The Geneva Papers on Risk and Insurance, Vol. 24 No. 1 (January 1999) 64–76

Age, Wealth Inequality and Life-Cycle Modelling

by James B. Davies*

1. Introduction

Most of the main observed features of the distribution of wealth have been studied in detail. We know much about the possible reasons for the long upper tail of the distribution, the age profile of mean wealth, portfolio composition, changes in the distribution over time, differences between countries, and the very low wealth-holdings observed for some people at all ages (see Davies and Shorrocks, forthcoming). Yet one of the most interesting stylized facts about the distribution of wealth – that wealth inequality follows a pronounced U-shaped age profile – has received hardly any attention and has not been satisfactorily explained. In all countries where data are available, wealth inequality declines at least until middle age, and in almost all datasets it increases from middle age through retirement.

The fact that the age profile of wealth inequality is U-shaped contrasts with the behaviour of inequality in earnings, income, and consumption. While inequality in earnings declines for the first five to ten years of the working lifetime, the dominant pattern is that inequality in all three of these variables rises over the lifetime. This is such a strong and well-known empirical regularity that it may have drawn attention away from the finding of a U-shaped profile for wealth inequality.

This paper explores alternative possible explanations for the U-shaped age profile of wealth inequality. Aside from satisfying our curiosity, such an investigation may help us understand processes of wealth accumulation better. Deaton and Paxson (1994) obtained valuable insights requiring just that a saving model should predict rising income and consumption inequality over the lifetime. But a satisfactory model must be capable of more. It should predict not only that inequality in income and consumption will rise with age but also that wealth inequality will initially decline with age before rising over the latter portion of the life-cycle. As argued by Friedman (1957), the ability to reproduce the stylized facts of wealth inequality is an important test of models of consumption.

The only previous attempt to model a U-shaped age profile of wealth inequality of which I am aware is Shorrocks (1975a). Shorrocks investigated a queuing model of wealth accumulation and illustrated the model through an application to the age profiles of mean wealth and wealth inequality. In his model there are two sources of shocks to wealth at a point in time. One is independent of current wealth, and arises due to earnings innovations. The other is proportional to current wealth (in expected value). The independent shocks act to reduce wealth inequality over time, whereas the proportional shocks do the opposite. Initially

* Department of Economics, University of Western Ontario, London, Canada N6A 5C2. email: davies@sscl.uwo.ca. The author thanks Rob Alessie, James Banks, Tony Shorrocks and participants in the Conference on Saving at Tilburg University in July 1996 and in the Conference on the Distribution of Wealth at University of Essex in June 1997 for their comments. The author also thanks the Social Sciences and Humanities Research Council of Canada for its support. Responsibility for any errors or omissions is the author's alone.

Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK.

the effect of earnings shocks is dominant, and wealth inequality tends to decline with age. However, as mean wealth rises the proportional shocks become more important, and as people begin to retire the earnings shocks clearly must become relatively less important. Thus, wealth inequality will begin to increase beyond some age.

In this paper the evolution of wealth inequality is initially modelled in a simple life-cycle model where all individuals earn the same rate of return on their assets, and are the same in all respects except for differences in earnings. In each period earnings are subject to transitory and permanent shocks. The transitory shocks, by themselves, would produce a "white noise" earnings process, while the permanent shocks would give a random walk in earnings. It turns out that the transitory shocks tend to make wealth inequality decline with age, as was true for earnings shocks in Shorrocks (1975a). However, the permanent shocks tend to make wealth inequality rise with age, so that whether wealth inequality falls or rises depends on the relative importance of transitory and permanent earnings shocks.[1]

The paper is organized as follows. In section 2 we look at empirical evidence on the age profile of wealth inequality. Section 3 examines the predictions of the simple life-cycle model (LCM) with uncertain earnings. This model assumes certain (and equal) lifetimes and rates of return, homogeneous preferences, and perfect capital markets. We move on, in section 4, to consider variations on the life-cycle model. We discuss briefly the effects of allowing (i) greater heterogeneity among individuals than in the simple LCM, (ii) uncertain lifetime and imperfect capital markets, and (iii) a bequest motive for saving. Section 5 concludes.

2. Empirical evidence

The stylized facts about the age profile of dispersion and inequality in non-human wealth holdings are simple. First, absolute dispersion (as measured, for example, by the standard deviation) increases monotonically with age. Second, relative inequality (as measured, for example, by the Gini coefficient or coefficient of variation) falls sharply over the first 10 to 15 years of the working lifetime. It then continues falling, or at least does not trend upward into middle age, and then at some point begins to rise with age. Typically, wealth inequality rises quite strongly through the retirement years.

These stylized facts can be seen in estimates of wealth distribution based on the estate mutliplier method. The U-shaped profile of relative inequality shows up, for example, in the upper quantile wealth figures for the U.K. over 1912–71 reported by Shorrocks (1975b). These estimates show relative convergence over the working lifetime and divergence in retirement. This pattern is most marked when corrections are made for changes in population composition with age. They are also evident in longitudinal as well as cross-section estimates.

