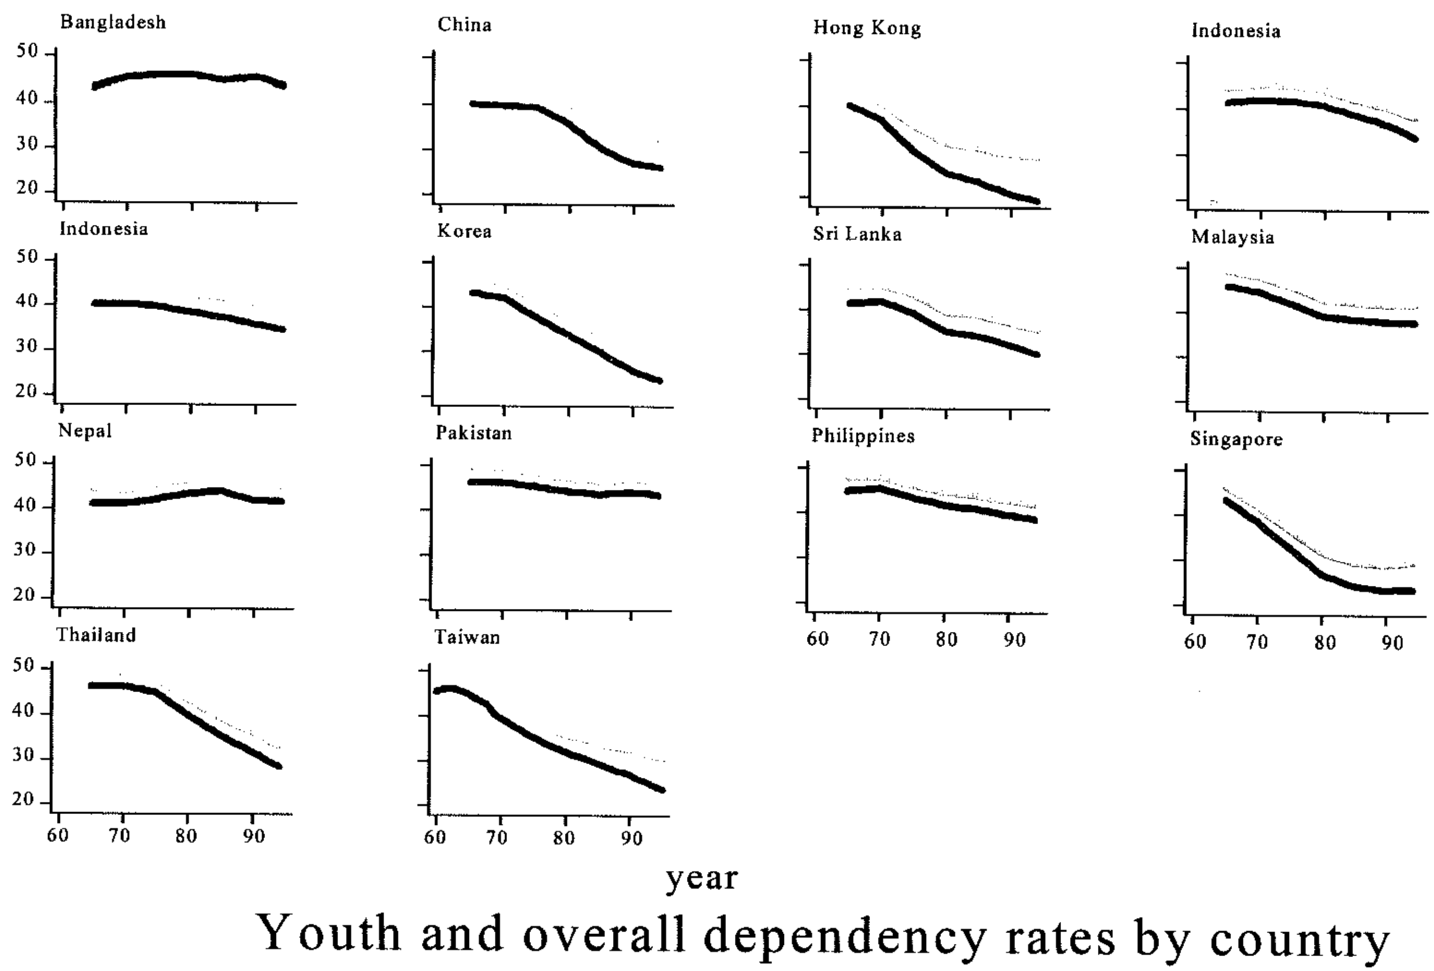

Figure 2.2: Dependency rates by country over time: lower line is fraction less than 15, upper line is fraction less than 15 or over 65.

