NHH
INSTITUTT FOR SAMFUNNSØKONOMI
DEPARTMENT OF ECONOMICS

SAM 6 2010
ISSN: 0804-6824
FEBRUARY 2010

Discussion paper

Baby Booming Inequality? Demographic Change and Inequality in Norway, 1967-2004

BY
INGVILD ALMÅS, TARJEI HAVNES, AND **MAGNE MOGSTAD**

Norges
Handelshøyskole
NORWEGIAN SCHOOL OF ECONOMICS AND BUSINESS ADMINISTRATION

Baby Booming Inequality? Demographic Change and Inequality in Norway, 1967–2004*

Ingvild Almås[†], Tarjei Havnes[‡] and Magne Mogstad[§]

February 22, 2010

Abstract

We demonstrate how age-adjusted inequality measures can be used to evaluate whether changes in inequality over time are because of changes in the age structure. In particular, we explore the hypothesis that the substantial rise in earnings inequality since the early 1980s is driven by the large baby boom cohorts approaching the peak of the age–earnings profile. Using administrative data on earnings for every Norwegian male over the period 1967–2004, we find that the impact of age adjustments on the trend in inequality is highly sensitive to the method used: while the most widely used age-adjusted inequality measure indicates that the rise in inequality in the 1980s and 1990s is indeed driven partly by the baby boom, a new and improved age-adjusted measure indicates the opposite, namely that the rise in inequality was even larger than what the inequality measures unadjusted for age reveal.

JEL: D31, D63, D91, E21.
Keywords: Inequality trend, age structure, age–earnings profile, Gini coefficient.

1 Introduction

The rise in earnings inequality in almost all developed countries since the early 1980s is one of the most extensively researched topics in economics. While there is substantial agreement about the facts, there is no consensus about the underlying causes. A number of explanations have been proposed and scrutinized, including skill-biased technical change,

*Thanks to seminar participants at "Inequality: New Directions" at Cornell University for useful comments and suggestions. Financial support from the Norwegian Research Council (194339) is gratefully acknowledged. This project is part of the research activities at the ESOP center at the Department of Economics, University of Oslo. ESOP is supported by The Research Council of Norway.

[†]Norwegian School of Economics and Business Administration, ingvild.almas@nhh.no.
[‡]University of Oslo, Dept. of Economics
[§]Statistics Norway, magne.mogstad@ssb.no.

international trade and globalization, and changes in labor market institutions (such as a decline in unionization and an erosion of the minimum wage).[1] In this paper, we investigate an alternative, demographic explanation: how much does the changing age structure matter for the trend in inequality? Is the substantial rise in inequality since the early 1980s driven by the baby boom cohorts approaching the peak of the age–earnings profile?

These questions arise from two stylized facts. First, there is a strong age–earnings relationship. Second, we have seen substantial changes in the age structure over time. Both theoretical models and empirical results suggest a strong relationship between age and earnings (see, e.g., Heckman et al., 2003). In particular, the age–earnings relationship is firmly established as increasing during the working lifespan and usually declining slightly when approaching retirement. This implies that inequality of earnings at a given point in time is likely to be present even in an economy where everyone is equal in all respects but age. Furthermore, almost all developed countries experienced a large increase in their population growth rate following World War II—more familiarly called the baby boom.[2] Since the 1970s, the baby boomers have gradually entered the labor market, and as their careers mature, they are making their way up the age–earnings profile. Together, the changing age structure and the strong age–earnings relationship may be an important determinant of the observed trends in earnings inequality over recent decades. Identifying the age effects on inequality and their trend over time is also of interest from a normative perspective. It has long been argued that inequality attributable to age should be of little concern for policymakers: differences arising from age disappear over time and are, therefore, irrelevant for the distribution of lifetime income (see, e.g., Atkinson, 1971).

In this paper, we examine empirically to what extent the changing age structure can explain the trends in earnings inequality in Norway during the period 1967–2004. Specifically, we adjust the trends in inequality for changes in the age composition of the working-age population, using data from administrative registers on earnings for every Norwegian. Our analytical sample is restricted to males, given their role of primary breadwinner over most of this period.

In some respect, our approach follows Paglin's (1975) pioneering paper, which first raised the question of the effects of the changing age structure on the trend in inequality. While the validity of Paglin's method for isolating the age effect on inequality has been questioned from a number of perspectives—which we address in our analysis—the issue of isolating the age effect on inequality remains an important research question. In fact, given the rise in inequality accompanying the aging of the baby boom cohorts over the last

[1] See e.g. Gottschalk and Smeeding (2000) and Lemieux (2008) for extensive reviews of the literature on earnings inequality.

[2] The baby boomers usually include children born from 1946 to about 1960. For example, The US Census Bureau considers a baby boomer to be someone born during the demographic birth boom between 1946 and 1964.

decades, the issue may be viewed as more important than in the earlier period (1947–1972) considered by Paglin and others.[3]

This paper proceeds as follows. Section 2 reports some stylized facts about the age–earnings profile and the age structure in Norway, linking them to the observed trends in earnings inequality. Section 3 sets out the methods used to identify and adjust for age effects. Section 4 describes the data and presents the empirical results. Section 5 concludes.

