



United Nations
University

WIDER

World Institute for Development Economics Research

Discussion Paper No. 2004/02

The Rise or Fall of World Inequality

A Spurious Controversy?

Bart Capéau¹ and André Decoster²

January 2004

Abstract

In the age of globalization, the question whether inequality in the world rose or fell down, is a hot topic. Leading scholars in the field of economic inequality measurement developed methods to estimate empirically the distribution of welfare (income) amongst world citizens. Despite their similar methodologies, they do not seem to agree about the conclusion. In the present paper we pinpoint what drives the two extreme positions apart. Sala-i-Martin (2002a, b), who claims that there can be no doubt that world inequality went down between the late 1970s and the late 1990s, has in fact calculated population weighted inequality between countries. Milanovic (2002a, b, c) does not deny this, but illustrates the empirical importance of divergent tendencies at the sub-national level (especially urban versus rural regions) for assessing true world inequality and comes to the reverse conclusion. Nevertheless, there seems to be unanimity, especially amongst the contributions quoted here, about the inequality measure(s) to be used for assessing world income distributions. We show that at least for international inequality, there is empirical evidence for rank reversals among the class of generalized entropy measures and expect the same to be true of world inequality. However, the normative debate about which inequality measure to use for assessing true world inequality has not yet begun.

Keywords: inequality, world income distribution

JEL classification: D31, D63, O15

Copyright © UNU-WIDER 2004

^{1, 2} Center for Economic Studies, KULeuven.

This study has been prepared within the UNU-WIDER project on Global Trends in Inequality and Poverty.

UNU-WIDER acknowledges the financial contributions to the research programme by the governments of Denmark (Royal Ministry of Foreign Affairs), Finland (Ministry for Foreign Affairs), Norway (Royal Ministry of Foreign Affairs), Sweden (Swedish International Development Cooperation Agency – Sida) and the United Kingdom (Department for International Development).

Acknowledgements

Part of the work was done while André Decoster was research fellow at UNU-WIDER. Bart Capéau gratefully acknowledges financial support from the KULeuven research fund, grant 02/07. We thank Tony Shorrocks, Philip Verwimp and an anonymous referee for comments on an earlier draft of this paper and André Watteyne for providing us with some excerpts of the World Development Indicators (World Bank). This paper is a slightly modified and revised version of an article which appeared in *Tijdschrift voor Economie en Management* 48(4): 547-572 under the title 'The Rise or Fall of World Inequality. Big issue or apparent controversy?'

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

www.wider.unu.edu

publications@wider.unu.edu

UNU World Institute for Development Economics Research (UNU-WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by Janis Vehmaan-Kreula at UNU-WIDER
Printed at UNU-WIDER, Helsinki

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

ISSN 1609-5774

ISBN 92-9190-570-4 (printed publication)

ISBN 92-9190-571-2 (internet publication)

1 Introduction

Assessing the welfare or income inequality amongst world citizens has been since at least as long as Theil's (1979) pioneering contribution, a major concern of welfare economists (see a.o. Whalley 1979; Berry *et al.* 1983; Chotikapanich *et al.* 1997; Schultz 1998; Bourguignon and Morrisson 2002; Milanovic 2002a). At least some of these contributions have fed the (political) world with alarming reports stating that:

'The richest 1% of people in the world receive as much as the bottom 57%, or in other words, less than 50 million income-richest people receive as much as 2.7 billion poor.'

'An American having the average income of the bottom US decile is better-off than 2/3 of world population.'

'The ratio between average income of the world top 5% and the world bottom 5% increased from 78 to 1 in 1988, to 114 to 1 in 1993.'

'75% of world population receive 25% of world \$PPP income; and the reverse.'

(Milanovic 2002a: 88-89).

More recently, the subject received attention from (endogenous) growth theorists (Quah 1996, 1997). Among them, a most provocative contribution came from Sala-i-Martin (2002a, b), challenging views as the ones quoted above as not being warranted by the empirical facts:

'The result is striking: rather than a steady increase in inequality, the population weighted variance has fallen during the last two decades. When considering individuals rather than countries, the evidence does not show "divergence, big time", but "convergence, period".'

(Sala-i-Martin 2002a: 7-8).

Despite a vigour rarely seen amongst academics,¹ there seems to be little or no divergence of opinion about the bare facts which gave rise to it, as we will argue below (section 5). Both report for example that the recent prolonged period of growth of China's mean income, had a major inequality reducing effect. On the other hand, the battery of inequality measures used by these authors do not rank always unanimously the same facts. Without claiming any original contribution in this sense, we simply establish that astonishingly little has been said about those normative issues in this application of inequality measurement, probably because of the hidden unanimity about the evidence. Rather than using similar value judgements to assess different data-sets, as is done now, conceptually it seems more fruitful to us, to investigate whether different measures of inequality rank the same couple of income vectors to be compared, differently and why this would be so.

In the present contribution we want to highlight some aspects of the debate between the 'convergence, period' protagonists and those who believe that the world income distribution is characterized more and more by 'twin peaks' (poor-rich) and a disappearing middle class. More specifically, we want to explain how two major recent

¹ See especially Milanovic (2002c).

In the next section we discuss three concepts of inequality central to the subsequent analysis:

- These concepts were thus baptized by Milanovic (2002b).

2 Which inequality?

From a welfare theoretic point of view however, this ‘Concept I inequality’ has the drawback of considering every country on an equal footing, irrespective of its size in terms of population. Consider the following two vectors of incomes, each consisting of only two different income values, €100 and €1000:

and assume the €100-earners in both cases belong to one and the same country X, while €1000 earners all live in another country Y.

Many people would find income distribution A, in which six out of ten persons earn an income of €100, while the other four earn €1000, as more unequal than vector B, irrespective of the information on the country to which the income earners belong. Indeed, in the latter distribution the bulk of people, say 99 out of 100, earn the same €100, while the remaining, now almost negligible (in population size but not in terms of inequality) part earns €1000. Yet, in ‘Concept I inequality’ both income distributions are described simply by means of the vector of two elements: (100, 1000). Something similar happens if one counts populated poor countries such as Bangladesh, India and China (who constitute respectively 2.2 per cent, 16.8 per cent and 21.1 per cent of total world population in 1999) on an equal footing, in terms of population, with, say for example, Belgium (0.17 per cent of world population in 1999).²

Of course, the example suggests that the deficiency of ‘Concept I inequality’ is easily cured by weighting the vector (100, 1000) by the respective populations receiving one of both incomes. This leads us to ‘Concept II inequality’, which is particularly well suited to measure inequality between groups of people. This is something different than inequality among nations or countries, since we take into account how densely populated those different nations are. In fact, contrary to what happens in ‘Concept I inequality’, every world citizen now counts for one. If we are interested, for example, in the inequality between Africa and Western Europe, this second concept is perfectly well suited, since it takes into account population size, which plays a role, as we saw by the example above.

