NBER WORKING PAPER SERIES

INTERTEMPORAL CHOICE AND INEQUALITY

Angus Deaton

Christina Paxson

Working Paper No. 4328

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1993

We are grateful to Orazio Attanasio and Richard Blundell for providing us with the cohort data for the US and Great Britain respectively. We should also like to thank Tim Besley, Anne Case, Ben Eden, Steve Davis, Marjorie Flavin, Alan Krueger, Robert Lucas, Joe Stiglitz and participants in the NBER's February 1993 Economic Fluctuations Research Meeting for helpful suggestions. We should like to thank the Lynde and Harry Bradley Foundation for financial support, and David Wise and the NBER Aging Program for starting us thinking about these issues. The first draft of this paper bore the title 'Intertemporal Choice and Consumption Inequality.' This paper is part of NBER's research programs in Economic Fluctuations and Aging. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4328
April 1993

INTERTEMPORAL CHOICE AND INEQUALITY

ABSTRACT

We show that standard models of intertemporal choice, including the permanent income hypothesis, imply that for any given cohort of people born at the same time, inequality in both consumption and income will grow with age. At any given date, each individual's consumption depends on the integral of unanticipated earnings shocks up to that date, so that consumption becomes more dispersed with time. If earnings are not themselves similarly dispersing, assets will do so, so that the dispersion of total income will increase, irrespective of the behavior of earnings. Because the result applies to an increase in inequality over time within a given age cohort, it has no immediate implications for the behavior of inequality in the economy as a whole, and is consistent with constant aggregate inequality over time. Cohort data are constructed from 11 years of household survey data from the U.S., 22 years from Great Britain, and 14 years from Taiwan. They show that within-cohort consumption and income inequality does indeed grow with age in all three economies, and that the rate of increase is broadly similar in all three.

Angus Deaton
Center of International Studies
Bendheim Hall
Princeton University
Princeton, NJ 08544
and NBER

Christina Paxon
Woodrow Wilson School
Princeton University
Princeton, NJ 08544

0. Introduction

This paper is concerned with the proposition that, for any fixed membership cohort of consumers, the cross-sectional dispersions of consumption and income are increasing over time. Such a proposition is perhaps not very surprising in a world in which the dispersion of earnings is increasing over time, but it holds more generally. The consumption of each individual depends on lifetime earnings, so that, even if individual earnings are stationary, consumption at any age depends on the integral of all unanticipated income shocks up to that age. As a result, if there is some degree of independence in the shocks experienced by different individuals, dispersion within the group will be increasing over time. While the dispersion of earnings does not change in such a case, the dispersion of assets will be increasing during the working years of the life-cycle, and this implies that there will also be widening inequality in the distribution of total income, defined as the sum of earnings and asset income. These results do not imply that consumption and income inequality will increase for society as a whole. Inequality is greater among older cohorts, and less among the young, but since young people are continually replacing the old, there is no presumption that overall dispersion should increase. However, if people live for ever, or if households do so through infinitely lived dynasties, then society should experience ever-increasing inequality.

The fact that the permanent income hypothesis implies that the variance of consumption should increase over time was originally proposed by Benjamin Eden in correspondence with one of us, and some of the more general propositions are implicit in Eden (1980). In Section 1, we start from Eden's insight and lay the theoretical groundwork, starting from the case of independent random walks. We show in what sense the result continues to hold when individual consumption innovations are not assumed to be independent, and we show that the result is more general than the certainty equivalence assumptions required to justify random walk consumption. We also discuss the role of preference shifters. Most importantly, we show that the examination of whether or not consumption dispersion increases with age can be useful for testing alternative theories of consumption. If there is perfect consumption insurance—effectively complete contingent markets—consumption dispersion within age cohorts will be constant over time, while if consumers are liquidity constrained, the cross-sectional dispersion of consumption should track the cross-sectional dispersion of income over time. These implications are quite different from those of unconstrained autarkic intertemporal optimization, and are readily examined on the data.

This taxonomy linking market structure to the dynamics of inequality is closely related to that in Lucas (1992) who explores a model in which endowments are constant, but in which consumers experience identical and independently distributed shocks to their marginal utilities. Although Lucas' framework is designed more to cast light on theoretical issues than are the empirically directed models of consumption and saving with which we are concerned here, the dynamics of inequality are essentially the same. In particular, Lucas shows that in a market version of his model where consumers have access to a credit market, there is ever growing inequality in consumption and wealth, while if there is a cash-in-advance (liquidity) constraint, there exists a unique stationary distribution of consumption, so that inequality is constant.

Section 2 is concerned with a more detailed examination of the special case of the permanent income hypothesis. The availability of closed-form solutions for the permanent income model allows us to be much more precise about the life-cycle dynamics of consumption and income inequality. We can also derive formulas for the behavior of saving and assets, and see how retirement saving influences the spread of inequality.

Section 3 contains empirical evidence for three countries using time-series of household survey data for each. We use data from the US Consumer Expenditure Survey (CEX) for 1980–90, from the Surveys of Personal Income Distribution in Taiwan for 1976–90, and from the British Family Expenditure Survey for 1979–90. Taiwan is included because it is a high saving and rapidly growing economy, with consistently high real interest rates, where household saving behavior is likely to be very different from the US or the UK, where saving rates are low, where few households possess significant financial assets, and where borrowing constraints are likely to be prevalent. For the three countries, we use constructed data on age-cohorts of individuals, using successive surveys to track the evolution of their consumption and income levels over time. The data for all three countries is in accord with the predictions of intertemporal optimization theory; within-cohort consumption becomes more dispersed over time for all three countries, and its rate of dispersion is similar in the three cases. These increases in consumption inequality are evidence against models of perfect insurance. However, we also find that the within cohort dispersion of earnings increases with age, so that it is difficult to rule out the possibility that increasing earnings dispersion, combined with liquidity constraints, is responsible for increases in the dispersion of consumption.

Section 4 concludes and draws out a number of further implications of our results.

1. Theory

The permanent income model with independent shocks

Suppose that intertemporal preferences are quadratic, that the real interest rate is constant and equal to the rate of time preference. Then optimal intertemporal choice implies the permanent income hypothesis, and consumption follows a random walk,

$$c_{it} = c_{it-1} + u_{it} \tag{1}$$

where i indexes the individual (or household) and t the time period. From (1), we can take variances *over any set of households in existence at both* $t-1$ *and* t, to give

$$var_t(c) = var_{t-1}(c) + \sigma_t^2 \tag{2}$$

where σ_t^2 is the period t variance of u_{it} and it is assumed that

$$cov(c_{it-1} u_{it}) = 0. \tag{3}$$

Note that (3) is the cross-sectional covariance, not the covariance over time. The latter has to be zero by the random walk property, that for each i, u_{it} is orthogonal to previously known information, including c_{it-1}. Intertemporal allocation theory by itself does not imply (3). For the moment we simply assume that it is true, that consumption innovations are unrelated to consumption levels in the cross-section, and return to the issue below.

Given the permanent income hypothesis, and the orthogonality restriction (3), equation (2) shows that the cross-sectional variance of consumption will grow over time. Indeed, (1) gives the much stronger result, that with independent innovations, the cross-sectional distribution of consumption at t is (second-order) stochastically dominated by the cross-sectional distribution of consumption at $t-1$. The Lorenz curve for consumption at t lies entirely outside of the Lorenz curve at $t-1$. The passage of time continuously spreads out the distribution of consumption, and any measure of consumption inequality that satisfies the principle of transfers will be increasing over time. Note that (1) holds for individuals who are present in both periods t and $t-1$, so that the proposition that consumption inequality increases is only true for a group with fixed membership, not for the population as a whole, where old people are constantly being replaced by new-borns. Only if people live for ever, or if families live for ever through eternal dynasties does (2) imply that consumption inequality should be increasing over time for society as a whole. The fact that most economies often experience substantial periods of unchanging inequality would therefore seem to be good evidence against the validity of the dynastic version of the permanent income hypothesis.

It is important to note that the increasing dispersion of consumption is independent of what is happening to labor income (earnings) and has no implications for earnings inequality. For example, if each household's earnings is a stationary stochastic process, typically differing from household to household, cross-sectional earnings inequality will remain constant while consumption inequality increases. By contrast, if household earnings are integrated processes, such as random walks, earnings inequality will increase over time. Matters are different for asset income and hence for total income, a topic to which we shall return below.

Non-independent consumption innovations

We now consider what happens when we drop the assumption that the covariance in (3) is zero. In much of the microeconometric work on the permanent income and life-cycle models, most notably Hall and Mishkin (1982), it was assumed that the orthogonality over time for each individual would imply a corresponding orthogonality over individuals at a moment in time. The fact that the theory has no such implications was noted by Chamberlain (1984) and emphasized by Hayashi (1987), and has recently been seen as a serious potential source of difficulty in using short panel data to test intertemporal allocation theory, see particularly Mariger and Shaw (1990) and Deaton (1992: Ch. 5).

To see the difficulty, consider a simple example in which the consumption innovation can be written in the form

$$u_{it} = \varepsilon_{it} + \alpha_i w_t \tag{4}$$

where w_t is an i.i.d. macro shock which affects different households differently according to the parameters α_i, and ε_{it} is a component that is i.i.d. both over time and individuals. Evaluating the covariance in (3) gives

$$cov(u_{it} c_{it-1}) = \sigma_\alpha^2 \sum_1^\infty w_t w_{t-k}. \tag{5}$$

This is a random variable that will take on different values in different time periods. However, the expectation of (5) over time is zero, and this result generalizes under appropriate assumptions.

Write θ_t for the covariance in (3), so that

$$\begin{aligned}\theta_t &= E\big((c_{it-1}-c_{t-1})(u_{it}-u_t)|t\big)\\ &= E(c_{it-1}u_{it}|t) - c_{t-1}E(u_{it}|t)\end{aligned} \tag{6}$$

where variables with only t subscripts are averages over households, and where the explicit conditioning on t is to emphasize the fact that these are cross-sectional expectations. Take expectations of (6) over time, and use the law of iterated expectations to give

$$E(\theta_t) = E\{E(c_{it-1}u_{it}|i)\} - E\{E(c_{t-1}u_{it}|i)\} \tag{7}$$

The terms inside the braces relate to individual behavior over time, about which the theory is informative. The first term is zero, because individual innovations are orthogonal to individual lagged consumption levels. The second term will be zero if we make the additional assumption that lagged *macro* consumption is known to each individual. This is a standard assumption in aggregating Euler equations, see particularly Grossman and Shiller (1982), but it is by no means obviously correct, and its failure can perhaps account for some of the rejections of intertemporal choice theory in the macro data, see Goodfriend (1992), Pischke (1991), and Deaton (1992, Ch. 5) for a review.