2 Stylized facts

Age–earnings profiles are widely used by economists, both to help forecast the course of future earnings and to depict how earnings typically change over the life cycle. Panel (a) in Figure 1 draws the age–earnings profile in 2001 for Norwegian males aged 25–59, confirming the picture from other developed countries: average earnings rise rapidly at younger ages, peak around age 50, and then decline slightly in the latter parts of the working life.[4] Similar profiles are evident throughout the period 1967–2006. The strong relation between age and earnings implies that earnings inequality in a given year may be present even in an economy where everyone is equal in all respects other than age, simply because individuals are at different stages in the life cycle. For example, the earnings of a 50-year-old are about 40 percent higher than that of a 30-year-old on average, but that does not mean that the lifetime earnings of 50-year-olds are any higher than the lifetime earnings of 30-year-olds.

Panel (b) in Figure 1 graphs the size of individual cohorts of Norwegian males from the total resident population aged 25–59 in 1970, 1980, 1990, and 2000. We immediately see the relatively small birth cohorts before and during World War II, and the subsequent boom in population growth. This demographic shift has manifested large changes in the age composition of the labor force over the course of the previous decades. In particular, the baby boomers have gradually entered the labor market since the 1970s, and as their careers have matured, they have made their way along the age–earnings profile. It follows that inequality in annual earnings may change over time simply because of changes in the age structure, affecting different parts of the age–earnings profile differently, even though inequality in lifetime earnings may be unchanged.

Figure 2 illustrates the large amount of ballast that may be embedded in snapshots of earnings inequality and its time trend, because of the changing age structure. As a

[3]Paglin's age adjustment of the Gini coefficient was subject to three rounds of comments and replies in the *American Economic Review* (Paglin, 1977, 1979, 1989), has numerous citations, and continues to be subject to controversy. For a review of the literature, see Almås and Mogstad (2009).

[4]Throughout this paper, earnings are adjusted for wage growth. This is implemented by using the basic amount thresholds of the Norwegian Social Insurance Scheme (used to define labor market status, determining eligibility for unemployment benefits as well as disability and old age pension). The basic amounts are adjusted for wage growth by parliament in the national budget each year.