But the issue of the present paper, and the debate it refers to, concern *the measurement of income inequality among world citizens*. This ‘Concept III inequality’ considers the distribution of income across individuals, irrespective of the country where people live. In this concept, every world citizen counts for one, and the income distribution is one long vector of (in 1999) about 6 billion entries.

2.1 Linking ‘Concept II’ and ‘Concept III’

In some views on economic inequality, ‘Concept II’- and ‘Concept III’-inequality are related. We already noticed that population-weighted international inequality counts every world citizen as one. But it does not (and does not intend to) take into account *within* country inequality, being the inequality within a country, *irrespective of the income that the rest of world’s citizens obtain*. Neglecting within country inequality means that every individual is treated as if they had the mean income of the country or region they belong to. Only differences in mean incomes *between* countries or regions and their respective population sizes play a role. If one accepts that total inequality can exactly be broken up in a *between group* (‘Concept II’-type) and a *within group* component, in the way just defined, one should add to ‘Concept II inequality’ (the between component) an aggregate of the inequality within countries or regions to arrive at ‘Concept III inequality’ (total inequality between all world citizens). Exactly this property was in a more general fashion developed in the literature on *decomposability* of inequality measures (Bourguignon 1979; Shorrocks 1980).

² All figures which are not stemming from quoted sources are based on the World Development Indicators (WDI), a database compiled by the World Bank. For more information, see <http://www.worldbank.org/data/wdi2003/index.htm>.

However, not all inequality concepts agree that global inequality can exactly be broken up in a between and a within part. The most famous example is the celebrated Gini-inequality ranking. Indeed, if one tries, as one does under decomposability, to build up total inequality between individuals from information on inequality within certain groups and inequality between those groups, one deliberately gives up the information about the place of an individual income in the global income distribution. For example, if one really wants to take into account where the top decile incomes of China are situated as compared to the bottom decile incomes in Western Europe, then a decomposable measure will not do the job. A decomposable measure only relates the regions through their mean incomes, and adds to this the relative position of the top decile incomes in China with respect to the other deciles in China. The lack of decomposability of the Gini exactly follows from the fact that each individual income in China is compared with all other individual incomes in the world.

We illustrate the three concepts central in the empirical literature on measuring world income inequality and their (possible) interrelationships for two frequently used inequality measures: the Gini coefficient and the Theil measure.³

Let us introduce some notation: y_{ih} denotes the income of an individual i belonging to a group h , containing n_h members; the income distribution in group h then is denoted by $\mathbf{y}_h = (y_{1h}, y_{2h}, \dots, y_{ih}, \dots, y_{n_h h})$. Let there be H groups. The mean income in group h is then $\mu_h \equiv \frac{\sum_{i=1}^{n_h} y_{ih}}{n_h}$. Let the global population size be equal to $n \equiv \sum_{h=1}^H n_h$. Mean

income in the global society then is $\mu \equiv \frac{\sum_{h=1}^H \sum_{i=1}^{n_h} y_{ih}}{n} = \sum_{h=1}^H \frac{n_h}{n} \mu_h$. The population share of group h in the total population will be denoted by $p_h \equiv \frac{n_h}{n}$, while the income

share of group h in the global income equals $s_h \equiv \frac{\sum_{i=1}^{n_h} y_{ih}}{\sum_{h=1}^H \sum_{i=1}^{n_h} y_{ih}} = \frac{n_h \mu_h}{n \mu}$. Finally, let $\bar{\mu} \equiv \frac{\sum_{h=1}^H \mu_h}{H}$ be the mean of group mean incomes.

‘Concept I inequality’ for the specific case of the Theil measure is then given by:

$$T_1 = \frac{1}{H} \sum_{h=1}^H \left(\frac{\mu_h}{\bar{\mu}} \right) \ln \left(\frac{\mu_h}{\bar{\mu}} \right), \quad (1)$$

where we use the subscript 1 to refer to ‘Concept I inequality’. The population sizes do not enter the picture. Each country is weighted equally, irrespective of its population size, and inequality is only built upon differences in average incomes. For ‘Concept II

³ Only for these two measures we will, in section 4, give numerical calculations from the literature.

inequality' we introduce the population shares, and replace the average of group mean incomes, $\bar{\mu}$, by the global average, μ :

$$T_2 = \sum_{h=1}^H p_h \left(\frac{\mu_h}{\mu} \right) \ln \left(\frac{\mu_h}{\mu} \right). \quad (2)$$

Finally, 'Concept III inequality' for the Theil measure equals:

$$T_3 = \frac{1}{n} \sum_{h=1}^H \sum_{i=1}^{n_h} \left(\frac{y_{ih}}{\mu} \right) \ln \left(\frac{y_{ih}}{\mu} \right). \quad (3)$$

If we now write within group inequality for group h as:

$$T_h = \frac{1}{n_h} \sum_{i=1}^{n_h} \left(\frac{y_{ih}}{\mu_h} \right) \ln \left(\frac{y_{ih}}{\mu_h} \right), \quad (4)$$

the reader can, after some tedious manipulations, verify the following crucial decomposability relationship:

$$\begin{aligned} T_3 &= \sum_{h=1}^H s_h T_h + T_2 \\ &= W + B. \end{aligned} \quad (5)$$

Total inequality between world citizens is decomposable in a *between* term (B), which is 'Concept II inequality' or population weighted international inequality, and a weighted sum of within country inequalities (W). In case of the Theil coefficient the weights are the income shares.

