

Baby booming inequality? Demographic change and earnings inequality in Norway, 1967–2000

Ingvild Almås · Tarjei Havnes · Magne Mogstad

Received: 1 February 2010 / Accepted: 20 January 2011 / Published online: 12 February 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract In this paper, we demonstrate how age-adjusted inequality measures can be used to evaluate whether changes in inequality over time are due to changes in the age-structure. To this end, we use administrative data on earnings for every male Norwegian over the period 1967–2000. We find that the substantial rise in earnings inequality over the 1980s and into the early 1990s, is to some extent driven by the fact that the large baby boom cohorts are approaching the peak of the age–earnings profile. We further demonstrate that the impact of age-adjustments on the trend in inequality during the period 1993–2000 is highly sensitive to the method used: While the most widely used age-adjusted inequality measure indicates little change in inequality over this period, a new and improved age-adjusted measure suggests a decline in inequality.

Keywords Inequality trend · Age structure · Age–earnings profile · Gini coefficient

JEL Classification D31 · D63 · D91 · E21

Thanks to seminar participants at “Inequality: New Directions” at Cornell University and two anonymous referees for useful comments and suggestions. Financial support from the Norwegian Research Council (194339) is gratefully acknowledged. The 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.

Havnes also gratefully acknowledges support from the GRASP Collaborative Project (EU-7FP, no. 244725).

I. Almås (✉)

Norwegian School of Economics and Business Administration, Bergen, Norway
e-mail: ingvild.almas@nhh.no

T. Havnes

Department of Economics, University of Oslo, Oslo, Norway
e-mail: tarjei.havnes@econ.uio.no

M. Mogstad

Statistics Norway, Oslo, Norway
e-mail: magne.mogstad@ssb.no

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, 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 over the 1980s and into the early 1990s driven by the baby boom cohorts approaching the peak of the age–earnings profile?

These questions spur from two stylized facts. First, there is a strong age–earnings relationship. Both theoretical models and empirical results suggest a strong relationship between age and earnings (see e.g., [14]). 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 completely equal in all respects but age. Second, almost all developed countries experienced a large increase in the 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 matured 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 the last decades. Identifying the age effect on inequality and its 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 even out over time and are, therefore, irrelevant for the distribution of lifetime earnings (see e.g., [6]).

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 1967–2000. 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 goes back to Paglin’s [23] 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

¹See e.g. [13] and [20] for extensive reviews of the literature on earnings inequality.

²The 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. Source: <http://www.census.gov/population/www/socdemo/age/general-age.html> Reading date: 2010/09/13.

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

Our 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 Section 5 presents the empirical results. Section 6 concludes.

2 Stylized facts

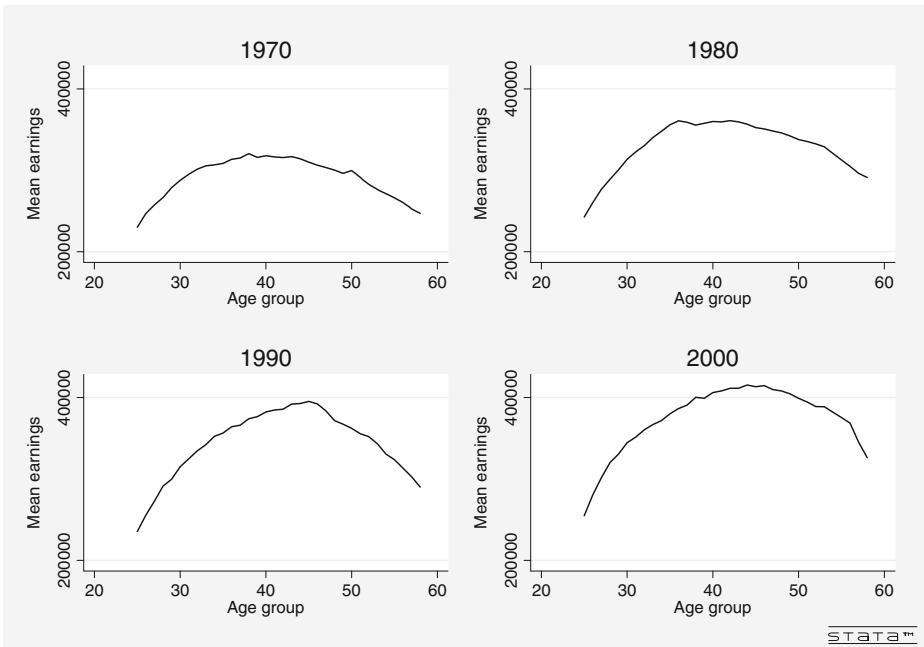
Age–earnings profiles are widely used by economists, both to help forecast the course of future earnings and to depict how earnings typically change over the life cycle. Panel A in Fig. 1 draws the age–earnings profiles in 1970, 1980, 1990, and 2000 for Norwegian males aged 25–59. This figure confirms the picture from other developed countries: Average earnings rise rapidly at younger ages, peak when individuals are in their 40s, and then decline slightly in the latest parts of the working life.⁴ 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 completely equal in all respects other than age, simply because individuals are at different stages in the life-cycle. In 2000, for instance, the average annual earnings of a 50-year-old is 40% higher than that of a 30-year-old, but that does not necessarily imply that the average lifetime earnings of 50-year-olds is any higher than the average lifetime earnings of 30-year-olds.

Panel B in Fig. 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 since the 1970s gradually entered the labor market, and as their careers have matured they have been making 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, weighing different parts of the age–earnings profile differently, even as inequality in lifetime earnings could be unchanged.


Figure 2 illustrates the large amount of ballast that may be embedded in snapshots of earnings inequality and its time trend, as a result of the changing age structure. As a benchmark, we compute the classical Gini-coefficient (G) in annual earnings

³Paglin's age-adjustment of the Gini coefficient was subject to three rounds of comments and replies in the *American Economic Review* [24–26], has numerous citations, and continues to be subject to controversy. For a review of the literature, see Almås and Mogstad [5].