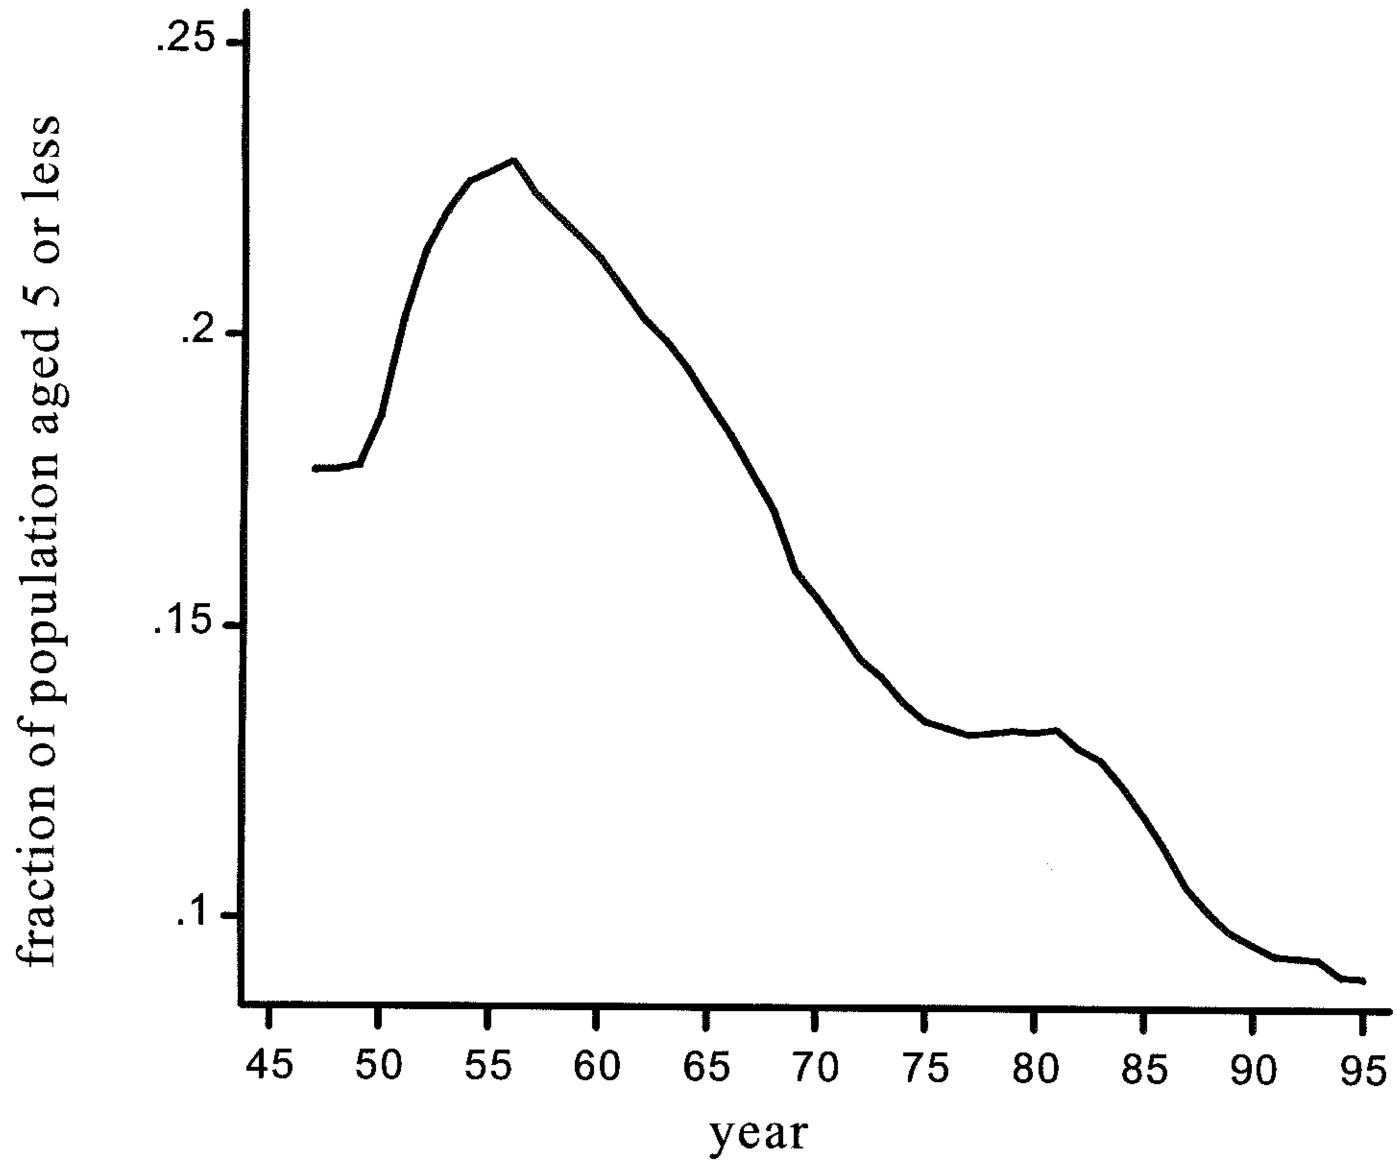

Figure 3.1: Fraction of the population aged 5 or less, by year