Let us now look at the case of the Gini coefficient, which is not decomposable. For the three concepts, we have respectively:

$$G_1 = \frac{1}{2\bar{\mu}H^2} \sum_{h=1}^H \sum_{g=1}^H |\mu_h - \mu_g|, \quad (6)$$

$$G_2 = \frac{1}{2\mu} \sum_{h=1}^H \sum_{g=1}^H p_h p_g |\mu_h - \mu_g|, \quad (7)$$

$$G_3 = \frac{1}{2\mu n^2} \sum_{h=1}^H \sum_{g=1}^H \sum_{i=1}^{n_h} \sum_{j=1}^{n_g} |y_{ih} - y_{jg}|. \quad (8)$$

Let us now define Gini-inequality within group h as:

$$G_h = \frac{1}{2\mu_h n_h^2} \sum_{i=1}^{n_h} \sum_{j=1}^{n_h} |y_{ih} - y_{jh}|. \quad (9)$$

Then the reader may verify that a decomposition, analogous as the one in equation (5), that is:

$$\begin{aligned} G_3 &= \sum_{h=1}^H p_h s_h G_h + G_2, \\ &= W + B \end{aligned} \quad (10)$$

can only be written if the income distribution and the group definition is such that there are no overlaps between incomes of the different groups in society.

If this does not hold, that means, if, for example, the richest of the poorest group is richer than the poorest of the next to poorest group, then (10), which writes the ‘global Gini’ as the sum of ‘within Gini’s’ and a ‘between Gini’, does not hold. Of course, one can always write:

$$\begin{aligned} G_3 &= \sum_{h=1}^H p_h s_h G_h + G_2 + R \\ &= W + B + R, \end{aligned} \quad (11)$$

where R is the residual.

Lambert and Aronson (1993) interpreted this residual term R as the contribution of the degree of overlap between the income distributions of different groups to overall inequality because it vanishes in case of absence of overlaps. Consequently, contrary to the Theil measure, the Gini inequality index does not allow to reconstruct ‘concept III inequality’ on the base of ‘concept II inequality’ and an aggregate of the within country inequalities.

On the other hand, *both* inequality rankings discussed so far (Gini and Theil) satisfy the *anonymity* principle (if any two people in the world exchange their income, global inequality remains unchanged), *scale invariance* (multiplying all incomes with a certain (positive) factor leaves inequality unchanged), *population invariance* (concatenating the current world income distribution with that of another world, of equal size and with the same income distribution, leaves inequality unchanged) and the *Pigou-Dalton principle* (a rich to poor transfer that does not lead to a rank reversal decreases inequality). Satisfying these four criteria is equivalent with what is known in the literature as Lorenz dominance: an income distribution is less unequal if the cumulative income share of the poorest $p\%$ is higher for all $p \in [0,100]$. Whenever the Gini and the Theil measure would give different results in ranking the same couple of income vectors, it could not be ascribed to any of these properties.

In case one wants to add on top of this the just discussed criterion that we must be able to write ‘Concept III inequality’ as a function of ‘Concept II inequality’ and a within component alone, we are restricted to the generalized entropy class, of which the Theil

measure is a special case.⁴ So, differences in opinion on the evolution of inequality could follow from the acceptance or rejection of this decomposability property.

2.2 Budget survey versus GDP per capita

Until so far, we concentrated on the issue of which inequality concepts could shed light on the question of global world inequality, but didn't specify – except for the general terms 'income' or 'welfare' – in what dimension one should measure inequality. Two conflicting views can be discerned in the literature, of which Sala-i-Martin (2002a) and Milanovic (2002a) can be considered as the respective proponents: either one uses income or expenditures from household budget surveys (which is mostly some kind of 'disposable income'), or one uses GDP as the appropriate income concept, the difference lying in the in- or exclusion of public goods.

Sala-i-Martin (2002a) defends that an individual welfare measure should include the contributions of *public goods* (especially health and education might be important). Therefore, measures of welfare based on household income or expenditures should be corrected for the level of public spending. This is usually done by applying information on the income distribution – which we discuss in the next section – on GDP *per capita* figures. There is a second reason to sustain this approach: as it turns out, GDP-based measures of income tend to be higher than what one obtains from income surveys, the reason not only being the public goods, but that people tend to underestimate their income.

For reasons of internal consistency, Milanovic (2002c: 14) uses the same source to determine overall mean income, and the distribution of income. Since GDP is by nature an aggregate measure, no distributional information can be retrieved from it, and household budget surveys with detailed micro-data on a representative sample of the population enter the scene. Moreover, he points out that the underestimation of income, referred to by Sala-i-Martin, comes in the first place from misreporting in the tails of the income distribution. Therefore using a GDP based measure would not be a good idea to correct for this, if one is interested in distributional issues. On top of this, and certainly for less recent years of observation, GDP-based measures of income mostly exclude an important source of welfare for poor (and hence many) people in developing countries: consumption from own produce. In most cases this home production is included in household surveys.

If anything, GDP-based measures of global inequality are expected to produce higher figures, than measurement on the basis of income or expenditures, and this for two reasons: first of all the level of public goods is higher in countries with higher GDP (see

⁴ The generalized entropy class of inequality measures of an income vector $\mathbf{y} = (y_1, y_2, \dots, y_i, \dots, y_n)$

with mean μ , is defined as: $GE(\gamma) = \frac{1}{\gamma^2 - \gamma} \left[\frac{1}{n} \sum_{i=1}^n \left(\frac{y_i}{\mu} \right)^\gamma - 1 \right]$, where $-\infty \leq \gamma \leq +\infty$ and

$\gamma \neq 0, 1$. As a special case, the Theil measure is the limit of this expression for $\gamma \rightarrow 1$. In general, the decomposition of the generalized entropy class of measures uses $p_h^{1-\gamma} s_h^\gamma$ as weights in the 'within'-term (see Cowell and Kuga 1981; Shorrocks 1984).

Milanovic 2002a: 64-65). Secondly, home production and auto-consumption tends to be higher in poor economies.

2.3 Purchasing Power Parity conversion

Data, whether it are GDP-figures or household surveys on income or expenditures, are expressed in local currencies. The dollar remains the standard for data covering the whole world. Yet, it may be clear that a conversion of these figures in a common currency by means of an ordinary exchange rate (from local currency to dollar at the moment of observation) cannot serve as a good basis for welfare comparisons. Indeed, the purchasing power represented by a US\$1 daily income (the poverty line in 1985 calculated by Ravallion *et al.* (1991) is simply not the same when the dollar has to be spent in the US as compared to in Ethiopia. In the US you could hardly survive with this little amount of money, whereas in Ethiopia this income represents about the average daily income of someone employed. The Penn World Tables (Summers and Heston 1991) contain time series of Purchasing Power Parity conversion rates into dollars (PPP\$) for a large set of countries, some of them dating back to the 1950s. The conversion of local currency observations in comparable purchasing parity units by mean of these data, is indispensable for suitable welfare comparisons and by now common practice.⁵

3 Data and methodology

Clearly, measuring world inequality ('Concept III inequality') requires a huge amount of data. Just as in the case of measuring inequality in a specific country, one ideally needs a representative sample of observations on people's income. However, the population now is world population. Contrary to what is the case in many countries, there is no statistical bureau which collects this type of information at the world level. The only solution is to make a *detour* and try to collect as much as possible information at national (or sub-national) levels, in order to compile on the basis of within country distributional information a large enough sample which covers a large enough fraction of the world in terms of population size.

