Saving, Inequality and Aging: an East Asian Perspective

Angus S. Deaton
Christina H. Paxson
Research Program in Development Studies
Princeton University

Executive summary

This paper is concerned with the relationship between population aging and inequality. Our research is motivated by the fact that the populations of many countries in the world, and in particular those of Asian countries, are rapidly aging. We ask whether there are good theoretical reasons to believe that aging will increase inequality. Our results are examined in the context of Taiwan, an economy that is rapidly aging and which experienced an increase in inequality in the 1980s.

We show that the life-cycle model of saving has implications for how aging and inequality are related. Consider two countries with different rates of population growth, and assume each is in demographic equilibrium, meaning that the age structure of each population is stable over time. Life-cycle theory implies that, all else equal, the country with the lower population growth rate (implying an older population) will have more inequality. We demonstrate this result by first considering how inequality within a cohort of same-aged people should, according to life-cycle theory, evolve over time as the cohort ages. The theory has a clear implication: inequality within cohorts should increase with age. The reason is that the consumption of life-cycle consumers at a specific date reflects the accumulated effects of random shocks to life-time wealth, positive or negative, up to that date. If one considers a fixed group of people, their consumption levels will 'fan out' over time. We have documented elsewhere that there are large increases in inequality with age within cohorts of same-aged people in Taiwan, the United States and Great Britain.

The result of widening within-cohort inequality with age has implications for the effects of population aging on inequality. Specifically, as the rate of population growth falls, there will be a higher fraction of people at the older ages for which within-cohort inequality is largest. This redistribution of the population towards older and more unequal groups will cause aggregate inequality to rise. Aggregate inequality is also affected by between-cohort inequality — that is, inequality in average consumption levels across groups of people of different ages — so that if old people consume much more or much less than the young, between-cohort inequality will be high. Under several of the standard assumptions about preferences, we show that life-cycle theory predicts that aging will either cause between-cohort inequality to either increase or remain unchanged, so that aggregate inequality will unambiguously increase with aging.

Realistic modifications of the theory to account for things such as the effects of life-cycle patterns in family size and the age composition of households on consumption can potentially undermine this result, and yield a negative relationship between aging and inequality. To better understand the predicted size of the effects of aging on inequality, and to see if a negative relationship is likely, we benchmark the model using Taiwanese household survey data. Our results indicate that aging potentially has a large and positive effect on inequality, especially if economic growth remains high. For example, given a per capita economic growth rate of 6 per cent, a decline in the rate of population growth from 3 per cent to 1 per cent would produce a substantial increase in the Gini coefficient from 0.309 to 0.352.

This paper is concerned with the relationship between population aging and inequality. We show that simple life-cycle models of saving and consumption predict that, in demographic equilibrium, economies with slower rates of population growth will have higher consumption and income inequality. Our results are consistent with recent trends in Taiwan, where the population is aging. Although Taiwan is one of the leading cases of 'growth with equity', inequality has been increasing for more than a decade. We use our model, together with household data from Taiwan, to quantify the effects predicted by the theory and to calculate measures of inequality associated with alternative rates of growth of population and of real incomes.

Although the implications of life-cycle theory for the relationship between aggregate saving and the rates of population and real income growth are well understood, no clear results have previously been derived for the effects of growth rates on inequality (Lam 1987). While there are a number of possible non-life-cycle mechanisms that could link population growth and inequality — differential fertility decline by income level is a leading contender, see Chu and Koo (1990) — the mechanism discussed in this paper is a direct implication of standard permanent-income or life-cycle theories of consumption and saving under uncertainty. In Deaton and Paxson (1994b) we show that if a finite life version of the permanent income hypothesis is correct, then consumption and income inequality within any given cohort of individuals will widen with age. This result, which can be generalised to a range of other life-cycle formulations, and which is supported by our empirical analysis of data from Taiwan, Great Britain and the United States, implies that in the absence of any cohort effects in inequality, income and consumption will be less equally distributed among older people than among younger people. In consequence, demographic change that increases the share of older people in the population will work to increase the inequality of consumption and income in the population as a whole. While overall measures of inequality depend on both between-cohort and within-cohort inequality, we show that the former will not be affected by changes in the rate of population growth in at least some leading versions of the life-cycle model.

Taiwan is a country that has experienced one of the most dramatic demographic transitions in history, and its population structure has been rapidly aging. Measures of inequality are very low in Taiwan in relation to those in many other countries, but there has been concern over rising inequality. The Gini coefficient for disposable income has risen from 0.277 in 1980 to 0.312 in 1990 (Republic of China 1990). Our calculations of consumption inequality indicate a similar increase over the same period, from 0.256 to 0.289. Although Gini coefficients can in theory range from zero to one, with higher values representing more inequality, these observed changes in Taiwan are actually quite large, and are certainly big enough to be a cause for concern among policymakers. The excellent household survey data from Taiwan, which are available annually since 1976, allow us to investigate the phenomenon of increasing inequality in some detail, and in particular, to look at inequality both within and between cohorts.

