Vienna Yearbook of Population Research 2014 (Vol. 12), pp. 193–209

How large are the effects of population aging on economic inequality?

Joshua R. Goldstein and Ronald D. Lee*

Abstract

The attention given to Piketty (2014) has renewed interest in the level and causes of inequality. In this paper, we look at the role that population aging plays in increasing economic inequality. We provide estimates of the magnitudes of the effects on inequality of three different factors related to population aging: capital intensification, changing population age structure, and increasing longevity. Changing age structure is found to have a small effect on aggregate inequality, while capital deepening and longevity-based life cycle savings are shown to be more important. Taken together, our findings suggest that aging has a substantial effect on economic inequality.

1 Introduction

Thomas Piketty's 2014 book, *Capital in the 21st Century*, has sparked an enormous resurgence in interest in inequality. Demography is one of the factors at the heart of Piketty's prediction of rising inequality. In this paper, we discuss several of the mechanisms through which population aging could influence economic inequality and try to provide estimates of the magnitude of each factor.

The three aspects of population aging we consider are the slowdown in population growth from fertility decline, the accompanying shift to older age structures, and increases in longevity. The rate of population growth influences inequality through its effect on the capital intensity of the economy (Piketty 2014, Solow 1956). The older age structure of the population has the potential to influence aggregate measures of inequality because of the tendency of inequality to increase with age (Paglin 1975, Lillard 1977, von Weizsäcker 1989). Longer life has its

* Joshua R. Goldstein (corresponding author), University of California, Berkeley, CA, USA
Email: josh@demog.berkeley.edu

Ronald D. Lee, University of California, Berkeley, CA, USA
Email: rlee@demog.berkeley.edu

DOI: 10.1553/populationyearbook2014s193

own effect on inequality through changes in the economic life cycle (Lee and Goldstein 2003). In considering these three factors, our goal is not to discuss all of the demographic factors that might influence inequality. For example, among the interesting aspects of population aging that we do not cover are the effects of longer life and lower fertility on inheritances. Other features like differential fertility, international migration, and the role of intergenerational transfers are also important potential mechanisms through which demographic change may influence economic inequality.

In this paper, we seek to connect demography with economic inequality. This approach is similar to those which have attempted to link population growth to economic growth and savings. For example, the Solow model shows how slower population growth with constant savings can lead to capital accumulation and higher living standards, and Piketty makes the further argument that higher capital intensity is associated with higher levels of inequality. Where Bloom et al. (2001) and others have argued that the demographic dividend in savings that accompanies the demographic transition will come to an end as populations age, we investigate the compositional consequences of these same age structure changes for inequality. Finally, as rising longevity increases life-cycle savings and thus the amount of capital in the economy, we explore here the consequences of longer time horizons for changes in the distribution of capital.

Methodologically, we use simple approaches to describe and quantify these different effects. In the first section on capital intensification, we show the compositional effects of changing the factors of the economy on the inequality of total income from assets and labor earnings. A good approximation of the compositional effect can be obtained simply by changing the shares in a weighted average of earnings and capital-based income. In the second section, we apply Lam's (1984) stable population theory results to the age profiles of inequality in income and net worth, constant at levels observed in the 2001 U.S. Consumer Expenditure Survey (Federal Reserve 2013). In the last section, we estimate the effect of increased life-cycle savings by stretching out current schedules in a stylized manner consistent with increases in longevity.

Among the many possible measures of inequality, we focus on the share of wealth held by the top decile of the population. This measure is readily interpretable, and makes comparison with the work of Piketty and his colleagues straightforward. The share of the top decile also turns out to be readily estimable from other information on means, medians, and variances, by assuming a log-normal distribution of income and assets. Although not applicable to the richest rich, the assumption of log-normality is reasonable enough for studying the holdings of the top 10 percent. The appendix provides a derivation of our estimator for the holdings of the top decile.

2 Kinds of inequality

Both cross-sectional and longitudinal measures of inequality are of potential interest. The most common measures of inequality are taken in the cross section, at a moment in time. Measures like the variance of assets or the share held by a given upper fraction are usually made in reference to the population as it is observed at a moment in time, across all ages.

From a welfare point of view, it can make more sense to compare inequality over the life cycle. Levels of inequality within a given age group or over the life cycle might be considerably lower than the levels observed in the cross section (Lillard 1977). On the other hand, cross-sectional inequality may have considerable salience from a psychological point of view. A 20-year-old may well feel disadvantaged relative to a 40-year -old, even if she knew with certainty that she will eventually reach age 40 herself. This feeling can be attributed partly to impatience with having to wait 20 years, partly to uncertainty, and partly to human psychology and the inability of people to make fully compensating comparisons over time.

Cross-sectional inequality may also be of substantial importance because of its role in determining power. In politics, individual votes, campaign contributions, and other influence is all cross-sectional; with the individual's relative power being determined by the amount of influence others have at that moment.

Finally, market prices are determined largely by supply and demand at a moment in time. Intertemporal substitution is costly and uncertain. Credit markets are not perfect. Thus, the distribution of resources at a moment in time will influence prices.

3 Capital intensification

Personal or household income is the sum of labor income and asset income. The distribution of income therefore depends on the size and distribution of both labor and asset income—and their covariance. According to Piketty, because wealth is far more unequally distributed than labor income, increasing capital intensity generally results in greater income inequality. "The most important factor [determining capital intensity] in the long run is slower growth, especially demographic growth, which, together with a high rate of saving, automatically gives rise to a structural increase in the long-run capital/income ratios, owing to the law $\beta = s/g$." (Piketty 2014:173) This assertion is based on the assumption that saving rates are constant while demography varies. The same long-run relationship or 'law' can be readily derived from the Solow growth model and is quite general.[1] In the long run the growth rate g

[1] For example, in the Solow models steady state, in order to keep capital per head constant, $sf(k) = gk$, where s is the savings rate, k is capital per worker, and $f(k)$ gives output per worker as a function of capital per worker. It follows that $k/f(k) = s/g$. $k/f(k)$ is beta.

of National Income equals the rate of productivity growth plus the rate of population growth, n.