⁴To make nominal figures comparable across different years, earnings throughout this paper 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. Specifically, nominal income in year t , Y_t , is adjusted such that $\tilde{Y}_t = Y_t G_{2006} / G_t$, where G_t is the basic amount threshold in year t and \tilde{Y} is the adjusted income measure.



(a) Panel A: Earnings profile in 1970, 1980, 1990 and 2000



(b) Panel B: Age composition in 1970, 1980, 1990 and 2000

Fig. 1 Earnings profile (*panel A*) and the changing age composition of the labor force (*panel B*) for males aged 25–59




Fig. 2 Trends in overall and between-group earnings inequality among males aged 25–59 over the period 1967–2000

of Norwegian males aged 25–59 over the period 1967–2000. 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. However, after the peak in inequality in 1993, G declines somewhat. The time trend in the classical Gini-coefficient is discussed in detail in Section 5.

Consider instead inequality in a hypothetical situation where everyone in the economy is completely equal in all respects other than age: While earnings vary over the life-cycle, every individual at a given age would have exactly the same earnings as others at that age. To illustrate that there could be substantial earnings inequality at a given point in time in this hypothetical situation, we compute the between-group Gini-coefficient (G_b) for every year in the period 1967–2000. Specifically, G_b replaces the earnings of each individual with his age-group mean, where each cohort is a separate age-group, and can therefore be viewed as a measure of inequality in the age-earnings profiles. We find that a substantial fraction of overall inequality is attributable to inequality between age-groups. In 1993, for example, G_b accounts for more than 20% of overall inequality in G . This illustrates that the age-earnings relationship may make us confuse older with richer, as G incorporates substantial cross-sectional inequality that might even out over time.

To get a sense of how inequality in this hypothetical situation might have evolved over time, Fig. 2 also displays the time trend in G_b . We can see that inequality between age-groups increased over the late 1980s and into the early 1990s, when it had almost doubled since 1967. However, after the peak in inequality between age groups in 1993, G_b declines steadily and in 2000 it reaches the levels observed in the early 1980s. Since the early 1980s, the time pattern in G_b mirrors well the time trend in G . This suggests that inequality between age-groups may have been an important determinant of the observed trend in earnings inequality. However, as

will be apparent in Section 3, this exercise is too stylized to draw inference about the age-effect on earnings inequality and its trend. Yet, it serves as a motivation for taking a closer look at how much the changing age structure matter for the trend in inequality. That is the focus of the rest of the paper.

3 Age-adjustment of the Gini coefficient

Empirical analysis of inequality in income or earnings distributions is 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. This measure is equal to twice the area between the Lorenz curve and its equality reference. In a seminal paper, Paglin [23] argues that G overspecifies the conditions of equality when applied to cross-sectional data: Assuming for the moment no economic growth, perfect equality requires not only equal lifetime earnings, 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 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 completely 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., [6]). Such an adjustment allows us to utilize the cross-sectional data at our disposal, while avoiding some of the pitfalls associated with its use.

In our empirical analysis, we use three different age-adjusted inequality measures. They all have the same objective, namely to purge the classical Gini coefficient applied to cross-sectional data of its inter-age or life-cycle component: In particular, the implicit assumption of a flat age-earnings profile is relaxed. Below, we first describe the new and improved method for age-adjustment of inequality, proposed by Almås and Mogstad [5]. Next, we discuss its relationship with the classical Gini coefficient as well as its relationship with two previous age-adjusted inequality measures.

3.1 The setup

Consider a society consisting of n individuals where every individual i is characterized by the pair (y_i, \tilde{y}_i) , where y_i denotes his actual earnings and \tilde{y}_i is the equalizing earnings in a given year. If actual and equalizing earnings are the same for all individuals and they live equally long, there is perfect equality of lifetime earnings. Roughly speaking, the equalizing earnings is the same for all individuals belonging to the same age group in this society; it is a function of individual i 's age, but not of any other individual characteristics. If no other earnings generating factor is correlated with age, the equalizing earnings is simply the mean earnings of each age group.

Furthermore, if there are no age effects on earnings, the equalizing earnings will be equal to the mean earnings in the society as a whole.

The joint cross-sectional distribution Y of actual and equalizing earnings is given by

$$Y = [(y_1, \tilde{y}_1), (y_2, \tilde{y}_2), \dots, (y_n, \tilde{y}_n)],$$

where n is the population size. Let Ξ denote the set of all possible joint distributions of actual and equalizing earnings, such that the sum of actual earnings equals the sum of equalizing earnings. Following Almås and Mogstad [5], suppose that the social planner imposes the following modified versions of the standard conditions on an inequality partial ordering defined on the alternatives in Ξ , where $A \preceq B$ represents that there is at least as much age-adjusted inequality in B as in A .⁵ Let μ denote the mean earnings of the population as a whole, and Δ_i represent the difference between individual i 's actual earnings y_i and equalizing earnings \tilde{y}_i . Let the distributions of such differences for the two distributions ($\Delta_i(A) = y_i(A) - \tilde{y}_i(A)$) and $\Delta_i(B) = y_i(B) - \tilde{y}_i(B)$) be sorted in ascending order.

Condition 1 (Scale Invariance) For any $a > 0$ and $A, B \in \Xi$, if $A = aB$, then $A \sim B$.

Condition 2 (Anonymity) For any permutation function $\rho : n \rightarrow n$ and for $A, B \in \Xi$, if $(y_i(A), \tilde{y}_i(A)) = (y_{\rho(i)}(B), \tilde{y}_{\rho(i)}(B))$ for all $i \in n$ then $A \sim B$.