To motivate the following sections of this paper, we begin with some evidence on within-cohort inequality. Figure 1 shows one measure of inequality, the variance of the logarithm of consumption, for 1976 through 1990 (excluding

Figure 1: Within-cohort inequality for selected cohorts

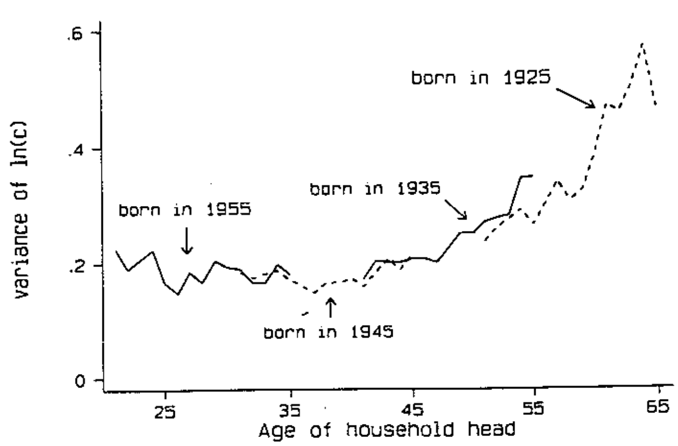

1978) for four groups of households, those whose heads were born in 1925, 1935, 1945 and 1955 respectively. The Taiwanese household data are not longitudinal, but we can use year of birth to identify the different representatives of the same cohort that show up in each of the surveys. We then plot inequality against age over time for each of the cohorts, with each connected segment showing the experience of one cohort. While it is possible to construct plots for 30–45 cohorts, we maintain clarity by showing just four; the results are very much the same for the other intermediate age groups.

The figure shows a clear effect of age on consumption inequality, flat (or even slightly declining) until around age 40, and rising thereafter. We also see that each cohort segment picks up near where the previous one left off, so that the level of inequality seems not to depend on the birth cohort, but only on calendar age. Furthermore, these effects are quite large: the increase in the variance of the logarithm of consumption from age 25 to age 55 that is shown in the figure corresponds to an increase in the Gini coefficient from about 0.23 to about 0.33, assuming consumption is lognormally distributed. In such circumstances, there is clear scope for aggregate inequality to be increased by a shift in the balance of the population towards older age groups such as that which has taken place in Taiwan.

Although Figure 1 is important to motivate our approach, we are less concerned in this paper with explaining what has happened in Taiwan — where there have been many other forces at work — than in using the Taiwanese numbers to fill out our theoretical formulae and thus to quantify the theoretical effects of fertility decline on inequality. We are particularly concerned with the interactions between population growth and economic growth, since increases in the latter affect the between-cohort distribution of income and consumption by redistributing lifetime resources to younger cohorts. The rate of population growth affects not only the relative numbers of old and young, but also the numbers of children in each household which, in turn, affects the life-cycle profile of household consumption. The effects on the between-cohort distribution of consumption, and thus on overall consumption inequality, depend on the distribution of resources between cohorts, and thus on the rate of economic growth. We can thus generate a rich set of interactions between economic growth, population growth, and aggregate inequality.

The paper is organised in the following sections: first there is a brief recapitulation of the results of Deaton and Paxson (1994b), sufficient only to make this paper self-contained and to motivate the basic results. Second, the life-cycle inequality theory is applied to a stable population growing at a fixed rate, and the circumstances under which a decrease in the rate of population growth will increase inequality are discussed. Finally, the last section is concerned with quantifying these results for Taiwan, and with the preparatory work that is necessary to convert the theory into usable form, particularly to make the transition from the individual agents of the theory to the households in the data and the actual measures of inequality.

Inequality and age

In the simplest version of the permanent income hypothesis, individual consumption is a martingale, so that, for individual i at time t, consumption c_{it} satisfies

$$c_{it+1} = c_{it} + \eta_{it+1} \qquad (1)$$

where η_{it+1} is an innovation whose expectation at time t is zero. Consider a group of individuals in existence at both t and t+1. Provided that the innovations are independent of current consumption in the cross-section, (1) shows that the cross-sectional distribution of the group's consumption at t+1 is the current distribution of consumption plus noise, so that the distribution of consumption at t+1 is (second-order) stochastically dominated by the distribution at t. As a result, consumption inequality at t+1 must be at least as large as consumption inequality at t, provided only that the inequality measure is quasi-concave. That the cross-sectional variance and coefficient of variation are increasing is immediate from (1), but the result applies more generally, for example to the Gini coefficient or to Theil's entropy measure.

Generalisations of this result are discussed in detail in Deaton and Paxson (1994) and can be summarised as follows. First, the theory of intertemporal allocation does not imply that the innovations η_{it+1} are independent of current consumption in the cross-section, only that they should be independent over time. As a result, the implication that inequality should increase in every year has to be weakened to the result that inequality must increase on average over a run of years. Second, the martingale theory of consumption is a special case of the general intertemporal allocation model with intertemporally additive preferences. In this more general case, (1) has to be replaced by the corresponding condition on the marginal utility of consumption

$$(1+r_{it+1})\lambda_{t+1}(c_{it+1}) = \lambda_t(c_{it}) + \varepsilon_{it+1} \qquad (2)$$

where r_{it+1} is the real interest rate linking t and $t+1$ and ε_{t+1} is an innovation. Equation (2) is equivalent to (1) under the conditions that validate the permanent income hypothesis, that the marginal utility function is linear and constant up to discounting by the rate of time preference, and that the real rate of interest is constant and equal to the rate of time preference. More generally, it is sufficient for (2) to give increasing inequality that the marginal utility function be concave and that the real rate be no less than the rate of time preference, conditions that are close to the opposite of those used in the precautionary saving literature. Note however that these conditions are sufficient but not necessary, and the permanent income hypothesis case (1) corresponds most closely to the flat life-cycle consumption profile of the simplest 'stripped-down' life-cycle model that is the basic workhorse for calculating the effects of growth on aggregate saving. It is thus of particular interest in the current context, where we are interested in extending that analysis to the relationship between growth and inequality.

The basic inequality result also has to be extended to deal with finite life-times, and to allow for changes in household characteristics over the life-cycle, particularly to do with the birth and rearing of children. It is straightforward to show that (1) remains true in a finite-life version of the permanent income model. The simplest life-cycle model has consumption constant: under uncertainty, the appropriate generalisation is that consumption is constant except in the face of new information. Changes in household characteristics can be taken into account in a number of ways. The simplest procedure, which we shall also use in the final section below, is to assume that there is a deterministic life-cycle path of household consumption, shaped by children and general taste change with age — for example the elderly may need more heat but less entertainment — and then to note that (1) applies to the non-deterministic part of life-cycle consumption. As a result, the within-cohort distribution of consumption will diverge or 'fan out' around this standard profile as the cohort ages. Such a modification means that the result of increasing inequality within cohorts holds for consumption relative to trend, not consumption itself, something that will have to be taken into account in the calculations.