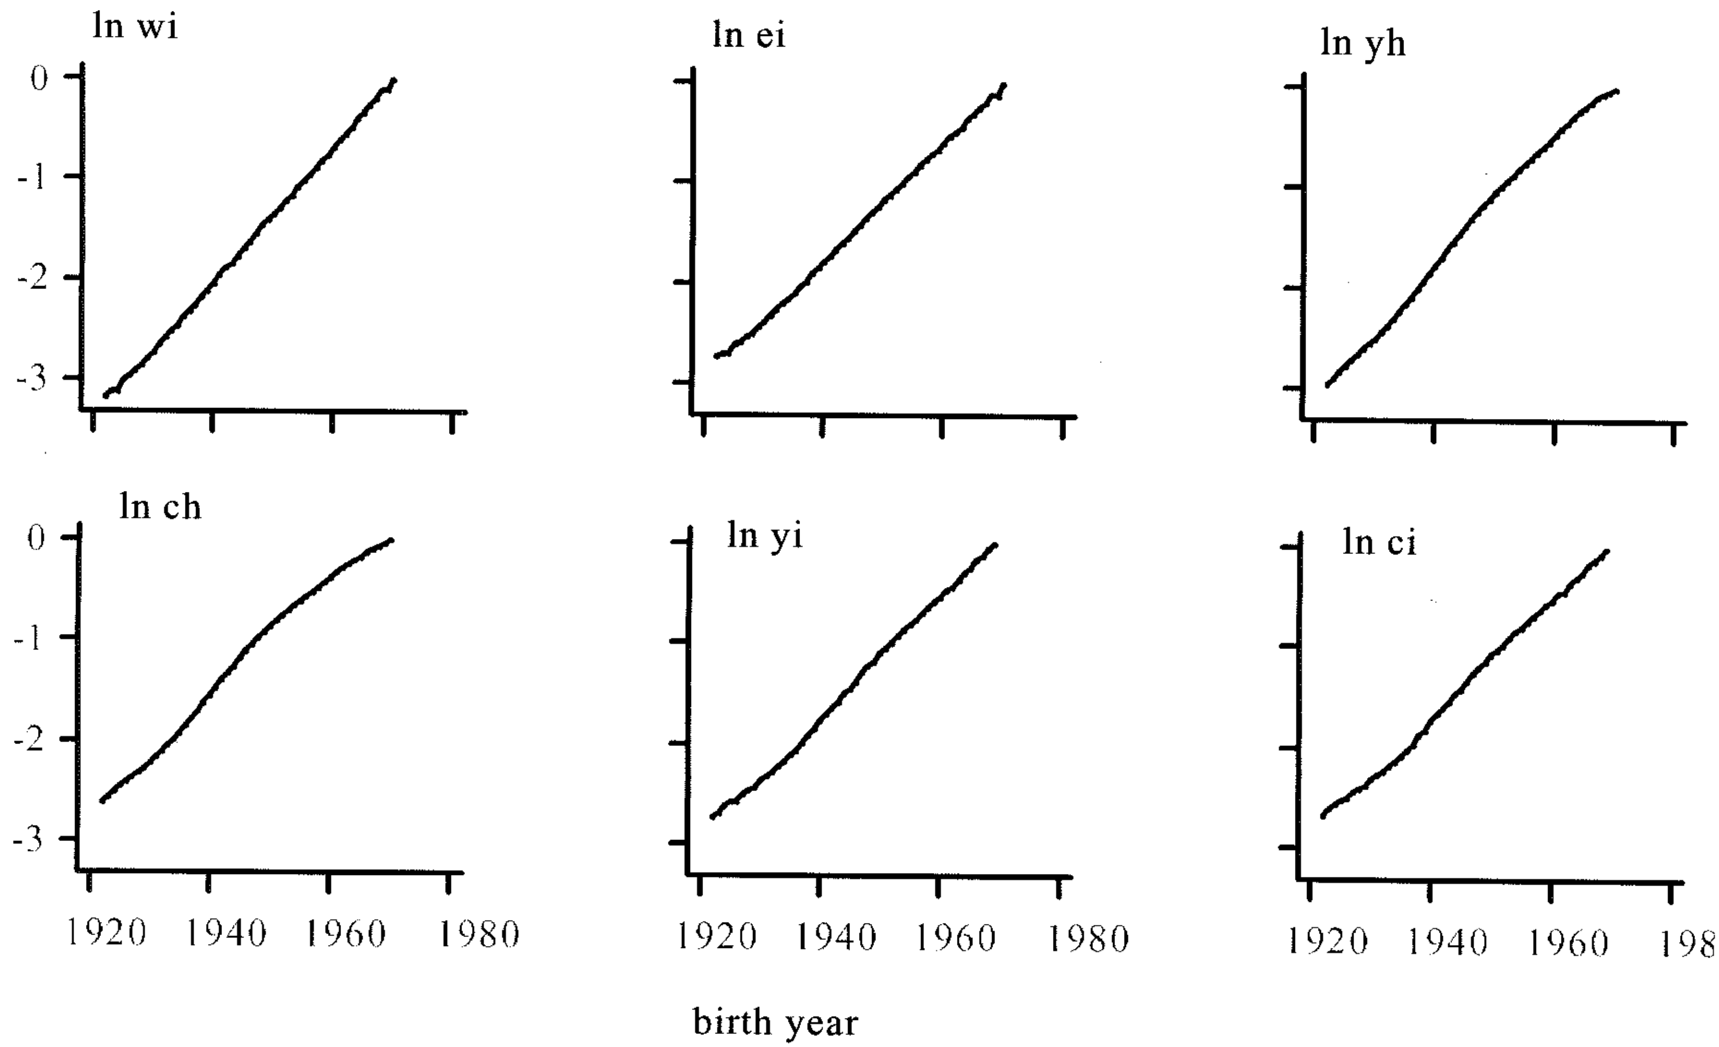

Figure 3.2: Cohort effects in wages, earnings, income and consumption

