

DESA Working Paper No. 26

ST/ESA/2006/DWP/26

August 2006

**Global Income Inequality:
What It Is And Why It Matters?***Branko Milanovic*

Abstract

The paper presents a non-technical summary of the current state of debate on the measurement and implications of global inequality among citizens of the world. It discusses the relationship between globalization and global inequality, shows why global inequality matters and proposes a scheme for global redistribution.

JEL Classification: D31

Keywords: globalization, global inequality, aid

Branko Milanovic is a lead economist in World Bank Research Department, unit dealing with poverty, income distribution and household surveys; senior associate at Carnegie Endowment for International Peace in Washington; adjunct professor at the School for Advanced International Studies at Johns Hopkins University, Washington, DC. He has a Ph.D. in economics (income distribution) from the University of Belgrade (1987). He writes on methodology and empirics of inequality; poverty and social policy in transition economies, and globalization and inequality. His recent publications include *Income and Influence: Social Policy in Emerging Market Economies* (co-authored with Ethan Kapstein, Upjohn, 2003); "True world income distribution 1988 and 1993: first calculations based on household surveys alone", *Economic Journal*, 2002; and *Income, Inequality, and Poverty during the Transition from Planned to Market Economy* (World Bank, 1998). His book *Worlds Apart: Measuring International and Global Inequality* was published by Princeton University Press in 2005.

Comments should be addressed by email to the author (bmilanovic@worldbank.org).

Contents

Some methodological issues	4
How great is global inequality?	7
Is there a link between globalization and global inequality?	11
Does global inequality matter?	13
References	15

UN/DESA Working Papers are preliminary documents circulated in a limited number of copies and posted on the DESA website at <http://www.un.org/esa/desa/papers> to stimulate discussion and critical comment. The views and opinions expressed herein are those of the author and do not necessarily reflect those of the United Nations Secretariat. The designations and terminology employed may not conform to United Nations practice and do not imply the expression of any opinion whatsoever on the part of the Organization.

Typesetter: *Valerian Monteiro*

United Nations
Department of Economic and Social Affairs
2 United Nations Plaza, Room DC2-1428
New York, N.Y. 10017, USA
Tel: (1-212) 963-4761 • Fax: (1-212) 963-4444
e-mail: esa@un.org
<http://www.un.org/esa/desa/papers>

Global Income Inequality: What It Is And Why It Matters?

*Branko Milanovic*¹

Global inequality is a relatively recent research topic. The first calculations of inequality across world citizens were done in the early 1980s (Berry, Bourguignon and Morrisson, 1983; Grosh and Nafziger, 1986). This is because in order to calculate global inequality, one needs to have data on (within-country) national income distributions for most of the countries in the world, or at least for most of the populous and rich countries. But it is only from the early to mid-1980s that such data became available for China,² Soviet Union and its constituent republics and large parts of Africa.

Before we move to an analysis of global inequality, however, it is useful to set the stage by delineating what topics we shall be concerned with and what not. This is necessary, precisely because of the relative underdevelopment of the topic, reflected in the fact that the same or similar terms are often used in the literature to mean different things. We need to distinguish between inequality among countries' mean incomes (inter-country inequality, or Concept 1 inequality, as dubbed by Milanovic, 2005), inequality among countries' mean incomes weighted by the countries' populations (Concept 2 inequality), and inequality between the world's individuals (global, or Concept 3 inequality).

Concept 1 inequality deals with convergence and divergence among countries, and although this line of work was at first couched in inequality terms (see Baumol, 1986), most of the later work used cross-country regressions and β convergence.³ In such regressions, each country/year is one observation. This line of research, which has generated a huge literature, is interesting for a number of reasons, but it has very little to tell us about income inequality among world citizens. This is basically because countries are of unequal population size. Thus, a fast increase in the income of a poor small country will not have the same effect on global inequality as the same per capita increase in a poor and populous country.

Concept 2 inequality tries to take this into account by weighing each country by its population. It is a low-cost approach since it requires knowledge of only two variables: mean income, which is approximated by gross domestic income (GDI) per capita, and population size. The first such calculations were done by Kuznets in 1954 (see Kuznets, 1965: 162ff).⁴ Some thirteen years later, as part of their first study of purchasing power parity (PPP), Kravis, Heston and Summers (1978) calculated Concept 2 inequality for the non-socialist world.