The final theoretical result in Deaton and Paxson (1994b) applies to the inequality of income. We show that, in the case of the permanent income hypothesis where (1) is true, total income, which is earnings plus any income from assets, also disperses within cohorts as they age. Perhaps the easiest way to see this is to consider two cases, Case A, where earnings are fanning out within the cohort, and Case B, where the distribution of earnings within the cohort is stationary with age.

Case A could occur where individual earnings have unit roots and are at least partially independent. In the extreme case of a random walk, consumption will equal earnings, and consumption inequality will increase with earnings inequality. With a unit root, consumption will be a smoothed version of earnings, but the long-run stochastic trends would match — as in the random walk which is all stochastic trend — and consumption and earnings will fan out together. In neither of these cases is there long-run accumulation of assets, so that earnings and income have the same trends, and both show increasing inequality within the cohort. Case A can be generalised to include retirement. Retirement introduces a motive for saving and asset accumulation during working years, so that the equality of consumption and earnings in the random walk case no longer holds. However, with inequality in earnings and assets increasing over time, inequality in income will also increase.

In Case B, earning inequality is constant within the cohort, but each individual's consumption is following a random walk. The increasing divergence between earnings streams and consumption streams are financed by asset accumulation and decumulation, and it is the integrated behaviour of asset income that finances the integrated behaviour of consumption in the face of stationary earnings. Hence, although the distribution of earnings in the cohort is stationary, the distribution of earnings plus asset income is integrated, and the total fans out with consumption.

The empirical results in the earlier paper use time-series of household surveys from Taiwan, Great Britain and the United States. With 48 surveys, and several thousand observations in each, the analysis of inequality is largely graphical and non-parametric. In all three countries, we found results similar to those in Figure 1, with inequality increasing during the working years of life in all three countries, continuing to increase among the elderly in the United States, but remaining constant in Taiwan and Britain.

On average, the variance of the logarithms of household consumption increases at between 0.0074 (United States) to 0.0102 (Britain) per year of cohort age with Taiwan in between. Over the thirty years from age 25 to age 55, these changes generate increases in the Gini coefficient of consumption of between 0.08 (United States) to 0.19 (Britain). Similar results are found for earnings and for total income.

To a certain extent, the results in the next two sections are of interest whether or not the theory of this section is responsible for the fact of increasing within-cohort inequality with age. Provided that we are prepared to believe that within-cohort inequality is invariant to changes in the rate of population growth, increasing the share of the elderly in the population will increase inequality whether or not the intertemporal allocation model is correct. Of course, it is the intertemporal theory that tells us what to hold constant and what to change, and which protects us from possible internal inconsistency, so that we do without it at our peril.

Inequality and growth

This section, which contains the main theoretical result of the paper, considers an economy in which the population is growing steadily at rate n, and in which consumption and consumption inequality are determined by life-cycle considerations. Since our results require the decomposition of inequality into between and within-cohort components, we must use an additively decomposable measure of inequality. We work here with the variance of the logarithm of consumption. Although it is possible for the variance of logs to fail to correctly indicate stochastic dominance relationships, it does so only in peculiar circumstances, and this flaw is not serious enough to outweigh its convenience.

Let V_t denote the variance of log consumption in the population at time t. By the decomposition of variance into its between and within components, we can write

$$V_t = \sum_{a=0}^{\infty} n_{at}\left[v_{at} + (x_{at} - \bar{x}_t)^2\right] \qquad (3)$$

where a is age, running from 0 to infinity, n_{at} is the fraction of the population of age a at time t (which will be zero above the maximum age), v_{at} is the within-cohort variance, that is the variance of log consumption for all people aged a, x_{at} is average log consumption of people aged a, and $\bar{x}_t$ is the grand mean of log consumption at t. The last is given by

$$\bar{x}_t = \sum_{a=0}^{\infty} n_{at} x_{at} \qquad (4)$$

In demographic equilibrium, when the population is growing at rate n and has been doing so for a long time, the share of people aged a is given by

$$n_{at} = \frac{b_0(1+n)^{t-a}p(a)}{b_0\sum_{\alpha=0}^{\infty}(1+n)^{t-\alpha}p(\alpha)} = \frac{(1+n)^{-a}p(a)}{\sum_{\alpha=0}^{\infty}(1+n)^{-\alpha}p(\alpha)} \qquad (5)$$

where the survival function $p(a)$ gives the probability that someone lives to age a, and b_0 is the number of births at time 0. Given that the lifetable is taken to be independent of time t, the advantage of comparing equilibrium states is that n_{at} is also independent of t. The disadvantage, of course, is that given the very long time periods required to move from one steady state to another, the comparison of steady states may tell us little about a demographic transition within any one country.

As first noted by Coale (1972), equation (5) yields very simple formulae for the derivative with respect to n of a mean or variance of any quantity in the population, see also Preston (1982) and Lam (1984). In particular, differentiation of (5) gives

$$\frac{\partial n_a}{\partial n} = \frac{-n_a}{1+n}(a - \bar{a}) \qquad (6)$$

so that the effects of changes in n on the average of a population characteristic will depend on the average age where that characteristic is located, a result that has been put to good use in the analysis of transfers by Lee and Lapkoff (1988) and Lee (1993).