Condition 3 (Unequalism) For any $A, B \in \Xi$ such that $\mu(A) = \mu(B)$, if $\Delta_i(A) = \Delta_i(B)$ for every $i \in n$, then $A \sim B$.

Condition 4 (Generalized Pigou–Dalton) For any $A, B \in \Xi$, if there exist two individuals s and k such that $\Delta_s(A) < \Delta_s(B) \leq \Delta_k(B) < \Delta_k(A)$, $\Delta_i(A) = \Delta_i(B)$ for all $i \neq s, k$, and $\Delta_s(B) - \Delta_s(A) = \Delta_k(A) - \Delta_k(B)$, then $A \succ B$.

Scale invariance states that, if all actual and equalizing earnings levels are rescaled by the same factor, then the level of age-adjusted inequality remains the same. Anonymity implies that the ranking of alternatives should be unaffected by a permutation of the identity of individuals. Unequalism entails that the social planner is only concerned with how unequally each individual is treated, defined as the difference between his actual and equalizing earnings.⁶ Finally, the generalized version of the Pigou–Dalton criterion states that any fixed transfer of earnings from an individual i to an individual j , where $\Delta_i > \Delta_j$, reduces age-adjusted inequality.

3.2 A new age-adjusted Gini coefficient

The method proposed by Almås and Mogstad [5] for age-adjustment of inequality may be described as a three-step procedure. First, a new age-adjusted Gini coefficient

⁵See Almås et al. [4] for analogous conditions imposed to study equality of opportunity.

⁶This condition may therefore be viewed as analogous to the Focus axiom in poverty analysis, stating that a poverty index should focus entirely on the earnings of the poor. See e.g. Foster and Shorrocks [12].

(AG) is derived. Second, a multivariate regression model is employed, allowing us to isolate the net age effects on earnings while holding other determinants of earnings constant. Third, the earnings distribution that characterizes perfect equality in age-adjusted earnings is determined.

Definition of inequality measure AG is based on a comparison of the absolute values of the differences in actual and equalizing earnings between all pairs of individuals, and is defined as

$$AG = \frac{\sum_j \sum_i |(y_i - \tilde{y}_i) - (y_j - \tilde{y}_j)|}{2\mu n^2}. \quad (1)$$

It is straightforward to see that AG satisfies Conditions 1–4. Note that these conditions are similar to those underlying G in all respects but one: The equalizing earnings is not given by the mean earnings in the society as a whole, but depends on the age of the individuals.

Because it is straightforward to construct age-adjusted Lorenz curves based on the distribution of differences between actual and equalizing earnings, it is by no means necessary to focus on the Gini coefficient: Other inequality indices that are based on the Lorenz curve can also rely on this method for age adjustments.

Identifying the net age effects Suppose that the earnings of individual i at a given point in time depends on his age a and a vector of individual characteristics X , such that $y_i = g(a_i, X_i)$. The functional form of g depends on the underlying model of earnings. Following standard practice in empirical economics,⁷ we assume that age and the individual characteristics are multiplicatively separable, $y_i = f(a_i)h(X_i)$. However, we will allow for a flexible functional form of f and h , yielding the following log-earnings equation

$$\ln y_i = \ln f(a_i) + \ln h(X_i) = \delta_i + X_i' B, \quad (2)$$

where δ_i gives the percentage earnings difference of being in the age group of individual i relative to some reference age group, holding all other variables constant. Equation 2 is estimated by OLS separately for each year. The key assumption underlying this estimation is that there are no omitted factors correlated with age that determine individual earnings. In this case, we obtain consistent estimates of the net age effects on earnings. It is important to emphasize that the objective of the estimation of Eq. 2 is not to explain as much variation as possible in earnings, but simply to get an empirically sound estimate of the effects of age on earnings.

Defining equalizing 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 concerned with the problem of how to adjust for some, but not all, earnings generating factors when the earnings function is not additively separable (see e.g., [8] and [16]). To eliminate earnings differences attributable to age but preserve inequality arising from all other factors, AG employs the so-called general proportionality principle

⁷See Heckman et al. [15] for a discussion of functional form assumptions in earnings regressions.

proposed by Bossert [7] and Konow [17], and further studied in Cappelen and Tungodden [9].⁸ Then, the absence of age-adjusted inequality requires that any two individuals belonging to a given age group have the same earnings level. Moreover, in any situation where everyone has the same earnings generating factors except age, there should be no lifetime earnings inequality.

Specifically, the equalizing earnings level of individual i depends on his age as well as every other earnings generating factor of all individuals in the society, and is formally defined as

$$\tilde{y}_i = \frac{\mu n \sum_j f(a_i)h(X_j)}{\sum_k \sum_j f(a_k)h(X_j)} = \frac{\mu n e^{\delta_i}}{\sum_k e^{\delta_k}}, \tag{3}$$

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 \tilde{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, \tilde{y}_i is equal to the mean earnings level in his age group, μ_i . And if there is no age effect on earnings, \tilde{y}_i is equal to the mean earnings level in the society, μ .

3.3 Relationship to the classical Gini coefficient

The classical Gini coefficient can be expressed as

$$G(Y) = \frac{\sum_j \sum_i |(y_i - \mu) - (y_j - \mu)|}{2\mu n^2}. \tag{4}$$

By comparing this expression to Eq. 1, we can see there is a very close link between G and AG . Both measures are based on a comparison of the absolute values of the differences in actual and equalizing earnings between all pairs of individuals. The distinguishing feature is how equalizing earnings is defined. For G , the equalizing earnings level is assumed to be μ : perfect equality requires not only equal lifetime earnings, but additionally that individuals of all ages must have the same earnings in any given year, which can be realized only if there is a flat age–earnings profile.