Information on national income distributions usually comes from household surveys conducted at the national level. In some Western countries, also the tax administration provides some information on the basis of administrative data, but we are far from having enough information from such sources. Moreover administrative data might, certainly in developing countries, miss some important parts of income: the virtual income represented by the value of home production and the incomes derived in the informal and black sectors of the economy (the latter might be poorly assessed through surveys too). Usually, the information on individual records of national level household surveys are not publicly available and come across by average income or income shares of certain quantiles or subgroups of the sample. One should therefore in one way or

⁵ This does not mean of course that the specific conversion into PPP\$ proposed by Summers and Heston (1991) is uncontroversial. Dowrick and Akmal (2003) question the supposed declining 'Concept II inequality' on the grounds of biases in the PPP conversion rates used in the Penn World Tables and propose an alternative conversion.

another construct an entire income distribution for the world from this sparse information.⁶

There is some choice on how to reconstruct this distribution: some authors estimate a parametric (Chotikapanich *et al.* 1997) or nonparametric (Sala-i-Martin 2002b) form on the basis of the data-points. Another possibility is to assume that all individuals from the population which are covered by the sample data-point receive the same income. Take the case where we combine GDP *per capita* with distributional data. If one observes that the poorest 20 per cent Indonesian people obtain about 9 per cent of total income (in 1999) and GDP *per capita* is 2790\$, then the best guess is to assume that all 20 per cent poorest Indonesians earn $0.09 \times 5 \times 2790 = 1255.5\$$. Out of necessity one neglects the inequality within that group. Irrespective of whether the decomposability property is accepted, one could easily derive from the previous section that this amounts then to an underestimation of true world inequality. In fact, starting from the constructed distribution, one can obtain the true world income distribution after some regressive Pigou-Dalton transfers.

Question is whether this underestimation is worrying. From the literature (Davies and Shorrocks 1989) and empirical practice it is well known that at least a Gini-index can be estimated quite accurately from 6 to 7 data-points, provided that these are optimally selected. Reconstructing the true world from available data is however more complex than assessing an income distribution from few, but optimally chosen data-points, for at least two reasons. Firstly, the recovering of the true world income distribution has to be accomplished by constructing first a true income distribution per country on the base of some sparse data-points, and there are about 100 such countries in the authors' respective datasets. The true world income distribution is then composed by putting together those estimated country-distributions. The second complication follows from the first one. Working with distributional information per country implies that some data-points, namely those stemming from highly populated countries, cover a significant fraction of the world income distribution⁷ and one risks, therefore, to lose too much information on the inequality within the group of people represented by those observations.

Milanovic's (2002a) 'objective is that the number of such data-points be at least ten in order to have a sufficiently precise description of a distribution' (p. 56). He obtained his distributional information from a set of household surveys which he collected from different sources (the Luxembourg Income Study for Western Europe and North America, the Living Standard Measurement and other surveys of the World Bank for third world countries, and for the transition economies in Eastern Europe and the former Soviet Union the Household Expenditure and Income Database for Transitional

⁶ One should distinguish this from scaling up the information from the sample level to the population level. One could reliably estimate the population parameters on the basis of sample information if the sample is representative for the population. Usually a sample contains relevant weighting factors for each observation unit, in order to perform the scaling up correctly.

⁷ Milanovic's (1993) largest data point stems from a rural region in China covering 180 million people, or about 3.3 per cent of world population at that time. In Sala-i-Martin (2002a, b), there must be at least five data-points with over 4 per cent of world population.

Economies was consulted).⁸ He was able to obtain such information for a sample of 91 ‘countries’ both in 1988 and 1993, covering about 84 per cent of the world population. For about three-quarters of the cases, Milanovic was able to obtain individual records. He converted these into decile data. Moreover, he decided to create some ‘sub-countries’ by considering the rural and urban part of populous Asian countries as separate observation units. He used the rural/urban divide for China, India, Bangladesh, and Indonesia in both 1988 and 1993, and for Pakistan in 1988. To be sure, sometimes Milanovic (2002a) had to revert to surveys which were not conducted in the two years which form the focus of his attention (1988 and 1993), but in a close neighbourhood. And in 12 exceptional occasions in 1988 (mainly for countries which originated from the former Soviet Union), he had to be satisfied with quintile data-points only.

Sala-i-Martin (2002a, b) wants to calculate world income inequality from 1970 onwards on a yearly base. Evidently, this is a much more ambitious project. He therefore reverts to the Deininger and Squire (1996) data-base, which contains for a large set of countries, information on income shares of quintiles for some specific years (not the same for all countries) between 1947 and 1993.⁹ So the price to be paid for this ambitious undertaking is not only sparser data (five data points per country instead of 11 on average in Milanovic 2002a), but also more fragmentary data (no information in specific years).¹⁰ To solve the latter problem, he estimated a linear time trend through the shares for countries for which several observations between 1970 and 1998 were available. He then used the estimated shares to calculate inequality.¹¹ For countries which were not in that case (only one observation on income distribution during 1970-1998), it was assumed that inequality did not change throughout. This was the case for 29 countries, comprising 6 per cent of the sample population. Finally, for another 28 countries (also representing 6 per cent of the sample population), there was no distributional information at all. In that case, as for example for the Democratic Republic of Congo, it was assumed that everybody had the same income. In this way, Sala-i-Martin compiled a data-set with distributional information on 125 countries, covering 88 per cent of the world population in 1998. A major difference with Milanovic’s coverage is that Sala-i-Martin does not include countries from Eastern Europe or the former Soviet Union.