Since we are also assuming that there are no extraneous causes of inequality change other than the aging of the population, the within-cohort variances will be the sum of the cohort variance at birth, denoted v_0, and an age-specific variance effect θ_a. Given this and the equilibrium growth assumption, the variance (3) is time invariant and can be rewritten dropping the time suffixes as

$$V = \sum_{a=0}^{\infty} n_a\left[v_0 + \theta_a + (x_a - \bar{x})^2\right] \qquad (7)$$

Substituting from (5), differentiating with respect to n, and using (6), we have

$$\frac{\partial V}{\partial n} = -(1+n)^{-1}\sum n_a(a-\bar{a})\left[v_0 + \theta_\alpha + (x_a - \bar{x})^2\right] - 2\frac{\partial \bar{x}}{\partial n}\sum n_\alpha(x_\alpha - \bar{x}) \qquad (8)$$

By the definition of $\bar{x}$ in equation (4), the last term is zero, so that the effect of changes in the rate of population growth on inequality can be written as

$$\frac{\partial V}{\partial n} = -(1+n)^{-1}\sum n_a(a-\bar{a})\theta_a - (1+n)^{-1}\sum n_a(a-\bar{a})(x_a-\bar{x})^2 \quad (9)$$

Equation (9) is the basic formula that we need, and to which the economic arguments can be applied. There are two terms. The first is the effect on which we have been focusing so far, and which comes from the fact that within-cohort inequality changes with age, while the second is the contribution to inequality of the between-cohort consumption differentials. If the inequality-age theory of the previous section is correct, θ_a is increasing with a, and the first term will be negative. More generally, if within-cohort inequality is positively correlated with age, the first term will be negative.

The sign and size of the second term depend on the shape of the cross-sectional age profile of consumption. Note that this term is a third moment, which will be positive or negative depending on how the squared deviation of log consumption from the grand mean, that is $(x_a - \bar{x})^2$, varies with age. It will be zero for any cross-sectional consumption profile that is symmetric with respect to age, for example, for a flat profile, or a symmetric hump-shaped profile. In practice, the actual amount of between-cohort inequality in an economy can be directly estimated from a single cross-section of data. However, we can also use the theory to say something about its determinants, and we must follow this route if we are to say something about how it will change with changes in the rates of economic and population growth in steady state.

According to the life-cycle model, the cross-sectional age-consumption profile is determined partly by cohort effects (technology) and partly by life-cycle effects (tastes), and we can use this decomposition to gain further insight, as well as to show how economic growth enters the picture. In the simplest life-cycle model under certainty, consumption is the product of a wealth term and an age term

$$c_a = k(r,a)W \quad (10)$$

so that lifetime resources W are allocated over the life cycle according to tastes as modified by the incentives provided by the real interest rate. Taking logs of (10) yields

$$x_a = \ln c_a = \ln W + \ln k(a,r) = \omega_0 - ga + \alpha_a \quad (11)$$

where the last expression comes from assuming that lifetime resources are constant for each cohort, but grow at the rate of per capita economic growth g across cohorts, an assumption that will be justified if bequests are unit elastic with respect to lifetime resources. In Equation (11), ω_0 is the average logarithm of lifetime wealth for the newly born generation. The growth rate times age is subtracted from ω_0, indicating that older and less wealthy cohorts of people will consume less on average, controlling for the pure age effects in consumption (and assuming g is positive.) The last term in Equation (11), α_a, denotes a set of age effects given by the logarithms of the taste and interest rate-determined terms in Equation (10). With uncertainty, Equation (11) will be modified for any individual by the sum of all innovations up to that age, so that the scatter of individual paths around Equation (11) will fan out over time, which is what generates the within-cohort variances. However, Equation (11) will hold for the average of each age cohort provided there are no macroeconomic components to the innovations, a reasonable enough assumption when thinking about long-run demographic equilibria.

If we use ξ to denote the last term in Equation (9), the between-cohort contribution to the derivative of the variance, and use (11), we have

$$\xi = -(1+n)^{-1}\sum n_a(a-\bar{a})\left[(\alpha_a-\bar{\alpha}) - g(a-\bar{a})\right]^2 \quad (12)$$

an expression that can, as noted above, be positive or negative depending on the age profile of consumption. However, an important special case is when the age profile is linear in age, that is when

$$(\alpha_a - \bar{\alpha}) = \eta(a-\bar{a}) \quad (13)$$

for some parameter η. The standard 'stripped-down' life-cycle model is where $\eta = 0$ whereas in the case of isoelastic preferences, we have

$$\eta = \sigma(r-\delta) \quad (14)$$

for intertemporal elasticity of substitution σ and rate of time preference δ. Deaton and Paxson (1993) find that (13) is a good approximation to the actual age profile in Taiwan, even when compared with non-parametric estimates. When (13) holds, (12) becomes

$$\xi = -\frac{(\eta-g)^2}{1+n}\sum n_a(a-\bar{a})^3 \quad (15)$$

which is negative provided the third moment of the age distribution is positive, which will be the case unless population is falling rapidly enough. Since the within-cohort contribution to (9) is negative, we have the result that if consumption age profiles satisfy (13), declines in the rate of population growth will lead to greater inequality.

This result may be undone if (13) is not satisfied. For example, if the cross-sectional age-consumption profile is much lower than average at young ages, the squared deviation of consumption from its average will be very large for the young. By (9), this is the case where the between-cohort contribution to the inequality derivative can be positive, and it is possible for it to be large enough to reverse the negative within-cohort effect. Although there is no theoretical reason why the cross-sectional age-consumption profile should have such a shape, it might be so if inequality is measured over all individuals, including children. Children consume less than adults, and very young children consume very much less, and so the consumption of children is likely to lie far below the population mean of consumption. The life-cycle model laid out earlier implies that high rates of economic growth g may counteract this effect, by increasing the wealth and consumption levels of the young relative to the old. However, it is not clear that life-cycle theory is adequate to explain the consumption levels of children, who do not live independently of adults or make independent consumption decisions.