If we accept that people are informed about aggregate consumption, this analysis shows that, although the dispersion of consumption need not increase in every period, it can be expected to do so on average.

Relaxing certainty equivalence

As always, the algebra of consumption is most tractable under certainty equivalence when the permanent income hypothesis holds, but the result that consumption dispersion will increase is more general than the permanent income hypothesis. With intertemporally additive preferences, and a constant rate of time preference, the evolution of consumption is controlled by the Euler equation

$$(1+r_{t+1})\lambda(c_{it+1}) = (1+\delta)\lambda(c_{it}) + u_{it} \tag{8}$$

where $\lambda(c)$ is the instantaneous marginal utility (felicity) of consumption function and r_{t+1} is the real rate of interest from t to $t+1$. If impatience is large enough so that $\delta \geq r_{t+1}$, then the marginal utilities of consumption become more dispersed over time, in every period if the innovations are independent of lagged values in the cross-section, otherwise on average. When the rate of interest is greater than the rate of time preference, the distribution of marginal utilities in the cross-section can either concentrate or disperse, depending on the cross-sectional dispersion of the innovations.

The relationship between the dispersion of consumption and the dispersion of marginal utilities depends on the function $\lambda(.)$. If the distribution of marginal utilities in t is stochastically dominated by the distribution of marginal utilities in $t-1$, the same will be true for the distribution of consumption if the inverse function $\lambda^{-1}(.)$ is *concave.* Since $\lambda(.)$ is monotone declining, $\lambda^{-1}(.)$ will be concave if and only if $\lambda(.)$ is concave. We therefore have the result that the distribution of consumption in t stochastically dominates the distribution of consumption in $t+1$ if $\delta \geq r_{t+1}$ and if the marginal utility of consumption function is concave. The condition that $\lambda(.)$ be *convex* is what generates a precautionary motive for saving, and has received a good deal of recent attention as a potential explanation for a number of empirical phenomena. The concavity of marginal utility generates the opposite of precautionary saving, an 'anti-precautionary' motive, and might seem to be implausible. That a strong enough precautionary motive should inhibit the spread of inequality should not be surprising; precautionary saving is motivated by the desire to minimize future consumption variability. As a result, 'prudent' consumers will avoid the risks that are the fundamental cause of increases in inequality.

Even so, the impatience and concavity conditions, although sufficient to guarantee increasing dispersion, are certainly not necessary. Consider, for example, the case of isoelastic utility where, with no uncertainty and identical preferences, each household's consumption will be growing at the same (positive or negative) rate. The variance of log consumption is then constant over time, while the variance of consumption levels is increasing or decreasing as the growth rate is positive or negative. With positive growth, the dispersion of consumption is increasing although that of marginal utility is decreasing. The presence of idiosyncratic uncertainty adds uncorrelated innovations to the Euler equation, and will make it more likely that both marginal utilities and consumption disperse over time.

The results obtained above hold if consumers have identical preferences and identical innovation variances in earnings. Of course, there may be heterogeneity across individuals in the rate of consumption growth, due to differences in the rate of time preference, the degree of risk aversion, or future uncertainty in consumption growth. Those with lower rates of time preference, higher degrees of risk aversion, and greater future uncertainty, will experience greater average consumption growth, see for example Deaton (1992, Ch. 2). It may be supposed that heterogeneity will tend to result in additional consumption dispersion over time, and it is certainly capable of doing so. However, it is also possible to construct examples where heterogeneity inhibits the spread of dispersion. For example, suppose that factors that result in higher consumption growth are inversely correlated with initial consumption levels. High average consumption growth among low-consumption individuals and low average consumption growth among high-consumption individuals could result in reductions in total consumption dispersion, even if the dispersion in consumption within each group were increasing.

Preference shifters

It is clear that all of the previous results can be undone if we allow arbitrary variations in subutility functions from one period to the next. The more interesting question is whether or not taste induced changes in the dispersion of consumption are likely to be large compared with those induced by the resolution of idiosyncratic uncertainty over time. Changes in household composition are the most obvious issue, and certainly so in the empirical work that follows. As household size expands and contracts through the life-cycle of the household, we might also expect the distribution of household compositions to expand and contract, and with it the distribution of consumption levels. Such effects can readily be allowed for informally in interpreting the largely graphical evidence that we shall present. Alternatively, simple parametric models can be used.

Consider, for example, modifying the subutility function $\upsilon(c)$ to allow for a vector of family composition variables a by writing household (sub)utility as $\phi(a)\upsilon\{c/\phi(a)\}$, so that total welfare is the utility of consumption per 'equivalent' multiplied by the number of equivalents. Then the Euler equation implies that

$$\lambda\left(\frac{c_{it}}{\phi(a_{it})}\right) = \sum_{0}^{t}\left(\frac{1+\delta}{1+r}\right)^{k} u_{it-k} \tag{9}$$

where $\lambda(.)=\upsilon'(.)$, and where we have specialized to the case of a constant real interest rate. In this case, all the previous results about consumption and its dispersion apply to consumption per equivalent and its dispersion. If we have a reasonable method for computing equivalents, the propositions can be checked directly. But consider again the certainty equivalent case, where λ is linear, so that (9) implies

$$\ln c_{it} = \ln\phi(a_{it}) + \ln z_{it} \tag{10}$$

where z_{it} is an integrated random variable whose variance is increasing over time. The variance of the demographic term on the right hand side of (10) is likely to increase during the child-rearing phase of the life-cycle and decrease thereafter, and there will generally also be time-varying covariance between the demographics and the integrated innovations. But the demographic variation follows some well-defined pattern, while the variance of $\ln z_{it}$ will be continuously increasing if the innovation variances are sufficiently large, and so will eventually swamp the other terms. Such examples suggest the empirical procedure that we follow, which is to regress measures of consumption inequality on age, but with allowance for the effects of changing demographic composition.

Alternative hypotheses

The proposition that optimal intertemporal allocation implies increasing consumption inequality would not be very interesting if other models of consumption had the same implication. Fortunately, such is not the case.

Consider first a model in which consumers are liquidity constrained. The simplest case is where consumption is equal to earnings, there are no savings and no assets. In this case, the cross-sectional distribution of consumption is the same as the cross-sectional distribution of earnings, and the two must move together with age. Of course, this is too simple. An inability to borrow does not imply an inability to save, and even those who cannot borrow and have no wish to accumulate will typically have a small working balance of cash on hand that acts as a reservoir for emergencies and smooths over short-term fluctuations in earnings. The literature on such 'buffer-stock' models is reviewed in Deaton (1991); as in Lucas (1978, 1992) liquidity constraints and stationary earnings will often guarantee the existence of an invariant distribution for consumption in which inequality is constant. Even without borrowing opportunities and with small amounts of assets on average, buffering can provide consumption streams that are a good deal less variable than earnings. Such models untie consumption from income in the short-run but not in the long-run, and they do not permit large deviations from income in the cross section. They are therefore consistent with the finding that the cross-sectional variance of consumption is less than the cross-sectional variance of earnings, because consumption will be higher than earnings for those with temporarily low earnings, and lower than earnings for those with temporarily high earnings, just as in the permanent income story. However, if the dispersion of earnings changes systematically with age, as for example in Mincer (1974) where different individuals choose different amounts of post-school training, or because the (constant) distribution of abilities is only slowly revealed through screening, then the dispersion of consumption will similarly change in a population of liquidity constrained consumers. Liquidity constraints thus have the implication that the consumption inequality should change with earnings inequality, so that if we can identify ages over which earnings inequality falls within cohorts, we have a sharp test of liquidity constraints.

Another model of consumption that has received a good deal of attention in the recent literature is one in which there exists a complete set of state-contingent securities so that agents can insure themselves and remove idiosyncratic consumption risk. Such models have been tested for the US by Mace (1991), Cochrane (1991), Altonji, Hayashi, and Kotlikoff (1992) and Hayashi, Altonji, and Kotlikoff (1991), with almost uniformly negative results—see Nelson (1992) and Attanasio and Weber (1992) for a contradiction of Mace's results using the same data—and rather more successfully for Indian villages by Townsend (1991). A complete set of contingent markets implies that the marginal utilities of money $\lambda(c_{it})$ have the factor structure

$$\lambda(c_{it}) = \omega_i \gamma_t \qquad (14)$$

see for example Altug and Miller (1990), Hayashi, Altonji and Kotlikoff (1991) or Deaton (1992, Ch. 1) for an exposition. Given (14), the distribution of log marginal utility is constant over time, as would be the distribution of log consumption if preferences are isoelastic. As one would expect, under complete insurance, individual consumptions move together, and although particular functional forms will imply particular patterns for specific measures of inequality, there is no general increase in dispersion over time. One case where this might not be true has been drawn to our attention by Steve Davis. If preferences are not strongly intra-temporally separable between leisure and goods, so that leisure also appears on the left-hand side of (14), and if individual wage rates (marginal productivities) diverge over time—for example because different people learn from experience at different rates—then it is socially efficient for the most productive to work more hours, and to be compensated by having greater consumption. Even in the full insurance equilibrium, consumption inequality will grow if wage dispersion does so.

Models of complete insurance take no account of the problems of information and moral hazard that are likely to prevent their practical realization, and there is a literature that enquires into optimal intertemporal consumption schemes with limited information. In these normative—and typically non-decentralizable—models, such as Thomas and Worrell (1990) and Atkeson and Lucas (1992), the (constrained) socially optimal arrangements involve increasing inequality over time, as in our descriptive model, and in contrast to the case of complete insurance.

2. Consumption inequality and the permanent income hypothesis

In the special case of the permanent income hypothesis, it is possible to derive explicit expressions for consumption and for consumption innovations, and we can use these formulas to be more precise about how consumption dispersion changes for this specific but important special case. The relative simplicity of the algebra also allows us to characterize the behavior of capital income, so that we can we can explore changes with age in the inequality of total income. In this case, there also exists a straightforward methodology for dealing with preference shifters.