Once they obtained (or constructed) this distributional information in the forms of income shares, both authors apply these shares to the welfare concept they have chosen: GDP per capita for Sala-i-Martin (2002a), income or expenditures from surveys for Milanovic (2002a). They both refrain from estimating a functional form of the distribution function, considering this a too constraining straitjacket from an empirical

⁸ More detailed information on the data sources of Milanovic (2002a) can be found on <http://www.worldbank.org/research/inequality/data.htm>.

⁹ Sala-i-Martin (2002a) mentions that he uses the extension of the Deininger and Squire (1996) database provided by the World Bank, known as the World Development Indicators, but it was not clear to us which version of this database provides distributional information.

¹⁰ The terminology ‘sparse’ and ‘fragmentary’ data stems from Milanovic (2002c).

¹¹ In a companion paper, Sala-i-Martin (2002b) estimated kernel densities on the base of quintile points to generate a more complete picture of the distribution, but the results of this approach do not differ substantially from Sala-i-Martin (2002a), and so we concentrate further exclusively on the latter.

point of view. Instead, they assign to every world citizen of a given income group the same income as the income observed in the data-point.

For the conversion into comparable monetary units, Sala-i-Martin (2002a) uses real PPP\$ figures as presented by Heston *et al.* (2001). Milanovic (2002a) uses the nominal PPP\$ conversion rates of the same source to convert his data points (mean incomes of groups or quantiles) into comparable units. Since both authors use ray invariant concepts of inequality, the difference between real and nominal is not important, once the figures have been converted into common dollars.

Table 1 summarizes the main differences in the data and methodology between the two authors.

Table 1
Data sources and coverage of Sala-i-Martin (2002a, b) and Milanovic (2002a, b)

	Sala-i-Martin	Milanovic
Time period	1970-1998	1988 and 1993
Observation unit	countries	countries, but rural/urban split up for – China, India, Bangladesh, and Indonesia in 1988 and 1993 – Pakistan in 1988
Number of units	125	91
% of world population	88% in 1998	84% in 1988 and 1993
Distributional information	Deininger and Squire (1996) and extension with World Development Indicators of the World Bank	household surveys
Number of datapoints	quintiles	varying, but ≥ 10 ; on average 10.8 in 1988 and 11.4 in 1993; 12 exceptions in 1988 (only quintiles)
Welfare concept	GDP per capita	income or expenditures (depending on the source)
Source	Penn World Tables of Heston, Summers and Aten (2001)	national household surveys collected by the author
Currency	real, 1996 PPP\$	nominal PPP\$

4 Results

Table 2 summarizes the results of both authors. For Sala-i-Martin we have taken the figures bearing on the same time period 1988-1993 of Milanovic.

Table 2
Gini and Theil indices for global world inequality

	Sala-i-Martin (2002a: 60-61, tables 1 and 2)			Milanovic (2002a: 78, table 19)		
	1988	1993	change	1988	1993	change
Gini						
Overall	0.627	0.615	↓	0.628	0.660	↑
Between (sub)countries	0.507	0.496	↓	0.551	0.578	↑
Within			↓	0.013	0.013	-
Overlap				0.064	0.068	↑
Theil						
Overall	0.767	0.729	↓	0.765	0.873	↑
Between (sub)countries	0.574	0.533	↓	0.571	0.649	↑
Within	0.193	0.196	↑	0.194	0.224	↑

Let us first look at the global picture by means of overall or ‘Concept III inequality’. The difference is striking. Milanovic finds a definite increase in inequality. For Sala-i-Martin the ‘disturbing rise’ in world inequality is more apparent than fact: he finds a clear decrease during this same period. This conflicting evidence holds both for the Gini and the Theil-index.

Note that there is a striking resemblance between both authors as far as inequality in 1988 is concerned. Yet, this is not that good news. Indeed, as we explained above, both authors use a different welfare concept, and we would expect Sala-i-Martin, using GDP *per capita*, to find higher inequality figures. We come back to this issue in the next section, where we try to explain the difference.

Stepping into the decomposition methodology, both authors agree that the largest part of ‘Concept III inequality’ is explained by the contribution of the *between* component. Both authors agree that the growth of mean income in China and, to a lesser extent, India tends to decrease world inequality. Indeed, faster growth of a poorer country (China grows faster than Western Europe/US/Japan) can be translated into a combination of a Pigou-Dalton transfer from richer to poorer countries and then rescaling. Since both authors use Lorenz-consistent inequality concepts, one can understand agreement on the conclusions in this respect. But they obtain opposite results on the *total change* of the between country inequality component. This divergence in results for the between component is most intriguing. The difference will be explained in the next section. Notice that the continuing growth of China and India, in the end would tend to increase inequality again (because the divergence with poor Africa and Asia will dominate then the convergence of both China and India to the richer part of the world). But apparently we have not yet reached this turning point.

For the within component, we have to make a distinction between the results obtained by means of the decomposable Theil and the non-decomposable Gini coefficient. According to Milanovic, also the within component increases if one agrees to use the decomposability property embedded in the Theil measure. Hence the rise in overall

inequality can be explained by both between and within components. But if one rejects the latter property (as in the case of the Gini) there is no evidence that within country inequality has contributed to the rise in equality. The rise in overall inequality is explained by an increase in the between component and the overlap of the distributions.

Only in case he uses the Theil index, Sala-i-Martin admits that within inequality has risen, but not enough to counterbalance the falling tendency in the between component. If he decomposes the Gini, we can infer that the sum of the within component and the rest term must have gone down (slightly), but we have no information on the relative contribution of overlap and within in Sala-i-Martin.

Summing up, the results of both authors are surprisingly divergent. The more because they use the same indices, more or less the same methodology, and roughly the same data. More than time, therefore, to step into the section that tries to explain where the divergence comes from.

5 Quarrelling about facts, silence about norms

5.1 Three suspects

The first suspect for the opposite results is of course the fact that average incomes (the size of the cake) are quite different when measured in household surveys or by GDP per capita. We will not repeat the discussions and arguments for preferring one above the other. But recently, Milanovic (2002b) himself offered convincing evidence that this first suspect is not to blame. He rescaled the average incomes of the surveys, used for the distributional information, to the level of the GDP per capita for each country in the years under consideration. We replicate his results in Table 3. Note first that the hypothesis, formulated in section 2, that using GDP per capita instead of household budget survey averages, would increase the *level* of inequality is confirmed for 1988. Compared to the result in Table 2, the Gini in 1988 goes up from 0.628 to 0.633. However, this is not the case for the Theil in 1993 (and Milanovic does not give an explanation for this).¹² But the main conclusion of this exercise is that inequality still increases, and the contradiction with Sala-i-Martin remains unexplained. Remarkably, the overlap term in the Gini coefficient now goes down. This might be explained by the fact that average GDP's *per capita* tend to lie further apart than mean household incomes, which causes a tendency towards less overlaps.