In practice, consumption data are collected at the level of the household, and inequality is measured as inequality across households, not individuals. This generates two modifications to our results, one of which is essentially the same as that discussed above, and one of which is more fundamental. The former comes from the fact that children tend to be located in households with younger heads, so household consumption may be lowest among the youngest households (where the age of a household is measured by the age of the household head). Whether or not children indeed have this effect is an empirical matter, and will be addressed below. The important point is that if the consumption of the youngest households lies far below average consumption, the last term in (9) can be positive. This effect will be strongest when economic growth rates are lowest, so that cohort effects work in the same direction as age effects.

A more fundamental problem is that, with household data, it is no longer realistic to treat the age effects in consumption levels (that is the x_a in Equation (7)) as if they are independent of the rate of population growth. Household consumption at any age is likely to depend on household composition — for example the ratio of children to adults — and household composition will vary with n. In this case, we can no longer hold constant the age effects when varying the rate of population growth, and Equation (9) no longer holds for household data. Under plausible specifications for the effect of children on consumption, the effect will again be for population growth to be positively associated with between-cohort inequality at low rates of economic growth. In subsequent work we plan to model this effect theoretically. For the moment, we incorporate it in the empirical results that follow.

Aging, population growth and inequality in Taiwan

In this final section, we show how to implement the results in the previous section using data from Taiwan. We discuss the practical changes that are required to work with the data, which come at the household rather than individual levels, and we calculate measures of inequality for different equilibrium rates of economic and population growth.

The calculations are all based on a household version of Equation (9), written as

$$V = \sum_{a=20}^{75} n_a^h\left[v_0 + \theta_a + (x_a - \bar{x})^2\right] \qquad (16)$$

where n_a^h is the fraction of households that are headed by people aged a, where we limit ourselves to households headed by people between 20 and 75, and where the within-cohort variances and means of logs now refer to household, not individual consumption. In the tables to follow, we will tabulate (16) as a function of the two growth rates n and g, but prior to doing so, we describe how each of the elements of (16) are obtained.

We start from the life-tables. We use the Taiwanese life-table for 1985 taken from Keyfitz and Flieger (1990), and use regressions on polynomials in age to calculate the survival probabilities for each age from the quinquennial survival probabilities in the tables. These life-tables show survival probabilities only up to age 85, and rather than attempt to extrapolate to survival probabilities at older ages we assume that no one lives past 85. The composite probabilities $p(a)$ are calculated from the gender-specific probabilities using

Table 1: Probabilities of survival and headship, and age distributions

age	$p^m(a)$	$p^f(a)$	$h^m(a)$	$h^f(a)$	cumulatives $\pi(5,a)$	$\pi(30,a)$	$\pi(50,a)$	$\pi(75,a)$
20	0.978	0.985	0.018	0.040	0.000	0.000	0.002	0.000
25	0.972	0.983	0.338	0.095	0.015	0.005	0.078	0.038
30	0.964	0.979	0.742	0.059	0.137	0.480	0.179	0.076
35	0.953	0.973	0.893	0.082	0.583	0.785	0.196	0.171
40	0.938	0.966	0.932	0.129	0.859	0.872	0.201	0.351
45	0.920	0.956	0.915	0.140	0.920	0.887	0.220	0.488
50	0.896	0.944	0.852	0.083	0.938	0.893	0.720	0.575
55	0.862	0.924	0.742	0.077	0.954	0.915	0.903	0.669
60	0.812	0.891	0.564	0.080	0.978	0.969	0.965	0.683
65	0.736	0.836	0.482	0.091	0.994	0.997	0.991	0.694
70	0.630	0.751	0.398	0.087	1.000	1.000	0.998	0.700
75	0.491	0.627	0.390	0.071	1.000	1.000	1.000	1.000

Notes: $p^m(a)$ is the probability that a male survives to age *a*. $p^f(a)$ is the probability that a female survives to age *a*. $h^m(a)$ is the fraction of males aged *a* who are household heads. $h^f(a)$ is the fraction of females aged *a* who are household heads. $\pi(\alpha,a)$ is the fraction of all people aged α who live in household with a head of age *a*. The last four columns show cumulative values of $\pi(\alpha,a)$ for α equal to 5, 30, 50 and 75.

Sources: The data for columns 1 and 2 come from Keyfitz, N. and Flieger, W., 1990, *World Population Growth and Aging: demographic trends in the late twentieth century*, University of Chicago Press, Chicago. The data for the remaining columns is from the 1991 Personal Income Distribution survey.

$$p(a) = p^m(a)p_m + p^f(a)p_f \qquad (17)$$

where the superscripts indicate the gender specific survival probabilities for males and females, and $p_m = 0.5162$ and $p_f = 0.4838$ are the probabilities that a birth is male or female respectively.

If $p^h(a)$ is the probability that someone survives to age *a* and is a household head at age *a*, we have

$$p^h(a) = p^m(a)p_m h^m(a) + p^f(a)p_f h^f(a) \qquad (18)$$

where $h^m(a)$ and $h^f(a)$ are the probabilities that males/females aged *a* are household heads. The age and gender-specific headship rates are estimated using the 1991 Survey of Personal Income Distribution, the latest of the sixteen annual household surveys. These rates are shown for selected ages in the third and fourth columns of Table 1 which also shows the gender-specific survival probabilities. The fractions of heads who are aged *a*, and thus the fraction of households headed by individuals aged *a*, are therefore given by

$$n_a^h = \frac{p^h(a)(1+n)^{-a}}{\sum_{\alpha=20}^{75} p^h(\alpha)(1+n)^{-\alpha}} \qquad (19)$$

The within-cohort variance terms v_0 and θ_a in (16) are taken from Deaton and Paxson (1994) and their values are shown in the top left-hand panel of Figure 2. These are obtained by using the 14 surveys from 1976, and 1978

Figure 2: Age and cohort effects in the variance and mean of ln(c)