Another early source of evidence was the 1962 Survey of Financial Characteristics of Consumers (SFCC) conducted for the Federal Reserve Board in the U.S. This reported an overall Gini coefficient of 0.76, and a Gini falling from 0.83 for households with head aged under 35 to 0.70 for those aged 45–54 and 55–64. The Gini for households aged over 65 was 0.76 (see Projector and Weiss, 1966, Table 8, p. 30).

More recent confirmation of the U-shaped profile of wealth inequality can be found for

[1] The impact of permanent earnings shocks here is not inconsistent with Shorrocks' analysis. Permanent earnings shocks cause continuing changes in saving, that is asset changes over following periods which are correlated with the level of wealth. Thus the permanent earnings shocks here are like the proportional shocks studied by Shorrocks rather than his independent earnings shocks.

Germany (Börsch-Supan, 1994, financial assets only), Italy (Jappelli, 1995), the Netherlands (Alessie et al., 1997), and the U.S. (Wolff, 1980; Greenwood, 1987). Canada provides the single exception to the U-shaped age profile of which I am aware. In each of the 1970, 1977, and 1984 Surveys of Consumer Finance (SCFs) wealth inequality shows a monotonic decline with age.[2] (See, for instance, Davies, 1979 and Magee et al., 1991. The 1984 SCF is the most recent wealth survey for Canada.)

While the U-shaped age profile of wealth inequality is quite well established, there are some data problems which should be considered. First, most of the studies mentioned above use cross-section data. Panel studies on wealth are beginning to become available, and will allow this deficiency to be corrected. Second, in most cases the published results present grouped data. In the earliest studies families or individuals were typically grouped in ten-year age ranges, but five-year ranges are now used in some cases. Grouping introduces potentially serious errors in estimated inequality where wealth is trending up or down strongly with age.

In a model like the simple LCM, which has a strongly peaked age-wealth profile, striking but spurious falling inequality during the working lifetime and rising inequality in retirement can be generated simply by this grouping error. Consider the following simple example. Suppose everyone has equal and constant earnings over a 35-year working lifetime beginning at age 20. In line with the current trend towards earlier retirement, suppose that everyone stops work at age 55 and enjoys a 20-year retirement before dying at age 75. Assuming a zero interest rate for convenience, and letting everyone consume at a constant rate over the lifetime, there will be a triangular-shaped age profile of wealth holding. While for any specific age there will be no wealth inequality at all, when members of this society are grouped into five-year age bands the following data would be obtained for coefficients of variation (CV) in these age groups.

Age Group	CV	Age Group	CV
20–24	0.577	50–54	0.044
25–29	0.193	55–59	0.083
30–34	0.116	60–64	0.116
35–39	0.083	65–69	0.193
40–44	0.064	70–74	0.577
45–49	0.053		

In practice, the rise and fall of wealth with age is much less sharp than in the simple LCM, reducing the possible importance of grouping error. And, fortunately, there are also studies which do not suffer from grouping error. An example is provided by Alessie et al. (1997) which looks at single-year age groups in the 1989 Dutch SEP data. Figure 1 shows the Alessie et al. kernel-smoothed age profile for the coefficient of variation of household net worth in the Netherlands. While this profile shows a small hump in the age range from 30 to 45, overall it shows a downward trend from age 25 to 50. There is then a plateau from 50 to 60, and only after age 60 (by which time most Dutch workers are retired) does an increase in wealth inequality begin.

[2] In order to investigate whether this simply resulted from the grouping of all households over age 65, I have examined the 1984 SCF microdata. There is a slight tendency for inequality, as measured by the coefficient of variation, to increase at around age 70 for several family size types. But the contrast with what is found in other countries remains strong.