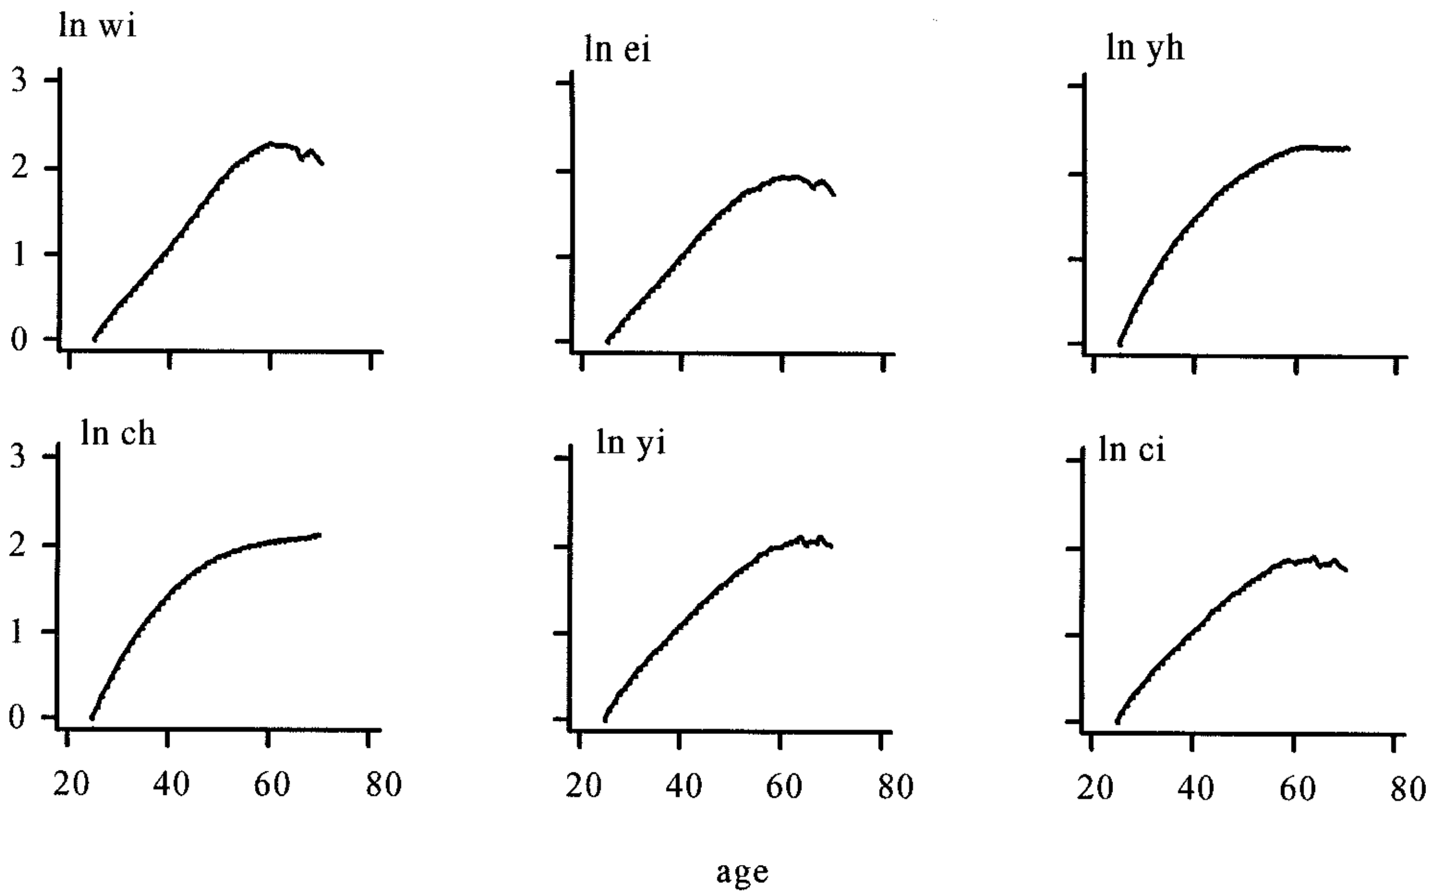

Figure 3.3: Age effects in wages, income and consumption