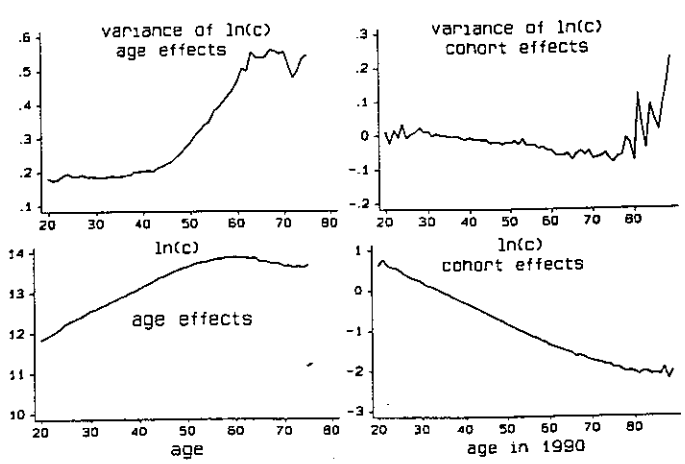

through 1990, calculating variances of the logarithms of consumption for each cohort at each age, and regressing the results on a set of age and cohort dummies to decompose the variances into age and cohort effects. The corresponding cohort effects — which play no role in the current analysis — are shown in the top right panel of Figure 2. These show very little secular trend in inequality. Although the younger cohorts appear to start their lives with slightly more inequality in log consumption than do their elders, a test that the 69 cohort effects in the variance of log consumption equal zero yields an F-value of only 1.66.

The between-cohort variance term in (16) requires the cross-section squared deviations by age $(x_a - \bar{x})^2$. In what follows, we compute between-cohort inequality at different rates of economic growth g and population growth n, holding fixed the pure age effects in consumption which we assume will remain constant over time. We isolate the pure age effects in consumption using two different strategies. In the first, we ignore the possible effects of household size and age composition on consumption. We simply regress the average of the logarithm of consumption for each cohort at each age on a set of age and cohort dummies. The regression, discussed in detail in Deaton and Paxson (1993) is

$$x_{ab} = \alpha_0 + \alpha_a + \lambda_b + \varepsilon_{ab} \qquad (20)$$

where x_{ab} is the average logarithm of consumption for households with heads aged a and born in year b, $\alpha_0+\alpha_a$ measures age effects in consumption, and λ_b measures cohort (that is, growth) effects in consumption. The age effects from (20) are then substituted into (11), and we calculate

$$x_a - \bar{x} = (\alpha_a - \bar{\alpha}) - g(a - \bar{a}) \qquad (21)$$

for different values of g.

Our second strategy takes explicit account of the fact that age-consumption profiles may depend on household size and the age composition of household members, both of which depend on the rate of population growth n. For example, when the rate of population growth rises, households will become larger on average and there will be a higher ratio of children to adults, and these changes in size and composition may affect consumption. Furthermore, these changes in size and composition are unlikely to affect households of all ages in a uniform manner. For example, suppose a higher ratio of children to adults (given household size) reduces consumption, and children tend to be concentrated among households with young heads. Then increases in the rate of population growth may depress consumption of younger households relative to old. This will cause the age profile of consumption to tilt counter-clockwise as the rate of population growth rises, with the rate of economic growth held constant. Changes in household size that differ across households of different ages may exert additional effects. The estimation of age effects with explicit allowance for household size and composition allows us to take these effects into account. In this second case (20) is modified to read

$$x_{ab} = \alpha_0 + \alpha_a + \beta_s \ln S_{ab} + \beta_k k_{ab} + \lambda_b + \varepsilon_{ab} \qquad (22)$$

where $\ln S_{ab}$ and $k_{ab} = (K/S)_{ab}$ are the average of the logarithm of family size (that is, number of household members) and the average ratio of numbers of children to family size in households headed by a person of age a born in year b. The age effects α_a and the coefficients on the demographic variables are then substituted into (11), and we calculate

$$x_a - \bar{x} = (\alpha_a - \bar{\alpha}) + \beta_s(\ln S_a - \overline{\ln S}) + \beta_k(k_a - \bar{k}) - g(a - \bar{a}) \qquad (23)$$

for different values of g and for values of S_a and k_a that are implied by different values of n. (Calculation of S_a and k_a is discussed below.)

The bottom left panel of Figure 2 graphs the estimates of age effects from (20). The age effects in consumption rise with age, at least up to age 60. However, because the Taiwanese economy has been growing so rapidly over the last fifteen years, the cohort effects — which are shown in the bottom right hand panel of the figure — are very much larger for the young than the old. Estimation of (20) with the cohort effects constrained to be linear yields an between-cohort growth rate of 0.049. In other words, controlling for age, the average consumption of each cohort is 4.9 per cent higher than the cohort born a year earlier. The combination of the age and cohort effects results in cross-sectional age profiles of consumption that are relatively flat.

This is best seen in Figure 3 overleaf, which gives two different decompositions of the same underlying data. The left-hand panel shows fifteen years of average log consumption for every (second) cohort, with points within each cohort connected. This shows the rapid growth of

Figure 3: Age-consumption profiles, by cohort and year

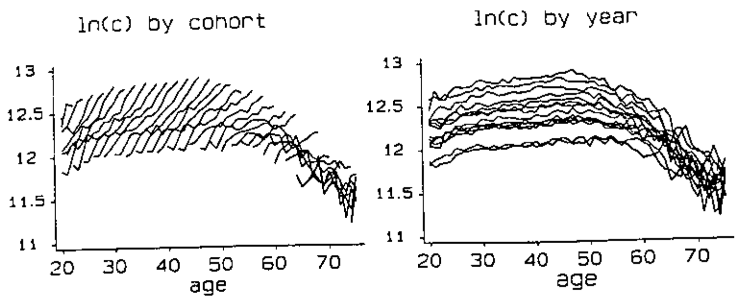

consumption for the younger cohorts, and much lower or even negative growth for older groups. The right-hand panel shows the corresponding cross-sections for each of the fifteen years. These have a quite different shape, rising gently with age until around age 50, and falling thereafter. The growth effects show up here in the upward movement of the cross-sections over time.