¹² Since the Theil measure is known to give some more attention to rather high incomes, we could infer that incomes *in this spectrum of the global distribution* have become closer than when constructed from mean expenditures/incomes.

Table 3
‘Concept III inequality’ from survey data with average rescaled to GDP/capita
(Milanovic 2002b: 88, Table 28)

	1988	1993	Change
Gini			
Overall	0.633	0.639	↑
Between subcountries	0.561	0.576	↑
Within	0.013	0.013	-
Overlap	0.059	0.050	↓
Theil			
Overall	0.778	0.847	↑
Between subcountries	0.584	0.621	↑
Within	0.194	0.226	↑

A second possible explanation might lie in the omission by Sala-i-Martin of the countries of Eastern Europe and the former Soviet Union. There is ample evidence that, definitely in the period 1988-1993, these countries have experienced a huge increase of inequality (see, for example, Milanovic 1998). But, of course, this increase in within country inequality cannot explain the divergence in the results for the between component. And as was clear from Table 2, this was by far the most important effect. Moreover, although the total effect on world inequality of the increase in inequality in these countries might be positive, the serious drop in the average income of these countries also exercised a reverse effect. The collapse of these ‘poorer rich countries’, and the fast growth of some richer poor countries, produced some convergence of mean incomes. More specifically the convergence of the mean income of China and India to the one of Russia and Ukraine pushed the Gini down with between 0.3 and 0.12 points (Milanovic 2002a: 84-85, table 22).

So, let us examine a third suspect: the distributional information itself. It is difficult to deny that Sala-i-Martin seems to be satisfied with less demanding distributional information. That he does not find an important contribution from within inequality undoubtedly is partly influenced by the fact that he only works with quintile shares. In general, quintiles give too poor information on the within country distribution. This holds all the more so for densely populated data-points like China, India and Bangladesh. Moreover the way he linearly interpolates the data for years where the information is lacking, in practice leads to a constructed lack of variation in his within component. As Milanovic (2002c: 16) shows, the constant character of Sala-i-Martin’s within component is not a coincidence, but points to the fact that his calculations practically reduce to the calculation of the between component (‘Concept II inequality’) augmented with a constant term. This is confirmed by Milanovic’s (2002b: 90) own calculations of ‘Concept II inequality’, based on GDP per capita: the Gini goes down from 52.5 to 51.2 and for the Theil, he obtains a decline from 0.623 to 0.565 (see Table 4).

5.2 A spurious controversy

None of the three suspects, examined so far, can explain the divergence in the conclusion about the between term of the inequality change. Yet, one other experiment by Milanovic provides the key. The information is in Table 4: instead of working with his own data, that subdivided China, India, Indonesia and Bangladesh into rural and urban subgroupings, he lumped them together into ‘whole’ countries, just as Sala-i-Martin did (and most other authors do). The result is striking, and confirms our dismissing of the other three possible explanations. When he calculates the between component on the distributional information of the surveys, with an average income rescaled to GDP per capita, *but with countries comparable to Sala-i-Martin’s definition*, then he obtains a negative contribution to overall inequality of the between component for the Gini (as compared to the positive contribution before!), and a much smaller positive contribution in case of the Theil.

Table 4
Between component of ‘concept III inequality’ from survey data, with average rescaled to GDP/capita and ‘concept II inequality’ (Milanovic 2002b: 90, table 29)

	With urban/rural divide of the data for China, India, Indonesia, Bangladesh			Whole countries			‘Concept II inequality’		
	1988	1993	change	1988	1993	change	1988	1993	change
Gini	0.561	0.576	↑	0.551	0.541	↓	0.525	0.512	↓
Theil	0.584	0.621	↑	0.558	0.567	↑	0.623	0.565	↓

This is revealing of course. The splitting up of countries in large regions, largely improves the within country distributional information of China, India and Bangladesh (Pakistan was split in 1988, not in 1983). Working with whole countries, it is not surprising that Sala-i-Martin misses a lot of within country inequality. The biggest contribution to the rise in inequality in Milanovic’s figures comes from the increasing divergence, not only between rural China and India on the one hand and urban China on the other, but also between the former and some fast growing richer countries (Japan, Germany and to a lesser extent France and the United States) (Milanovic 2002a: 86, table 23).

The fact that both authors seem to agree that it is the ‘between’ component which is the driving force behind the change, is therefore non telling at all. Their between component is simply not comparable. The increase in inequality, found by Milanovic, and stemming in the first place from a ‘between’ component, is, at least in terms of the ‘normal’ definition of countries, a ‘within’ component. And the between component turns out to have a negative or *a much smaller* positive contribution to overall inequality, if one lets the groups coincide with the ‘normal’ country definition. We therefore consider the controversy (or at least this aspect of it) to be more apparent than real.

5.3 On ‘Concept III’, the controversy is empirical

More generally stated, our conclusion would be that, if one wants to measure true world inequality on the basis of sparse data, it is *not* a good idea to start from ‘Concept II-inequality’. Surely not if it is calculated on the basis of the largely arbitrary definitions of countries as states. The problem does not seem to be sparse data as such. Sparse data do allow you to compute overall inequality rather exactly. But one does need a critical amount of data-points; and one has to examine first critically the data-points representing large parts of the world population. Using accessible empirical information on the urban/rural divergence in, for example, China and India, allows Milanovic to derive that representing the whole of China or India’s distribution by means of quintiles is unwarranted. In this respect, it seems that Milanovic gets a better, more accurate, estimation of true world inequality.

There is one minor *caveat* here. The finer distributional information of Milanovic is not constant over time. Between 1988 and 1993, the average number of data-points increases from 10.8 to 11.4, and, maybe more importantly, for 12 countries (mainly from Eastern Europe and the former Soviet Union) with only poor quintile information in 1988, at least ten data-points were obtained in 1993. This might have pushed Milanovic’s results a little bit in the upward direction, revealing however not a true increase in inequality, but a better perception over time of possibly unchanged inequality.