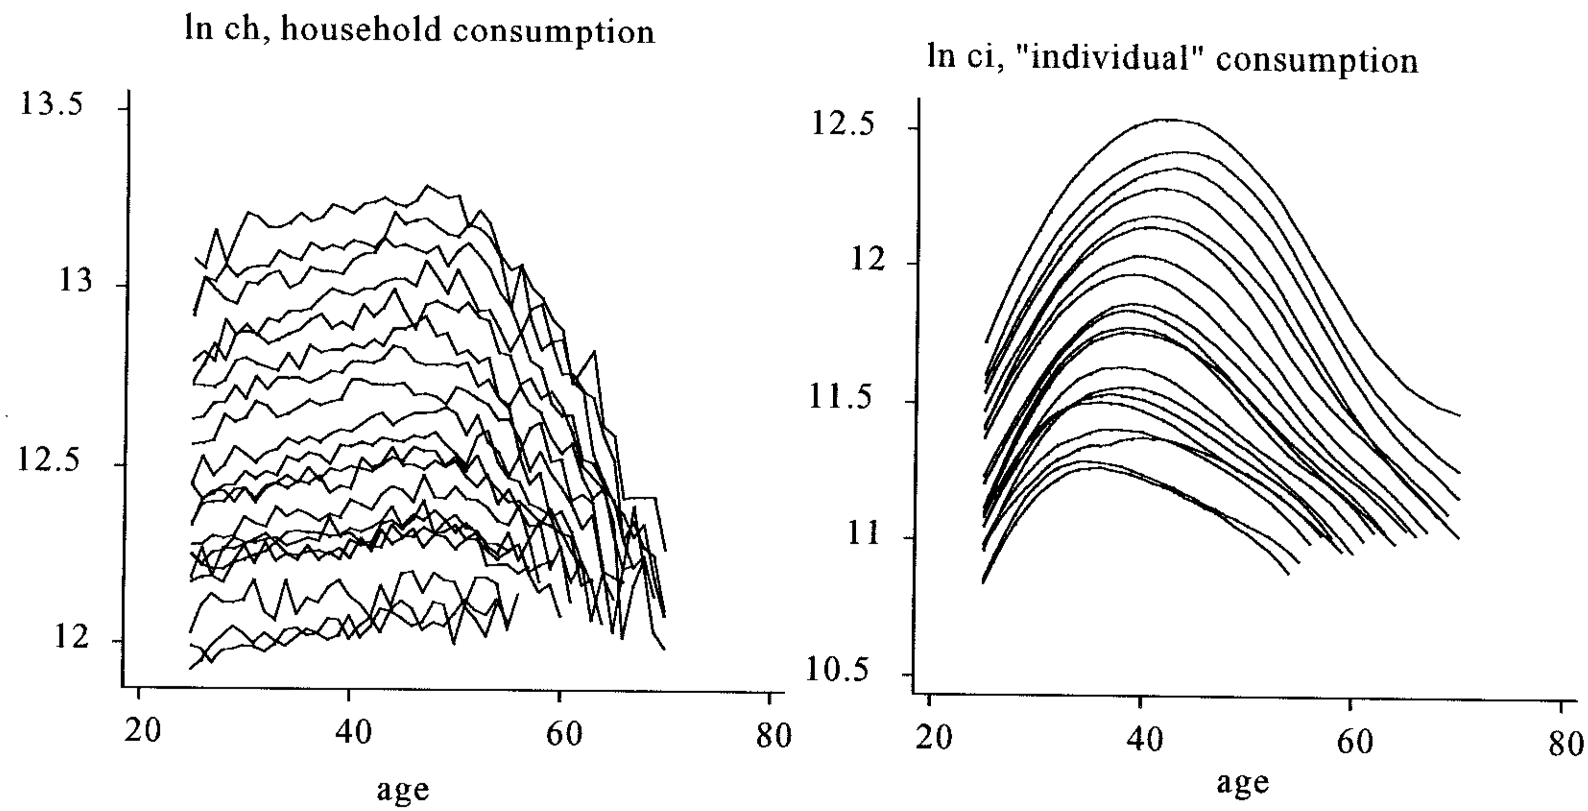

Figure 3.4: Age-consumption profiles for each year of survey data