There are two reasons for the enduring popularity of this approach (for recent examples, see Schultz, 1988; Boltho and Toniolo, 1999; Firebaugh, 2003). First, Concept 2 inequality is the largest component of global inequality. Global inequality is, by definition, composed of population-weighted international

1 The views expressed are personal and should not be attributed to the World Bank or its affiliated organizations.

2 The first post-Cultural Revolution household survey in China was conducted in 1978. The first available rural and urban surveys are from 1980 and 1981 respectively (see Ravallion and Chen, 2006: 3).

3 Some of the initial emphasis on inequality, rather than on β coefficients, can still be seen in the use of the sigma convergence where sigma is the standard deviation of income logs.

4 For the year 1949, Kuznets calculated a Concept 2 inequality that covered around a third of the world's population.

inequality (Concept 2 or between-inequality), and inequality due to income differences within countries. The relationship is shown in equations (1) and (2) for Gini and Theil coefficients respectively, where y_i = per capita income of i -th country, p_i = population share of i -th country in total world population, π_i = share of i -th country in total global income, n = number of countries, G_i = Gini coefficient of national income distribution, T_i = Theil coefficients of national income distribution, and L = the so-called overlapping component.⁵ Since the ‘between’ component is by far the larger, accounting for between $\frac{2}{3}$ and $\frac{3}{4}$ of global inequality (depending on what inequality measure one chooses), Concept 2 inequality can be used as a lower-bound proxy to global inequality. Moreover, its movements can be presumed to track changes in global inequality. Second, the data requirements for the calculation of Concept 2 inequality are modest.

$$\text{Concept}_3\text{-Gini} = \sum_{i=1}^n G_i p_i \pi_i + \frac{1}{\mu} \sum_i^n \sum_{j>i}^n (y_j - y_i) p_i p_j + L \quad (1)$$

Concept 2 Gini

$$\text{Concept}_3\text{-Theil} = \sum_{i=1}^n p_i T_i + \sum_{i=1}^n \left(p_i \frac{y_i}{\mu} \right) \ln \frac{y_i}{\mu} \quad (2)$$

Concept 2 Theil

Of course, what Concept 2 inequality does not take into account are within-country inequalities. In calculating Concept 2 inequality, we implicitly assume that each individual within a country has the same per capita income (and thus $G_i=T_i=L=0$). This last assumption needs to be abandoned if we want to calculate ‘true’ global inequality across individuals. But in order to abandon it, one must have access to national income distributions which are available only from household surveys. It is this ‘jump’ that makes such a big difference in data requirements between Concept 2 and Concept 3. From being ‘modest’, the data requirements now become huge, since, ideally, we should have access to national income distributions from all the countries in the world.

This leads to a very important, albeit not sufficiently appreciated, difference between Concepts 2 and 3. This is not a conceptual difference, but rather the difference in what is a commonly used metric of welfare. To be sure, Concept 2 inequality can be calculated using either GDIs per capita or mean disposable incomes retrieved from countries’ household surveys (HS). It is however almost never calculated using the latter because HS means are much more difficult to obtain than from national accounts (NA) data. In contrast, Concept 3 inequality *must* be based on household surveys because the only sources of distributional data are, as said before, household surveys. Because there is no world-wide household survey, this means that the best one can do is to combine individual countries’ surveys, and use disposable per capita income or personal per capita consumption as welfare indicators.

Now, the first problem is that there is a definitional difference between GDI which comes from national accounts, and disposable income which comes from surveys. Second, there has been a recent tendency for these two measures not to move in unison in several important countries (see Deaton, 2005). Thus, even if everything else were fully comparable, a commonly calculated Concept 2 measure that uses national accounts data will differ from an equivalent Concept 2 measure calculated using household surveys because welfare

5 The overlapping component accounts for the fact that somebody who lives in a richer country may have an income lower than somebody from a poorer country (and the converse). L is calculated as a residual, and this is why the Gini index is, unlike the Theil index, not exactly decomposable.

indicators are different and because they have recently diverged for reasons that are not yet quite clear.⁶

This review will deal only with studies of global or Concept 3 income inequality.⁷ The way to estimate global inequality is to calculate Concept 2 inequality using nation accounts data, and to combine it with the empirical observation that within-country income distributions tend to follow a log-normal pattern. Then, the only additional piece of information needed is a Gini coefficient, or some other summary inequality statistic describing national income distributions. They are published in various compendia of Gini coefficients such as the WIDER and Deininger-Squire databases, etc.