Accordingly, changing population growth rates alter the capital intensity. The annual population growth rates of the rich industrial nations have varied between −0.5% and +1% in recent decades, and are projected to remain at these levels in the coming decades. The annual productivity growth rates of these countries are expected to be around 1.5%. Piketty assumes that the average saving rate will be 10%. Thus, with $n = 1\%$, $\beta = s/g$ will be $4 = 10/(1 + 1.5)$. With $n = 0$, β will be $6.7 = 10/(1.5)$. And with $n = -0.5\%$, $\beta = 10$. Thus, we see that the population growth rate can play a very important role. Because many of the observed variations in the population growth rate are relatively short-term fluctuations attributable to, for example, baby booms and busts, their effects will be muted. But over the long term, population growth rates may also be expected to vary by around one percent.

These results, which link slower growth to capital intensification, are straightforward. Current wealth is accumulated from past savings. If economic growth (from productivity and population) is rapid, then these savings will have been accumulated starting from a base that is low relative to current income. Thus, current wealth levels will be low relative to current income levels. However, this tells us nothing about issues such as why people save or how wealth is transmitted. If people have very short lives, then wealth will be largely inherited. If people have infinitely long lives, then there will be no inherited wealth.

The results also rest on assumptions that may or may not be true. One portion of savings consists of life-cycle savings, or funds that are accumulated to provide income for retirement. This portion of savings should be strongly influenced by fertility, longevity, and population age distribution. Furthermore, what matters in this context is net savings, after allowing for the portion of savings that is needed to maintain or replace capital that wears out; that is, gross savings minus depreciation. The rate of depreciation should depend on the age of the capital stock. When the economy is growing rapidly (g is high), then the capital stock will be young and – all other things being equal – have a lower depreciation rate; and when g is low, the capital stock will be old and have a higher rate of depreciation. These considerations cast some doubt on the assumptions underlying the law of long-term capital intensity.

Others have discussed the assumptions underlying Piketty's formulation elsewhere. Here, we set these complications aside and instead focus on estimating the magnitude of Piketty's compositional effect of capital intensity on inequality.

Consider the hypothetical case in which population growth in the United States falls from its recent historic rate of about $n = 1$ percent per year to zero percent. As our calculation above showed, this increases the steady-state capital/income ratio from its current level reported by Piketty of about four to close to 6.7.

To estimate the effect of this change in the capital labor ratio on income inequality, note that the share of income attributable to returns on capital (Piketty's α) is $r\beta$, the product of the rate of return to capital times the capital/income ratio (Piketty).

Taking $r = 0.05$, this gives us a share of capital income to total income of 0.2 for $\beta = 4$ and 0.34 for $\beta = 6.7$.

If we assume the same ranking of labor earnings and capital earnings (i.e. a perfect correlation), then the share of total income held by the top 10 percent will be a simple weighted average of the shares held by the top 10 percent of labor earners and the top 10 percent of earners from capital. We denote the holdings of the top decile as $H_{.1}$, with superscripts l referring to labor income, k capital income and $l + k$ referring to total income. The relationship between these when the correlation is perfect is

$$H_{.1}^{l+k} = (1 - r\beta)H_{.1}^{l} + r\beta H_{.1}^{k}, \tag{1}$$

where $r\beta$ is the share of total income from capital and $1 - r\beta$ is the share from labor.[2] Differentiating with respect to β, assuming, as per Piketty, a constant rate of return, then gives us

$$\frac{dH_{.1}^{l+k}}{d\beta} = r(H_{.1}^{k} - H_{.1}^{l}). \tag{2}$$

Table 1 provides the top decile shares of labor income, asset income, and total income reported by Piketty for different stylized inequality regimes. It also shows the effect on the top decile share of total income of a unit increase in β, using the above result. We see that the effects are larger for more unequal societies, reflecting the tendency for capital income to concentrate more than labor income.

The value of this derivative for the United States allows us now to state the estimated effect of a one percent slowdown in population growth. As we saw above, this change in population growth increased β from 4.0 to 6.7; and we can now say that this would increase the share held by the top decile by about five percentage points, or $(6.7 - 4.0) * 0.0175$.

This kind of calculation represents an upper bound for the composition effect, because we have assumed a perfect correlation between labor income and asset income, as well as a perfect correlation between the existing asset income at a given β and any new asset income implied by a higher value of β. The correlations are likely to be high; but the lower they are, the smaller the compositional effect will be. Simulation with uncorrelated labor and asset income suggests that the effect on the share held by a top decile of a unit increase in β would be about one percentage point, or slightly more than half of the 1.75% found for the case of perfect correlation. [3] As Piketty (pages 244–246) has argued that the correlation between labor and asset income is quite high in modern industrial societies, we believe that a reasonable estimate of the derivative of the top decile's share with respect to β would be close to the case of perfect correlation, or about 1.5 percent.

---

2 These accounting identities are discussed by Piketty on page 52. The addition of top deciles of labor and asset income holds because of the assumption of perfect correlation in the two types of income.

3 Our simulation was based on log-normal distributions of labor and capital, with the appropriate ratio of mean labor income to mean asset income.