Consumption inequality

When subutility functions are quadratic, and when the real interest rate is constant and equal to the rate of time preference, the Euler equation implies that

$$E_t c_{t+k} = c_t. \tag{11}$$

If terminal assets are zero, then *ex post* consumption must also satisfy the stochastic intertemporal budget constraint

$$\sum_0^{T-t}(1+r)^{-k}c_{t+k} = A_t + \sum_0^{R-t}(1+r)^{-k}y_{t+k} \tag{12}$$

where A_t is the current value of assets, and y_{t+k} is *labor income* (earnings) in period $t+k$, R is the date of retirement, and T the date of death. To deal with the annuity formulas that repeatedly occur in this finite life version of the PIH, it is useful to introduce the quantity

$$\beta_t = 1 - \left(\frac{1}{1+r}\right)^{T-t+1} \tag{13}$$

β_t will be unity when T is infinite, but in finite life problems is a concave monotone declining function of t, with values close to unity at the beginning of life and declining to $r/(1+r)$ at death

If the expectations operator is passed through (12), an explicit solution for c_t can be obtained, which is the finite life version of the permanent income hypothesis (or equivalently the stochastic version of the 'stripped-down' life-cycle model, Modigliani, 1986),

$$\beta_t c_t = \frac{r}{1+r} A_t + \frac{r}{1+r} \sum_{k=0}^{R-t} (1+r)^{-k} E_t y_{t+k} \tag{14}$$

which, apart from the factor β_t, is the standard PIH, see Flavin (1981). Equation (14) can be recast into 'innovation' form by substituting for assets using the budget constraint

$$A_t = (1+r)(A_{t-1} + y_{t-1} - c_{t-1}) \tag{15}$$

and subtracting (14) lagged one period. The result can be rewritten as

$$\beta_t \Delta c_t = \eta_t \tag{16}$$

where η_t is an innovation, related to the innovations in earnings by

$$\eta_t = \frac{r}{1+r} \sum_{k=0}^{R-t} (1+r)^{-k} (E_t - E_{t-1}) y_{t+k}. \tag{17}$$

It is useful to rewrite (16) in terms of the history of the innovations, so that

$$c_t = c_0 + \sum_{\tau=0}^{t} \beta_\tau^{-1} \eta_\tau, \tag{18}$$

and, since the innovations are serially uncorrelated, we have

$$var(c_t) = var(c_0) + \sum_{\tau=0}^{t} \beta_\tau^{-2}\sigma_{\eta\tau}^2 \tag{19}$$

where $\sigma_{\eta t}^2$ is the variance of age t's consumption innovation.

This analysis allows us to deduce a number of useful facts. First, if there are no earnings after retirement, there can be no further innovations in consumption, and there will be no further dispersion in the cross-section distribution of consumption. According to (19), the variance of consumption will grow until retirement, but remain constant at $var(c_R)$ thereafter. Of course, in any given cohort, not everyone will retire on the same date, and even with no earnings uncertainty, there are likely to be idiosyncratic surprises in rates of return that are not captured in the permanent income model, and which will result in idiosyncratic innovations to consumption. Even so, it seems reasonable to expect that the rate of growth of consumption inequality will at least slow down among old cohorts.

Equations (16) through (19) also characterize how consumers at different ages respond to innovations, and thus how the dispersion of consumption across members of a cohort changes with the age of the cohort. Equation (16) shows that, for any given innovation η_t, the size of the consumption response Δc_t will increase with age, and, since β_t^{-1} is convex in t, will do so at an increasing rate. However, the relationship between η_t and the underlying innovations in earnings is also a function of age, and will depend on the persistence of shocks to earnings. For example, if earnings are white noise, $\eta_t = r\varepsilon_t/(1+r)$, the consumption innovation is the annuity value of the income innovation independently of age, so that, according to (19), the cross-sectional variance of consumption will be a convex function of age. With white noise earnings, shocks have no implications for the future, but are spread over a larger number of years for younger consumers. If by contrast, earnings shocks are persistent, a given earnings innovation may imply a larger consumption innovation for younger consumers, since any effects of a current shock that would be received after retirement are lost. A case with a good deal of empirical relevance is that where the *first difference* of earnings is a first-order moving average $\varepsilon_t - \theta\varepsilon_{t-1}$, with $1 > \theta > 0$. which occurs if earnings is the sum of a random walk and white noise. In this case (16) and (17) imply that

$$\Delta c_t = \beta_t^{-1}\left(\beta_t^R(1-\theta) + \theta r(1+r)^{-1}\right)\varepsilon_t \qquad (20)$$

where β_t^R is given by (13) with R in place of T. The multiplier in (20) decreases with age if

$$\frac{1}{\theta} > 1 + \frac{r}{(1+r)\left((1+r)^{T-R} - 1\right)} \qquad (21)$$

which will hold for reasonable values of r, T, and R unless θ is very close to 1. Hence, unless the unit root has a very low share of the innovation variance in earnings, we can expect the variance of consumption to increase more rapidly at younger ages, and thus to be concave in age. Older households are closer to retirement and so get the benefits of an earnings innovation for fewer periods; they therefore spend less of it and save more.

Income inequality

The assumptions of the permanent income hypothesis also enable us to be explicit about the behavior of capital income and thus of total (or disposable) income, the sum of earnings and capital income. In the infinite horizon permanent income model, consumption and disposable income are cointegrated under the assumption that the first-difference of earnings is stationary, an assumption that allows for earnings to be either stationary or $I(1)$, see Campbell (1987) and Stock and West (1988). When earnings are non-stationary with idiosyncratic innovations, there will be increasing dispersion in earnings over time. However, the cointegration result shows that, at least in the infinite horizon context, the inequality of *disposable* income must increase over time, whether or not earnings are dispersing. In the finite-life case considered here, matters are a little more complicated.

Define total (or disposable) income y_t^d by

$$y_t^d = \frac{r}{1+r} A_t + y_t, \tag{22}$$

which is the sum of asset income and earnings. It is convenient to define 'saving,' not as the difference between disposable income and consumption, but as

$$s_t = y_t^d - \beta_t c_t. \tag{23}$$

We are interested less in saving than in the behavior of total income, and nothing would be lost by using some other label for (23). If (22) is used to substitute for disposable income in (23), and (14) for $\beta_t c_t$, some algebra gives the finite-life version of Campbell's (1987) 'rainy day' equation

$$s_t = \sum_{k=1}^{R+1-t} (1+r)^{-k} E_t(-\Delta y_{t+k}) \tag{24}$$

whereby saving is the discounted present value of expected future falls in earnings. Note that the expected falls in (24) include the fall in earnings that will take place immediately after retirement; indeed even if earnings are constant or expected to remain so (a random walk) there will still be a role for saving. The combination of (24) and (18) yields

$$y_t^d = \beta_t [c_0 + \sum_{\tau=1}^{t} \beta_\tau^{-1} \eta_\tau] - \sum_{k=1}^{R+1-t} (1+r)^{-k} E_t \Delta y_{t+k} \tag{25}$$

which, with (17), links disposable income to the underlying earned income process.

The simplest case is again where earnings is white noise with mean μ and innovations ε_t, so that, by (24), saving is $\mu/(1+r)^{R+1-t} + \varepsilon_t/(1+r)$. The first term, which is saving for retirement, is the same for all members of the cohort with the same μ, so that the cohort variance of disposable income is the cohort variance of consumption plus a constant. If there is variance in μ in the cross-section, the variance of saving and of disposable income will contain a term $(1+r)^{-2(R+1-t)} var\mu$ which also grows with t. This analysis clearly extends to any other case where the earnings process is stationary, and in all these cases, income inequality grows with consumption inequality, although the inequality of earnings does not change over time.

When earnings is an integrated process, so will be the time t expectation of earnings in R, of the post-retirement fall in earnings, and thus savings as given by (24). Disposable income will then be the sum of two integrated processes, each with positive weights on the earnings innovations, so that, once again the variance of disposable income will grow until retirement, but now more rapidly than the rate at which the variance of consumption expands. Assets are the sum of previous saving and so will be an $I(2)$ process, whose cross-sectional variance will therefore expand more rapidly than that of either consumption or disposable income.

After retirement, saving as defined above is always zero—see again (24)—although there is dissaving as conventionally measured. In consequence, disposable income is equal to consumption multiplied by β_t, so that, post-retirement we have

$$vary_t^d = \beta_t^2 varc_t = \beta_t^2 varc_R \tag{26}$$

so that the variance of income is declining, although the variance of the logs of income will be constant. Consumption inequality is predicted to increase until retirement and then remain constant, while income inequality increases with consumption inequality during the working life, and declines thereafter. In these simple models where everyone in the cohort is assumed to retire at the same date, there is a discontinuity in the variance of income (but not of consumption) at R. In reality, not everyone retires at the same age, and during the transition there is likely to be an increase in variance as people switch from being earners to non-earners. Ultimately however, once all members of the group are retired, and provided assets are dissaved during retirement, income inequality should decrease.

Allowing for life-cycle patterns in earnings and consumption

We note finally that the permanent income hypothesis also permits a straightforward method for dealing with deterministic life-cycle fluctuations in tastes and in earnings. The following is an adaptation of the procedures used in a similar context by Hall and Mishkin (1982).

Suppose that in period t there are deterministic components ϕ_t and ψ_t in consumption and earnings respectively. The former are supposed to capture life-cycle variations in needs and in the composition of the household, and the latter to recognize the occupation-specific lifetime shapes of earnings. We define

$$\tilde{c}_t = c_t - \phi_t; \quad \tilde{y}_t = y_t - \psi_t \tag{27}$$

Note that it makes more sense to work with an indeterministic stochastic process for $\tilde{y}_t$ than for y_t, and to assume that the Euler equation generates a random walk for $\tilde{c}_t$ rather than for c_t. The latter implies that additional unanticipated earnings are allocated equally over all remaining periods in the life-cycle once the basic requirements ϕ_t have been met. If (27) is substituted into the budget constraint (12), and the Euler equation used as before, we get a version of (14) that applies to the corrected consumption and earnings quantities which appropriate allowance for the deterministic factors. However, the innovation equation (16) for the change in consumption now holds for $\tilde{c}_t$, so that we finally reach a version of (18) which is

$$c_t = \phi_t + c_0 + \sum_{\tau=1}^{t} \beta_\tau^{-1} \eta_\tau . \tag{28}$$

In the empirical work, we can therefore either 'purge' consumption of the influence of compositional and other variables, as do Hall and Mishkin, and then examine the variances of the residual, or we can take variances of (28), which would suggest allowing for the variance of household types in assessing the evolution of consumption variance.