Under the assumption of a log normal distribution of income, the inequality statistics allow us to derive an estimate of the variance of each national distribution. Once we know the variance and the mean, and given the assumption of log-normality, we can estimate the entire distribution, that is, each fractile's income. It is then a relatively simple task to combine these national distributions into a single world-wide income distribution, particularly so if one uses an exactly decomposable measure of inequality like the Theil coefficient or the variance of logs. This was precisely the approach followed by many early and some recent studies of global inequality (Berry, Bourguignon and Morrison, 1983; Grosh and Nafziger, 1986; Quah, 1999; Shultze, 1998; Chotikapanich, Valenzuela and Rao, 1997). At times, this approach can be refined by using a bit more information than what is contained in a Gini or Theil index. Sala-i-Martin and Bhalla have used quintiles of income distribution to get a better handle of national distributions and thus a more precise estimate of global distribution. All these methods can be considered *tatonnements*, groping for the global distribution.

These methods are quite ingenious, given their rather minimal information requirements. But they are also very 'costly', because it is often the numerous assumptions, rather than the data, that drive the results. A lot of assumptions are made simultaneously (e.g. that each country's distribution is log normal; that GDP per capita gives the correct mean income, and that its under or over-estimation, compared to household surveys, is constant across poor and rich alike), and it is nearly impossible to tell the impact of each separate assumption on the results. Further, since even the minimal data requirements (national Ginis) are not satisfied annually, authors are led to make additional assumptions (for example, that national inequality does not change or changes in an assumed fashion), so that in the end, the part of the results driven by various assumptions may vastly outstrip the part based on actual data. The best recent examples of such approaches, which are often thinly disguised Concept 2 inequality calculations, are Bhalla (2002) and Sala-i-Martin (2002a; 2002b). In Bhalla's (2002) calculations of global inequality, only one out of his 24 distributions is based on actual data, while 23 are 'derived' through assumptions; in Sala-i-Martin's (2002a) paper, the ratio is one actual to four 'derived' distributions.⁸

6 This area—understanding why national accounts and household survey averages move differently—represents one of the most important areas for further research. Bhalla (2002) must be credited through his, at times single-minded, insistence on using national accounts data for highlighting this issue.

7 And with global inequality conventionally defined as inequality in relative, not absolute, incomes, and using conventional measures of inequality like the Lorenz curve, Gini coefficient or Theil index. The focus on absolute inequality, however, has its own uses (see Atkinson and Brandolini, 2004; Svedberg, 2003; Ravallion, 2004). Similarly, relative income inequality, with the use of different inequality aversion parameters (reflecting, in principle, different welfare judgments), will produce ambiguous results, even where conventional statistics yield a clear outcome (see Capeau and Decoster, 2004: Table 5).

8 For a critique, see Milanovic (2005: 119-127).

Compared to this method, the quantum leap is to directly use household surveys from as many countries as possible (ideally, all). This was done by Milanovic (2002; 2005) and by the World Bank (2005). Here, in principle, global inequality is calculated the same way one calculates within-country inequality, using not national accounts data, but survey data. Another quantum leap in this line of research will occur when these, so far disparate, national surveys are standardized, or a single world-wide household survey is conducted. We shall come back to this in the third section. Now, we have to briefly cover some methodological issues that are very important in this type of research although they seldom receive the attention they deserve.

Some methodological issues

We start with the simplest question of all: what is ‘income’ in calculations of global inequality? As we have seen, most of the early work used national accounts data, that is GDI per capita expressed in the same currency (international or PPP dollars). This is because household survey data for many countries were simply unavailable (and even when they existed, researchers could not get them because the statistical agencies refused to release them).⁹ There are currently three main sources of world wide data on GDI per capita across time and across countries. They are World Bank data, available in World Bank World Development Indicators (WDI), the Penn World Tables, and Angus Maddison’s data. The advantage of using GDI per capita as ‘income’ is that these numbers are relatively uncontroversial, even if the three sources do not always agree among themselves. We know what is meant by GDI per capita and we know that these values do give some generally accepted mean incomes of all nations.¹⁰

The main drawback of this approach is that GDI per capita is not ‘income’ in any recognizable sense to any individual or household. Gross domestic income includes components such as corporate investment from retained profits, build-up of stocks, government spending on defence, etc., which are not part of even broadly defined household disposable income. In addition, publicly financed health and education are part of GDI per capita, but not included in household per capita disposable income unless one is able to impute—which is quasi impossible in a multi-country context—these values back to individual households based on survey data on school attendance and use of medical services.