Table 1:
Piketty's estimates of labor and asset income received by the top decile for various inequality regimes with our estimate of the effect on total income of a unit increase in the capital/income ratio β

	Low inequality (Scandinavia, 1970s)	**Medium inequality (Europe, 2010)**	**High inequality (US 2010, Europe 1910)**	**Very high inequality (US 2030?)**
Labor income (H_l)	20%	25%	35%	45%
Asset income (H_k)	50%	60%	70%	90%
Total income (H_{l+k})	25%	35%	50%	60%
Effect of β increase ($dH_{l+k}/d\beta$)	1.50%	1.75%	1.75%	2.25%

Note: For example, if β were to increase by 2.0 from a 'low inequality' baseline, then the top decile share of income (H_{l+k}) would increase from 25% to 28% (2.0 × 1.50%). The first three lines of this table are from Piketty (p. 247-249). The derivative is our calculation based on change in weighted average of top decile share of labor earnings and capital earnings, assuming new capital earnings are perfectly correlated with existing capital earnings.

Our calculation suggests that the capital intensification accompanying a hypothetical end to population growth in the United States would increase the total income of the top decile by about four to five percentage points. This increase in inequality is substantial, but it is still smaller than the differences across inequality regimes, which are on the order of 10 to 15 percentage points of total income held by the top decile (shown in the third row of the table). Our finding that slow growing Europe would have lower levels of inequality than the more rapidly growing United States tells us that population growth is not the overwhelming determinant of inequality. Our calculations do, however, indicate that changes in population growth within countries will, if Piketty's formulation holds, indeed result in quantitatively important increases in inequality.

4 Shifting age structure

Over the course of the demographic transition, the population initially gets younger as population growth accelerates. Then, as fertility falls, there is a transitional period during which there are relatively few children and elderly people, and many people of working age. As fertility remains low, the people who had been of working age grow older, and the population rapidly ages. At the end of the transition, the age structure of the population becomes similar to that of a stable population with low fertility.

The period in which the share of the population who are of working age is growing and the share of the population who are dependents (young and old) is shrinking gives rise to the ‘demographic dividend’. This dividend in dependency rates can also be seen when we attempt to measure population-level inequality. Because inequality increases as cohorts age, a population with a relatively large share of young people will tend to be more equal. As population aging implies that a greater share of the population will progress to ages characterized by more inequality, aggregate measures of inequality may be expected to increase as the population grows older.

In considering the magnitude of the effect of shifts in age structure on population inequality, it is useful to begin with the standard decomposition of the population variance into the between and within variances of subpopulations, which in our case are defined by age groups. For an age-structured population with a share $c(x)$ at age x, the population variance decomposition is

$$\sigma^2_{\text{pop}} = \sum_x c(x)\sigma^2_{\text{within}}(x) + \sum_x c(x)(\bar{\mu} - \mu_x)^2, \tag{3}$$

where $\sigma^2_{\text{within}}(x)$ is the variance within age group x, $\bar{\mu}$ is the population mean, and μ_x is the mean for each age group.[4] The standard deviation is the square root of the variance.

Lam (1984) applied the formula above to stable populations and showed that the derivative of the log of the population variance is given by

$$\frac{d \log \sigma^2_{\text{pop}}(n)}{dn} = \alpha(\bar{x} - \bar{x}_b) + (1 - \alpha)(\bar{x} - \bar{x}_w), \tag{4}$$

where Lam’s α (unrelated to Piketty’s α) is the share of the total variance in equation (3) above that is between groups, $\bar{x}$ is the mean age of the economic quantity of interest (e.g. the mean age of log-earnings), $\bar{x}_b$ is the mean age weighted by each age group’s share of between-age variance, and $\bar{x}_w$ is the mean age weighted by each age group’s within-group variance.

Lam’s results show the role of two offsetting effects. Consider income profiles for which the mean and the variance both rise with age. Younger people tend to have relatively low incomes. If we increase population growth, the share of the population who are young will rise, and the share of the population who are of the ages at which incomes tend to be far below average will also increase. Thus, the between-group variance will increase. In this case, the between component of variance will increase

---

[4] This standard decomposition of the variance of a mixture of subpopulations has been applied to the age composition of population inequality by Lam (1984) and von Weizsäcker 1989).

the variance if population growth increases, and it will decrease the variance if population growth declines.

The variance within age groups will be in the opposite direction. If the variance within groups increases with age, then reducing population growth will result in a concentration of the population at older, higher within-group variance ages. Thus, a decline in population growth will cause this component of the variance to increase.

5 U.S. age profiles of inequality

Lam's approach can be applied to contemporary age profiles of income and asset accumulation. We first describe the profiles currently observed in the United States and then analyze the effect of declining population growth. For comparability, we consider the same scenario used for studying capital deepening, reducing the growth rate by one percent.

We obtain age profiles for inequality in the United States from the published tabulation of the Survey of Consumer Finance. These tabulations report the mean and the median pretax family income and family net worth by age group of the family head.[5] Using the log-normal approximation allows us to estimate all of the moments from the reported means and medians. In order to convert family income to individual income by age, we multiply each family-level quantity by the age-specific headship rate that we tabulated from the 2014 Current Population Survey.

Several features of Figure 1 are worth noting. First, the top row shows the life-cycle patterns we would expect to see from increasing earnings and savings during the working years and declining earnings and assets after retirement. Assets continue to increase longer than earnings, reflecting the returns on capital and the delays in drawing down during retirement. Substantial assets are left at age 80, which suggests that large bequests are likely. This figure does not show the much lower medians observed at each age.