Estimates of Equation (22) yield age and cohort effects that are similar to those shown in Figure 2. The coefficients $\beta_s \equiv 0.238(t = 3.49)$ and $\beta_k \equiv -0.205(t = 1.20)$ show that, controlling for age and cohort — and note that the former captures average fertility — cohorts with larger than average household sizes consume more, while those with lower proportions of children per household consume (insignificantly) less. Although the inclusion of these demographic variables makes very little difference to the shape of the age-consumption profile shown, it will be seen below that whether or not we make explicit allowance for the effects of family size and composition has a major impact on the calculations of inequality.

To use the results of (22) to calculate between-cohort inequality, we need a procedure for calculating S_a and K_a, the average number of household members and children in households headed by persons aged a, for different rates of population growth. To compute S_a, we use

$$S_a = \frac{\sum_{\alpha=0}^{85} \pi(\alpha,a)p(\alpha)(1+n)^{-\alpha}}{p^h(a)(1+n)^{-a}} \qquad (24)$$

where $\pi(\alpha,a)$ is the probability that person aged α lives in a household with head aged a. Equation (24) is simply the number of people aged α divided by the number of households with head aged a multiplied by the probability that a person aged α lives in such a household. The probabilities $\pi(\alpha,a)$ are estimated from the 1991 survey data from the actual distribution of people of different ages across households with different aged heads. Exactly parallel formulae and parallel calculations are used to calculate the numbers of children. Equation (24) is simply modified to sum over the age range 0 to 20 (the Taiwanese statistical authorities define children as all those aged 20 or less.) The final columns of Table 1 show cumulative values of $\pi(\alpha,a)$ for selected values of α, and indicate that, as expected, children are concentrated among relatively young households.

Table 2 shows the final results of calculations of (16), the relationship between inequality as measured by the variance of logarithm of consumption and the rates of growth of population and per capita income. The upper panel shows the results when household composition is not explicitly modelled, and the middle panel the results when the effects of population growth on family structure are taken into account. The last three rows of Table 2 show for each rate of population growth the average household size, the average fraction of household members that are children, and the average age of household heads. Rates of per capita income growth from 0 per cent to 6 per cent are in the rows, and population growth rates from 0 per cent per annum to 4 per cent per annum are shown in the columns. In each case the total variance is decomposed into its between and within-cohort components. [For those who prefer to think of inequality in terms of Gini co-efficients, there is a precise one to one correspondence when log consumption is normally distributed: a variance of logs of

Table 2: Within-cohort, between-cohort and total inequality

		$n = 0.00$	$n = 0.01$	$n = 0.02$	$n = 0.03$	$n = 0.04$
	No demographic effects on consumption levels					
Growth (g)						
0.00	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.273	0.286	0.293	0.297	0.296
	total inequality	0.562	0.560	0.555	0.547	0.536
0.02	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.093	0.099	0.104	0.108	0.110
	total inequality	0.383	0.374	0.366	0.358	0.351
0.04	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.047	0.042	0.037	0.033	0.031
	total inequality	0.337	0.316	0.299	0.284	0.271
0.06	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.136	0.112	0.091	0.073	0.057
	total inequality	0.425	0.387	0.353	0.323	0.297
	Demographic effects on consumption levels					
0.00	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.261	0.287	0.309	0.329	0.347
	total inequality	0.550	0.562	0.572	0.580	0.588
0.02	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.082	0.096	0.112	0.127	0.142
	total inequality	0.371	0.372	0.374	0.377	0.382
0.04	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.037	0.035	0.035	0.038	0.042
	total inequality	0.326	0.310	0.297	0.288	0.283
0.06	within cohort inequality	0.289	0.275	0.262	0.250	0.240
	between-cohort inequality	0.126	0.101	0.080	0.062	0.049
	total inequality	0.416	0.376	0.342	0.313	0.289
	average household size ($\bar{S}$)	3.99	4.40	5.00	5.85	6.98
	average of (kids/size) ($\bar{k}$)	0.259	0.338	0.419	0.497	0.569
	average age of heads ($\bar{a}$)	45.70	44.06	42.53	41.09	39.76

0.26 is a Gini of 0.28, a variance of 0.43 is a Gini of 0.36, and the Gini increases by 0.005 (at low values) and 0.004 (at high values) for a 0.01 increase in the variance.]

Note first that although the Taiwanese population structure is far from a stable growth equilibrium — there was a baby boom in the early 1950s followed by very little growth in births — the variance of log consumption in 1990 was 0.299, which decomposes into an within-cohort component of 0.260 and a between-group variance of 0.039, figures that are close to those given in the top panel of Table 2 for $n = 0.02$ and $g = 0.04$. Although 4 per cent is a good deal less than the average rate of growth in per capita GDP in recent years, it is closer to the 4.9 per cent per annum at which the cohort effects in consumption are growing (Figure 2).

In the top panel of Table 2, we see the expected result that total consumption inequality falls as the rate of population growth rises, and the size of the effect is greater the larger the rate of growth of per capita income. When income is growing slowly, the cross-sectional age profile of consumption is dominated by the pure age effects, with the largest deviations from average (log) consumption occurring at the youngest ages. This is sufficient to make the final term in (9) positive, so that the between-cohort contribution to the derivative offsets the within-cohort contribution. As the growth rate rises, the age profile of consumption pivots clockwise and flattens out. As a result, the between-cohort component of the variance becomes smaller, and its derivative with respect to n changes sign.