Another drawback is that the combination of GDI per capita with some distributional statistics (to reflect national income distribution) mixes two aggregates, calculated from different sources, and this ‘mixing’ is not distribution-neutral. On the one hand, we are using country GDI per capita, and on the other, we are applying to this mean, distributional parameters obtained from surveys of household disposable income.¹¹ It was already explained that the two instruments are different by definition. But in addi-

9 This is still the case with many countries that refuse to release micro data to any institution or individual. Such countries are quite diverse, ranging from Japan to Algeria. There is thus a paradoxical situation that a number of expensive instruments like household surveys exist in the world; yet, they cannot be used because of misplaced policies of some countries’ statistical offices. The issue of confidentiality, with which they sometimes defend their practice, is clearly bogus since no researcher can ever identify the participating households. Are we to believe that a researcher in 2006 is going to identify the Japanese households that took part in a survey conducted in 1973?

10 It has not always been like that. Consider the problems of converting Communist countries’ national accounts methodology to the national accounts statistics (NAS) system, and of course, the issue of the deliberate falsification of national accounts.

11 The difference in coverage and definition between national accounts and surveys means that, even if everything were perfectly measured, it would be incorrect to apply inequality or distributional measures—which are derived from surveys—which measure one thing, to means—which are derived from national accounts—which measure another (Deaton, 2003: 35).

tion, such an adjustment is not distribution-neutral. We know that surveys tend to under-represent capital income or under-survey rich people (see Mistiaen and Ravallion, 2003).¹² Thus, a simple scaling-up of all survey incomes by a given parameter will reduce poverty below what it really is and underestimate inequality.¹³ When a different mean (GDI per capita) is applied to a survey-based distribution, we implicitly allocate the difference between GDI per capita and disposable household per capita income across all households (more exactly, in proportion to reported household income). Poor people's incomes are increased in the same proportion as rich people's incomes. But if most of the difference between the two concepts is due to the unreported income of the rich, then this approach wrongly inflates the incomes of the poor.

Now, the income of the rich, which accounts for the bulk of the difference between GDI and disposable income, is of two types: first, the definitionally different part, which consists of publicly-financed health and education, corporate profits, etc., that are consumed by the rich, but also by the middle-classes and some poor in *rich* countries;¹⁴ second, the income of rich people in each individual country that is missed out by surveys (e.g. property incomes). All actual recipients of these incomes are globally-rich because the middle-class and even the poor from the rich world are in the top quintile of the global income distribution, but the difference between the two aggregates is spread much more widely: some of it is imputed to the poor in the poor countries which, we know, receive none of it.

Consider the following example. Let the poverty line \$PPP1 per capita per day. Let the average per capita disposable income from surveys of several groups of the poor in India be \$0.75, \$0.8, \$0.85. Now, suppose (very realistically) that India's GDI is some 35 per cent higher than disposable income. We know that most of this 35 per cent is received by the rich, either because they benefit more from publicly-funded public services, or because they fail to report their property incomes. What the authors (in particular, Sala-i-Martin, 2002a; Bhalla, 2002) then do is to multiply the incomes of the poor by the factor of 1.35. Then, suddenly, none of them is poor any longer: they have all crossed the poverty threshold. Pure magic!¹⁵

But if we decide that 'income' in global studies should be the same concept as in national studies of inequality—that is per capita disposable household income—the problem is not solved yet. This is because national definitions of survey income are very different, and the more countries we include, the more different they become. A huge effort, conducted by the Luxembourg Income Study, has gone into standardization of national definitions. A similar project is underway at the World Bank using Living Standards Measurement Surveys (LSMS). Yet, the standardization covers only a small portion of all countries and surveys.

12 Income from property is notoriously underestimated in household surveys (even leaving aside the fact that most surveys do not include capital gains at all). Concialdi (1997: 261) claims that the best available French household surveys underestimate capital incomes by about 40 per cent. Wagner and Grabka (1999) estimate German property income to be underestimated by almost one-half compared to national accounts data. In Japan, according to Ishizaki (1985), only 12 per cent of property income is 'captured' by household surveys (quoted in Bauer and Mason, 1992: 407).