Second, the pattern of variability of income and assets shown by the standard deviation in the second row is remarkably similar to the life-cycle profile average levels shown in the first row. In part, these are mechanical effects of scale. As income and assets increase, their variability also increases.

The third row shows the standard deviation of the log of pre-tax income. This is the σ parameter in the log-normal distribution. This measure of relative variability is closely related to the coefficient of variation. It reveals that the increase in variability is not just a function of a rising mean but also exists in relative terms. Indeed, the pattern for income is nearly linear and increases steadily up to ages 65–75. The pattern for net worth is more complicated, as net worth rises rapidly between ages

5 The micro data are available for more detailed tabulations and other definitions of stock and flow of wealth.

Figure 1:
Age profiles of income and assets estimated from the Survey of Consumer Finances, triennially from 1989 to 2013. The solid black lines are from 2001. All amounts are in 2013 constant dollars. Estimates are made from the reported means and medians reported in Federal Reserve (2014), assuming a log-normal distribution at each age

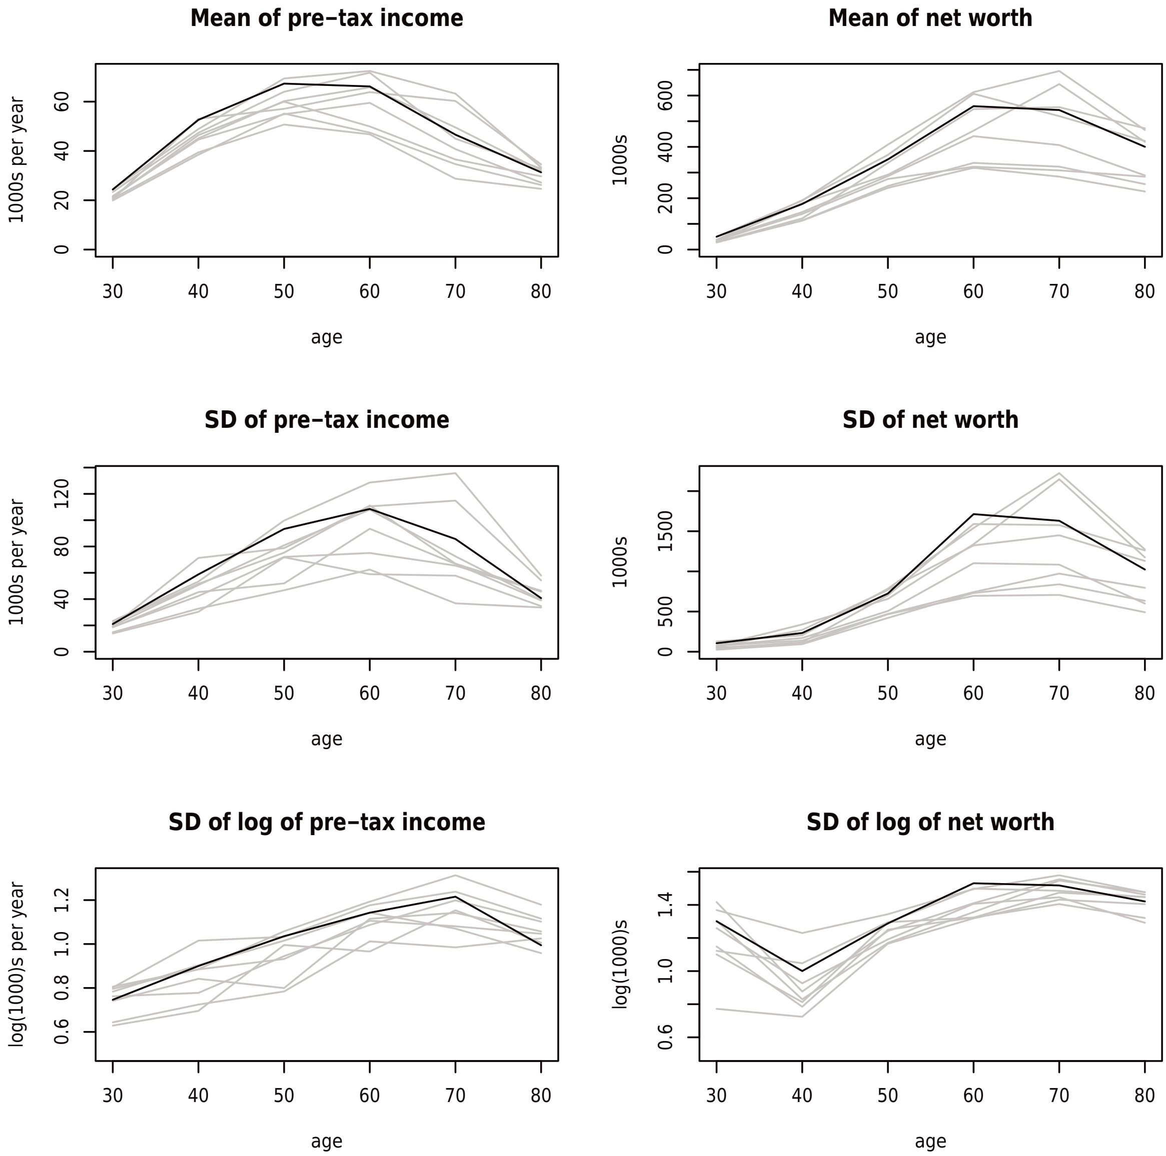

40 and 60, and then plateaus thereafter. Among the youngest age group net worth varies considerably, but this is relative to a near-zero base.