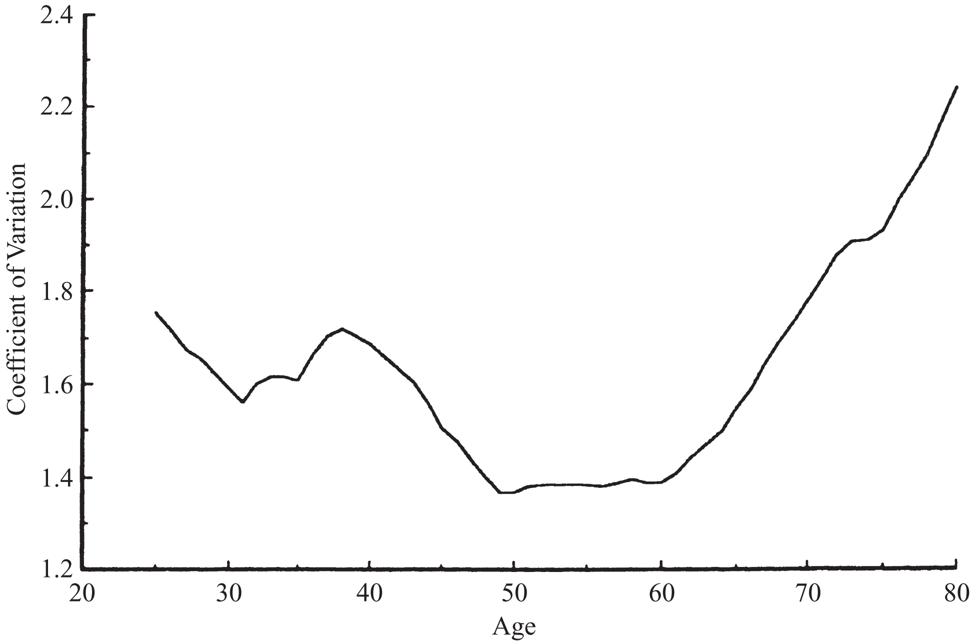

Figure 1: Coefficient of variation of net worth, the Netherlands, 1989
Source: Alessie et al. (1997).

A further concern about the wealth data is that it generally excludes social security and pension wealth. Although several studies have imputed state and occupational pensions, I am not aware of any available estimates of wealth inequality by age group making use of such adjustments. It is likely that imputing pension wealth would mute the tendency for wealth inequality to rise toward the end of the lifetime. However, even so the remarkable rise in inequality in non-pension wealth at these ages remains an interesting fact which needs to be explained.

3. Life-cycle model

Our attention is confined here to the simplest version of the "stripped down" LCM. The units considered are families, of constant (adult) composition from "birth" at age 20 (a typical age of labour force entry) to death T years later. All families receive labour income, E_t, from age 20 until retirement after R years of employment. Consumers display certainty equivalence behaviour. I assume that the interest rate equals the rate of time preference, and that there is a perfect capital market. Thus, all families desire constant consumption over the lifetime. I will refer to this specification as the simple LCM.

Carroll (1992) and Hubbard et al. (1993) find that a realistic approximation to the earnings process is provided by a combination of white noise and a random walk. Throughout this paper the following process is assumed for the earnings of an individual family, E_t:

$$E_t = \begin{cases} E_{t-1} + \varepsilon_t - \beta\varepsilon_{t-1} & \varepsilon_t \sim (0, \sigma^2) \quad t \leqslant R \\ 0 & t > R \end{cases} \tag{1}$$

where $\beta = 1$ gives the case of pure white noise and $\beta = 0$ a random walk. When $0 < \beta < 1$, of course, we have a case which is intermediate between the pure white noise and random walk extremes.

The earnings process assumed implies constant mean earnings, E_t. The following analysis can readily be generalized to allow the age profile of mean earnings to have a realistic hump shape, but this comes at the cost of tedious algebra, and yields little further insight.[3] I am also ignoring state and occupational pensions. Their inclusion would reduce the strength of the saving motive, producing effects on wealth accumulation in the simple LCM over the working life similar to those of shortening the retirement period. One impact of using additive rather than multiplicative shocks here, which should be kept in mind, is slightly less of a tendency for wealth inequality to decline over the working lifetime in the models considered.[4]

I have specified an earnings process with additive shocks in order to be able to get analytical results on the behaviour of inequality. Multiplicative shocks fit the data better, and conveniently accommodate a lognormal earnings distribution (which is fairly realistic). The formulation (1) might be taken to suggest that, in contrast, I will work with symmetric, or possibly even normal, distributions. This is not the case. The shapes of the distributions of E_0 and the ε_i's do not have to be unduly restricted, and may be strongly positively skewed.

Notice that in the white noise version of (1), absolute dispersion of E_t, as measured by the variance, $V(E_t)$, or standard deviation will be constant over the lifetime. If we have $\beta < 1$, however, the variance of earnings will rise with age. Since (1) features constant mean earnings, *relative* earnings inequality – which can be readily measured here by dividing the standard deviation by the mean to get the coefficient of variation, $CV(E_t)$ – moves in exactly the same way as absolute dispersion. So, if $\beta = 1$ relative as well as absolute earnings dispersion are constant, and with $\beta < 1$ both relative and absolute dispersion are increasing.

In what follows, we find that the behaviour of relative and absolute consumption dispersion is also the same, since C_t is constant over the lifetime in the simple LCM. However, the situation is very different for wealth, which is not constant. This illustrates the general point that we must distinguish carefully between the behaviour of relative and absolute measures of dispersion. For simplicity, the term "dispersion" will be used as synonymous with "absolute dispersion". Dispersion may be measured using the variance or standard deviation. The term "inequality" will be reserved for relative dispersion, and may be measured with such scale independent measures as the coefficient of variation or the Gini coefficient.

---

[3] It can be argued that unless account is taken of the consumption needs impact of the hump-shaped age profile of family size, introducing a hump-shaped earnings profile makes the exercise *less* rather than more realistic. See Davies (1988).