13 We speak of scaling-up, rather than scaling-down, because GDI per capita is normally greater than household per capita disposable income.

14 Disposable household income, retrieved from surveys in West European nations, amounts to about 60 per cent of GDI. A bulk of that difference is explained by health and education consumption. Publicly-financed health and education as a share of GDI is much less in poor countries.

15 On the additional pitfalls caused by the use of averages from national accounts and distributions from household surveys, see Ravallion (2000), Deaton and Dreze (2002), and Deaton (2003).

The main differences arise in the treatment of self-employed income (what are business expenditures for the self-employed?), valuation of home-consumption, including owner-occupied housing,¹⁶ treatment of publicly-provided health and education benefits, and the use of top-coding of high incomes (where all incomes above a certain ceiling are coded as equal to that ceiling).¹⁷ For different countries and at different levels of development, differences in the treatment of these categories are not equally important. For poor countries, it is the problems of valuation of own consumption and self-employed income that are the most difficult and that can make individual incomes often move up by a factor of 2 or more; for the rich countries, it is the treatment of publicly-provided health and education benefits that is of most concern. Swedish disposable income with them or without them is quite different. For countries with extravagantly rich individuals, it is the underestimation of capital incomes which is of concern.

But there is no agreement that ‘income’ in global inequality studies should be income at all. Many people think that rather than income, one should look at consumption or expenditures as the true indicator of the standard of living. This debate mirrors the debates in individual countries since obviously inequality can be measured using either income or consumption. What lends this debate an added importance in the case of global inequality is that in many countries, household surveys ask questions about income only, while in other countries, they ask for both, or for expenditure only. Then, a global study of inequality has to do what all national studies try to avoid, that is to mix household survey data that use two different concepts of ‘income’: disposable income and consumption. This introduces an error whose direction and magnitude cannot be estimated.

Although in the last few years, there has been a trend toward the use of consumption measures (not the least through the efforts of the World Bank and the influence it has exerted on the choice of survey instruments in the former communist countries and in Africa), we are still far from unanimity on this issue. In the study of global inequality based on 1998 benchmark data, Milanovic (2005: 104) used 63 consumption instruments and 59 income instruments. This represents a significant increase in the number of consumption instruments compared with ten years before (80 income-based and 22 consumption-based distributions), but for some of the most important countries (like China), one still depends on income data alone. If a guess had to be ventured, it could be said that there is likely to be a tendency toward greater use of household per capita consumption as the main welfare indicator. While quite defensible from a strictly methodological perspective, this will open up a number of issues of comparability since most historical income distribution statistics (e.g. in the United States, UK, France, Germany) are income-based. One thus needs to weigh methodological and quality improvements in snapshots of recent income distribution against the breaks in historically existing series.

Considering the problems of the appropriate welfare indicator, other methodological issues are easier to deal with. Whatever ‘income’ is, that ‘income’ should be expressed in per capita terms, and should be equal for the members of a household. This means that the two issues often debated in national inequality studies are ‘solved’ here: the issue of equivalence scales, and intra-household inequality. They are ‘solved’ because at the current level of statistical development, there is simply no way to account for economies of scale and size across different countries.

16 For example, in 1990, the Chinese statistical office changed the valuation of grain output produced by rural households from state-mandated to market prices. This generated a large change in calculated poverty rates and a break in the rural mean income series (see Ravallion and Chen, 2005).

17 For example, the US Current Population Survey ‘top-codes’ all very high wage and capital incomes. Similarly, the maximum capital gain that can be recorded in the survey is \$149,999 per household annually.

The main reason for this is that economies of scale and size depend on the relative prices of public and private goods (if housing is very cheap, economies of size will be small) and they systematically differ between poor and rich countries (see Lanjouw and others, 2004). Until we have a better handle of the relative prices of public and private goods,¹⁸ we cannot adjust internationally for equivalent units. The use of a given equivalence scale for all countries in the world would be much more arbitrary than the use of per capita calculations. Similarly, we lack information about within-household inequalities.¹⁹

Another issue on which there is agreement is that whatever ‘income’ is, it needs to be adjusted using a country’s relative price level. In other words, we need to use PPP exchange rates to translate domestic currencies into international dollars. Ideally, of course, one would like to move toward a better adjustment where, at least for some large countries with less than fully integrated markets, PPP exchange rates would differ between different parts of the country (e.g. the price level in the richest Chinese province is estimated to be 76 per cent higher than in the poorest; see Brandt and Holz, 2006).