3. *Evidence*

Introduction

Our general procedure for all three countries is to use successive years of cross-sectional household survey data to follow cohorts of individuals or individual households through time. While none of the survey data are longitudinal, so that it is never possible for us to follow individuals or households over time, we can follow cohorts of people through their randomly selected representatives in each survey. Cohorts are identified by their year of birth, or equivalently, by their age in a base year. Hence, to take Taiwan as an example, we can take the cohort born in 1945, who were 31 years old in our first survey, that of 1976, and we can examine the distribution of consumption, earnings, and incomes for all 31 year olds in the 1976 survey. This distribution is then compared with the distributions of consumption, earnings, and incomes for all 32 year olds in 1977, those for all 33 year olds in 1978, and so on, ending with the distribution for 45 year olds in 1990, the last survey year. Such procedures are for many purposes superior to the use of panel data. Sample attrition is not an issue, and summary statistics from the cohort distributions are likely to be more accurately measured than the individual data that would be used in panels. The propositions of this paper concern the dynamics of the distributions of consumption in successive years, and do not require knowledge of the joint distributions over several years, information that could only come from panel data.

The major difficulty with our procedures is that while income and earnings data are available at the level of the individual, consumption data are available only at the household level. We must therefore examine the evolution in the distribution of consumption over households, and convert income and earnings to the same basis. The difficulty lies in the fact that, unlike individuals, households form and dissolve over time, so that when we track households labelled by the age of the household head, we cannot always be sure that we are sampling from the same population in successive years. In practice, we believe that the difficulties are only likely to be severe among older households, where death, living with children, or institutionalization mean that surviving households are increasingly selected—and presumably increasingly unrepresentative—as the age of the household head increases. There are similar difficulties among very young household heads, who are also unlikely to be representative of their cohort, most of whom are in full-time education or (in Taiwan for males) in the military. Even among middle-aged heads, divorce and remarriage will imply that the population of households with heads aged 40 in one year is not the same as the population of households with heads aged 41 in the next.

An alternative to defining cohorts by the birth year of the household head is to track individuals (whether household heads or not) born in the same year, and to assign to each individual the (per capita or per equivalent adult) consumption and income levels of the household to which he belongs. This alternative has problems of its own. Although selectivity into and out of headship is no longer an issue, individuals who live with family members from different generations may be assigned inappropriate consumption and income levels. Again, these problems are likely to be most severe for young people living with parents, or old people living with children. For Taiwan, where these problems are likely to be more serious, we have the micro data that allow us to construct cohorts based on both household heads and individuals, and we present results for both. However, faced with selection and measurement problems at young and old ages, our reporting emphasizes results that do not depend on these age groups, even though in some cases, the exclusion decreases the power of our tests.

In this paper, we give greatest emphasis to the Taiwanese case. Taiwan is inherently interesting in its own right. Household saving rates are very high, over 20% in the 1980s, and it is therefore implausible that many households are liquidity constrained. With such high saving rates, and the wide ownership of assets, capital income is an important component of household disposable income, and increasingly so with age, see Deaton and Paxson (1992, Figs. 5 and 6). It is therefore *possible* for earnings and disposable income to behave differently, as should be the case if the theory is correct, and if the earnings distribution is not dispersing with age. However, we are also emphasizing Taiwan because we have the individual data that enable us to perform more elaborate calculations. The US data, from 11 years (44 quarters) of the Consumer Expenditure Survey (CEX), were supplied to us by Orazio Attanasio in cohort form, and thus have the advantage of being consistent with his previous work on US saving behavior, Attanasio (1991) and Attanasio and Weber (1992a). Similarly, for Great Britain, we have used the cohort data supplied to us by Richard Blundell based on 22 years (88 quarters) of the Family Expenditure Survey.

Data and construction of cohorts

The Taiwanese data come from the *Personal Income Distribution Surveys* which have been collected each year since 1976. In this paper we use fourteen surveys, covering 1976 through 1990 but excluding 1978. In 1976 and 1977, the sample sizes are a little over 9,000 households, but from 1978, there are over 14,000 households in each survey; there are 50,000 or so persons in the first two years, and around 75,000 later. We exclude data from 1978, which was the first year of the expanded survey, because the income data contain a number of outliers. These outliers do not appear to affect the means, but inflate measures of dispersion to levels that are clearly incompatible with the adjacent years. The survey design is described in Republic of China (1989). For income and consumption there is a single interview, at which questions are asked about major items of income and expenditure in the past year. A separate control sample of a small number of households keeps diaries of all items of income and expenditure throughout the year. These households are regularly visited by field inspectors to control quality, and the results are used to monitor (but not to alter) the results from the main survey. Consumption is defined as total expenditure on all goods (both durable and nondurable); income is after-tax income from all sources. Earnings (measured before tax) are defined to include business and farm income. Note that although this is the best that can be done, and certainly preferable to assuming that all business income is the return to physical assets, the procedure introduces a component of capital income into our measure earnings, something that must be borne in mind when interpreting the results. The age profiles of household incomes, earnings, consumption and saving are presented in Deaton and Paxson (1992), who also describe the changing demographic structure of households in Taiwan.

The Taiwanese surveys are sufficiently large for us to define cohorts for each age. Table 1 lists the number of households in each of five cohorts (those who were 20, 30, 40, 50, and 60 in 1976) that appear in each of the survey years. The 1976 and 1977 figures are low because these surveys are smaller, but for the young and old households also show the selection effects. In the empirical analysis we restrict ourselves to households with heads between the ages of 20 and 75 inclusive. After excluding 1978, we have in total 784 'observations' of cohort-year pairs, each 'observation' consisting of a distribution over the households in that cohort-year.

The US data we use are from the 1980 through 1990 surveys of the CEX. These surveys are smaller than those from Taiwan, covering between 5,700 and 8,300 households per year, with roughly a fourth of households surveyed in each quarter. There are ten cohorts defined by five year age bands, from those aged 21–25 in 1980 to those aged 66–70 in 1980. Households are allocated to the quarter in which they were interviewed so that there are 44 observations on each of the cohorts, or 440 'observations' in all. The numbers of households in each of cohort and (the first quarter of each) year are given in Table 2. These are urban households only; the CEX in 1982 and 1983 had no rural households, and for consistency; we work with urban households in all periods. We are less concerned in this paper with explaining inequality in the US as a whole than with testing the implications of the theory, and for this any fixed group of households will do. There is a good deal less attrition among older households in the US than in Taiwan, because the elderly in the US are more likely to continue as independent households than to move in with their eldest son, as is typical (although decreasingly so) in Taiwan. For the US, household consumption refers to nondurable expenditure, disposable income to the total after-tax income of the household, and earnings to before-tax labor income. The 'age' of a cohort is defined as the midpoint of the 5-year band, and our sample includes people with an age of 23 (e.g. 21-25) in

1980 to those with an age of 78 (e.g. 76-80) in 1990.

The British data come from the Family Expenditure Survey, an annual survey that has been in continuous operation since 1954 collecting data on some 7000 households each year. The sample is representative of Great Britain, which is the United Kingdom excluding Northern Ireland. We use 22 years of data, from 1969 to 1990, split into quarters. We use information for 11 cohorts, again in five year age bands, from those aged 5-9 in 1969 to those aged 55-59 in 1969. We restrict our analysis to those between the ages of 22 and 77 (where 'age' is defined as the mid-point of the 5 year group.) So, for example, the youngest cohort of those aged 5-9 in 1969 is not included in the sample until 1984, when the individuals in this cohort are between the ages of 20 and 25, and are assigned an age of 22. Allowing for the fact that the youngest cohorts are not observed in the earliest years, nor the oldest in the latest, there are 844 quarterly 'observations' on the cohorts. The number of households in each of these cohorts is listed in Table 3. For the British data, consumption excludes durables, earnings is the answer to a question about 'normal weekly earnings', and income is total disposable income.

Results: age profiles of consumption and income inequality

There are a number of different ways of looking at the evidence, but given the very large numbers of observations involved—the Taiwanese data alone cover nearly a million individuals—we make heavy use of graphics, both for the underlying data, and for many of the estimated parameters.

We begin by examining the lifetime profiles of the variances of the logarithms of consumption and total income. There are a number of reasons to work with the variances of logs. Although the theory under the permanent income hypothesis yields predictions about variances, not variances of logs, use of the former will give increasing dispersion even when the consumption level of everyone in the cohort was expanding proportionately, as would be the case in a growing economy with liquidity constraints or as might occur with perfect insurance. A model of autarkic intertemporal allocation with isoelastic preferences also predicts that the variance of log consumption will be constant in the absence of idiosyncratic shocks, and this is a natural baseline from which to look for the dispersion that should occur under individual uncertainty. Note finally that even if the permanent income hypothesis is true so that individual consumption levels follow independent random walks, the distribution of consumption at t will stochastically dominate that at $t-1$, so that any summary measure of inequality that respects the principle of transfers will be increasing over time. Although there are examples where the variance of logs violates the principle, see Atkinson (1970) and Sen (1973), these are theoretical curiosities rather than good reasons to abandon such a convenient measure.

Figure 1 for Taiwan, Figure 2 for the US, and Figure 3 for Britain show life-cycle profiles of consumption and income inequality. The horizontal axis in all these graphs shows age of the head of household, and the vertical axis the variances of logs, and each panel shows the experience of a single cohort. For Taiwan, we show the graphs for every fifth cohort, rather than attempting to show the full set of results. For the youngest cohorts, for example those aged 10 in 1976 in Figure 2, the graphs do not cover the full span of years because such cohorts are only observed when they reach the cutoff age. Similarly the oldest cohorts are not observed once they pass the upper age cutoff. The figures shown in the graphs are not the raw variances, but the fitted values from the regressions (run separately for consumption and for income and for each country) of the variances on dummies for cohorts, ages, and years, or more precisely on a maximally linearly independent set of such dummies.

In all three countries, the inequality of consumption increases with age, as predicted by the theory. Further, and again as the theory requires, the increase in inequality is concentrated in the working years, slowing down or ceasing after retirement. Inequality of total income follows much the same pattern, increasing with age during the working years and remaining constant or falling thereafter. In Taiwan, income variance falls sharply among the oldest households, but given the serious selectivity issues among the old in Taiwan, we do not wish to attribute any great significance to the finding. In the US and Taiwan, but surprisingly not in Britain, income is more unequally distributed than is consumption. There is a good deal of similarity in the life-cycle evolution of inequality in total income and consumption. While such a result is consistent with the theory, it also would occur if consumption were closely tied to income, as in a model of liquidity constrained consumers. The British results, where the two variances are close, and where they track each other closely, are particularly suggestive of such an interpretation. The correlation between the two sets of variances is much weaker in Taiwan, and very much weaker in the US. Note finally that the variances of log income tend to be themselves more variable than the variances of log consumption, particularly in the US, where there seems to have been a sharp increase in income inequality—although not in consumption inequality—in the single year 1986.