[4] With the usual formulation in which shocks are multiplicative and lognormal, the variance of earnings rises at an increasing rate, in contrast to the linear behaviour with (1). The convexity of the age profile of the variance in the lognormal case is, however, not particularly pronounced for realistic parameter values. Convexity of this profile should strengthen the relative tendency coming from permanent earnings shocks towards rising wealth inequality at later ages.

A family of working age has the budget constraint:

$$\sum_{i=t}^{T} \frac{C_t}{(1+r)^{i-t}} \leqslant (1+r)W_t + \sum_{i=t}^{T} \frac{\mathrm{E}(E_i)}{(1+r)^{i-t}} \qquad t \leqslant R \tag{2}$$

where expected future earnings equal the permanent component of current earnings. This gives rise to the consumption rule:

$$C_t = \frac{[(1+r)W_{t-1} + \beta\varepsilon_t] + (E_t - \beta\varepsilon_t)D_t^R}{D_t^T}: \qquad D_t^K = \sum_{i=t}^{K} \frac{1}{(1+r)^{i-t}}. \tag{3}$$

This rule provides much of the intuition behind the positively sloped age profile of consumption inequality, when taken together with:

$$\overline{C}_t = \overline{E}\frac{D_1^R}{D_1^T}; \qquad \overline{E} = \mathrm{E}(E_1) = \ \ldots \ = \mathrm{E}(E_R) \tag{4}$$

which indicates that mean consumption is constant. Whereas mean consumption is constant over the lifetime, the variance of W_{t-1}, $V(W_{t-1})$, grows with age (even if wealth inequality is trending downwards), and so does $V(E_t)$ unless we have the pure white noise model, in which case it is constant. Since the dispersion in these determinants of consumption is rising, it is not surprising that the dispersion in C_t is increasing with age, as demonstrated by Deaton and Paxson (1994). But, given that mean consumption is constant this also means that consumption *inequality* is rising with age.[5]

In contrast to consumption, mean wealth rises very quickly up to the retirement age in the simple LCM, and falls quite sharply afterwards. We have:

$$\overline{W}_t = \begin{cases} \overline{E}\left(\dfrac{D_1^T - D_1^R}{D_1^T}\right)\displaystyle\sum_{i=1}^{t}(1+r)^{t-i} & t \leqslant R \\ \left(\dfrac{D_{t+1}^T}{D_{R+1}^T}\right)\overline{W}_R & R < t < T \\ 0 & t = T \end{cases}$$

The sharply peaked mean age–wealth profile is not echoed in real-world data. While hump-shaped age profiles of wealth are observed in cross-section, the rise and fall of wealth is much less sharp than in the simple LCM, and, as has been widely commented, (non-pension) wealth appears to be decumulated fairly slowly in retirement.

Through successive substitution expressions for current wealth in the working and retirement periods are readily obtained, and analytic results on the variance of wealth also result. The following results emerge (see Davies, 1996, for proofs):

Result 1: During the retirement period in the simple LCM, consumption and wealth inequality are equal and constant.

[5] The prediction that consumption inequality rises with age is very robust, and is obtained in all the models considered in this paper. In order to economize on algebra, then, while numerical results are presented for consumption inequality the corresponding equations are omitted. They are provided in an appendix which may be obtained from the author.

The explanation for this result is that in the simple LCM consumption in retirement is proportional to wealth, and wealth is in turn proportional to net worth at the retirement date, which is the same for everyone. This result implies that the simple LCM is not capable of explaining why wealth inequality is typically observed to rise in retirement. (As we see below, however, fairly straightforward extensions to the model make such a prediction possible.)

Result 2: In the absence of earnings shocks, but with some dispersion in initial earned incomes, inequality in W_t would be constant over the entire lifetime, and, together with consumption, would equal earnings inequality.

This result follows from the fact that lifetime earnings would be proportional to initial earnings. Consumption and saving would, of course, be proportional to lifetime earnings in this simple case where the distribution of earned income would not change as a cohort aged.

Result 3: If the earnings shocks were purely white noise ($\beta = 1$), $CV(W_t)$ would decline over the working lifetime. On the other hand, if the earnings process were a pure random walk, $CV(W_t)$ would increase over the working lifetime.

To provide intuition for the last result, think first of what happens in the first period in the white noise model. Families save a very large portion of positive shocks, and dis-save the same proportion of negative shocks. In the first period, when mean wealth is low, this produces large relative inequality of wealth. As we move to the second and subsequent periods two things happen to moderate this inequality. The first is that the wealth (or debt) incurred in response to first-period shocks is consumed (paid off). The second is that there are new shocks, uncorrelated with previous ones. Thus, wealth inequality declines.