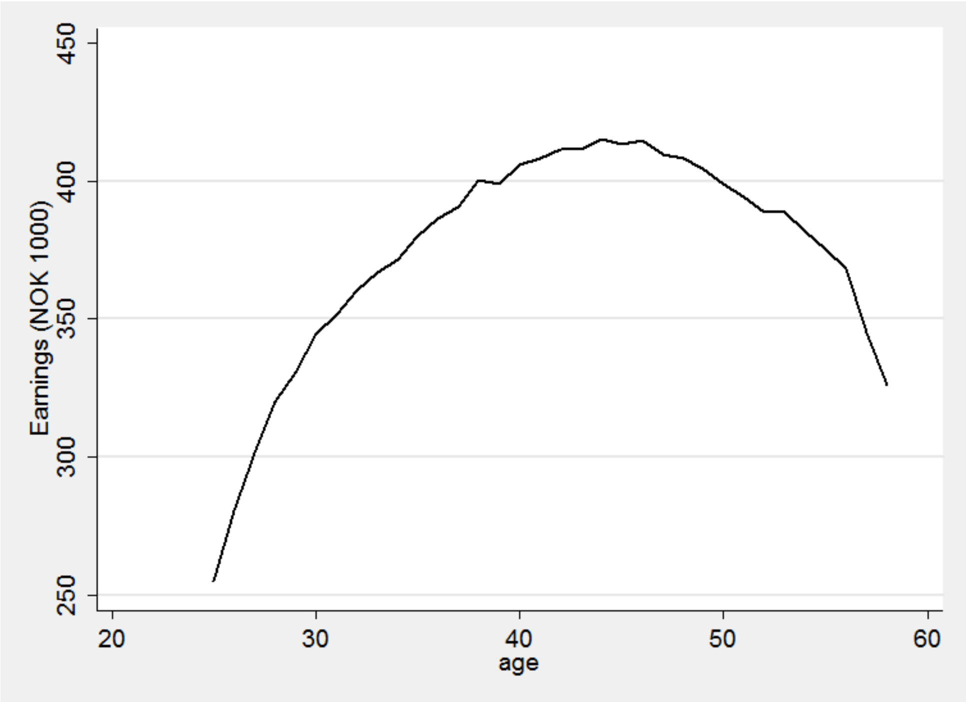

(a) Earnings profile in 2000

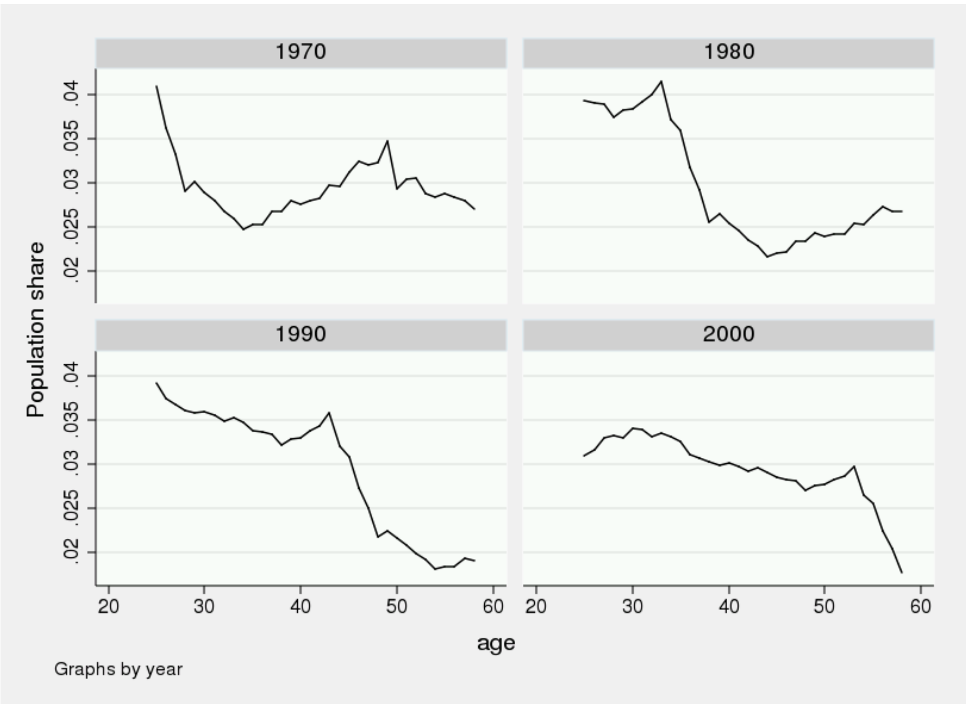

(b) Age composition in 1970, 1980, 1990 and 2000.

Figure 1: Earnings profile (panel (a)) and the changing age composition of the labor force (panel (b)) for males aged 25–59.

benchmark, we compute the classical Gini coefficient (G) in annual earnings of Norwegian males aged 25–59 over the period 1967–2004. Similar to the situation in most other developed countries, inequality fell slightly during the 1970s before increasing over the 1980s and into the early 1990s. It should be noted that the spike in inequality in the early 1990s was associated with tax reform, and that inequality would be likely to have increased more steadily in the absence of this reform (see, e.g., Fjærli and Aaberge, 2000).

Consider instead inequality in a hypothetical situation where everyone in the economy is equal in all respects other than age: while earnings vary over the life cycle, every individual at a given stage of the life cycle would have exactly the same earnings as others at that stage; i.e., everybody would follow the same age–earnings profile and thus have equal lifetime earnings. First, we compute the between-group Gini coefficient (G_b) in 2000, which captures the inequality in the age–earnings profile in Figure 1. Specifically, G_b replaces the earnings of each individual with his age-group mean. We find that G_b makes up a little under 25 percent of overall inequality in G in 2000. This illustrates that the age–earning relationship may make us confuse older individuals with richer individuals; G incorporates substantial cross-sectional inequality that evens out over time.

To get a sense of how inequality in this hypothetical situation would have evolved as the age structure changes, suppose that the age–earnings profile was the same in each year as in 2000. Let $\widetilde{G_b}$ denote the between-group Gini coefficient when we replace the earnings of each individual in every age group with the relevant age-group mean in 2000. We can see that the trend in $\widetilde{G_b}$ mirrors well the shape in G. This indicates that the changing age structure may have been an important determinant of the observed trend in earnings inequality. However, as will be apparent in Section 3, this exercise is far too stylized to draw inference about the age effect on earnings inequality and its trend, both because inequality within and across age groups is likely to be changing over the period, and because overlap in earnings between age groups complicates the picture. Yet, it serves as motivation to take a closer look at how much the changing age structure matters for the trend in inequality. That is the focus of the rest of the paper.

3 Age adjustment of inequality

Empirical analyses of inequality in earnings distributions are conventionally based on the Lorenz curve. To summarize the information content of the Lorenz curve and to achieve rankings of intersecting Lorenz curves, the classical Gini coefficient (G) is often used, which is equal to twice the area between the Lorenz curve and its equality reference. In a seminal paper, Paglin (1975) argues that G overspecifies the conditions of equality when used with annual data: assuming for the moment no economic growth, perfect equality requires not only equal lifetime income, but also that individuals of all ages must have equal earnings in any given year, which can be realized only if there is a flat age–earnings