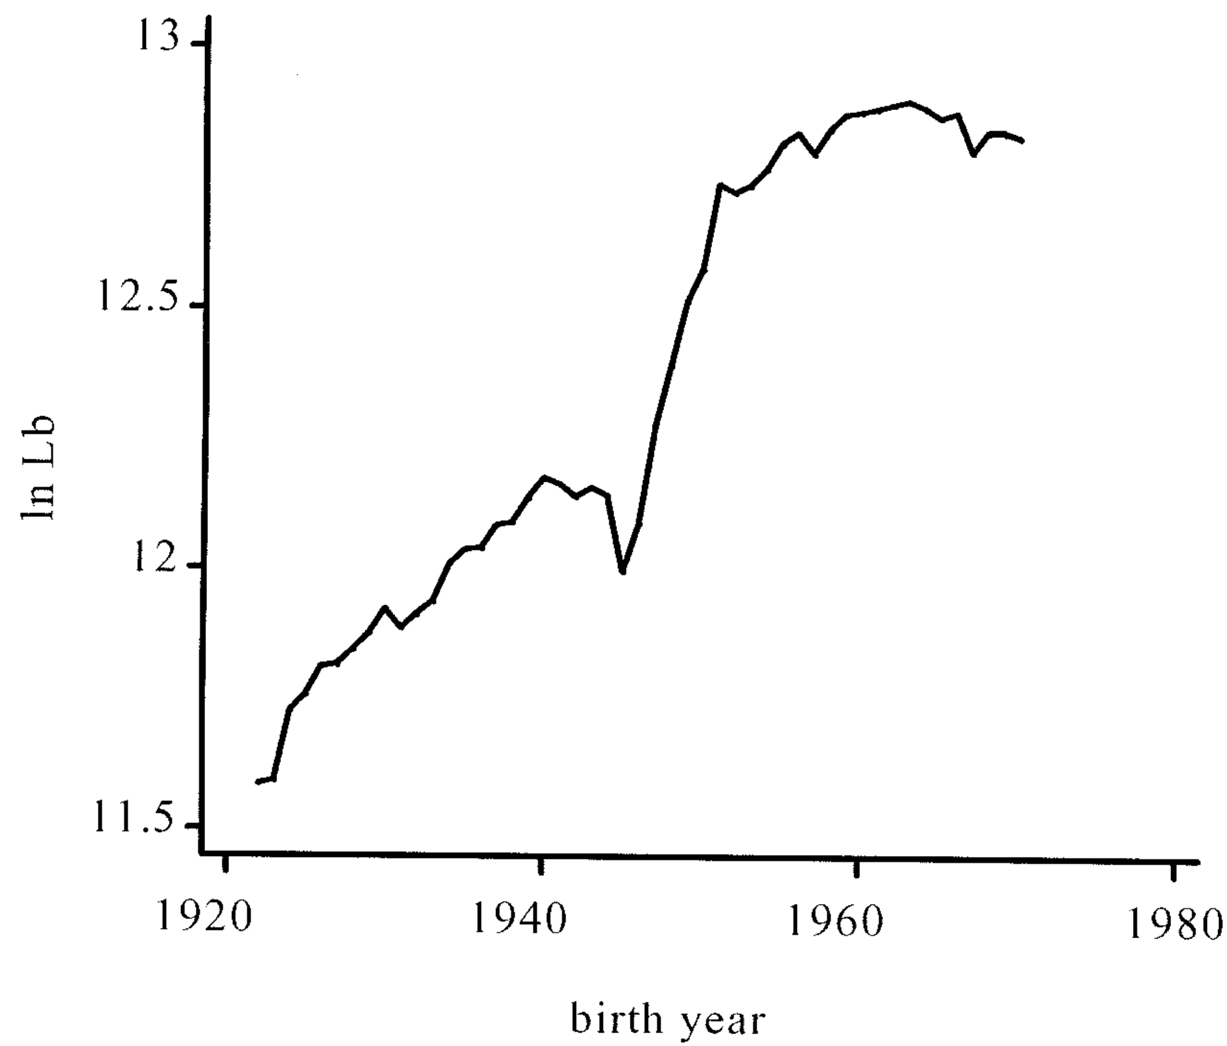

Figure 3.5: Log cohort size, measured at age 25