However, a flat age–earnings profile runs counter to both consumption needs over the life cycle and productivity variation depending on human capital investment and experience. Indeed, the relationship between earnings and age can produce earnings inequality at a given point in time even if everyone is completely equal in all respects but age. As transitory earnings differences even out over time, a snapshot of inequality produced by G runs the risk of producing a misleading picture of actual variation in lifetime earnings. In comparison, AG abandons the assumption of a flat age–earnings profile and allows equalizing earnings to depend on the age of the individuals. In doing so, AG purges the cross-sectional measure of inequality of its

⁸Note that the generalized proportionality principle is compatible with any specification of the earnings equation, and not only our multiplicative earnings generating function. However, we would not need to introduce this principle if the earnings generating function was assumed to be additively separable.


inter-age or life cycle component. If $\tilde{y}_i = \mu$ for all individuals in every age group, the age-earnings profile is flat and AG coincides with G .

To get further intuition on the similarities and differences between G and AG , it is helpful to see the correspondence between the standard representation of the Lorenz curve and a Lorenz curve expressed in differences between actual earnings and mean earnings in the society as a whole. Figure 3 displays standard and difference based Lorenz curves for the same earnings distribution. The area between the standard Lorenz curve and the diagonal of the upper diagram (the line of equality) is identical to the area between the difference based Lorenz curve and the horizontal axis (the line of equality) in the lower diagram. In both cases, G is equal to twice the area A , between the Lorenz curve and the line of equality.

Along the same line, we can draw the age-adjusted Lorenz curve underlying AG , expressing the differences between actual earnings and the equalizing earnings in the population. And just as for G , AG is equal to twice the area between this difference based Lorenz curve and the horizontal axis (line of equality). When drawing age-adjusted Lorenz curves, however, individuals are ordered not by their earnings per se, as in Fig. 3, but according to the difference between actual and equalizing earnings.

Both G and AG reach their minimum value of 0, if 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,

Fig. 3 The standard and difference based representation of the Lorenz curve



G reaches its maximum value of 1, if one individual has all earnings in the society. In comparison, AG takes its maximum 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 zero 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 individuals' actual earnings, but also a result of differences in equalizing earnings between individuals at different points in the life cycle.

By the same token, AG will be smaller (greater) than G whenever the differences in individuals' earnings because of age is positively (negatively) correlated with differences in individuals' earnings attributable to other earnings generating factors.⁹ For example, an individual with zero earnings will contribute less to inequality in AG than in G whenever his equalizing earnings level is lower than the mean earnings in the society.

3.4 Relationship to previous age-adjusted inequality measures

There are two distinguishing aspects of age-adjusted inequality measures. First, they differ in the way they aggregate up the differences between actual and equalizing earnings. Second, they hold different views on how equalizing earnings should be measured. In this paper, we consider two alternative age-adjusted inequality measures: Paglin's Gini (PG) and Wertz' Gini (WG). They both have the same objective as AG , namely to purge G applied to snapshots of earnings inequality of its inter-age or life cycle component. In particular, the condition of a flat age-earnings profile is relaxed. Below, we use Conditions 1–4 to assess the properties of PG and WG , and to characterize their relationship to AG .

Because of its close relationship to AG , it is convenient to first consider WG , which was proposed by Wertz [28]. He claims that PG fails to adjust properly for age effects, but his comment has been largely ignored, perhaps because Wertz does not put up conditions that allow a formal assessment of the properties of PG and WG . Let WG be defined by

$$WG(Y) = \frac{\sum_j \sum_i |(y_i - \mu_i) - (y_j - \mu_j)|}{2\mu n^2}, \tag{5}$$

where μ_i and μ_j denote the mean earnings of all individuals belonging to the age group of individual i and j , respectively. Like AG , WG is based on a comparison of the absolute values of the differences in actual and equalizing earnings levels between all pairs of individuals and ranges over the interval $[0, 2]$. It is also straightforward to see that it satisfies Conditions 1–4.

The distinguishing feature between AG and WG is that the latter measure defines the equalizing earnings of an individual i as the unconditional mean earnings in his age group, μ_i , whereas the former measure defines his equalizing earnings as the

⁹To see this, let $\epsilon_i = y_i - \tilde{y}_i$ for any individual i , and note that AG and G have the same denominator. While the numerator of AG aggregates $|\epsilon_i - \epsilon_j|$ over all pairs of individuals, the numerator of G aggregates $|(\tilde{y}_i + \epsilon_i) - (\tilde{y}_j + \epsilon_j)|$ over all pairs of individuals. Hence, $G > AG$ whenever $cov(\tilde{y}, \epsilon) > 0$.

net age effect of belonging to his age group after integrating out the effects of other earnings generating factors correlated with age, \tilde{y}_i . Any differences between AG and WG is therefore a result of omitted variables bias in using μ_i to measure equalizing earnings. As is well known, the omitted variables bias in μ_i depends on the effects of the omitted variables on earnings and the effects of the omitted variables on age. For example, education is correlated with both age and earnings. When using WG to evaluate the influence of age-adjustment on the time trend in inequality, we may therefore confuse the effects of changes in the age-structure with the impact of more people taking higher education than before. The omitted variables bias formula tells us that WG will be equal to AG whenever age is uncorrelated with omitted earnings generating factors. Hence, AG may be viewed as a generalization of WG , and is important in situations where omitted variables bias is a major concern.

Next, consider the much used PG , which can be expressed as

$$PG(Y) = \frac{\sum_j \sum_i (|y_i - y_j| - |\mu_i - \mu_j|)}{2\mu n^2}. \tag{6}$$

Applying the standard Gini decomposition, PG can be rewritten as

$$PG = G - G_b = \sum_i \theta_i G_i + R, \tag{7}$$

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 earnings μ_i , G_i represents 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 is 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., [18]).¹⁰

Both WG and PG define the equalizing earnings of an individual as the mean earnings of the age group he belongs to, disregarding that other earnings generating factors are correlated with age. Unlike AG , they may not only eliminate inequality due to age but also inequality because of these other factors.