We now replicate Lam's analysis for the SCF profiles using the 2001 profiles shown in bold in Figure 1.[6] We decompose the variance of the logarithm of income

[6] We note that our analysis here does not account for the trend toward increasing inequality at a given age that can be seen in these schedules, as well as in those in other countries. See, for example, Bönke, Corneo, Lüthen (2015) for cohort changes in Germany.

Table 2:
The effect of a one percent decline in population growth on the variance of the logarithm of income and net worth and accompanying quantities from Lam's stable population analysis and our analysis using the 2001 U.S. Survey of Consumer Finances

	Total variance $n = 0$	Within age group component	Between age group component	Mean ages $\bar{x}$	x_w	x_b	Effect of 1% less pop. growth on total variance	on top decile share ($H_{0.1}$)
Income	1.11	1.01	0.10	51.9	55.8	48.8	+3.3%	+0.7%
Net worth	2.40	1.79	0.61	51.9	54.7	37.5	−1.5%	−0.3%

Note: The effect of 1% less population growth on total variance is obtained by calculating the derivative of log variance with respect to n using Lam's formula and then multiplying by $\Delta n = -1\%$. The effect on the top decile is obtained by estimating $\Delta\sigma$ as half of the change in the total variance and applying the log-normal approximation of the top decile shown in the appendix.

(and of wealth). Measures of welfare and utility are often more closely related to the logarithm of earnings or wealth. Moreover, a convenient property of the standard deviation of the logarithm is that it can be easily converted into the share held by the top decile, assuming a log-normal distribution.

For income, we find (see first line of Table 2) that a one percent reduction in the population growth rate increases the variance by about 3%. In terms of the share held by the top decile, this implies an increase of about 0.6 percentage points (e.g. from 50% to 50.6%).[7] The magnitude of this effect is detectable, but is still quite small. It is much smaller than the four to five percentage point increase in the top decile share of income that we estimated for the capital-deepening response to the same decline in population growth.

For net worth, we find (see second line of Table 2) that the aging of the population due to slower population growth actually decreases the variance by about 1.5%. This small negative effect results from a much stronger compressing effect of the between-age-group component of variance. Intuitively, we would expect that the population-level variability of assets would be compressed, because there would be fewer young people with very low asset levels. This effect can be seen in the low average age (37.5) of the between-age-group variance weighted mean age and in the larger share of the 'between' component of variance relative to the total.

Although there is intuitive appeal to the idea that shifting the population to ages of greater within-age-group inequality should increase aggregate inequality, the actual effects turn out to be quite small. The direction of the effect also appears to be highly

[7] Increasing the variance by 3% increases the SD by about 1.5%. The relationship between the share held by the top decile and the SD of the log of a log-normally distributed quantity is approximately $0.4 \times SD$. If the SD has a value of one, then the total effect on the holdings of the top decile will be an increase of 0.6 percentage points.

sensitive to the details of the age profiles of the mean and the variance. Although income and assets have the same general age profile (in Figure 1), the profiles differ enough that the relative importance of the within-group and the between-group variances in formula (3) can change substantially. Lam (1984) also found that the sign and the magnitude of the effect of changing population growth rates on inequality are susceptible to small differences in the earnings profiles and the age ranges under consideration. The older age structure of aging populations appears to be at most a minor driver of population-level measures of inequality.

6 Stretching the economic life cycle

Although the major cause of population aging is declining fertility, increases in longevity are a key factor in the individual economic life cycle. Life expectancy at birth is increasing by about 0.15 years per year in the United States, and life expectancy at age 60 is increasing at about 0.1 years per year. In Japan, life expectancy is increasing much more rapidly, with period life expectancy increasing by about 0.25 years per year. In the next half-century, we can expect to see adult longevity increase by some five to 10 years.

As a cohort ages, there is a more opportunity for their incomes and assets to fluctuate randomly. Thus, a cohort tends to see increases in within-age-group variance as it gets older. As increases in longevity extend the period of time over which fluctuations can occur, they are generally associated with increases in both life-cycle and population-level inequality.

The effects of increasing longevity on the life-cycle patterns of earnings, consumption, and savings are complex. It is, however, easy to perform a simulation in which we modify our observed schedules of mean and variance by linearly extrapolating the extra five years of working life that we expect to gain over the next half-century.

Figure 2 shows the result of such a hypothetical stretching of the economic life cycle. We apply the stretching to both the means and the standard deviation. For income, the result of stretching is to slow the decline in the mean and the SD. For net worth, the result is to extend the period of accumulation, leading to a higher peak mean net worth and to higher peak variation.

We simulate the consequences of these new schedules by calculating the aggregate inequality implied by the original and the stretched schedules. The results of this analysis are shown in Table 3.[8]

We can see from this simulation that the effect of stretching the schedule is considerably greater for assets than for income. This makes sense given the profiles shown in Figure 2, in which the extrapolated increase beyond age 60 is much larger

[8] Our analysis here does not take into account the induced changes in interest rates or the wage rates; this simplification is consistent with treating the U.S. economy as open to international forces.