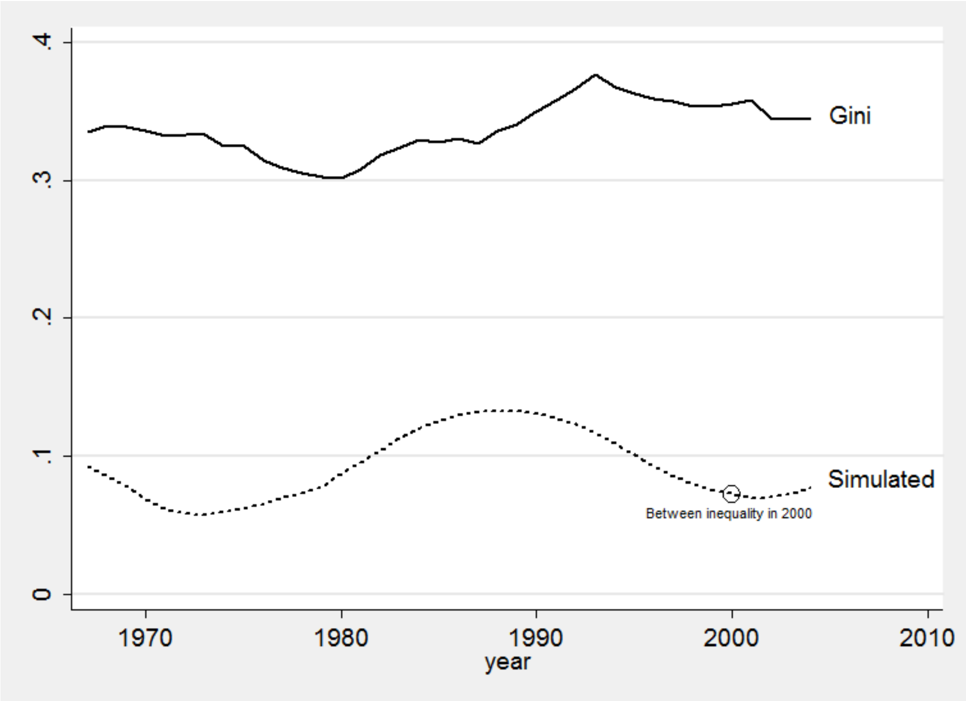

Figure 2: Trends in earnings inequality among males aged 25–59 over the period 1967–2004
Notes: Simulated = $\widetilde{G_b}$, between inequality in 2000 = G_b in 2000.

profile.

However, a flat age–earnings profile runs counter to consumption needs over the life cycle as well as productivity variation depending on human capital investment and experience. As illustrated above, the relationship between earnings and age can produce inequality at a given point in time even if everyone is equal in all respects but age. Moreover, inequality in annual earnings may change over time simply because of changes in the age structure: a change in the age structure changes the weights we give to the different parts of the age–earnings profile, and may subsequently change the measured inequality even if inequality in lifetime earnings is unchanged. For this reason, it has long been argued that age adjustments of cross-sectional measures of inequality are necessary (see, e.g., Atkinson, 1971). Such an adjustment allows us to utilize the cross-sectional data at our disposal, while avoiding some of the pitfalls associated with its use.

3.1 Age-adjusted inequality measures

In this paper, we use two different approaches to age adjustment. Both have the same objective: namely, to purge the classical Gini coefficient applied to cross-sectional data of its interage or life cycle component. In particular, perfect equality corresponds to everyone receiving the mean earnings of their age group, thereby relaxing the assumption of a flat age–earnings profile.

Paglin–Gini. The most common age-adjusted inequality measure, the Paglin–Gini (PG), was proposed by Paglin (1975) and can be expressed as:

$$PG = \frac{\sum_j \sum_i (|y_i - y_j| - |\mu_i - \mu_j|)}{2\mu n^2}, \tag{1}$$

where μ_i and μ_j denote the mean earnings of all individuals belonging to the age group of individual i and j, respectively. Applying the standard Gini decomposition, PG can be rewritten as:

$$PG = G - G_b = \sum_i \theta_i G_i + R, \tag{2}$$

where G_b represents the Gini coefficient that would be obtained if the earnings of each individual in every age group were replaced by the relevant age group mean μ_i, G_i is the Gini coefficient of earnings within the age group of individual i, θ_i is the weight given by the product of this group's earnings share $\frac{n_i \mu_i}{\mu n}$ and population share $\frac{n_i}{n}$ (n_i being the number of individuals in the age group of individual i), and R captures the degree of overlap in the earnings distributions across age groups (see, e.g., Lambert and Aronson, 1993).[5]

A main controversy surrounding the PG is whether or not R should be treated as a between-group or a within-group effect. Nelson (1977) and others argue that R is part of between-group inequality and should thus be netted out when constructing age-adjusted inequality measures. Paglin (1977), however, maintains that R is capturing within-group inequality and that PG is accurately defined. Until recently, the issue was unsettled simply because little was known about the overlap term. Shorrocks and Wan (2005), for example, refer to R as a "poorly specified" element of the Gini decomposition. However, Lambert and Decoster (2005) provide a novel characterization of the properties of R, showing first that R unambiguously falls as a result of a within-group progressive transfer, and second, that R increases when the earnings in the poorer group are scaled up, reaching a maximum when means coincide. This makes Lambert and Decoster (2005, p. 378) conclude that "The overlap term in R is at once a between-groups and a within-groups effect: it measures a between-groups phenomenon, overlapping that is generated by inequality within groups." Consequently, it appears that neither Paglin nor Nelson was right, as R in part, but not only, reflects earnings differences attributable to age. Therefore, if PG is to eliminate earnings differences attributable to age in a consistent way, there must be no overlap between the earnings distributions of different age groups.

As pointed out by Jenkins and O'Higgins (1989), the debate over how R should be

[5]Overlap implies that the earnings of the richest person in an age group with a relatively low mean earnings exceeds the earnings of the poorest person in an age group with a higher mean earnings; that is, $y_i < y_j$ and $\mu_i > \mu_j$ for at least one pair of individuals i and j.

interpreted conflates the separate issues of inequality measurement per se and the subsequent decomposition of measured inequality. The fact that we cannot separate out the interage component of R, does not imply that age-adjusted inequality measures are futile. Instead of trying to adjust for differences attributable to age at the aggregation stage of inequality measurement, Wertz (1979) advocates that one should transform the measure of earnings at the individual level, to purge it of its interage component before aggregating up earnings differences.