In addition, PG stands out in the way it aggregates up the differences in actual and equalizing earnings. Specifically, PG is based on a comparison of differences in the absolute values of actual and equalizing earnings levels between all pairs of individuals, $|(y_i - y_j)| - |(\mu_i - \mu_j)|$. This runs counter to the Unequalism condition, because $|(y_i - y_j)| - |(\mu_i - \mu_j)| = 0$ does not necessarily imply that $|(y_i - \mu_i) - (y_j - \mu_j)| = 0$.

The following numerical example shows that PG violates the Unequalism condition. Consider two distributions A and B with two age groups, each consisting of two individuals. Suppose that A ’s distribution of actual and equalizing earnings, $(y_i(A), \mu_i(A))$, is given by

$$A = [(20, 60), (100, 60), (60, 80), (100, 80)],$$

¹⁰Overlap implies that the earnings of the richest person in an age group with a relatively low mean earnings level 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 .

whereas B 's distribution of $(y_i(B), \mu_i(B))$ is given by

$$B = [(0, 40), (80, 40), (80, 100), (120, 100)].$$

In both cases, the distribution of differences between the actual and equalizing earnings, $y_i - \mu_i$, is given by $\{[-40, 40], [-20, 20]\}$. According to the Unequalism condition, age-adjusted inequality measures should be the same when the distributions of differences between actual and equalizing earnings are the same. While WG satisfies this condition, PG violates it.¹¹

Arguably, the Unequalism condition is an intuitively appealing condition as it ensures that age-adjusted inequality measures follow G in measuring inequality according to the differences in actual and equalizing earnings, between all pairs of individuals, rather than the aggregated differences in actual earnings minus the aggregated differences in equalizing earnings.¹²

As $|(y_i - y_j) - (\mu_i - \mu_j)|$ provides an upper bound for $|(y_i - y_j)| - |(\mu_i - \mu_j)|$, it follows that $WG \geq PG$. This begets the question: Under which conditions will WG be equal to PG , and subsequently, can we be sure that the two measures produce the same inequality ranking? As proved by Almås and Mogstad [5], PG will differ from WG if there is any age effect on earnings, provided that there is some within age group earnings variation. In particular, overlap in the earnings distributions across age-groups, that is, $R > 0$, is a sufficient, but not a necessary, condition for $WG > PG$. This implies that PG is likely to yield a different ranking than WG in situations where earnings distributions differ substantially in the degree of overlap.

This result relates to a major controversy surrounding PG , namely whether or not R should be treated as an inter-age or a within age-groups component.¹³ Until recently, the issue was unsettled simply because little was known about the overlap term; Shorrocks and Wan [27], for example, refer to R as a “poorly specified” element of the Gini decomposition. However, Lambert and Decoster [19] 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 is scaled up, and reaches a maximum when means coincide. This makes Lambert and Decoster [19, 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”. Therefore, $R = 0$ is necessary (although not sufficient) for PG to net out the inter-age component, and nothing but the inter-age component, from cross-sectional inequality measures.

¹¹Specifically, $WG(A) = WG(B) = 0.25$, whereas $PG(A) = 0.179 \neq PG(B) = 0.107$.

¹²Our numerical example illustrates the difference. Consider distribution A and the contribution to age-adjusted inequality from the comparison of the richest individuals in the two age-groups, for which $(y_i(A), \mu_i(A))$ is given by $(100, 60)$ and $(100, 80)$. Paglin advocates that perfect equality corresponds to everyone receiving the mean earnings of their age-group. An earnings comparison of this pair of individuals should thus contribute with 20 to age-adjusted inequality, which is captured by the numerator of WG . By contrast, the numerator of PG records a -20 contribution to age-adjusted inequality—the rationale for which is hard to grasp.

¹³Nelson [22] and others argue that R is part of inter-age inequality and should thus be netted out when constructing age-adjusted inequality measures. Paglin [24], however, maintains that R is capturing within-group inequality and that PG is accurately defined.

4 Data

Our data are based on administrative registers from Statistics Norway covering the entire resident population in Norway between 1967 and 2000. The unique individual identifier allows merging information about individual characteristics, like age and education, with data on annual 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 1967–2000, 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 (reported in 2006-NOK) 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 Fig. 1. At the same time there are several changes in the labor force, both in individual characteristics and in family background. For instance, the level of education increases over the whole period, especially during the 1970s and the early 1980s. In addition, we can see that the immigrant share of the labor force more than doubles. Furthermore, family size and parental age is decreasing somewhat, particularly between 1990 and 2000.

To calculate AG , we first estimate the net age effects running OLS on the regression equivalent of Eq. 2, controlling for education, birth order, family size and immigration, as well as parental education and age at birth. To allow for non-linearities in the covariates, we include six dummies for education groups, as well as four dummies each for number of siblings and birth order. We then calculate equalizing earnings as \tilde{y}_i by applying the transformation in Eq. 3, and then estimate AG from Eq. 1. PG is estimated from Eq. 8, while WG is estimated from Eq. 5.

Table 1 Descriptive statistics

	1970	1980	1990	2000
<i>Earnings, 2006-NOK</i>				
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
<i>Individual characteristics</i>				
Age	42.34	40.73	40.30	41.46
Education (years)	9.79	10.67	11.35	11.88
Immigrant (share)	0.06	0.09	0.13	0.14
No. of children	2.53	2.55	2.47	2.31
<i>Family characteristics</i>				
Mother's education (years)	7.82	8.18	8.59	9.19
Father's education (years)	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

5 Empirical analysis

Main findings Figure 4 displays the evolution over time in age-adjusted and age-unadjusted earnings inequality. The insights may be summarized in two conclusions. First, between 1980 and 1993, when the baby-boomers were approaching the peak of their age-earnings profile, AG shows a more modest increase in inequality than G . This finding suggests that the large increase in earnings inequality over this period was partly driven by changes in the age-structure. Second, the trends in age-adjusted inequality during the periods 1967–1980 and 1993–2000 are quite sensitive to the method used. In the former period, AG shows a more moderate decrease compared to G and PG . And in the latter period, AG follows G in suggesting a decline in inequality, whereas PG indicates little change in inequality.