Note that our conclusion of ‘Concept II’ as a nonstarter for assessing true world inequality, is mainly an empirical statement, and has nothing to do with the acceptance or rejection of decomposability as an attractive property. Both authors give results for decomposable and non-decomposable inequality concepts, though Milanovic clearly prefers the non-decomposable Gini above the decomposable Theil coefficients. But decomposability *per se* is a property of the inequality ranking, not of the facts. Indeed, even decomposable inequality concepts can rank the world distribution in 1993 and 1988 differently, but apparently this is not the case here.

5.4 An almost forgotten normative rejoinder

The negative contribution of the between component in the change of overall inequality, becomes even clearer if one looks at the last three columns of table 4, which gives Milanovic’s results for population weighted international inequality (‘Concept II inequality’) on the basis of a much larger sample (122 countries instead of 87). Population weighted international inequality definitely has gone down, a result perfectly in line with Sala-i-Martin’s findings. If one adds to this the fact that Sala-i-Martin essentially adds a constant to this information, we can understand his results.

Hence, one thing seems clear: ‘Concept II inequality’ (between nations) seems to have unambiguously fallen the last couple of decades (and especially between 1988 and 1993). Many, though not all contributions in the literature seem to confirm this finding of both Milanovic (2002b) and Sala-i-Martin (2002a, b). A noteworthy exception is Bourguignon and Morrisson (2002). They obtain increasing between country inequality from 1980 to 1992 for the Theil index. However, once again the definition of the groups might be important: their countries are in fact groups of countries (larger units than Sala-i-Martin), and they defined groups in such a way that the member states are

roughly homogeneous, which might explain the observed increase in heterogeneity (inequality) between regions.

Remarkably, the same source (Bourguignon and Morrisson 2002: 734, table 2) reports decreasing international inequality for the *mean logarithmic deviation*.¹³ Precisely this kind of rank reversals when comparing the same couple of income distributions with different inequality indices, have been a prominent issue in the literature of the past decades (Atkinson 1970; Sen 1973; Kakwani 1984, among many others). Therefore it is all the more remarkable that, as far as this topic is concerned, so little can be learnt from the literature reviewed here. What's more, the data even seem to suggest that welfare economists would *not* agree about the evolution of international inequality over time, even if they would agree about the figures. In Table 5 we give the results from our own calculations of 'Concept II inequality' between 1991 and 1999 measured by the generalized entropy indices with $\gamma \in \{-5, -4, \dots, -1, 0, 1, \dots, 4, 5\}$. We used a sample of 158 countries representing 96.2 per cent of the world population in 1999.¹⁴ We obtain essentially the same results as Milanovic (2002b) and Sala-i-Martin (2002a, b): decreasing population weighted international inequality for $\gamma = 0, 1, 2$. But, widening the normative scope, this unanimous picture disappears for $\gamma \leq -3$ and $\gamma > 3$. Normative differences in opinion determine whether one concludes that population weighted international inequality has risen or the other way around.

Table 5
Population weighted international or 'concept II' inequality 1991-99 on the base of World Development Indicators

Year	value of the parameter γ in the generalized entropy class										
	-5	-4	-3	-2	-1	0	1	2	3	4	5
1991	156.56	26.14	5.77	1.85	0.88	0.61	0.59	0.78	1.32	2.75	6.70
1992	190.01	29.31	5.96	1.80	0.85	0.59	0.58	0.78	1.34	2.83	6.98
1993	158.89	25.35	5.34	1.66	0.80	0.57	0.57	0.77	1.33	2.83	7.06
1994	197.56	29.45	5.77	1.68	0.79	0.56	0.56	0.77	1.34	2.87	7.20
1995	205.11	30.22	5.80	1.65	0.76	0.54	0.55	0.75	1.30	2.77	6.87
1996	217.49	31.71	5.97	1.66	0.75	0.53	0.54	0.74	1.28	2.73	6.77
1997	248.15	34.89	6.33	1.70	0.75	0.53	0.54	0.73	1.28	2.73	6.82
1998	268.16	37.18	6.60	1.72	0.75	0.52	0.53	0.73	1.29	2.79	7.02
1999	312.66	41.50	7.05	1.77	0.75	0.52	0.53	0.73	1.29	2.78	7.05

Which normative opinions lie behind the choice of the parameter γ ? A higher γ tends to give more weight to transfers at the top of the income distribution and *vice versa*. At the lower end of the world income distribution, one might consider the growth of

¹³ The mean logarithmic deviation corresponds to the limit of the generalized entropy index for $\gamma \rightarrow 0$.

¹⁴ We used current PPP\$ GDP *per capita* and population figures from the World Development Indicators 2003 (WDI) database (<http://www.publications.worldbank.org/WDI/>).

(urban) China and India *vis-à-vis* the rest of the poor as a regressive Pigou-Dalton transfer. At the other end, divergence between Eastern European countries and the rich countries is again an anti-Dalton transfer. Putting both these changes in the distribution together, might be interpreted as a convergence of the rich poor and the poor rich. However, an additional Pigou-Dalton transfer between (urban) China and India on the one hand and Eastern European countries on the other hand, might overrule the negative effect of those anti-Dalton transfers. If the evolution of the world income distribution can indeed be sketched by a combination of anti-Dalton transfers at the top and bottom, counterveiled by a Pigou-Dalton transfer in the middle, this might explain the observed fall in inequality for moderate values of γ , while one would perceive an increase in inequality for extreme values of γ (as in Table 5).

Finally, what do we conclude about the evolution of true world inequality (‘Concept III’) in this respect? Comparing two 6 billion long income vectors is a complex matter. We would have expected therefore that the data would leave ample space for normative differences of opinion. This is also what one learns from a closer look at the results of Sala-i-Martin’s (2002a: 61 table 2): despite a lack of emphasis for this finding from the author’s side (and even less an attempt to explain it), at least some rank reversals were obtained. For the period 1980-1989, for instance, the mean logarithmic deviation goes down, but the square of the coefficient of variation goes up.¹⁵

Unexpectedly, Milanovic (2002a: 72-73) suggests that this problem would not occur with his data. Indeed, he obtains Lorenz dominance for the 1988 income distribution as compared to the 1993 distribution. In other words, all Lorenz consistent inequality indices (among which the Generalized Entropy class, and the Gini) would yield a higher value for 1993 as compared to 1988. The results of Table 5 reveal that Lorenz curves of population weighted inequality do intersect between 1991 and 1999. So, we feel justified to expect that Milanovic’s dominance result is vulnerable to the coarse grid he chose: while he has about 1000 observation points, dominance is only checked at percentiles. Many possibly statistically relevant crossings might remain hidden in that way.