Adding a random walk element to earnings creates permanent earnings shocks. These shocks tend to make wealth inequality rise with age for the following reason. While the *immediate* impact of these permanent shocks on saving is smaller than that of transitory shocks, the changes they produce in wealth do not fall off with time, as in the white noise model. A permanent negative earnings shock, for example, means lower expected lifetime wealth, lower consumption, and lower wealth throughout the rest of the lifetime. Permanent earnings shocks thus not only produce rising earnings inequality, they also generate rising wealth inequality.

So transitory and permanent shocks have opposite influences on the evolution of wealth inequality over the life-cycle. The white noise aspect of the earnings process is tending to make wealth inequality fall, even while earnings inequality is rising. The random walk aspect is tending to make wealth inequality rise. The remainder of this section asks which effect will dominate when plausible parameter values are assigned to the simple LCM.

Table 1 shows $CV(W_t)$ and $CV(C_t)$ at five-year intervals through the lives of a population of simple life-cycle savers who work for 35 years and retire for 20 years. In view of declining retirement ages and increasing longevity, it would appear that these figures are more appropriate for current work than the 40-year working lifetime and 10-year retirement that featured, for example, in Atkinson (1971) and some other earlier illustrative calculations. Against this, recall that a longer retirement period implies a greater need for life-cycle saving, which has the opposite influence to including state and occupational pensions. Thus, it is important to check the sensitivity of the results to alternative assumptions on the length of the retirement period. This is done in Table 2, which increases the length of the working period, and reduces longevity, alternatively. A further assumption in the calculations reported here is a zero interest rate.

The earnings process in Table 1 is parameterized to mimic the age profile of earnings

Table 1:
Simple LCM: inequality by age

I. $CV(W_t)$ Age	$\boldsymbol{\beta = 0}$ **(Random walk)**	**0.25**	**0.5**	**0.75**	**1 (White noise)**
20	0.4001	0.4292	0.4658	0.5082	0.5550
25	0.4456	0.4319	0.4200	0.4103	0.4030
30	0.4944	0.4585	0.4276	0.4031	0.3862
35	0.5439	0.4879	0.4397	0.4026	0.3797
40	0.5949	0.5195	0.4542	0.4044	0.3763
45	0.6482	0.5535	0.4706	0.4075	0.3741
50	0.7052	0.5908	0.4894	0.4116	0.3727
54	0.7677	0.6324	0.5064	0.4170	0.3716
55–74	0.7677	0.6324	0.5064	0.4170	0.3716
II. $CV(C_t)$ Age	$\boldsymbol{\beta = 0}$ **(Random walk)**	**0.25**	**0.5**	**0.75**	**1 (White noise)**
20	0.4035	0.3909	0.3815	0.3756	0.3733
25	0.5612	0.4901	0.4311	0.3904	0.3738
30	0.7492	0.6174	0.5006	0.4126	0.3744
35	0.9833	0.7829	0.5970	0.4462	0.3752
40	1.2867	1.0028	0.7315	0.4971	0.3763
45	1.6971	1.3048	0.9225	0.5754	0.3776
50	2.2822	1.7394	1.2036	0.6987	0.3795
54	2.9633	2.2478	1.5368	0.8521	0.3823
55–74	2.9633	2.2478	1.5368	0.8521	0.3823

Labour force entry age = 20; retirem#ent age = 55; age of death = 75.

inequality in the Netherlands. It is appropriate to look at family income, since it is best to think of the decision unit as a married couple, or family. Also, it is best to include all forms of non-capital income, given the considerable importance of transfer payments in the real world. Both these requirements are met with the help of cross-section data from the Dutch SEP survey of 1987, 88, and 89. Aldershof (1994, Figure 29) shows the age profile of the coefficient of variation of non-capital income for Dutch families. The pooled data indicate $CV(E_t)$ of about 0.4 around age 25, and a figure of about 0.6 at age 55.[6] Accordingly, I have set up the earnings process to make $CV(E_t)$ rise from 0.4 at the start of the working lifetime to 0.6 at the end when β takes on its "best guess" value of 0.5. (The implied value of σ^2 is held constant when β is varied from 0.0 to 1.0 in Table 1. This means that in the current version of

[6] $CV(E_t)$ is a little higher for those aged 20, but the value for those aged 25 is used instead. A decline in earnings inequality up until the late 20s is a typical finding, and could be accommodated in the earnings process used in this paper. However, that would require us to allow the distribution of ε_t to vary with t, which is not convenient for the derivation of the analytical results we are attempting to illustrate.

the paper the age profile of earnings inequality departs from its benchmark slope when $\beta \neq 0.5$.)