Another concern is that the relative prices faced by different parts of the income distribution are not the same. According to Pogge and Reddy (2002), relative food prices faced by the poor in poor countries are higher than implied by the use of a single all-consumption PPP.²⁰ Food prices are what really matters for the poor, and the use of a lower overall price index will artificially boost poor people’s incomes in India and elsewhere in poor countries.²¹ Pogge and Reddy advocate a cross-country project akin to the one currently conducted by World Bank’s International Comparison Project which would generate PPPs relevant for the very modest, principally food, basket consumed by the poor across the world. But so long as within-country (e.g. provincial) PPPs and PPPs differentiated by income class are not available, we are obliged to use a single PPP exchange rate per country.

How about the use of market (rather than PPP) exchange rates in global inequality calculations? This is a useful complement because it gives us a different insight into inequalities. If one is interested in global purchasing power or ability to affect the world economy, then conversion of local incomes into actual US dollars makes sense. But there are relatively few instances where we are interested in this, and most of our interest in global inequality is really based on the desire to compare living standards of different people. For that purpose, PPP exchange rates are, of course, better.

How great is global inequality?

There is general agreement about the size of global inequality, and there is general disagreement about the recent direction of change of global inequality. Table 1 shows the results for global inequality obtained by a

18 And also relative prices of child vs. adult goods if we are to adjust for household composition, and not only for household size.

19 Schultz (1998) however tries to account for intra-household inequality by using gender gaps in schooling (for each country) to estimate gender gaps in income at the household level.

20 This means that the ratio of food prices consumed by the poor in (say) India compared to the food prices of the same goods in the US is higher than the ratio between the overall price level in India and that in the US.

21 The reason behind such income overestimation is as follows. The weights in the ‘world’ consumption basket of goods and services are decisively influenced by the prices and structure of consumption in rich countries since they are obviously the largest consumers. Then, a relatively high consumption of services in poor countries -- which are cheap there, but are assessed at much higher ‘world’ prices -- tends to show poor countries’ (and poor people’s) incomes to be higher than they ‘really’ are. One possible way to adjust for this is to move from the commonly used Geary-Khamis index, which has this property, to more ‘neutral’ price indices (Afriat or EKS) where the weight of rich countries is less (see Dowrick and Akmal, 2001).

number of authors using quite different techniques: most of them mix national accounts information (using GDI per capita as mean income) and household survey information, and only a few use household surveys directly. In all the studies however, the recipients are individuals (inequality is expressed on a per capita basis), and national incomes are converted into international (PPP) dollars although the PPP exchange rates may be drawn from different sources. All Gini values for the 1990s, with the exception of the two extremes (61 and 71), lie within a relatively narrow range between 63 and 66. The similarity in the results is even more remarkable when one realizes that the standard errors of these estimates are between 2 and 3 Gini points,²² and that most of the estimates are consequently within one standard error of each other.

Table 1.

Global inequality (in Gini points) in the 1990s, according to various authors

Author	Year	Gini value	National mean incomes from:	National income distributions from:
Milanovic (2005)	1993	66	Household surveys	Household surveys
Milanovic (2005)	1998	65	Household surveys	Household surveys
Bourguignon and Morrison (2002)	1990s	66	GDI (Maddison)	Household survey estimates
Sala-i-Martin (2002a)	1998	61	GDI (Penn World Tables)	Ginis and quintiles from HS
Bhalla (2002)	2000	65	GDI (Penn World Tables and WDI)	Ginis and quintiles from HS
Dikhanov and Ward (2001)	1999	68	National consumption (WDI)	Ginis and quintiles from HS
Dowrick and Akmal (2001)	1993	71	GDI	Ginis and quintiles from HS
Sutcliffe (2003)	2000	63	GDI (Maddison)	Ginis and quintiles from HS
Chotikapanich,				
Valenzuela and Rao (1997)	1990	65	GDI (Penn World Tables)	Ginis for HS

Key: HS: household survey; GDI: Gross Domestic Income; WDI: World Development Indicators (World Bank).