Figure 4 presents the same results in a summary form that focusses more precisely on the theoretical predictions, and that allows easy comparison of the three countries. These graphs are derived by regressing the variances of log income and log consumption on age and cohort dummies, again separately for each country, and then graphing the estimated age effects. Unlike the regressions used to derive Figures 1 through 3, no year effects are included. Year, cohort, and age dummies are linearly dependent, and while the dependence has no consequence for deriving fitted values, the identification of any one set requires a normalizing assumption. The issue is discussed in some detail in Deaton and Paxson (1992), who argue for a normalization that, in the current application, generates results that are almost identical to those reported here, which were obtained by the more transparent procedure of excluding the year dummies.

That consumption inequality rises with age is even more obvious in Figure 4 than it was in Figures 1–3. In Taiwan, the profile is convex until age 60 or so, and constant thereafter, suggesting that it is in middle age that most information is received about lifetime prospects, with information accruing only slowly at the beginning of the life-cycle. In Britain too, consumption inequality increases faster in middle age than in youth, although here the spread slows down before ceasing at around age 60. The profile in the US is somewhat different, with an approximately steady increase in consumption inequality throughout life. Perhaps the relatively widespread ownership of assets and of private pension schemes in the US allows dispersion in yields to continue to add to inequality even among elderly households. Because these graphs show age effects from regressions that are normalized to be zero at age 20, they cannot be used to compare the levels of variances across countries. However, the scales are the same, and it is clear that the *increase* in consumption inequality with age is very similar in the US, Taiwan, and Britain. Table 4 shows regression coefficients of the estimated age effects on age; the variance of logs within each cohort increases by 0.07 every decade in the US, by 0.08 in Taiwan, and by 0.10 in Britain. These are large effects relative to differences in inequality over time or between countries. For those used to thinking about inequality in terms of gini coefficients, the Table converts the log variances to ginis on the assumption that the distributions are lognormal. In the US, where consumption inequality is highest at all ages, the age-inequality effect increases the gini from 0.339 at age 25 to 0.411 at age 55, a change that is comparable in magnitude to the recent increase in the inequality in earnings that has attracted a good deal of attention, see Cutler and Katz (1992). In Taiwan and in Britain, the increases are much larger, from 0.234 at age 25 to 0.337 at age 55 in Taiwan, and from 0.234 at age 25 to 0.418 at age 55 in Great Britain. Indeed, it seems probable that the much discussed increase in consumption inequality in Taiwan can be largely accounted for by the aging of a population that has experienced one of the most rapid demographic transitions in history.

Figure 4 also shows the age effects in income inequality. In Britain, we see again what looks like the effects of liquidity constraints, with income inequality and consumption inequality very closely linked. The Taiwanese data again show the dip in income inequality among older households, and only the US shows the rise and fall that would come from the simplest permanent income hypothesis with no bequests. Table 4 shows the estimated time trends, as well as the implied increases in income inequality. Note in particular the very large effects on gini coefficients in Britain. Whether these effects come from the theory as presented here, or from some other cause—such as an increase in earnings inequality with age—the data show a strong relationship between inequality and age, for both consumption inequality and income inequality.

Further results for Taiwan: Lorenz curves and household composition

Figures 5 through 8 all concern Taiwan, where we have direct access to the micro data, and explore the results in more depth. Figures 5 and 6 present Lorenz curves for consumption and income for every fifth cohort in the three years 1976, 1983, and 1990. Lorenz curves are the obvious way of examining stochastic dominance, and show the whole distribution, not just a summary statistic such as the variance of logs. The appearance of Lorenz curves is notoriously insensitive to changes in distribution, but the graphs do show some movement, and in all cases except for income in the oldest cohort, the later years have the Lorenz curves that are further away from the 45–degree line, as predicted by the theory. Again we see that throughout the distribution there is little increase in consumption inequality with age at young ages, although the Lorenz curves are visibly different for cohorts born in 1941 and earlier. It is worth noting that the increase in consumption inequality appears to begin at the bottom and middle of the distribution, but moves up with time, so that among those born before 1921, increases in dispersion between 1976 and 1990 were limited to the upper part of the distribution of consumption.

Figures 7 and 8 turn to the effects of allowing for variations in household size and composition. The left hand panel of Figure 7 shows the age profile of the number of adult equivalents, defined as the number of household members over 14 years of age plus half of the number aged 14 or less. Each continuous line shows the average numbers for a given cohort over time, and the figure shows both the characteristic life-cycle pattern of household size together with the substantial cohort effects associated with the decline in fertility over time. Note that the apparent variability among the older cohorts is a sampling effect, since there are few households in these cohorts over which to calculate the means. The right hand panel shows the coefficient of variation of adult equivalents for the same cohorts. The point to note here is that from ages 40 to 60, although household size is falling for most of the period, the dispersion of household sizes is rising. Some of the rising variability of total household consumption that we have already documented may therefore be attributable, not to the accumulated effects of idiosyncratic uncertainty, but to the fact that, in the critical age range, the dispersion of household sizes is also increasing. Figure 8 shows that this conjecture is correct, in that the rise in consumption inequality with age in Taiwan is less when household size is taken into account. Even so, and apart from a small decline from 40 to 45—which might not be statistically significant, or which might be removable by a more sophisticated treatment of equivalents—the dispersion of 'corrected' consumption still increases with age. The Figure shows three different correction methods. The upper (solid) line shows the variance of the logs of consumption per equivalent. The lower (broken) line shows the variance of logs of consumption per equivalent where cohorts are defined, not by household heads, but by individuals. Finally, the middle (dashed) line shows the variance of the residuals of regressions of the logarithm of household consumption on a range of demographic variables. Although the rate of increase in the variances differs between the methods, the general shapes are very much the same, and once again there is increasing consumption inequality with age.

Earnings dispersion and alternative hypotheses

We turn finally to an examination of the life-cycle behavior of the dispersion of earnings. The theory of this paper predicts increasing inequality of consumption and income independently of whether earnings inequality is increasing, so that there is a sense in which the behavior of earnings is irrelevant. However, if we consider alternative explanations the fact that inequality increases with age, the inequality of earnings becomes a crucial part of the story. An ideal environment for testing the theory would be one in which individual earnings are stationary, so that earnings inequality remains constant over the working life. Intertemporal allocation theory would then predict that consumption and income inequality would increase with age, a prediction that is in sharp contrast to the fixed consumption and income inequality that would result from liquidity constrained behavior, which is the leading alternative hypothesis. However, if earnings inequality increases with age, as would occur if individual earnings processes are nonstationary with independent innovations, consumption and total income inequality will also increase, whether or not consumers are liquidity constrained. Perfect insurance implies that consumption inequality should be constant whatever happens to earnings, but this is not a very interesting hypothesis, and in any case, has already been rejected by the data.

Figure 9 shows the life-cycle patterns of the variances of log earnings and log income for the three countries. As before, these figures plot the estimated age effects from regressions including both age and cohort dummies. Since many households, particularly older households, have income but no earnings, we cannot compute variances of logs for all households, and must restrict our calculations to the subsample of households with at least one earner. As a result, the variances of log earnings for cohorts with older heads are dominated by selectivity, and we display results only up to age 65. Even so, the Taiwanese results suffer from the fact that business income is important, and cannot be satisfactorily split into its capital and labor components. In our calculations we have treated business income as earnings, which is almost certainly its largest component, but hardly provides an adequate basis for testing whether inequality in earnings and total income behave differently. In consequence, although the Figure shows that the variance of the log of earnings is increasing with age, we cannot be sure that the result is not due to the mislabelling of some capital income as earnings.

These problems are less for severe for the US and Britain, and both countries show an increased dispersion of earnings with age during normal working years. The results for the U.S. are consistent with previous examinations of individual earnings variances. Dooley and Gottschalk (1984) use the Current Population Surveys from 1968–79 to show that variances within experience cohorts of men fall for a few years and then rise, a result that is confirmed by our own calculations using later years from the CPS, and that is not inconsistent with the results on the variances of household earnings in Figure 9. Since cohort earnings variances are increasing with age, and while the predictions of increases in inequality of consumption and total income are borne out for all three countries, we cannot use that fact to conclude in favor of life-cycle allocation theory and against the hypothesis of liquidity constraints. Indeed, a simple model of buffering with liquidity constraints would tie the cross-sectional distribution of consumption to the cross-sectional distribution of earnings, although the buffering would mean that the absolute level of dispersion of the former would be less than the latter. The proposition can be examined by testing whether the age effects in the variances of log consumption and log earnings are equal, allowing them to differ in levels, but not in profile. The F-tests for the absence of age effects in the *difference* between the variance of log earnings and log consumption are 3.72 (Taiwan), 1.64 (US) and 10.82 (Britain). For Taiwan and the U.S., these F-tests are unimpressive relative to the sample sizes, although the p-value for Taiwan is very small. The British results appear to provide the most evidence against liquidity constraints. However, the shape of the consumption and earnings profiles differ substantially only at the youngest and oldest ages, indicating that selectivity may be driving the result, and when we restrict the analysis to 25 to 60-year-olds, the F-test is only 3.77. Although the liquidity constraints story makes very little sense for Taiwan, where households save more than a fifth of their incomes, it is not inconsistent with other evidence for the US and Britain.

4. Conclusions and further implications

In this paper we have examined the proposition that consumption inequality should widen over time for a group of individuals who are making optimal intertemporal allocations of resources in an uncertain environment. Although the proposition is most obvious when each individual's consumption follows a random walk with innovations that are independent across people, we have demonstrated that the result will hold under a much wider set of circumstances. Under the special case of the permanent income hypothesis, we have also shown that income inequality should increase with age, even when there is no change in the distribution of earnings. These predictions are borne out in data from the US, Britain, and Taiwan, where both income and consumption inequality increase up to around the age of retirement, and where the latter is constant thereafter (or increasing, for the US) The absolute size of these effects is large, so that the gini coefficient of consumption inequality in the US is 25% larger among 55-year olds than among 25-year olds, and that for income inequality more than 40% larger. The effects appear to be even larger in Britain.

The main contribution of this paper is not the construction of another test of life-cycle theory or of the permanent income hypothesis. The rather limited evidence we have suggests that the dispersion of earnings also increases with age, so that the dispersion of consumption and total income will also increase under a wide range of hypotheses about behavior. Instead, our main purpose is to draw attention first, to the implications of life-cycle theory for inequality, and second, to the fact that inequality does indeed increase with age, at least for the three countries we have examined. Even if we remain uncertain about the connection between the former and the latter, the fact that inequality increases with age has a number of implications, implications that are made the more interesting by their possible basis in intertemporal theory.