Given the empirical finding that transitory and permanent shocks to earnings are of about equal importance, the central case in Table 1 would appear to be where $\beta = 0.5$. In this case we find that not only does consumption inequality increase with age throughout, as Deaton and Paxson (1994) emphasize is required for realism, but so also does wealth inequality after a very brief decline at the youngest ages. The weight on the white noise component of the earnings process has to be raised to 0.75 before what might be described as realistic behaviour of $CV(W_t)$ is obtained. With $\beta = 0.75$, $CV(W_t)$ declines from age 20 to 35, and thereafter rises mildly through to the retirement age.[7]

Table 2 investigates sensitivity to *R* and *T*, holding β at 0.5 throughout. The first column repeats the results we have already seen. The second column raises the retirement age to 60, holding the age of death at 75. The last column keeps the retirement age at 60, and reduces the age of death to 70. Interestingly, as the length of retirement is reduced in turn to 15 and then to 10 years from the original 20, the performance of the model improves. Wealth inequality at age 50 is little affected – rising slightly. But $CV(W_t)$ becomes substantially higher for those aged 20, and also becomes quite a bit higher for those in retirement. The result is a longer period of declining $CV(W_t)$ – up to age 28 in the third column – followed by gently rising inequality up until retirement.

Table 2 is of interest not only from the viewpoint of testing sensitivity to *R* and *T*, but, as mentioned earlier, because a reduction in the strength of the life-cycle saving motive caused by shortening retirement has a roughly analogous effect on saving to introducing state and occupational pensions. The table therefore suggests that the result of modelling the latter could be to improve the performance of the model significantly.

Table 2:
Simple LCM: sensitivity of $CV(W_t)$ to R and T

Age	$\boldsymbol{R = 55}$ $\boldsymbol{T = 75}$	$\boldsymbol{R = 60}$ $\boldsymbol{T = 75}$	$\boldsymbol{R = 60}$ $\boldsymbol{T = 70}$
20	0.4658	0.5108	0.5842
25	0.4200	0.4329	0.4534
30	0.4276	0.4366	0.4506
35	0.4397	0.4471	0.4583
40	0.4542	0.4604	0.4703
45	0.4706	0.4760	0.4854
50	0.4894	0.4939	0.5037
54	0.5064	0.5104	0.5212
60 to *T*	0.5064	0.5343	0.5481
Age of Min. W_t	26	28	30
W_R/E_0	12.7	10.9	8.0

[7] Note that, unlike the shape of the age profile, the *values* obtained for $CV(W_t)$ differ considerably from what is observed in the Dutch data (see Figure 1). The observed values are much higher. This suggests the importance of taking into account aspects of heterogeneity among families aside from earnings differences.

4. Variations on the life-cycle model

In this section we first ask what difference it would make to the above results if we allowed greater heterogeneity among households. We then discuss the role of uncertain rates of return, uncertain lifetime and imperfect capital markets.

Heterogeneity in ages of labour force entry, retirement, and death

In section 2 we noted that, because there is a systematic hump-shaped relationship between age and wealth, grouping households into five- or ten-year age bands itself would tend to make apparent wealth inequality decline with age during the working lifetime and rise during retirement. These were spurious effects. Similar effects, which however are *not* spurious, occur when not all individuals' periods of work and retirement are the same. Consider the following differences in turn.

People begin work at different ages

For those in about the first half of the working period, this will produce effects similar to those of grouping in the simple LCM, once the bulk of individuals are in the labour market. Prior to this, wealth inequality will rise for a few years as current members of the labour force become wealthier and they keep being augmented by new arrivals with zero wealth. Once labour force entry is complete, say around age 30, we will have high wealth individuals who entered the labour force when they were 20 contrasting with low wealth individuals who just entered in their late 20s and still have very low net worth. Wealth inequality will be high. As we move forward through people's 30s and 40s these differences will be reduced and relative wealth inequality will fall – especially if those who enter the labour force later are more highly qualified and earn more.

People retire at different ages

Suppose that everyone started work at 20, but there was a uniform distribution of retirement ages from 55 to 65 (and everyone died at age 75). Continue to assume that all the other simplifying assumptions of the LCM remain in place. Then relative wealth inequality would be unchanging up to age 55, when the first people begin to retire. (Wealth holdings in this initial age range would be proportional to target wealth at the retirement age.) It would then increase quite rapidly, as new retirees run down their wealth but those who are still working continue to accumulate. It seems likely that this captures a relevant real-world phenomenon, since there is quite a bit of variation in retirement ages in practice, and it is also known that much saving for retirement is concentrated in the final 10 to 20 working years.

People differ in lifespan

Assume that everyone starts work at age 20 and retires at 55. Instead of assuming that everyone lives until 75, however, assume that there is a uniformly distributed set of (perfectly anticipated) ages at death ranging from 70 to 80. Although the age of labour force entry and the length of the working lifetime are the same for everyone here, note that the variation in length of retirement introduces a source of wealth inequality among those of working age. Since the wealth of each working person at a particular age is proportional to their target wealth on retirement, however, wealth inequality is constant over the working period, and differences in lifespan alone therefore do not lead to the pattern of declining wealth inequality in the working lifetime that is observed in the real world.