Another common criticism of PG is that it assumes that the unconditional distribution of mean earnings by age represents the age effects, and it does, therefore, not only eliminate inequality attributable to age but also inequality because of factors correlated with age, such as education.[6] For example, when using PG to evaluate the influence of age adjustment on the time trend in inequality, we may confuse the effects of changes in the age structure with the impact of more people taking higher education than before.

Age-adjusted Gini. In the spirit of the critiques of PG, Almås and Mogstad (2009) propose an alternative age-adjusted Gini coefficient (AG).[7] Their method for age adjustment of inequality may be described as a three-step procedure. First, a multivariate regression model is employed, allowing us to isolate the net age effects while holding other determinants of earnings constant. Second, the earnings distribution that characterizes perfect equality in age-adjusted inequality is determined by using the principle of proportionality. Third, AG is derived and axiomatically justified as an age-adjusted inequality measure. To see how AG identifies the net age effects on earnings, suppose that the earnings of individual i in a given year depends on the age group a that he belongs to and a vector X of individual characteristics, such that:

$$y_i = f(a_i)h(X_i). \tag{3}$$

For simplicity, f is specified as a function of age alone but could in general also reflect other aspects affecting individuals' life cycle behavior, such as their time preferences, interest rates, and the rules governing retirement. In our empirical analysis, we allow for a flexible functional form of f, yielding the following earnings-generating function:

$$\ln y_i = \ln f(a_i) + \ln h(X_i) = \delta_i + X_i' B, \tag{4}$$

where δ_i gives the percentage earnings difference associated with being in the age group of individual i compared with some reference age group, holding all other variables constant. Equation 4 is estimated by OLS separately for each year. The key assumption underlying

[6]See, for example, Danziger et al. (1977) who in an early comment about Paglin's approach pointed out that "no single indicator is sufficient to capture trends in normatively relevant inequality without a well specified multivariate model."

[7]Almås and Mogstad (2009) use this measure to study wealth inequality across countries.

this estimation is that there are no omitted factors correlated with age that determine individual earnings. In that case, we obtain consistent estimates of the net age effects on earnings. It is important to emphasize that the objective of the estimation of equation 4 is not to explain as much variation as possible in earnings, but simply to obtain an empirically sound estimate of the effects of age on earnings.

Identifying the net age effect is only part of the job; it is also necessary to find a consistent way of adjusting for age effects when there are other earnings-generating factors. There is a considerable literature concerning the problem of how to adjust for some, but not all, earnings-generating factors when the earnings function is not additively separable (see, for example, Bossert and Fleurbaey (1996) and Kolm (1996)). To eliminate differences in earnings attributable to age but preserve inequality arising from all other factors, AG employs the so-called general proportionality principle proposed by Bossert (1995) and Konow (1996), and further studied in Cappelen and Tungodden (2007). In particular, the earnings level of individual i that makes him not contribute to age-adjusted inequality depends on his own age as well as every other earnings-generating factor of all individuals in the society, and is defined as:

$$\widetilde{y}_i = \frac{\mu n \sum_j f(a_i)h(X_j)}{\sum_k \sum_j f(a_k)h(X_j)} = \frac{\mu n e^{\delta_i}}{\sum_k e^{\delta_k}}, \tag{5}$$

where e^{δ_k} gives the net age effect of belonging to the age group of individual k after integrating out the effects of other earnings-generating factors correlated with age. No age-adjusted inequality corresponds to every individual i receiving $\widetilde{y}_i$, which is the share of total earnings equal to the proportion of earnings an individual from his age group would earn if all earnings-generating factors except age were the same for everyone in the population. If age is uncorrelated with all other earnings-generating factors, $\widetilde{y}_i = \mu_i$, such as in PG. Furthermore, if there is no age effect on earnings, $\widetilde{y}_i = \mu$, such as in G.

In order to generalize the standard approach in a way that is useful for age adjustment, it is helpful to see the correspondence between the standard Lorenz curve and Lorenz curves expressed in differences. Figure 3 displays standard and difference-based Lorenz curves, in the special case where the equalizing earnings are equal to the overall mean in society. Both summarize the area between the line of perfect equality and the Lorenz curve, and hence, the two areas marked "A" are of equal size. As is well known, G is equal to $2A$.

When drawing age-adjusted Lorenz curves, however, the population is ordered not by their earnings per se, but according to the difference between their individual earnings and their equalizing earnings. Perfect equality is then characterized as a situation where each individual receives that share of total earnings that he would hold if all earnings-generating factors, except age, were equal. Because of the analogy to the classical case, we can draw Lorenz curves based on the distribution of differences in actual and equalizing

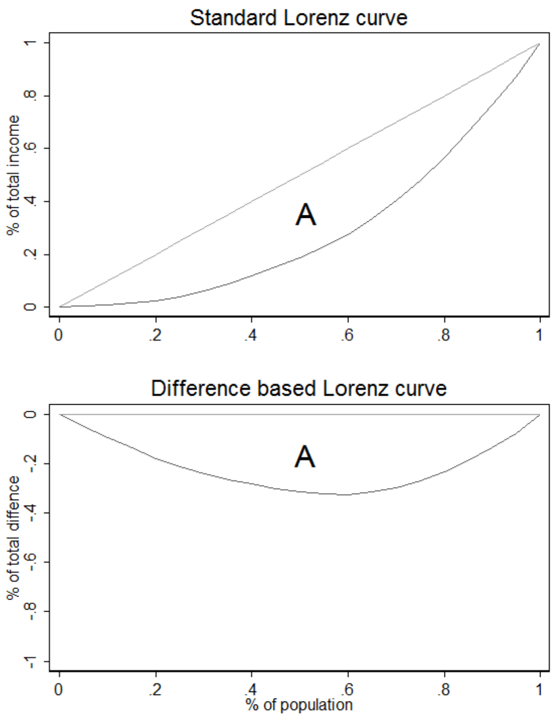

Figure 3: The standard and difference-based representation of the classical Lorenz curve.

earnings, $(y_i - \widetilde{y}_i)$. AG equals twice the area between this curve and the line of perfect equality (the x axis in Figure 3).