First, consider the implications of the results for inequality in society as a whole. If there are no links between successive generations, and if the inequality in the distribution of earnings remains constant, consumption and income inequality will also remain constant in the society as whole, even though inequality is increasing for each age group. Old, unequal, cohorts are constantly dying and being replaced by young cohorts among whom inequality is much less. The theory and evidence are therefore consistent with unchanging inequality in society as a whole, and the data do indeed show extended periods during which inequality is constant, see Levy and Murnane (1992) for the US. If there are strong intergenerational links, where bequests are important, or in the limit where households are organized into eternal dynasties, the theory predicts widening inequality for society as a whole, at least if no new dynasties are ever created. The failure of such predictions would seem to be evidence against the most extreme versions of such models, but there still remains an important research agenda of building bequest motives into the sort of models examined in this paper and drawing out the implications for the transmission of inequality from older to younger cohorts.

Second, the relationship between age and inequality forges a powerful link between demographic change and the distribution of resources, a link that is largely independent of the theoretical basis of the result. Countries such as Taiwan that have experienced a rapid demographic transition are aging rapidly, so that there are now many more 55-year olds relative to 25-year olds than was the case 25 years ago. Since income and consumption are more unequally distributed among the former than the latter, the overall distribution of income and consumption will become more unequal, and this is exactly what has happened in Taiwan over the last decade or so. We plan to examine this question in a good deal more detail in another paper, but our preliminary results suggest that much of the increase in consumption inequality in Taiwan can be accounted for by population aging. Although there has also been aging in the US, we do not believe that the same phenomena can account for very much of the current increase in inequality, though once again, the detailed research remains to be done.

Third, consumption inequality has implications for aggregation, especially for those attempts to model average consumption as the behavior of an intertemporally efficient representative agent. One defence of applying the theory to aggregate data is that, in the absence of distributional change, functions of averages will behave similarly to averages of functions. The ratio of the average of the logarithms of consumption to the logarithm of the average of consumption is a measure of consumption dispersion, so that when dispersion is increasing, the average of logs will behave quite differently from the logarithm of the average. Indeed, results by Attanasio and Weber (1992b) using the British data show that estimates of intertemporal models are quite sensitive to whether or not the aggregation is correctly done.

Fourth, our results have implications for the cross-country relationship between inequality and economic development, first investigated by Kuznets (1965). Kuznets suggested that inequality should rise in the early stages of economic development, but should eventually fall as incomes rose, a result that he explained in terms of the urbanization that typically accompanies economic growth. The empirical validity of this inverted U-shape has been subject to a great deal of controversy, and is far from widely accepted, largely because of the extremely poor quality of international measures of inequality. The results of this paper suggest a different basis for a Kuznets type relationship. Economic development is eventually accompanied by a demographic transition, from high fertility to low fertility, a transition that eventually redistributes population from young to old. Such a redistribution will tend to lead to widening inequality, at least until the new stable population distribution is established, at which point inequality growth will cease. Of course, income growth may itself affect inequality independently of these demographic effects, with higher growth rates resulting in greater inequality across cohorts.

5. List of works cited:-

Altonji, Joseph, Fumio Hayashi, and Laurence J. Kotlikoff, 1992, 'Is the extended family altruistically linked? Direct tests using micro data,' *American Economic Review*, **82**, 1177–98.

Altug, Sumru, and Robert A. Miller, 'Household choices in equilibrium,' *Econometrica*, **58**, 543–70.

Atkeson, Andrew and Robert E. Lucas, 'On efficient distribution with private information,' *Review of Economic Studies*, **59**, 427–53.

Atkinson, Anthony B., 1970, 'On the measurement of inequality,' *Journal of Economic Theory*, **2**, 244–63.

Attanasio, Orazio P., 1991, 'A cohort analysis of saving behavior by US households,' Stanford University (Sept.), mimeo.

—— and Guglielmo Weber, 1992a, 'Consumption growth and excess sensitivity to income: evidence from US micro data,' Stanford University and University College, London (Apr.), mimeo.

—— and —— 1992b, 'Consumption growth, the interest rate and aggregation,' Stanford University and University College, London (June), mimeo. (*Review of Economic Studies*, forthcoming).

Campbell, John Y., 1987, 'Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis,' *Econometrica*, **55**, 1249–73.

Chamberlain, Gary, 1984, 'Panel data,' in Zvi Griliches and Michael D. Intriligator (eds.), *Handbook of Econometrics*, Vol. 2, North-Holland. Amsterdam, 1247–318

Cochrane, John H., 1991, 'A simple test of consumption insurance,' *Journal of Political Economy*, **99**, 957–76.

Cutler, David M. and Lawrence F. Katz, 1991 'Macroeconomic performance and the disadvantaged,' *Brookings Papers on Economic Activity*, 2, 1–74.

Deaton, Angus, 1991, 'Saving and liquidity constraints,' *Econometrica*, **59**, 1221–48.

—— 1992, *Understanding consumption*, Oxford. Oxford University Press.

—— and Christina H. Paxson, 1992, 'Saving, growth, and aging in Taiwan,' Research program in development studies, Princeton (June), mimeo.

Dooley, Martin D., and Peter Gottschalk, 1984, 'Earnings inequality among males in the United States: trends and the effects of labor force growth,' *Journal of Political Economy*, **92**, 59–89.

Eden, Benjamin, 1980, 'Stochastic dominance in human capital,' *Journal of Political Economy*, **88**, 135–45.

Flavin, Marjorie, 1981, 'The adjustment of consumption to changing expectations about future income,' *Journal of Political Economy*, **89**, 974–1009.

Goodfriend, Marvin, 1992, 'Information aggregation bias,' *American Economic Review*, **82**, 508–19.

Grossman, Sanford J. and Robert J. Shiller, 1982, 'Consumption correlatedness and risk measurement in economies with nontraded assets and heterogeneous information,' *Journal of Financial Economics*, **10**, 195–210.

Hall, Robert E. and Frederic S. Mishkin, 1982, 'The sensitivity of consumption to transitory income: estimates from panel data on households,' *Econometrica*, **50**, 461–81.

Hayashi, Fumio, 1987, 'Tests for liquidity constraints: a critical survey and some new observations,' in Truman F. Bewley (ed.), *Advances in Econometrics: Fifth World Congress*, Vol. 2, 91–120.

——, Joseph Altonji, and Laurence J. Kotlikoff, 1991, 'Risk-sharing, altruism, and the factor structure of consumption,' NBER working paper no. 3834 (Sept.), mimeo.

Levy, Frank and Richard J. Murnane, 1992, 'U.S. earnings levels and earnings inequality: a review of recent trends and proposed explanations,' *Journal of Economic Literature*, **30**, 1333–81.

Lucas, Robert E., 1978, 'Equilibrium in a pure currency economy,' *Economic Inquiry*, **18**, 203–20.

—— 1992, 'On efficiency and distribution,' *Economic Journal*, **102**, 233–47.

Mace, Barbara J., 1991, 'Full insurance in the presence of aggregate uncertainty,' *Journal of Political Economy*, **99**, 928–56.

Mariger, Randall P., and Kathryn Shaw, 1990, 'Unanticipated aggregate disturbances and tests of the life-cycle model using panel data,' Board of Governors of the Federal Reserve and Carnegie Mellon University, mimeo.

Mincer, Jacob A., 1974, *Schooling, experience, and earnings*, New York. Columbia University Press.

Modigliani, Franco, 1986, 'Life cycle, individual thrift, and the wealth of nations,' *American Economic Review*, **76**, 297–313.

Nelson, Julie, 1992, 'On testing full insurance using consumer expenditure survey data,' University of California, Davis, processed. (Jan.)

Pischke, Jörn-Steffen, 1991, 'Individual income, incomplete information, and aggregate consumption,' Industrial Relations Section working paper no. 289, Princeton University, mimeo.

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

Sen, Amartya K., 1973, *On economic inequality*, Oxford. Clarendon Press.

Stock, James, and Kenneth D. West, 1988, 'Integrated regressors and tests of the permanent income hypothesis,' *Journal of Monetary Economics*, **21**, 85–95.

Thomas, Jonathan, and Tim Worrell, 1990, 'Income fluctuation and asymmetric information: an example of a repeated principal–agent problem,' *Journal of Economic Theory*, **51**, 367–90.

Townsend, Robert M., 1991, 'Risk and insurance in village India,' NORC discussion paper no. 91–3 (May), mimeo.

Table 1: Numbers of households in each cohort, by year, Taiwan Survey of Personal Income Distribution

	cohort 1	cohort 11	cohort 21	cohort 31	cohort 41
age in 1976	20	30	40	50	60
1976	26	203	304	275	113
1977	35	274	280	237	71
1978	70	401	380	344	106
1979	133	352	383	300	88
1980	188	393	359	329	82
1981	261	375	372	318	87
1982	363	372	370	327	74
1983	394	364	376	278	71
1984	463	410	347	237	70
1985	519	400	296	242	48
1986	543	398	303	222	53
1987	563	367	281	191	65
1988	604	390	260	208	49
1989	636	364	256	184	41
1990	640	299	228	196	59

Table 2: numbers of households in each cohort, by year, US Consumer Expenditure Survey

	Cohort									
age in 1980	21–5	26—30	31–5	36–40	41–5	46–50	51–5	56–60	61–5	66–70
1980	355	378	357	261	246	221	245	280	219	212
1981	478	459	376	324	247	260	247	233	255	211
1982	451	458	450	324	235	241	285	244	242	214
1983	476	486	412	312	244	259	297	271	263	205
1984	484	494	443	359	270	246	297	268	247	207
1985	517	465	465	342	263	258	256	283	241	199
1986	584	530	448	373	311	309	308	298	285	191
1987	534	516	442	340	305	325	281	304	232	200
1988	453	418	412	300	226	229	257	250	223	170
1989	491	421	415	303	266	252	237	275	208	147
1990	515	440	417	320	275	198	263	236	199	137

Table 3: Numbers of households in each cohort by year, Great Britain, Family Expenditure Survey

age in 1969	5–9	10–4	15–9	20–4	25–9	30–4	35–9	40–4	45–9	50–4	55–9
1969	.	.	.	72	127	159	147	180	170	156	168
1970	.	.	.	93	139	140	180	153	174	119	159
1971	.	.	.	144	158	163	150	157	169	127	174
1972	.	.	.	143	191	144	173	153	192	138	162
1973	.	.	44	163	143	136	141	130	198	136	171
1974	.	.	58	143	144	115	134	115	135	119	167
1975	.	.	97	174	148	154	161	149	167	153	187
1976	.	.	115	185	134	137	157	149	190	132	149
1977	.	.	123	184	161	147	146	141	159	147	189
1978	.	54	149	197	141	150	130	170	153	138	146
1979	.	70	172	181	146	138	129	131	159	120	160
1980	.	98	148	192	173	141	126	115	161	136	142
1981	.	102	149	221	180	149	148	149	155	159	138
1982	.	134	176	172	142	170	148	156	178	139	164
1983	56	116	173	201	149	141	121	117	148	130	127
1984	81	139	182	162	156	127	135	146	176	124	133
1985	69	142	160	189	142	159	135	144	150	134	144
1986	124	167	171	159	151	138	124	116	177	126	127
1987	159	185	185	196	144	154	136	155	172	115	106
1988	169	151	189	197	162	133	125	143	164	121	.
1989	168	155	164	179	147	142	133	156	172	134	.
1990	162	192	149	182	147	117	139	151	142	121	.