The behaviour of wealth inequality in the retirement period is non-monotonic. Initially it increases as the least wealthy individuals (those with the lowest life expectancy) decumulate at the fastest rate. However, once the low life expectancy people start to die off, wealth inequality begins to moderate, and there will be a period of declining inequality at the end of the lifetime.

An alternative way to introduce differences in the length of life would be to shrink the length of the working lifetime along with the overall length of life. This would make the period of rising wealth inequality start earlier – when those with lowest life expectancy began to retire. Since there is some evidence that those with lower earnings do indeed both retire and die earlier, this is likely a point with real-world relevance.

In the real world there is an additional effect which should not be overlooked. This is that there is a positive correlation between measures of resources like income or wealth and life expectancy. Thus the wealthiest members of society should decumulate most slowly in retirement. Starting at age 65, say, wealth inequality will as a result tend to increase as the less wealthy run down their assets faster. Beyond a certain point, however, as the less wealthy begin to die off at a high rate, the population will become more homogeneous, and wealth inequality would be expected to decline.

Summing up this discussion, it appears that heterogeneity has quite a bit of potential to help to explain the U-shaped age profile of wealth inequality. Differences in the age of labour force entry can help to explain declining wealth inequality during the working lifetime. Differences in age of retirement or length of life, on the other hand, may have a role in explaining the increase of wealth inequality from late middle age onwards. Note, however, that brief initial and final periods during which behaviour does not follow the U-shaped pattern may occur due to changes in population composition.

Uncertain rates of return

The model considered in section 3 incorporated uncertain earnings but not uncertain rates of return. Suppose that all individuals must save using the same risky asset. Then there will clearly be a tendency for uncertain rates of return to make wealth inequality rise with age. If we return to the simplest life-cycle model with no earnings differences, certain rates of return will yield zero wealth inequality, controlling for age. As soon as there are differences in rates of return individuals will begin to separate in terms of their lifetime opportunities as time goes on. The luckiest will get a string of high returns, and become quite wealthy, and of course the opposite will also occur for the least lucky. In principle this effect could be offset by behavioural influences – say for example if saving rates declined as people became better off. But with the standard assumptions economists make about intertemporal preferences (additive separability, homotheticity) this will not occur. And if greater "realism" is aimed for it is likely to lead to higher saving rates for individuals who are better off, reinforcing the tendency for unequal rates of return to cause wealth inequality to rise with age.

In a world with more than one kind of asset, portfolio choice affects how risky rates of return affect the evolution of wealth inequality over the life cycle. Under reasonable conditions people with higher real lifetime incomes will put a larger share of their portfolio in riskier assets, once again reinforcing the tendency for differences in rates of return to lead to cumulative effects which make wealth inequality rise with age. Progressive income taxation may offset this tendency to some extent.

Uncertain lifetime

In a world without state or occupational pensions the impact of uncertain lifetime is likely to be to increase saving rates, and reduce the rate of decumulation in retirement (Davies, 1981). This could tend to offset the influence of factors like unequal retirement ages and differing life expectancy, which we saw above could make wealth inequality rise in retirement. If no one is dis-saving very fast, relative differences in wealth will not change much. On the other hand, in the presence of pensions it is possible that uncertain lifetime will lead to a higher rate of decumulation than is seen under certain lifetime (Davies, 1981). Thus, given the increasingly generous pensions observed in advanced Western countries in recent decades, it may be that uncertain lifetime is responsible for more rapid decumulation of non-pension wealth, tending to allow population heterogeneity in factors like retirement age and life expectancy to cause a more rapid increase of wealth inequality in retirement.

Bequests

Bequests may arise in a pure life-cycle model if annuity market imperfections result in the elderly holding assets in bequeathable form. The relative size of such "accidental" bequests and intentional ones as well, which can be brought in with the addition of a bequest motive, tends to be positively correlated with the resources of both decedents and heirs. As people move through the life cycle more and more of the higher earners will inherit, and they will see a relative increase in their wealth as a result. This is an effect which one would expect to be particularly strong as we move through middle age. It may help to explain why rising wealth inequality is observed in several datasets, including the Dutch data we have looked at here, starting in middle age rather than at retirement.

Another interesting possibility is that, although very few young people will inherit, the fact that a few receive large bequests may have a strong impact on measured wealth inequality. Thus it is possible that the bequest process may have something to do with the decline of wealth inequality at low ages, as well as with the increase later on.

5. Conclusion

This paper has asked what may explain the well-established empirical regularity that wealth inequality tends to decline in cross-section at least until middle age and then to increase. We have found that the life-cycle model, even in quite a simple form, provides several reasons to expect such a pattern. This provides some assurance that economists' continued attachment to this workhorse model makes sense.