In sum, our results conform well to Paglin's [23] 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 have significant impact on the trend in earnings inequality. However, Paglin's suggested age-adjustment performs quite differently from the adjustment implemented by AG , both when controlling only for individual background characteristics (AG_0) and when we add controls also for family background characteristics (AG_1). This illustrates that properly accounting for changes in age composition can be crucial for interpreting changes in the distribution of earnings.




Fig. 4 Trends in age-adjusted and unadjusted earnings inequality among males aged 25–59 over the period 1967–2000. Notes: G = Classical Gini coefficient, AG_0 = Age-adjusted Gini controlling for individual background characteristics (education, immigration status, number of siblings, birth order, and number of children), AG_1 = Age-adjusted Gini controlling for individual and family background characteristics (mother's education, mother's age at birth, father's education, and father's age at birth), PG = Paglin-Gini, and WG = Wertz-Gini

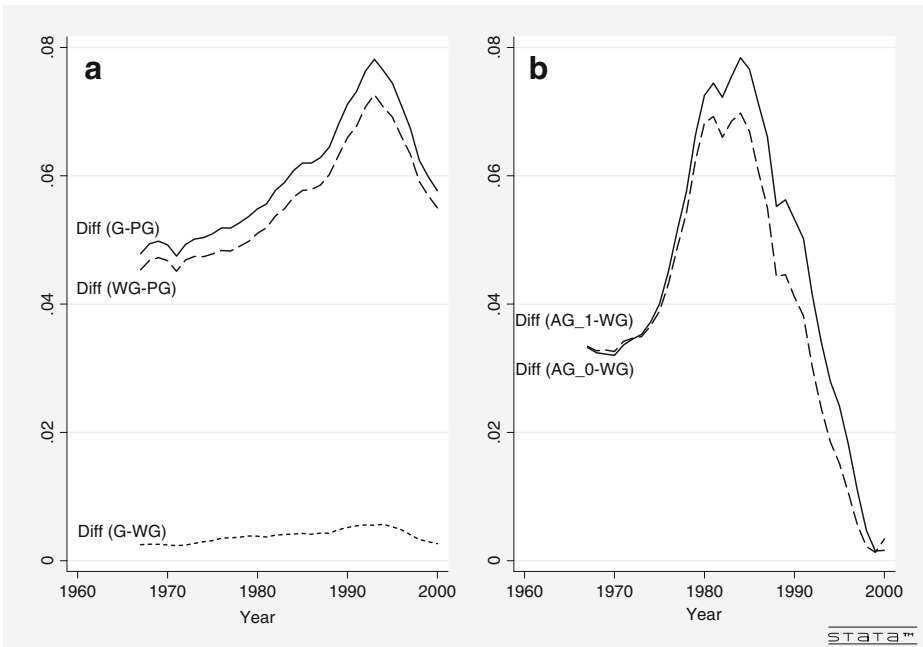


Fig. 5 Differences between inequality measures. Notes: This figure is based on estimates reported in Fig. 4. *Panel A* displays the difference between G and PG , and the difference between G and WG . *Panel B* displays the difference between WG and the two AG measures

Classical Gini coefficient Before turning to a more detailed investigation of the different age-adjusted inequality measures, let us first consider the time trend in G . The measured trend in unadjusted inequality does not only serve as a benchmark, but it is important in its own right, providing first evidence on the time trend in earnings inequality of the male labor force in Norway over the last few decades.¹⁴ We can see that G decreased substantially between 1967 and 1980, dropping by about 10%. Earnings inequality then rebounded between 1980 and 1993, to surpass previous levels already in the late 1980s, before dropping again between 1993 and 2000. Overall, the period under study saw a slight rise in inequality, estimated at about 2.5%. The much studied rise in inequality during the 1980s, is quite apparent, however, with an increase of about 25% from a bottom of 0.30 in 1980 to a peak of 0.38 in 1993.

It should be noted that the increase in inequality in the early 1990s was associated with a tax-reform, and that inequality would be likely to have increased more steadily in the absence of this reform. In particular, the increased inequality can, at least partly, be explained by the high earners' response to large reductions in marginal tax rates (see Aarbu and Thoresen [3]). However, the spike in inequality in 1993 is

¹⁴See Almås et al. (2011) for a discussion of the income inequality and fairness in Norway, Aaberge and Mogstad (2011) for a dominance analysis of income inequality in Norway, and Aaberge et al. (2010) for an analysis of equality of opportunity in Norway.

most likely a result of changes in the income reporting behavior, rather than factual changes in the distribution of income (see [10]).

The Paglin-Gini Panel A of Fig. 5 displays the difference between PG and G over the period 1967–2000. We can see that the two measures of inequality diverge at an increasing rate until 1993, after which the difference declines. As shown in Eq. 7, PG yields a different time trend in inequality insofar as there is significant time-variation in between-group inequality, G_b . Because G_b is a population share weighted average of the different age-group means, it increases with (i) the disparity in mean earnings across age groups, and with (ii) the number of people in the age groups with relatively low or relatively high mean earnings levels.

Figures 6 and 7 look into the first possible explanation for the difference between G and PG . The first figure displays the age-group mean earnings (divided by the mean earnings in the population as a whole) in different years, suggesting that the age-earnings relationship is strongest in the early 1990s. In line with this result, the latter figure shows that the coefficient of variation in mean earnings across age groups is relatively large in the late 1980s and in the early 1990s. These results suggest that increased disparity in mean earnings may help explain why PG differs most from G during this period.