As for the direction of change—comparison between 1990s and 1980s—there is no unanimity. Sala-i-Martin and Bhalla, using very similar methodologies, argue that global inequality has declined by between 3 and 4 Gini points. Dikhanov and Ward as well as Bourguignon and Morrison find an increase of about 1 Gini point. Sutcliffe concludes that there was no change, and Milanovic finds an increase of 3 Gini points between 1988 and 1993, followed by a decline of 1 Gini point in the next five years.²³ His most recent (and unpublished) calculations for 2002 show another small increase of about 1 Gini point. Thus, according to Milanovic, there are zigzags. They are explained by the slow growth of rural incomes in India and China as well as the economic collapse of Eastern Europe in the early 1990s, both of which contributed to global inequality. When both developments were reversed in the next five-year period, global inequality decreased. But these are zigzags caused by specific economic events in large countries, not a trend.

This lack of unanimity on changes, and disagreements on whether there is any trend at all, stem not only from the differences in methodology, but paradoxically, also from the very similar results that all authors obtain regarding the overall *level* of inequality. The reason is as follows. Different methodologies yield similar inequality levels, but they do so with quite a lot of noise caused by measurement problems. Mean incomes, whether obtained from surveys or national accounts, are not consistently calculated, and key data sources disagree among themselves. The computation of Concept 2 inequality using GDI per capita—a metric on

22 One has to be careful in the interpretation of the standard error. The standard errors are obtained using simple 'bootstrapping' techniques, so they basically show how sensitive the estimated Gini coefficient is to any single observation (Milanovic, 2002). These results do not include any information about the reliability of the underlying national income distributions (viz, how correctly incomes are measured).

23 The World Bank's *World Development Report, 2006* uses mean log deviation as the measure of global inequality. It finds that it has decreased between 1994 and 2000 from 0.87 to 0.82 (see World Bank, 2005: 64).

which there is apparently least dispute—from World Bank or Maddison’s data series will differ by several Gini points. This is because, as pointed out by Sutcliffe (2003), Maddison’s data include estimates for a number of worn-torn or otherwise ‘excluded’ countries like Sudan, Afghanistan, Somalia, the Congo, Cuba, North Korea, etc., that are almost invariably poor and not included in the World Bank database. In addition, Maddison’s growth rates for China are less than the official rates and those used by the World Bank. It is then not surprising that when one superimposes estimates of national distributions on one or another set of GDIs per capita to generate global inequality, the choice of the GDI database will clearly influence the end result.²⁴

Income distribution data, particularly when extrapolated from quintiles or from Gini coefficients, are even noisier. Furthermore, due to the absence of income distribution data for many countries, some authors (e.g. Bhalla and Sala-i-Martin) resort to very dubious assumptions, e.g. that income distributions do not change over time or change in a certain (linear) fashion or everyone in a country has the same income. In most cases, this myriad of assumptions and measurement errors tend not to bias the results in one direction only, but probably to offset each other, producing relatively similar levels of inequality. But when one re-estimates global inequality for another year, while the level hardly changes, the result is (on account of the measurement error, if nothing else) likely to be slightly different. It is that slight difference that is then interpreted as the evidence of a change, or in some cases, even of a trend.

How big is a Gini of around 65? It is larger than the inequality found in any single country including South Africa and Brazil, two of the most unequal countries in the world, whose Ginis are in the upper fifties or low sixties. The Gini coefficient however does not give an intuitive feel for how large global inequalities are. A better way to look at it is to consider how the overall pie is distributed among different fractiles of the distribution. Thus, for example, the top 5 per cent of individuals in the world receive about 1/3 of total world (PPP-valued) income, and the top 10 per cent get one-half. If we take the bottom 5 and 10 per cent, they receive 0.2 and 0.7 per cent of total world income respectively. This means that the ratio between the average income received by the richest 5 per cent and the poorest 5 per cent of people in the world is 165 to 1 (Milanovic, 2005). The richest people earn in about 48 hours as much as the poorest people earn in a year.