If there are no age effects on earnings, the age–earnings profile is flat and AG coincides with G as $\widetilde{y}_i = \mu$ for all individuals in every age group. If there is a relationship between age and earnings, AG will in general differ from G. Both G and AG reach their minimum value of 0 when everyone receives their equalizing earnings. Moreover, both measures take their maximum when the difference between actual and equalizing earnings is at its highest possible level. Specifically, G reaches its maximum value of 1 if one individual receives all the earnings. In comparison, AG takes its maximum value of 2 in the hypothetical situation where the equalizing earnings of the individual who has all the earnings is zero, and the equalizing earnings of one of the individuals with no earnings is equal to the aggregate earnings in the economy. The fact that AG and G range over different intervals is therefore a direct result of their different views of perfect equality: age-adjusted inequality is not only a result of differences in individual earnings, but also a result of differences in equalizing earnings between individuals of different ages.

Formally, AG is based on a comparison of the absolute values of the differences in actual and equalizing earnings between all pairs of individuals, defined by:

$$AG = \frac{\sum_j \sum_i \left|(y_i - \widetilde{y}_i) - (y_j - \widetilde{y}_j)\right|}{2\mu n^2}. \tag{6}$$

AG also stands out from PG in the way it aggregates earnings differences: in general, $|(y_i - \widetilde{y}_i) - (y_j - \widetilde{y}_j)|$ will differ from $|(y_i - y_j)| - |(\mu_i - \mu_j)|$, even when $\widetilde{y}_i = \mu_i$.

To justify their proposed adjustment, Almås and Mogstad (2009) offer a set of conditions that are similar to those underlying the classical Gini coefficient in all respects but one: equalizing earnings are given not by the mean earnings in society as a whole, but rather depend on age. First, they impose scale invariance, stating that if all actual and equalizing earnings are rescaled by the same factor, then the level of age-adjusted inequality remains the same. Next, they impose anonymity, implying that the ranking of alternatives should be unaffected by a permutation of the identity of individuals. Third, they put forward a generalized version of the Pigou–Dalton criterion, stating that a transfer of earnings from an individual i to an individual j, where $(y_i - \widetilde{y}_i) > (y_j - \widetilde{y}_j)$ both before and after the transfer, reduces age-adjusted inequality. Finally, they impose a condition called *unequalism*, implying that the measure only focuses on how unequally each individual is treated, defined as the difference between his actual and equalizing earnings: age-adjusted inequality measures should be the same when the distributions of differences in actual and equalizing earnings are the same.[8] As proven by Almås and Mogstad (2009), AG satisfies all four conditions, while PG fails to meet the unequalism condition.

4 Empirical analysis

Data and descriptive statistics. Our data are based on administrative registers from Statistics Norway covering the entire resident population in Norway between 1967 and 2004. The unique individual identifier allows us to merge information about individual characteristics, such as education, with data on annual earnings and earnings taken from tax registers in each year. From the individual identifiers, we are also able to link individuals to their parents, allowing the inclusion of controls for family background. In the analysis, we employ a measure of earnings including all market income from wages and self-employment. In each year from 1967 to 2004, we include the entire population of males aged 25–59 who were alive and resident in that year.

Table 1 reports descriptive statistics for selected years. We see that while real earnings are increasing over the period, the relation between age and earnings is present in all years. The demographic wave induced by the baby boomers was presented in detail in panel (b) of Figure 1. At the same time, there are several changes in the labor force, both in individual characteristics and in family background. In particular, the level of education increases steadily over the whole period, for both workers themselves and their parents, while the immigrant share of the labor force more than doubles. Family size and parental

[8] This condition may be viewed as analogous to the focus axiom in poverty analysis, stating that a poverty index should focus entirely on the incomes of the poor (see, for example, Foster and Shorrocks, 1991).

Table 1: Descriptive statistics

	1970	1980	1990	2000
Earnings				
Mean	289,707	325,998	342,462	373,084
Age 26–29	248,581	266,849	263,080	289,886
Age 30–39	302,713	336,944	342,187	367,602
Age 40–49	311,392	355,340	386,041	409,659
Age 50–59	275,295	319,540	335,314	380,551
Individual characteristics				
Age	42.34	40.73	40.30	41.46
Education (years)	9.79	10.67	11.35	11.88
Immigrant	0.06	0.09	0.13	0.14
No. of children	2.53	2.55	2.47	2.31
Family characteristics				
Mother's education	7.82	8.18	8.59	9.19
Father's education	8.42	8.96	9.41	9.96
Mother's age at birth	29.34	29.55	29.04	28.12
Father's age at birth	32.67	33.02	32.60	31.53
Observations	810,643	892,038	989,901	1,118,735

age decreases somewhat, particularly between 1990 and 2000.