We have seen that the initial decline of wealth inequality with age could arise, in the simple LCM framework as a result of transitory shocks to earnings or differences in the age of labour force entry. The rise in wealth inequality at later ages may be due to a decline in the importance of transitory relative to permanent earnings shocks (the latter tending to make wealth inequality rise with age), as well as to population heterogeneity in the form of differences in age of retirement and life expectancy. Greater life expectancy among higher income individuals can also contribute to rising inequality in the initial retirement years. While all these effects tend to be offset at the very youngest and oldest ages when changes in population composition are occurring, the potential of the LCM to explain the U-shaped profile of wealth inequality over most of the lifespan is quite striking. As we have briefly discussed, bequest behaviour may also contribute to the explanation.

REFERENCES

ALDERSHOF, T., 1994, "Income and Wealth Holdings Over the Life-Cycle", University of Tilburg, mimeo.

ALESSIE, R., LUSARDI, A. and ALDERSHOF, T., 1997, "Income and Wealth over the Life Cycle: Evidence from Panel Data", *Review of Income and Wealth*, 43, pp. 1–32.

ALESSIE, R., LUSARDI, A., ALDERSHOF, T. and KAPTEYN, A., 1995, "Saving and Wealth Holdings of the Elderly", *Ricerche Economiche*, 49, pp. 293–315.

ATKINSON, A.B., 1971, "The Distribution of Wealth and the Individual Life-cycle", *Oxford Economic Papers*.

ATKINSON, A.B. and HARRISON, A.J., 1978, *Distribution of Personal Wealth in Britain*, Cambridge University Press, Cambridge.

ATTANASIO, O.P. and JAPPELLI, T., 1996, "The Life Cycle Hypothesis and Consumption Inequality", IFS, University College London, and University of Salerno, mimeo.

BÖRSCH-SUPAN, A., 1994, "Savings in Germany – Part 2: Behavior", in James Poterba (ed.), *International Comparisons of Household Saving*, University of Chicago Press, Chicago.

CARROLL, C.D., 1992, "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence", *Brookings Papers on Economic Activity*, 2, pp. 61–135.

DAVIES, J.B., 1979, "Life-Cycle Savings, Inheritance, and the Distribution of Personal Income and Wealth in Canada", PhD thesis, London School of Economics.

DAVIES, J.B., 1981, "Uncertain Lifetime, Consumption and Dissaving in Retirement", *Journal of Political Economy*, 89, pp. 561–577.

DAVIES, J.B., 1988, "Family Size, Household Production, and Life Cycle Saving", *Annales d'Economie et de Statistique*, (9), Jan.–Mar., pp. 141–165.

DAVIES, J.B, 1996, "Wealth Inequality and Age", University of Western Ontario, Department of Economics, Research Report 9613.

DAVIES, J.B. and SHORROCKS, A.F., forthcoming, "The Distribution of Wealth", in A.B. Atkinson and F. Bourguignon (eds.), *Handbook of Income Distribution*, North-Holland, Amsterdam.

DEATON, A. and PAXSON, C., 1994, "Intertemporal Choice and Inequality", *Journal of Political Economy*, 102(3), pp. 437–467.

FRIEDMAN, M., 1957, *A Theory of the Consumption Function*, NBER, Princeton University Press, Princeton.

GREENWOOD, D.T., 1987, "Age, Income, and Household Size: Their Relation to Wealth Distribution in the United States", in Edward N. Wolff (ed.), *International Comparisons of the Distribution of Household Wealth*, Oxford University Press, Oxford, pp. 121–140.

HUBBARD, R.G., SKINNER, J. and ZELDES, S.P., 1993, "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving", *Carnegie–Rochester Conference Series on Public Policy*, 40, June, pp. 59–125.

HURD, M.D., 1995, "The Predictive Validity of the Subjective Probabilities of Survival in the Health and Retirement Survey", SUNY Stony Brook, mimeo.

JAPPELLI, T., 1995, "The Age–Wealth Profile and the Life-Cycle Hypothesis: A Cohort Analysis with a Time Series of Cross-Sections of Italian Households", CEPR Working Paper 1251.

MAGEE, L., BURBIDGE, J. and ROBB, L.A., 1991, "Computing Kernel-Smoothed Conditional Quantiles from Many Observations", *Journal of the American Statistical Association*, 86, pp. 673–677.

PROJECTOR, D.S. and WEISS, G.S., 1966, *Survey of Financial Characteristics of Consumers*, Federal Reserve Board, Washington.

SHORROCKS, A.F., 1975a, "On Stochastic Models of Size Distributions", *Review of Economic Studies*, pp. 631–641.

SHORROCKS, A.F., 1975b, "The Age–Wealth Relationship: A Cross-Section and Cohort Analysis", *Review of Economics and Statistics*, pp. 631–641.

WOLFF, E.N., 1980, "Estimates of the 1969 Size Distribution of Household Wealth in the U.S. from a Synthetic Data Base", in James D. Smith (ed.), *Modelling the Distribution and Intergenerational Transmission of Wealth*, NBER Studies in Income and Wealth, 46, University of Chicago Press, Chicago, pp. 223–271.