Figure 2:
Original and longevity-stretched age profiles of income and assets estimated from the Survey of Consumer Finances, 2001 (in 2013 constant dollars)

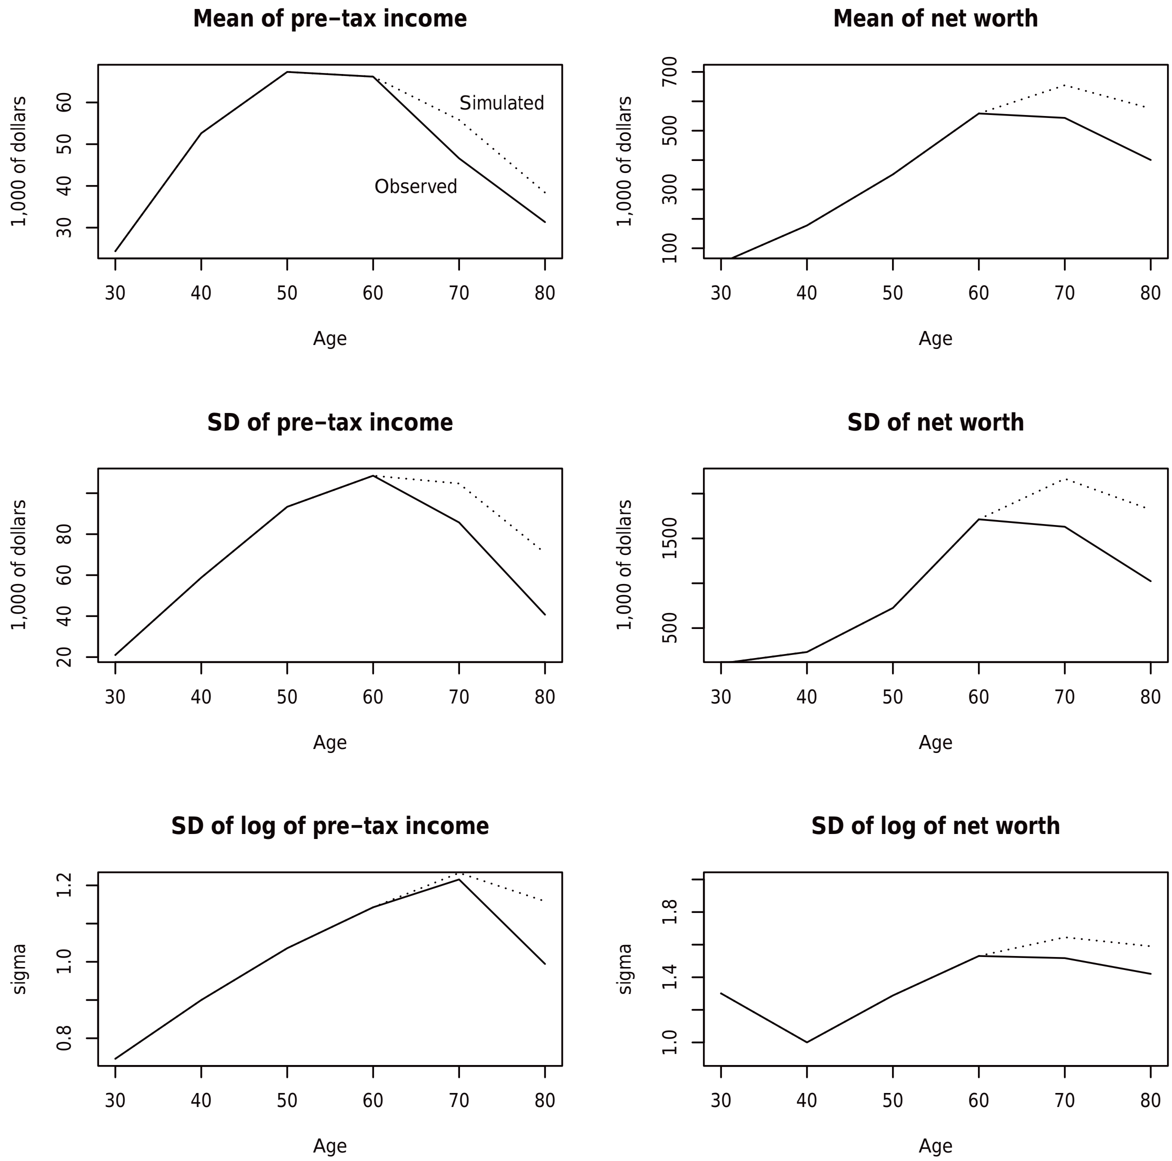

Note: Stretched profiles with an additional five years of longevity are obtained by extrapolating the trend from ages 50 to 60 and an additional five years to age 65. See the text for details.

than the original pattern of increase predicted for assets. This result is plausible, as we can imagine that living longer will have a greater effect on inequality through the compounding of random shocks to assets than through the compounding of shocks to earnings.

In Table 3, we see that the effect of the approximately five-year increase in longevity is a one to two percent increase in inequality, as measured by the share held by the top decile. While this is not a trivial effect, it is considerably smaller than the four to five percent increase in the share held by the top decile that we simulated

Table 3:
Inequality (top decile share) implied by the original and the stretched schedules of income and net worth

	Pre-tax income	Net worth
Original	44.6%	62.8%
Stretched	45.3%	64.9%
Difference	0.7%	2.1%

Note: The original schedule is based on the 2001 Survey of Consumer Finance, as shown in Figure 2. The stretched schedule extrapolates five additional years of the trend linearly from the observed values at ages 50 and 60, as shown in Figure 2. Inequality is estimated by applying these schedules to the stationary populations using the life tables of the United States in 2000, and a simulated version of the 2050 life table (note that the life expectancy in the 2050 life table is five years longer).

from falling fertility and the capital deepening that would result from a 1% decrease in the population growth rate.

7 Discussion

In this paper we have estimated the magnitudes of some of the important effects of demographic change on aggregate economic inequality. Our purpose has been to gain insight into the relative importance of several different mechanisms.