Table 4: Summary statistics for age effects

Taiwan	Coefficient on age	Variance of log Age=25	Variance of log Age=55	Gini coefficent Age=25	Gini coefficent Age=55
Consumption	.0084 (18.86)	.173	.368	.234	.337
Income	.0027 (3.46)	.385	.445	.344	.368
Consumption/adult equivalent (household data)	.0021 (14.74)	.184	.225	.241	.266
Consumption, demographics regressed out	.0015 (14.81)	.099	.133	.178	.206
Consumption/adult equivalent (person data)	.00068 (6.63)	.188	.211	.244	.258
United States					
Consumption	.0074 (30.02)	.373	.564	.339	.411
Income	.0126 (9.72)	.598	1.370	.422	.599
Britain					
Consumption	.0102 (22.95)	.173	.586	.234	.418
Income	.0136 (25.81)	.115	.640	.191	.435

Notes: The first column shows the regression coefficient of the estimated age effects on age (t-statistics are in parentheses). The variances of logs shown in the second and third columns are predicted values from regressions of variances on age and cohort effects. The predicted variances at the two ages (25 and 55) hold cohort constant. The Gini coefficients shown in the fourth and fifth columns are based on the variances in columns 2 and 3.

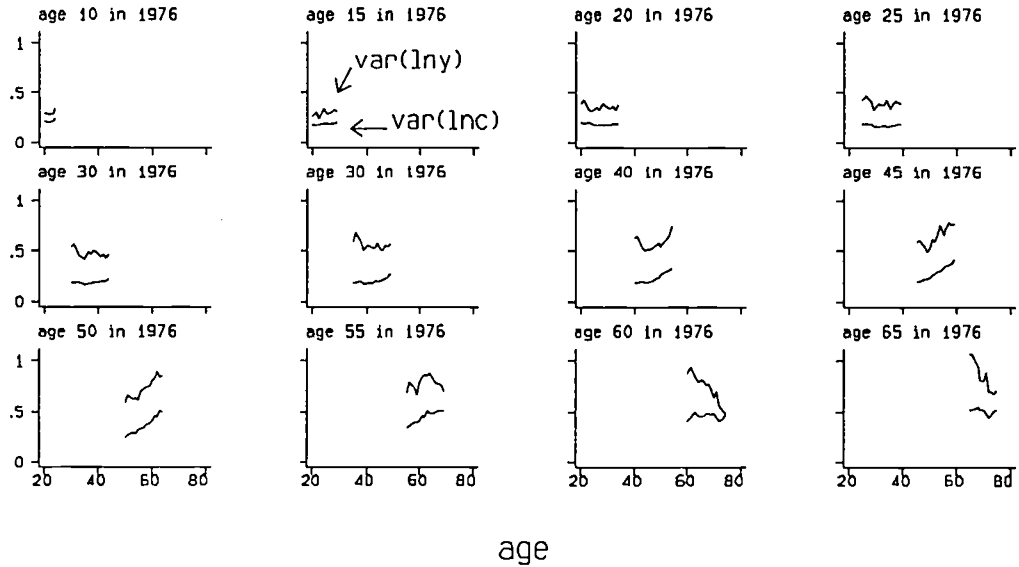

Taiwan: variances of ln(c) and ln(y).

Figure 1

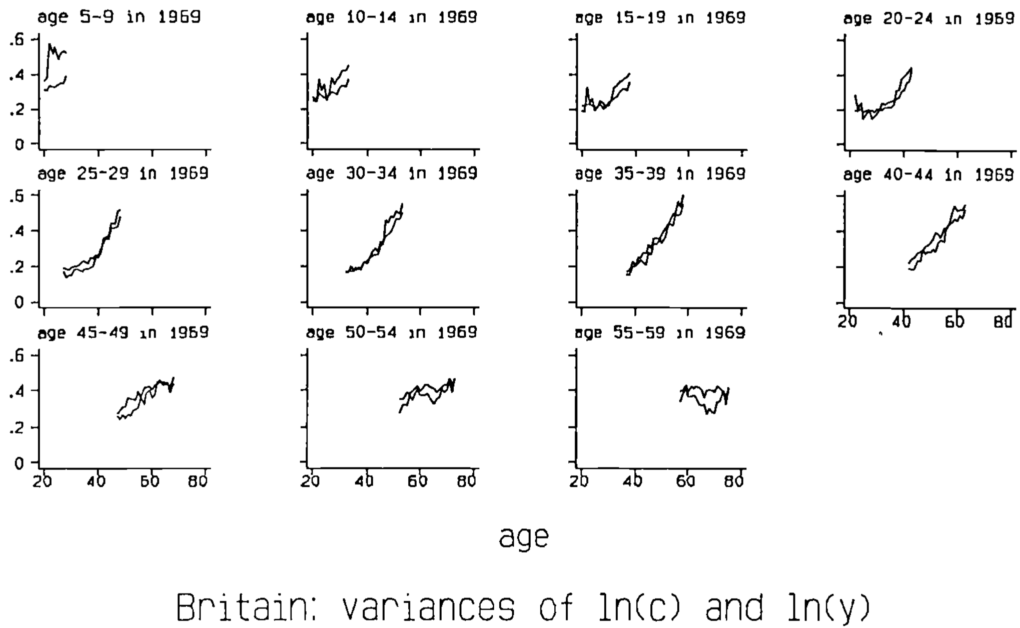

age 5-9 in 1969
age 10-14 in 1969
age 15-19 in 1969
age 20-24 in 1969
age 25-29 in 1969
age 30-34 in 1969
age 35-39 in 1969
age 40-44 in 1969
age 45-49 in 1969
age 50-54 in 1969
age 55-59 in 1969
.6
.4
.2
0
20
40
60
80
age
Britain: variances of ln(c) and ln(y)

Figure 2

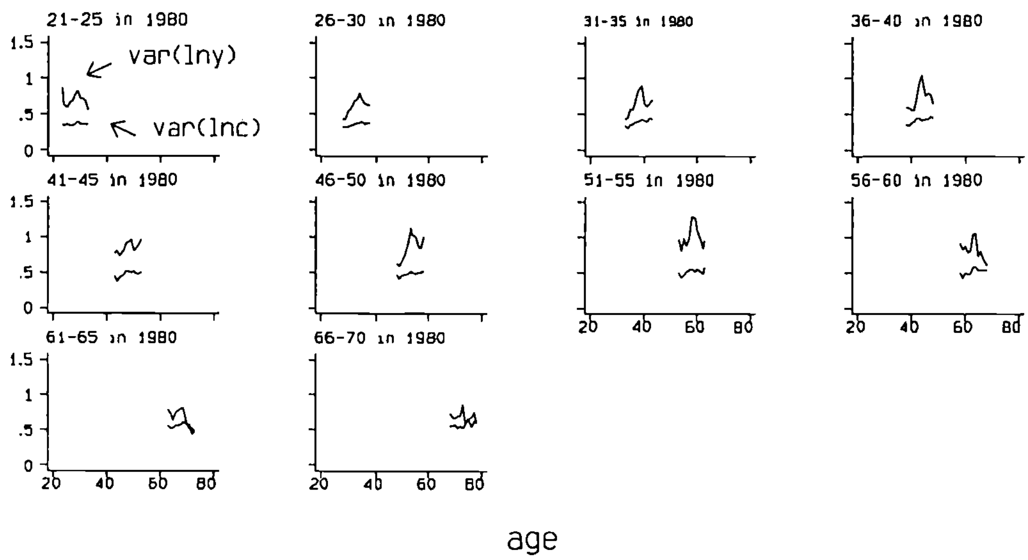

U.S.: variances of ln(c) and ln(y)

Figure 3

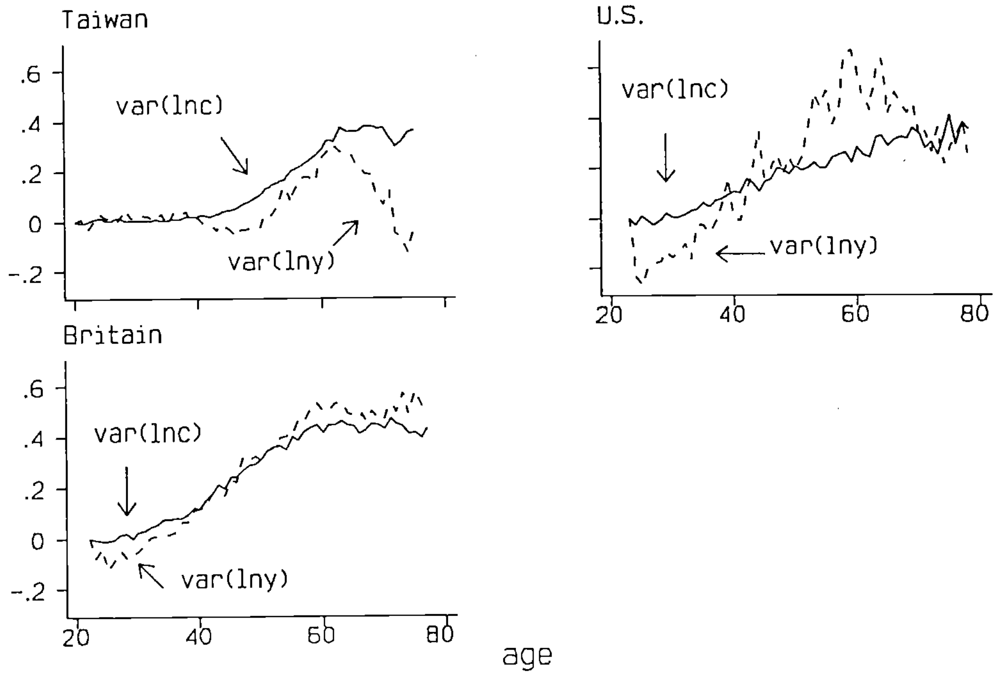

age effects in var(lnc) and var(lny)