To shed light on the second possible explanation for the difference between G and PG , Fig. 8 graphs the population share in different age groups when the age groups are sorted by their rank in the earnings distribution. We see that more people are belonging to the age-groups with relatively high and low mean earnings in the early 1990s than in the other years. This illustrates that the change in G_b over time is not only a result of changes in the age-earnings profile, but also driven by changes in the demographic composition.




Fig. 6 Relative age-group mean earnings for different years. Notes: This figure displays the age-group mean earnings, divided by the mean earnings in the population as a whole




Fig. 7 Coefficient of variation in mean earnings across age-groups, 1967–2000



Fig. 8 Population share in age-groups sorted by their rank in the earnings distribution. Notes: This figure displays the population share of the age groups, when the age-groups are sorted in an ascending order by their rank in the earnings distribution. The lines have been smoothed using a running line least squares technique with bandwidth equal to 0.8

A new age-adjusted Gini coefficient As discussed above, the age-adjusted Gini coefficient proposed by Almås and Mogstad [5] addresses two common criticisms of *PG*: (i) the way the differences between actual and equalizing earnings are aggregated, and (ii) how equalizing earnings is measured. To examine the first issue, we focus on the comparison between *PG* and *WG*. The reason is that *WG*, on the one hand, aggregates up the differences between actual and equalizing earnings in the same way as *AG*, but on the other hand, conforms with *PG* in measuring equalizing earnings of an individual as the earnings of the age group he belongs to.

Panel A of Fig. 5 displays the difference between *PG* and *WG* over the period 1967–2000. We can see that the time trend in *WG* differs substantially from that produced by *AG*, especially in the period 1980–2000. As discussed in Section 3.2, *WG* will differ from *PG* if there is any age effect on earnings, provided that there is some within-group earnings variation. In particular, overlap in the earnings distributions across age-groups is a sufficient, but not a necessary, condition for *WG* to exceed *AG*. Figure 9 shows that the inequality due to overlap in the earnings distribution makes up a large fraction of overall earnings inequality. It is also clear that the overlap term increases from 1980 until the early 1990s, during which the difference between *WG* and *PG* also increases. This indicates that the changes over time in the degree of overlap in the earnings distributions may help explain the differences in the age-adjusted time trends in inequality according to *WG* and *PG*.

To examine the second issue, we turn the attention to the comparison between *WG* and *AG*. From Panel B of Fig. 5, we can see that the time trend produced by *AG* is quite different from that of *WG*. This holds true for *AG*₀, in which case we have only controlled for individual background characteristics, as well as for *AG*₁, where we also control for family background characteristics. In particular, we can see that *AG* suggests a smaller decline in inequality over the period 1967–1980, and




Fig. 9 Degree of overlap in the earnings distribution across age-groups for different years. Notes: This figure displays the overlap term of the Gini-coefficient over time

a less pronounced increase in inequality during the late 1980s and early 1990s. After the peak in earnings inequality in 1993, AG suggests a stronger decline in inequality.

As discussed in Section 3, any difference between WG and AG is due to omitted variables bias in using μ_i to measure equalizing earnings. In particular, the omitted variables bias in μ_i depends on the effects of the omitted variables on earnings and the effects of the omitted variables on age. The relatively large differences between WG and AG_0 suggest that controlling for individual background characteristics is quite important for the age adjustment of inequality, whereas the small difference between AG_0 and AG_1 implies that the controls for family background characteristics matter little (conditional on individual background characteristics). When regressing the individual background characteristics on earnings as well as on age, we find that education stands out as a key source to omitted variable bias in WG and its trend. First of all, education has a strong positive correlation with both earnings and age, which helps to explain why AG exceeds WG in a given year. In addition, when using WG to evaluate the influence of age-adjustment on the time trend in inequality, we confuse the effects of changes in the age structure with the impact of higher education among the new cohorts that entered the workforce. This is mirrored in an increasing variation across age-groups in years of education over the 1970s and into the early 1980s. After peaking in 1983, this variation decreased throughout the rest of the period, reflecting that the older cohorts with low education were retiring. As a consequence, the contribution from education to omitted variables bias in WG is expected to increase until the early 1980s and then decrease, which is consistent with Panel B of Fig. 5. Along the same lines, other individual background characteristics contribute to explaining the omitted variable bias in WG and its trend, such as immigrant status and family size.

Additively decomposable inequality measures Finally, we follow Mookherjee and Shorrocks [21] in using members of the family of generalized entropy inequality measures to compute age-adjusted and unadjusted trends in earnings inequality. In the spirit of Paglin, they make age adjustment by subtracting between-group inequality from overall inequality. The Paglin–Entropy (PE) can be expressed as:

$$PE^\alpha = E^\alpha - E_b^\alpha = \sum_i \zeta_i E_i^\alpha, \quad \alpha > 0, \quad (8)$$

where E^α denotes the entropy measure of overall inequality, E_b^α is the entropy index that would be obtained if the earnings of each individual in every subgroup were replaced by the relevant age group means μ_i , E_i represents the entropy index within the age group of individual i , and ζ_i denotes the associated weight which is a function of this group's relative mean earnings $\frac{\mu_i}{\mu}$ and population share $\frac{n_i}{n}$. The α parameter reflects the sensitivity to changes in different parts of the distribution, with $\alpha = 0$ giving the most weight to the lower tail. Following common practice in empirical analysis using the entropy measures, we compute PE for α equal to zero, one, and two.