The first pathway we examined was the increase in capital intensity that accompanies a slowdown in population growth. To estimate the magnitude of this 'Piketty' effect, we calculated the increase in capital his model would predict for the United States if population growth were to slow by about one percent. We found that this effect would produce a significant increase in income inequality, raising the share of income held by the top decile by about four to five percentage points. To put this increase in perspective, we should note that the top decile in the United States now has about 50 percent of income. An increase of five percent in the earnings of the top decile would take the country about a third of the way to what Piketty calls "very high inequality," in which 65% of income is held by the top decile.

The second pathway of demographic change we looked at was the shift in the age structure toward the older and more unequal ages that may be expected to accompany a slower population growth rate. This effect is, as David Lam found in his development of the subject, less clear-cut than it might at first seem. This is partly because the increases in inequality with age are not so enormous that a change in population composition has a large effect and partly because changes in population-level inequality in the within-age-group and the between-age-group components partially offset one another. As there are more old people, the population shifts toward ages at which there are higher levels of within-age-group inequality. But at the same time, the presence of fewer young people pushes down the numbers of those with earnings that are far below average, which tends to lessen between-age-group inequality. Applying Lam's results for stable populations to the United States schedules of income and assets in 2001, we find that the net result of a one percent decline in population growth would increase the share of income held by the top decile by something like one-half of one percent. This 'Lam' effect might be detectable, but it appears that it would be much smaller than the 'Piketty' effect.

Finally, we consider the consequences of increased longevity on longer periods of life-cycle savings (Lee and Goldstein 2003). We crudely simulate the effect of extending life by the five years that are forecast over the next half century by stretching out the schedules we observe for 2001. We extrapolate the age trend observed from 50 to 60 an additional five years out to age 65, inserting this additional period of earnings and capital accumulation into the economic life cycle. We find that the aggregate inequality implied by these longer life profiles is an increase in the share held by the top decile of around one to two percent. The size of this longevity effect is in-between the 'Lam' age structure effect and the 'Piketty' capital intensity effect, both of which can be seen as consequences of changing fertility.

Together, the three mechanisms we explored could account for around seven percent of the increase in the share of income held by the top decile. This would be a substantial increase in inequality.

Income inequality expanded between 1970 and 2010, as the share of income held by the top decile of the population increased by about 20% in the United States and by about 5% in Europe (Piketty p. 324, Figure 9.8). We know that demography cannot explain the differences between Europe and the United States, since the United States has faster growth rates and a younger population than Europe. However, if we apply our results of the potential impact of demographic change within these two regions, we would expect to observe that population aging in the United States will lead to substantial increases in inequality in that country. In Europe, the same magnitude of change would be even more dramatic, more than doubling the increase in inequality seen in recent decades.

In this analysis, we have considered each of three factors independently. We considered capital deepening without taking age structure into account. We applied age structure profiles to changing populations without taking macroeconomic constraints into account. Finally, we considered age profiles of inequality without taking macroeconomic constraints into account. A more complete modeling approach would consider capital deepening in the context of the population age structure and the economic life cycle, rather than in the ageless context of Solow's neo-classical growth model.

Steps toward the development of more integrated models have been made by Lee, Mason, and Miller (2003) and by Romero-Sanchez (2013). Their approaches include the consideration of intergenerational transfers (notably, public pension programs), general equilibrium interest rate effects, and the combined forces of fertility decline and longevity increase. A key assumption in each of these models is how the age at retirement reacts to increases in longevity. Without increases in the retirement age, the increase in life-cycle savings needed for retirement is quite large, and produces even larger increases in the capital output ratio (β) than those considered here. Another important consideration is the role played by public pay-as-you-go transfer systems. These benefits can have the effect of replacing life-cycle savings and lessening β and are redistributive (or, at most, proportional to labor earnings). Thus, transfer programs can reduce the inequality-increasing effects of population aging.

We expect that the resurgence of interest in the topic of economic inequality sparked by Piketty will inspire many studies of the demographic influences on economic inequality. Our initial foray into this field suggests that there is indeed room for a substantial compositional increase in inequality, assuming Piketty is right about increased capital intensification and its compositional effect on inequality. However, the direct age structure effect appears to be of minor importance. Finally, the increase in the average lifespan also appears to have a substantial impact on life-cycle savings. Considered together, these findings give us reason to suspect that population aging will exacerbate the increases in inequality seen in recent years, strengthening the case that we should be concerned about who owns what in the 21st century.

References

Bloom, D. E., D. Canning and J. Sevilla 2001. Economic growth and the demographic transition. NBER Working Paper 8685.

Bönke, T., G. Corneo and H. Lüthen 2015. Lifetime earnings inequality in Germany. *Journal of Labor Economics* 33: 171–208.

Cowell, F. S. 2009. *Measuring Inequality.* LSE Perspectives in Economic Analysis, http://darp.lse.ac.uk/MI3, accessed May 2015.

Federal Reserve 2014. *2013 SCF Chartbook*, http://www.federalreserve.gov/econresdata/scf/files/BulletinCharts.pdf, accessed May 2015.

Lam, D. 1984. The variance of population characteristics in stable populations, with applications to the distribution of income. *Population Studies* 38: 117–127.

Lee, R. D. and J. R. Goldstein 2003. Rescaling the life cycle: longevity and proportionality. *Population and Development Review* 29(Supplement): 183–207.

Lee, R. D., A. Mason and T. Miller 2003. Saving, wealth, and the transition from transfers to individual responsibility: the cases of Taiwan and the United States, *The Scandinavian Journal of Economics* 105(3): 339–357.

Lillard, L. A. 1977. Inequality: Earnings vs. human wealth. *American Economic Review* 67: 42–53.

Paglin, M. 1975. The measurement and trend of inequality: A basic revision. *American Economic Review* 65: 598–609.