Figure 4

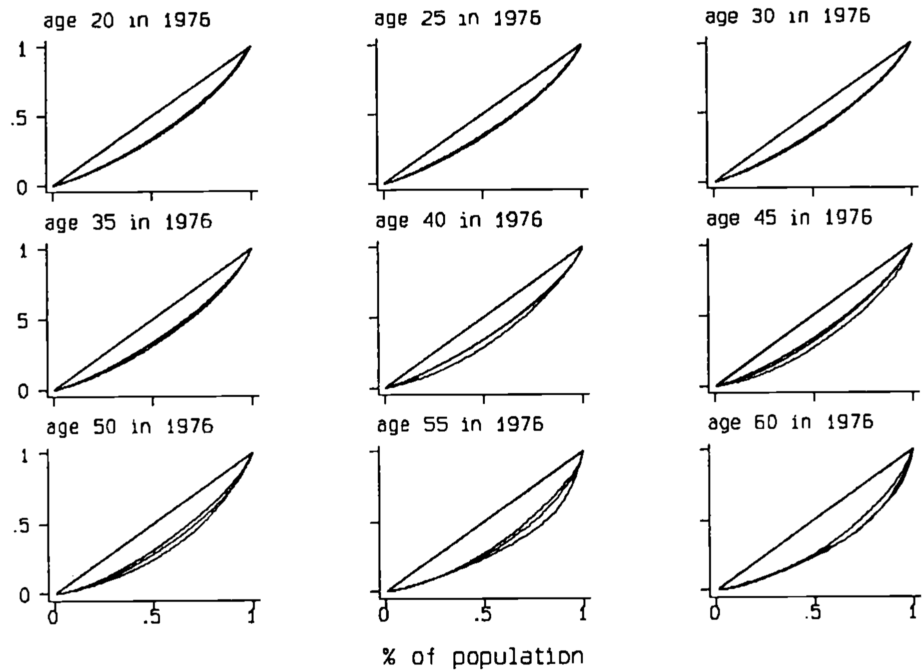

Lorenz curves for consumption: 1976, 1983, and 1990

Figure 5

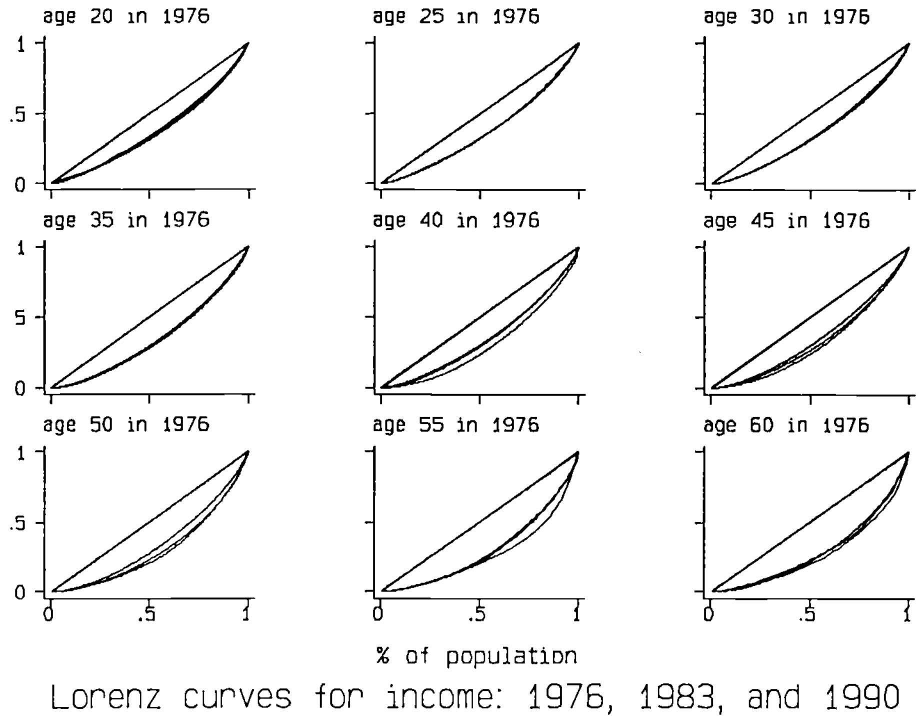
age 20 in 1976
age 25 in 1976
age 30 in 1976
age 35 in 1976
age 40 in 1976
age 45 in 1976
age 50 in 1976
age 55 in 1976
age 60 in 1976
1
.5
0
.5
% of population
Lorenz curves for income: 1976, 1983, and 1990

Figure 6

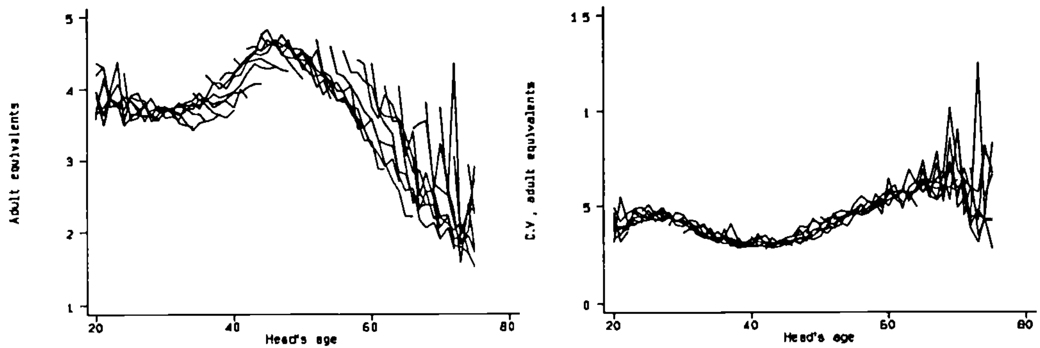
Adult equivalents
5
4
3
2
1
20
40
60
80
Head's age
C.V., adult equivalents
1.5
1
.5
0
20
40
60
80
Head's age

Figure 7

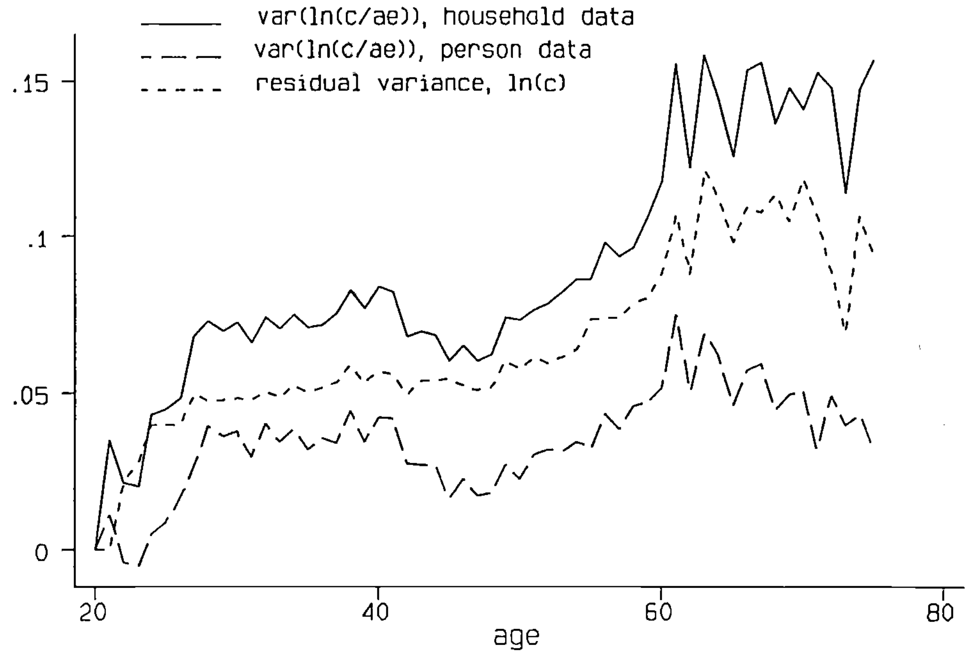

Taiwan: var(lnc), adjusted for demographics

Figure 8

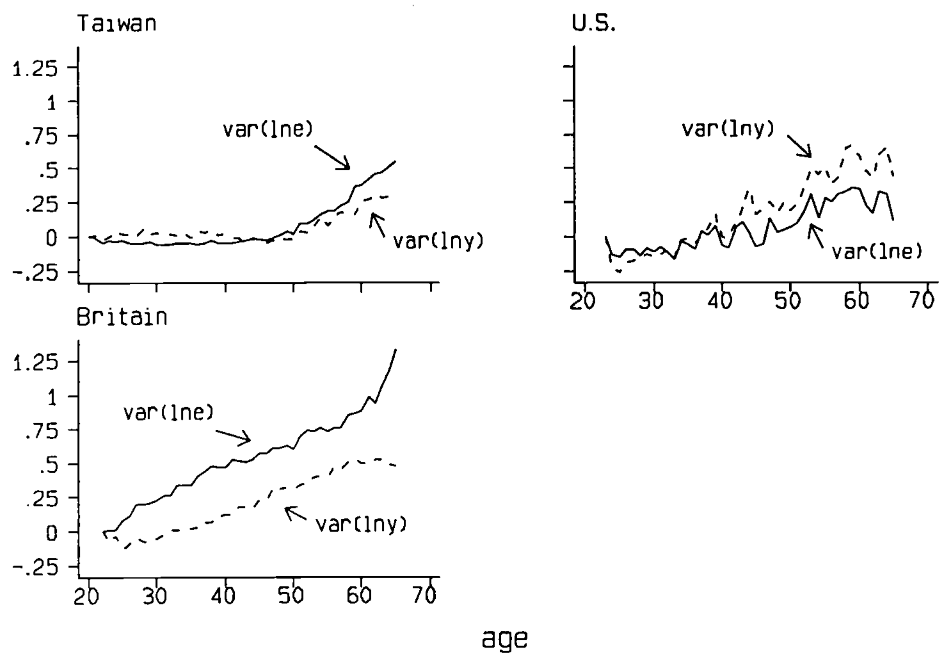

Age effects in var(lny) and var(lne)

Figure 9