To calculate AG, we first estimate the net age effects running OLS on the regression equivalent of equation 4, controlling for education, birth order, family size, and immigration, as well as parental education and age at birth. We then calculate equalizing earnings as $\widetilde{y}_i$ by applying the transformation in equation 5, and then estimate AG from equation 6. PG is estimated from equation 1.

It should be noted that the age-adjusted inequality measures, like the classical Gini coefficient, are ordinal in nature and any monotonic transformation of such a measure will preserve its ranking of distributions. This means that they are only of interest as a way of comparing and ordering the distributions by inequality. The fact that these inequality measures range over different intervals is, therefore, beside the point.[9]

Empirical results. The graph on the left in Figure 4 reports the development in AG, G, and PG for earnings, estimated separately for each year between 1967 and 2004. Three main findings appear. First, in the period of convergence between 1967 and 1980, AG shows a more moderate decrease compared with G. This conforms well to the prediction from the exercise in Section 2. PG, on the other hand, suggests that the changing age

[9]Specifically, G can range from 0 to 1, PG from 0 to G, and AG and WG from 0 to 2. Normalizing these measures so they range over the same interval is possible, but it will not affect the ranking of the earnings distributions.

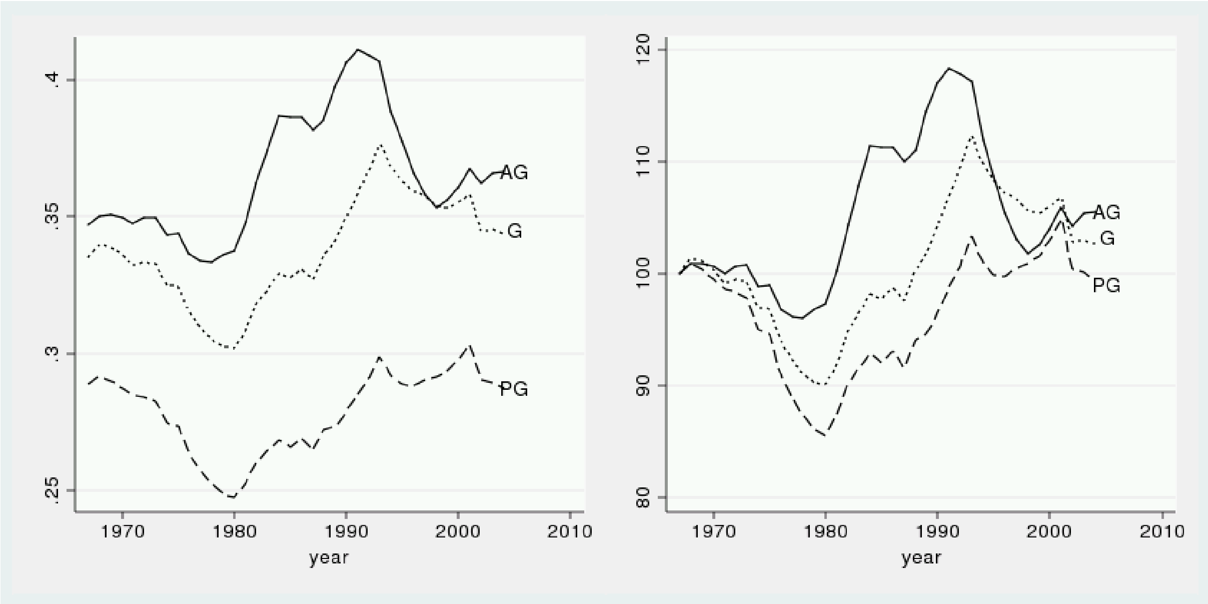

Figure 4: Trends in earnings inequality among males aged 25–59 over the period 1967–2004 (left panel), and relative change in inequality from the base year (right panel).
Notes: G = Gini, AG = Age-adjusted Gini controlling for education, PG = Paglin–Gini, see Section 3.

composition contributed to reduce the earnings inequality over this period, as reflected in a more pronounced drop in inequality in PG than in G.

Second, between 1980 and 1993, when the baby boomers were approaching the peak of their age–earnings profile, PG shows a slightly more modest increase in inequality than G. Relying on this most common approach to age adjustment, we might therefore conclude that the large increase in inequality in the 1980s was driven at least in part by changing demographics. However, the new and improved approach suggested in Almås and Mogstad (2009) concludes the opposite: the measured increase in inequality is even larger than classical inequality measures would suggest.

Finally, Figure 4 shows that the slight drop in inequality in the early 2000s, suggested by both G and PG, is not manifested in the estimated AG. If anything, inequality seems to be rising in this period, illustrating again that properly accounting for changes in age composition can be crucial for interpreting changes in the distribution of earnings.

The measured trend in unadjusted inequality is important in its own right, providing first evidence of the time trend in earnings inequality of the male labor force in Norway over the last few decades. We can see that classic inequality decreased substantially between 1967 and 1980, dropping by about 10 percent. Earnings inequality then rebounded between 1980 and 1993, to surpass previous levels in the late 1980s, before dropping again between 1993 and 2004. Overall, the period under study saw a slight rise in inequality, estimated at about 2.5 percent. The much-studied boom in inequality during the 1980s is quite apparent, however, with an increase of about 25 percent from a bottom of 0.30 in 1980 to a peak of 0.38 in 1993.