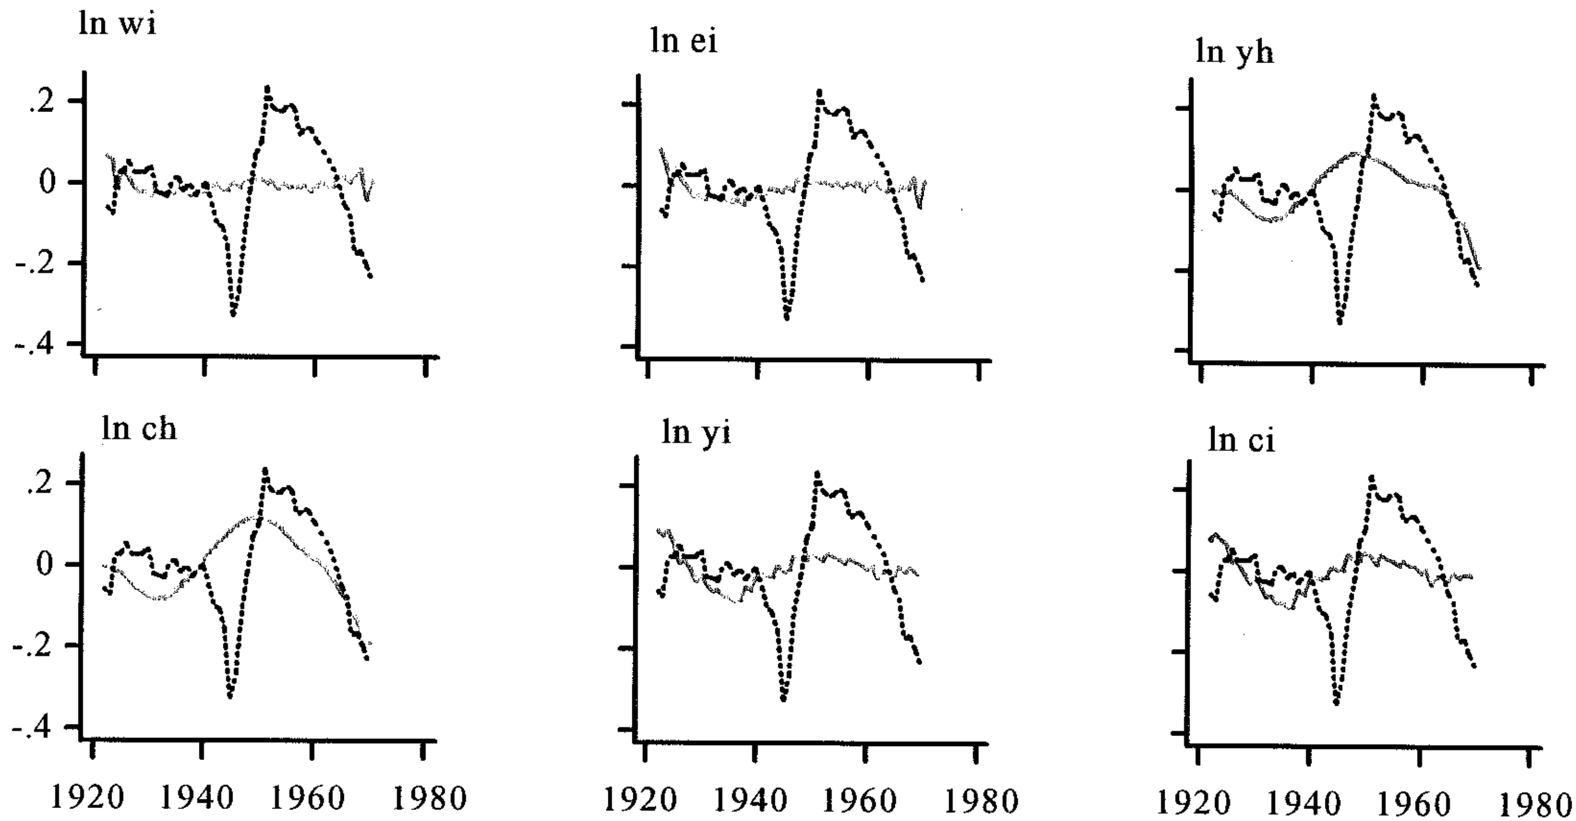

detrended logarithm of cohort size

detrended cohort effects in wage, income or consumption measure

Table 3.6: Cohort size and cohort effects in wages, income and consumption