Figure 10 displays the trends in overall inequality and within-group inequality according to the entropy measures. We can see that the entropy measures produce quite similar trends in overall inequality compared to the classical Gini coefficient. But more importantly, it is evident that between-group inequality makes up little of overall inequality. Hence, PE suggests that age adjustment is of minor importance. However, we need to be cautious in drawing this conclusion. While it is often argued




Fig. 10 Trend in overall inequality and within-group inequality according to the entropy measures. Notes: This figure displays overall inequality and within-group inequality according to the entropy measures, with α equal to zero, one, and two. When α is equal to zero or one, we have to exclude the observations with zero earnings, as these entropy measures are not defined for zero earnings

that the absence of an overlap term makes the entropy measures more suitable for decomposition analysis, the overlap term does have value from another perspective in giving useful information that additively decomposable indices must, by definition, ignore.¹⁵ In addition, *PE* suffers from the problem of omitted variables bias in the definition of equalizing earnings, just as *PG* and *WG*.

6 Concluding remarks

This paper demonstrates how age-adjusted inequality measures can be used to evaluate whether changes in inequality over time are due to changes in the age-structure. In particular, we have investigated to what extent age-adjustments affect the trend in earnings inequality in Norway between 1967 and 2000. We find that it does, and further that the impact of the adjustment depends crucially on the method applied.

Before adjusting for age, the classical 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 2000. Overall, our estimates reveal a slight increase in inequality from 1967 to 2000. Our preferred measure of age-adjusted inequality reveals, however, a more modest decline of inequality in the first period, a smaller increase in the intermediate period, and a steeper decrease in inequality in the latest period. These findings stand in contrast to the results from the much used approach proposed by Paglin [23], which has been questioned from a number of perspectives.

¹⁵See Foster and Sen [11] for a discussion.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Aaberge, R., Mogstad, M.: Robust inequality comparisons. *J. Econ. Inequal.* (2011, forthcoming)
2. Aaberge, R., Mogstad, M., Peragine, V.: Measuring long-term inequality of opportunity. *J. Public Econ.* doi:[10.1016/j.jpubeco.2010.11.023](https://doi.org/10.1016/j.jpubeco.2010.11.023) (2010)
3. Aarbu, K.O., Thoresen, T.O.: The Norwegian tax reform; distributional effects and the high-income response. Discussion Papers no. 207, Statistics Norway, Research Department (1997)
4. Almås, I., Cappelen, A., Thori Lind, J., Sørensen, E., Tungodden, B.: Measuring unfair (in)equality. *J. Public Econ.* (2011)
5. Almås, I., Mogstad, M.: Older or wealthier? The impact of age-adjustments on the wealth inequality ranking of countries. *Scand. J. Econ.* (2011, forthcoming)
6. Atkinson, A.: The distribution of wealth and the individual life cycle. *Oxf. Econ. Pap.* **23**, 239–254 (1971)
7. Bossert, W.: Redistribution mechanisms based on individual characteristics. *Math. Soc. Sci.* **29**, 1–17 (1995)
8. Bossert, W., Fleurbaey, M.: Redistribution and compensation. *Soc. Choice Welf.* **13**, 343–355 (1996)
9. Cappelen, A., Tungodden, B.: Fairness and the proportionality principle. NHH, Department of Economics Working Paper no. 31 (2007)
10. Fjørli, E., Aaberge, R.: Tax reforms, dividend policy and trends in income inequality empirical evidence based on norwegian data. Statistics Norway Discussion Paper no. 284 (2000)
11. Foster, J., Sen, A.: On economic inequality after a quarter century. Annex in *On Economic Inequality*. Sen (1997)
12. Foster, J., Shorrocks, A.: Subgroup consistent poverty indices. *Econometrica* **59**(3), 687–709 (1991)
13. Gottschalk, P., Smeeding, T.M.: Empirical evidence on income inequality in industrialized countries. In: Atkinson, A.B., Bourguignon, F. (eds.) *Handbook of Income Distribution*, pp. 261–308. North Holland, Amsterdam (2000)
14. Heckman, J., Lochner, L., Todd, P.: Fifty years of mincer earnings regressions. NBER Working Paper no. 9732 (2003)
15. Heckman, J., Lochner, L., Todd, P.: Earnings functions and rates of return. *Journal of Human Capital* **2**(1), 1–31 (2008)
16. Kolm, S.: *Modern theories of justice*. MIT Press, Cambridge (1996)
17. Konow, J.: A positive theory of economic fairness. *J. Econ. Behav. Organ.* **46**, 13–35 (1996)
18. Lambert, P., Aronson, J.: Inequality decomposition analysis and the Gini coefficient revisited. *Econ. J.* **103**, 1221–1227 (1993)
19. Lambert, P., Decoster, A.: The Gini coefficient reveals more. *Metron—International Journal of Statistics* **63**, 373–400 (2005)
20. Lemieux, T.: The changing nature of wage inequality. *J. Popul. Econ.* **21**(1), 21–48 (2008)
21. Mookherjee, D., Shorrocks, A.: A decomposition analysis of the trend in UK income inequality. *Econ. J.* **92**, 886–902 (1982)
22. Nelson, E.: The measurement and trend of inequality: comment. *Am. Econ. Rev.* **67**, 497–501 (1977)
23. Paglin, M.: The measurement and trend of inequality: a basic revision. *Am. Econ. Rev.* **65**, 598–609 (1975)
24. Paglin, M.: The measurement of inequality: reply. *Am. Econ. Rev.* **69**, 673–677 (1977)
25. Paglin, M.: The measurement of inequality: reply. *Am. Econ. Rev.* **69**, 673–677 (1979)
26. Paglin, M.: On the measurement and trend of inequality: reply. *Am. Econ. Rev.* **79**, 265–266 (1989)
27. Shorrocks, A., Wan, G.: Spatial decomposition of inequality. *Journal of Economic Geography* **5**, 59–81 (2005)
28. Wertz, K.: The measurement of inequality: comment. *Am. Econ. Rev.* **69**, 670–672 (1979)