In the lower panel, this effect is strengthened by the explicit modelling of the effects of population growth on household composition. Increasing the rate of population growth raises the ratio of children to adults in the typical household and thus tips the age consumption profile counter-clockwise. The effects of economic growth on the between-cohort contribution to inequality are therefore even stronger than in the top panel. As a result, when economic growth rates are low — the 0 per cent and 2 per cent cases in the lower panel — lower rates of population growth decrease overall inequality. At the higher rates of economic growth that are closer to those actually experienced by Taiwan, the within and between-cohort effects work in the same direction, and declining rates of population growth are a powerful force for increased inequality. If we think of these numbers in terms of Gini coefficients of consumption, at 6 per cent annual growth or per capita income, the difference between a 3 per cent population growth rate and a 1 per cent population growth rate is the difference between a Gini of 0.309 and a Gini of 0.352, a difference that would typically be regarded as of the first importance.

Conclusions

This paper demonstrates that life-cycle models of consumption predict that, in demographic equilibrium, there should be a negative relationship between population growth and inequality, at least at high rates of economic growth. The model fits the facts of the Taiwanese economy, where economic growth rates are high, population growth has declined, and inequality is increasing. The model also predicts increases in inequality in other fast-growing Asian countries, such as Thailand and Indonesia, where population growth rates have also slowed.

Several caveats must be kept in mind when interpreting our results. First, the results pertain to changes in inequality that result from movements from high to low steady state rates of population growth. Taiwan does not currently have a stable age distribution: the large in-migration and the baby boom of the early 1950s, and the subsequent rapid fertility decline, have resulted in an age distribution that has changed and will continue to change over time, see Deaton and Paxson (1993). This paper has not analysed what will happen to inequality during the transitional period as the effects of these particular demographic events are played out.

Second, a critical assumption in all we have done is that the pure age effects in the level and variance of consumption are stable over time, and will not change with n and g. The proposition that inequality within cohorts increases with age in a systematic way appears to be fairly robust, given that we find similar increases in within-cohort inequality with age in Britain, the United States and Taiwan. However, we are less convinced that the age-consumption profiles we observe in Taiwan would remain stable at different values of n and g. For example, the upward-sloped age profiles of consumption shown in the bottom left-hand panel of Figure 2 might be attributed to high real interest rates in Taiwan. If steady-state interest rates are determined by n and g, as is implied by closed economy growth models, then our method of holding the age-consumption profile fixed while changing n and g is invalid.

Third, it should be kept in mind that the theory in the second section strictly applies to inequality across individuals, not inequality across households. Because data on consumption come at the household level, and disentangling who gets what within households is no easy task, we benchmark our model using household-level data. The jump from inequality across individuals to inequality across households requires some strong assumptions about living arrangements and headship rates. Our approach has been to use data from Taiwan in 1991 to calculate age-specific headship probabilities (that is $h^m(a)$ and $h^f(a)$) and probabilities that individuals aged α live with heads aged *a* (that is $\pi(\alpha,a)$), and to assume that these probabilities will remain fixed over time. However, there is evidence that the traditional Taiwanese pattern of co-residence of older people with their adult children is already breaking down, see for example Lo (1987). Greater independence among the elderly could be due to increased wealth, which makes independent living possible, or to lower fertility rates, which imply that elderly people have fewer adult children with whom they can live. Changes in headship and living arrangements may alter observed age-consumption profiles, affecting measures of inequality between cohorts. Without a theory of how living arrangements are determined, it is difficult to say whether changes in living arrangements will strengthen or weaken the negative correlation between population growth and inequality that we predict.

Acknowledgements

We are grateful to Tim Besley, Colin Hargreaves, James Trussell, Charles Westoff and an anonymous reviewer for helpful discussions and comments, and to the Russell Sage Foundation and the National Institute of Aging for their support.

References

Chu, C.Y.C. and Koo, Hui-Wen, 1990. 'Intergenerational income-group mobility and differential fertility', *American Economic Review*, 80(5):1125–38.

Coale, A.J., 1972. *The Growth and Structure of Human Populations: mathematical investigation*, Princeton University Press, Princeton.

Deaton, A.S. and Paxson, C.H., 1993. *Saving, Growth, and Aging in Taiwan*, National Bureau of Economic Research Working Paper No. 4330, NBER Cambridge, Massachussets.

— and —, 1994. 'Intertemporal choice and inequality', *Journal of Political Economy*, 102(3):437–67.

Keyfitz, N. and Flieger, W., 1990. *World Population Growth and Aging: demographic trends in the late twentieth century*, University of Chicago Press, Chicago.

Lam, D., 1984. 'The variance of population characteristics in stable populations, with applications to the distribution of income', *Population Studies*, 38(March):117–27.

—, 1987. 'Distribution issues in the relation between population growth and economic development', in D.G. Johnson and R.D. Lee (eds), *Population Growth and Economic Development: issues and evidence*, University of Wisconsin Press, Madison:589–627.

Lee, R.D. and Lapkoff, S., 1988. 'Intergenerational flows of time and goods: consequences of slowing population growth', *Journal of Political Economy*, 96(3):618–51.

—, 1993. The formal demography of population aging, transfers, and the economic life cycle, University of California, Berkeley (unpublished).

Lo, J. Chi-chiung, 1987. 'The changing patterns of household structure and economic status of the elderly: 1976 to 1985', *Conference on Economic Development and Social Welfare in Taiwan* (I), The Institute of Economics, Academia Sinica, Taipei.

Preston, S., 1982. 'Relations between individual life cycles and population characteristics', *American Sociological Review*, 47(2):253–64.

Republic of China, 1990. *Report on the Survey of Personal Income Distribution in Taiwan Area of The Republic of China*, Directorate General of Budget, Accounting and Statistics Executive Yuan, Taipei.

—, various issues. *Survey of Personal Income Distribution*, Directorate General of Budget, Accounting and Statistics, Taipei (data tapes).