Piketty, T. 2014. *Capital in the Twenty-First Century* (translated by A. Goldhammer). Cambridge, MA: Harvard University Press.

Romero-Sanchez, M. 2013. The role of demography on per capital output growth and savings rates. *Journal of Population Economics* 26(4): 1347–1377.

Solow, R. M. 1956. A contribution to the theory of economic growth. *Quarterly Journal of Economics* 70(1): 65–94.

von Weizsäcker, R. K. 1989. Demographic change and income distribution. *European Economic Review* 33: 377–388.

Appendix: Estimation of top decile share of log-normal distribution

Figure A.1:
Accuracy of linear approximation of top decile share of a log-normally distributed variable by value of σ

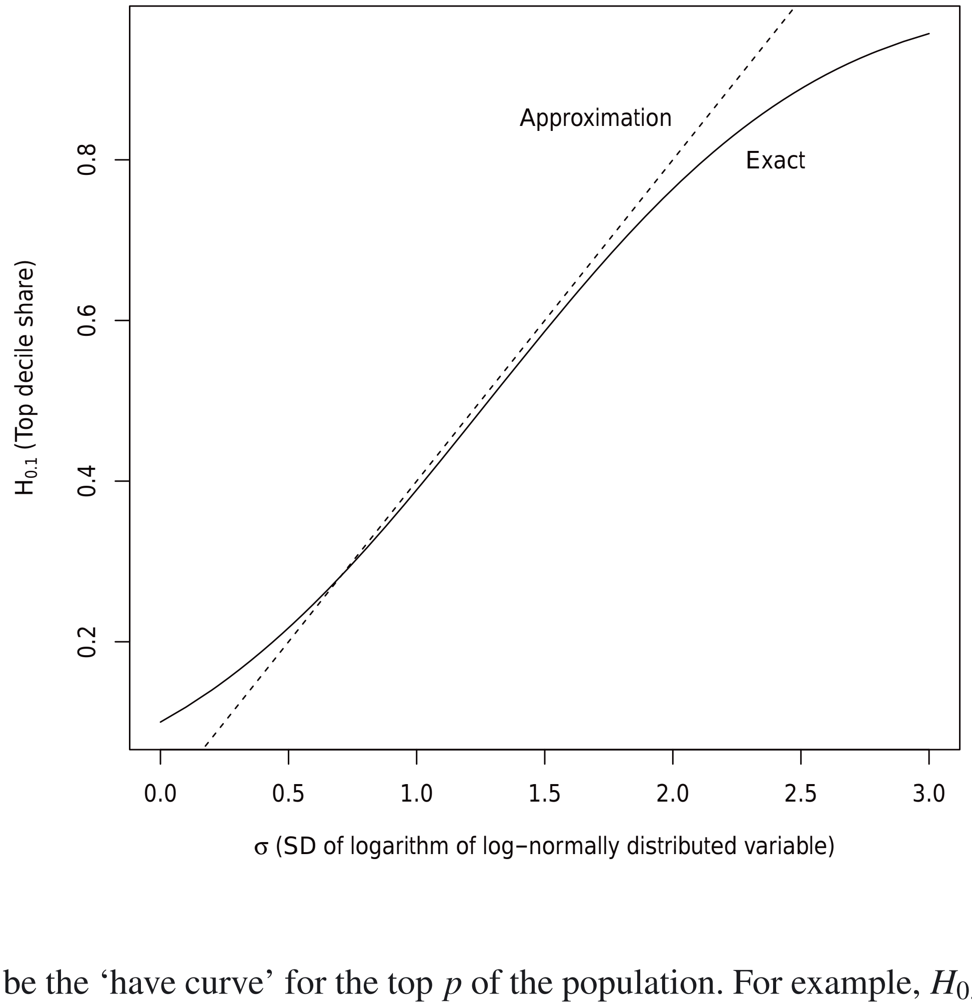

Let H_p be the 'have curve' for the top p of the population. For example, $H_{0.1}$ would be the share held by the top 10 percent.

When a quantity is log-normally distributed, we can approximate the share held by the top decile quite accurately for values of σ ranging from 0.5 to 2.0, the range of economic inequality seen in many populations. The approximation is

$$H_{0.1} \approx \frac{1}{\sqrt{2\pi}}\sigma = 0.40\sigma. \tag{A.1}$$

The accuracy of this approximation can be seen in Figure A.1.

To derive, we first express the 'have' function H in terms of the more common Lorenz function L for the share held by the bottom fraction of the population. This gives us $H_{0.1} = 1 - L_{0.9}$, where $L_{0.9}$ is the share held by the bottom 90 percent. The Lorenz curve for the log normal is known to be

$$L_{0.9} = \Phi(\Phi^{-1}(0.9) - \sigma), \tag{A.2}$$

where Φ is the cumulative distribution function of the standard normal (Cowell 2009, p. 154).

The linear approximation of the cumulative distribution of the standard normal for values x above zero is

$$\Phi(x) \approx \frac{1}{2} + \frac{x}{\sqrt{2\pi}}. \tag{A.3}$$

In our case, we use the approximation (A.3) to estimate the Lorenz function (A.2). Substituting $x = \Phi^{-1}(0.9) - \sigma = 1.28 - \sigma$ into (A.2) gives

$$L_{0.9} \approx \left(\frac{1}{2} + \frac{1.28}{\sqrt{2\pi}}\right) + \frac{\sigma}{\sqrt{2\pi}} \approx 1 + \frac{\sigma}{\sqrt{2\pi}}. \tag{A.4}$$

Substituting back into H gives the desired result in (A.1).