6 Conclusion

In the present contribution, we tried to pinpoint down what drives the so widely divergent results in two recent contributions on the measurement of true world inequality. The first one claims that the frequently heard opinion about the ‘disturbing rise’ of inequality is unwarranted by the facts (Sala-i-Martin 2002a, b). At the other extreme, Milanovic (2002a, b) claims that the 1993 income distribution is Lorenz dominated by the 1988 world income distribution.

We showed that it was neither the difference in welfare concept (GDP *per capita* versus income or expenditures) nor the sample (inclusion or not of the transition economies in Eastern Europe) that can be held responsible for the difference.

¹⁵ The squared coefficient of variation equals the generalized entropy index for $\gamma = 2$.

In fact both authors agree that the growth in China, and to a lesser extent in India, had a big, inequality reducing effect on population weighted international inequality. But neglecting the big divergences within China and India between still very poor rural populations, the very fast growing urban areas and some good performing rich countries in terms of growth (Japan and Germany), may have led Sala-i-Martin to a perhaps too hasty conclusion that global inequality went down.

And what about Africa? The continual low (and even negative) growth of many African countries gives rise to the suspicion that inequality will continue to rise. Moreover, the available information on Africa is still poor today, so that, on top of that, we also expect to observe better existing inequality on that continent in the near future.

At least in the eyes of some. Because after all, we do not expect there to be unanimity among scholars in the field on such a complex issue as comparing two 6 billion long income vectors, with unequal means, and unequal length. Nevertheless the normative positions to tackle that job remained largely undiscussed up to now.

Finally, the ongoing debate on the direction of (minor) changes in income inequality during the last decades should not make us blind for the fact this inequality remains to reach a disturbingly high level.

References

- Atkinson, A. B. (1970). 'On the Measurement of Inequality'. *Journal of Economic Theory*, 2, 244-63.
- Barro, R. J. and X. Sala-i-Martin (1995). *Economic Growth*. New York: McGraw-Hill.
- Berry, A., F. Bourguignon, and C. Morrisson (1983). 'Changes in the World Distribution of Income Between 1950 and 1977'. *Economic Journal*, 93, 331-50.
- Bourguignon, F. and C. Morrisson (2002). 'Inequality Among World Citizens: 1820-1992'. *The American Economic Review*, 92, 727-44.
- Bourguignon, F. (1979). 'Decomposable Income Inequality Measures'. *Econometrica*, 47, 901-20.
- Chotikapanich, D., R. Valenzuela, and D. S. P. Rao (1997). 'Global and Regional Inequality in the Distribution of Income: Estimation with Limited and Incomplete Data'. *Empirical Economics*, 22, 533-46.
- Cowell, F. A. and K. Kuga (1981). 'Additivity and the Entropy Concept: An Axiomatic Approach to Inequality Measurement'. *Journal of Economic Theory*, 25, 131-143.
- Davies, J. B. and A. F. Shorrocks (1989). 'Optimal Grouping of Income and Wealth Data'. *Journal of Econometrics*, 42, 97-108.
- Deininger, K. and L. Squire (1996). 'A New Data Set Measuring Inequality'. *World Bank Economic Review*, 10, 565-91.
- Dowrick, S. and M. Akmal (2003). 'Contradictory Trends in Global Income Inequality: A Tale of Two Biases. Australian National University.

- Foster, J. E. and A. Sen (1997). *On Economic Inequality after a Quarter Century*. Oxford: Clarendon Press.
- Heston, A., R. Summers, and B. Aten (2001). Penn World Table Version 6.0, Center for International Comparisons at the University of Pennsylvania (CICUP), December.
- Kakwani, N. C. (1984). 'Welfare Ranking of Income Distributions', in R. L. Basman and G. F. Rhodes (eds), *Advances in Econometrics*. Greenwich, CT: JAI Press, pp. 191-213.
- Lambert, P. and R. Aronson (1993). 'Inequality Decomposition Analysis and the Gini Coefficient Revisited'. *The Economic Journal*, 103, 1221-27.
- Milanovic, B. (1998). 'Income, Inequality, and Poverty During the Transition from Planned to Market Economy'. Washington, DC: World Bank.
- Milanovic, B. (2002a). 'True World Income Distribution, 1988 and 1993: First Calculation Based on Household Surveys Alone'. *Economic Journal*, 112, 51-92.
- Milanovic, B. (2002b). 'Worlds Apart: Inter-national and World Inequality 1950-2000'. Washington, DC: World Bank.
- Milanovic, B. (2002c). 'The Ricardian Vice: Why Sala-i-Martin's Calculations of World Income Inequality Cannot be Right'. Washington, DC: World Bank.
- Quah, D. T. (1996). 'Twin Peaks: Growth and Convergence in Models of Distribution Dynamics'. *Economic Journal*, 106, 1045-55.
- Quah, D. T. (1997). 'Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs'. *Journal of Economic Growth*, 2, 27-59.
- Ravallion, M., G. Datt, and D. van der Walle (1991). 'Quantifying Absolute Poverty in the Developing World'. *Review of Income and Wealth*, 37, 345-61.
- Sala-i-Martin, X. (2002a). 'The Disturbing 'Rise' of Global Income Inequality'. NBER Working Paper w8904.
- Sala-i-Martin, X. (2002b). 'The World Distribution of Income' (estimated from individual country distributions), Columbia University.
- Schultz, T. P. (1998). 'Inequality in the Distribution of Personal Income in the World: How it is Changing and Why'. *Journal of Population Economics*, 11, 307-44.
- Sen, A. (1973). *On Economic Inequality* (1st edition). Oxford: Clarendon Press.
- Shorrocks, A. F. (1980). 'The Class of Additively Decomposable Inequality Measures'. *Econometrica*, 48, 613-25.
- Shorrocks, A. F. (1984). 'Inequality Decomposition by Population Subgroups'. *Econometrica*, 52, 1369-85.
- Summers, R. and A. Heston (1991). 'The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950-1988'. *Quarterly Journal of Economics*, 106, 327-68.
- Theil, H. (1979). 'World Income Inequality'. *Economics Letters*, 2, 99-102.
- Whalley, J. (1979). 'The Worldwide Income Distribution: Some Speculative Calculations'. *Review of Income and Wealth*, 25, 261-76.