Overall, our findings conform well to Paglin's (1975) study of the effects of age adjustment on earnings inequality in the US over the period 1947–1972: they suggest that changes in the age structure had major consequences for the trend in earnings inequality. However, Paglin's suggested adjustment performs quite differently from the adjustment implemented by AG, which addresses two common criticisms of PG.

5 Concluding remarks

This paper demonstrated how age-adjusted inequality measures can be used to evaluate whether changes in inequality over time are because of changes in the age structure. In particular, we investigated to what extent age adjustments affect the trend in in earnings inequality in Norway between 1967 and 2004. We found that it does, and further that the impact of the adjustment depends crucially on the method applied.

Before adjusting for age, the classic Gini shows a substantial decrease in inequality between 1967 and 1980, a sharp increase between 1980 and 1993, and then a drop from 1993 to 2004. Overall, our estimates reveal a slight increase in inequality from 1967 to 2004.

Our preferred measure of age-adjusted inequality revealed, however, a more modest decline in inequality in the first period, a somewhat larger increase in inequality in the intermediate period, and a steeper decline in inequality in the latest period. Overall, this measure suggested an increase in inequality over the whole period that is modestly higher than what is revealed by the classical Gini coefficient.

These findings stand in sharp contrast to the results from the much-used approach proposed by Paglin (1975), which has been questioned from a number of perspectives. Using Paglin's measure, the estimates show a steeper decline of inequality in the first period, a less steep increase in inequality in the intermediate period, and no change in inequality in the latest period.

References

[1] Almås and M. Mogstad (2009): "Older or Wealthier?: The Impact of Age Adjustments on the Wealth Inequality Ranking of Countries", *SOEPpapers* 181, DIW Berlin, The German Socio-Economic Panel (SOEP).

[2] Atkinson, A. (1971): "The distribution of wealth and the individual life cycle", *Oxford Economic Papers*, 23: 239–254.

[3] Bossert, W. (1995): "Redistribution mechanisms based on individual characteristics", *Mathematical Social Sciences*, 29: 1–17.

[4] Bossert, W. and M. Fleurbaey (1996): "Redistribution and compensation", *Social Choice and Welfare*, 13: 343–355.

[5] Cappelen, A. and B. Tungodden (2007): "Fairness and the Proportionality Principle", *NHH, Department of Economics Working Paper*, No. 31.

[6] Danziger, S., R. Haveman and E. Smolensky (1977): "The measurement and trend of inequality: comment", *American Economic Review*, 67, 505–512.

[7] Fjærli, E. and R. Aaberge (2000): "Tax Reforms, Dividend Policy and Trends in Income Inequality Empirical Evidence based on Norwegian Data", *Statistics Norway Discussion Paper*, No. 284.

[8] Gottschalk, P. and T. M. Smeeding (2000): "Empirical evidence on income inequality in industrialized countries", in A. B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*, Amsterdam: North Holland, pp. 261–308.

[9] Jenkins, S.P. and O'Higgins, M. (1989): "Inequality measurement using 'norm incomes': were Garvy and Paglin onto something after all?", *Review of Income and Wealth*, 35(3), 265–82.

[10] Kolm, S. (1996): *Modern Theories of Justice*, Cambridge, MA: MIT Press.

[11] Konow, J. (1996): "A positive theory of economic fairness", *Journal of Economic Behavior and Organization*, 46: 13–35.

[12] Lambert, P. and J. Aronson (1993):"Inequality decomposition analysis and the Gini coefficient revisited", *Economic Journal*, 103: 1221–1227.

[13] Lambert, P. and A. Decoster (2005): "The Gini coefficient reveals more", *Metron-International Journal of Statistics*, 63: 373–400.

[14] Foster, J. and A. Shorrocks, (1991): "Subgroup consistent poverty indices", *Econometrica*, 59(3): 687–709.

[15] Heckman, J., L. Lochner and P. Todd, (1991): "Fifty Years of Mincer Earnings Regressions", *NBER Working Paper*, 9732.

[16] Lemieux, T. (2008): "The changing nature of wage inequality", *Journal of Population Economics*, 21(1): 21–48.

[17] Nelson, E. (1977): "The measurement and trend of inequality: comment", *American Economic Review*, 67: 497–501.

[18] Paglin , M. (1975): "The measurement and trend of inequality: a basic revision", *American Economic Review*, 65: 598–609.

[19] Paglin, M. (1977): "The Measurement of Inequality: Reply", *American Economic Review*, 67:520–531.

[20] Paglin, M. (1979): "The Measurement of Inequality: Reply", *American Economic Review*, 69: 673–677: 673-677.

[21] Paglin, M. (1989): "On the measurement and trend of inequality: reply", *American Economic Review*, 79: 265–266.

[22] Shorrocks, A. and G. Wan (2005): "Spatial decomposition of inequality", *Journal of Economic Geography*, 5: 59–81.

[23] Wertz, K. (1979): "The measurement of inequality: comment", *American Economic Review*, 69: 670–672.

Norges Handelshøyskole

Norwegian School of Economics and Business Administration

NHH
Helleveien 30
NO-5045 Bergen
Norway

Tlf/Tel: +47 55 95 90 00
Faks/Fax: +47 55 95 91 00
nhh.postmottak@nhh.no
www.nhh.no