Another important question is to ask how much of global inequality is due to differences in the mean incomes of countries and how much is due to income differences within countries. Some 70 per cent of global inequality is ‘explained’ by differences in countries’ mean incomes. This is a sharp reversal from a situation which existed around the time of the Industrial Revolution when more than half of the (admittedly very rough) estimate of global inequality was due to income differences within nations (see Bourguignon and Morrison, 2002).²⁵ Then, in contrast to today, the differences between countries’ mean incomes were relatively small. For example, in 1870, the average (un-weighted) GDI per capita of the ten richest countries was 6 times greater than the average (un-weighted) GDI per capita of the ten poorest countries. In 2002, the ratio was 42 to 1.²⁶

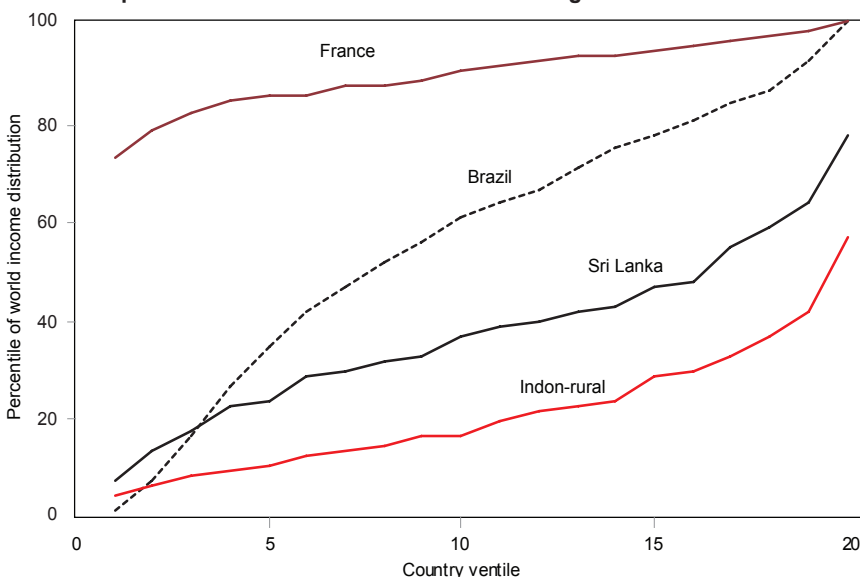
24 For example, Concept 2 inequality calculated using World Bank GDI per capita data from 138 countries shows a decrease of some 3 Gini points between 1985 and 2000. The same concept calculated using Maddison’s data over the same period for about 160 countries shows a decline of only 1 Gini point (author’s unpublished calculations).

25 This is an estimate based on the Theil decomposition between inequality due to the differences in incomes between six country groups, namely ‘Africa’, ‘Asia’, ‘Japan, Korea and Taiwan’, ‘Latin America’, ‘Eastern Europe’, and ‘Western Europe and its offshoots’, and inequality within country groups. Since there are no data on income distributions for most countries in the world prior to 1950, Bourguignon and Morrison use estimates for a few select countries to ‘impart’ the same distributions to other countries in the group. Their between-component accounts for some 30 per cent of global inequality. Obviously, if they had data on all countries’ distributions, the between-component would have been larger. However, it is unlikely to have exceeded one-half of global inequality.

26 Both calculated from Maddison (2004) data.

While income inequality between countries is the largest component of global inequality, overlaps between countries' distributions (that is, some people from a poor country being better off than some people from a rich country) are not negligible either.²⁷ We illustrate this in Figure 1, which plots the position of each 5 per cent (ventile) of different countries' distributions in the global distribution. Consider the line for France. We calculate the mean income (in international dollars) of each French ventile from the lowest (first) to the highest—arrayed on the horizontal axis—and then find their positions in the global income distribution. As can be seen, the poorest 5 per cent of Frenchmen have a mean income which places them in the 72nd percentile of the world income distribution; the richest 5 per cent have incomes which place them in the top percentile of the world. Hence, French income distributions span the range between the 72nd and 100th percentiles in the world. Consider now rural Indonesia at the bottom of the figure. Here, the range is from the 4th percentile to the 56th percentile in the world. The two distributions (France and rural Indonesia) do not overlap at all.²⁸ But this is not the case if we compare Brazil and France: more than a third of all Brazilians are richer than the poorest 5 per cent of the French.²⁹

Figure 1:
The position of different countries' ventiles in global income distribution



Source: Calculated from World Income Distribution (WYD) data. Available at: <http://econ.worldbank.org/projects/inequality>.

The figure illustrates not only that inequality due to within-country distributions is still significant, and that countries are not homogeneous entities composed of either rich and poor people only, but will have practical implications when we discuss global transfers (see the fifth section). In short, if transfers were to flow from mean-income rich to mean-income poor countries, and we do not a priori know who their beneficiaries are, a glance at Figure 1 immediately convinces us of the need to take recipient countries' income distributions seriously. This is because of the probability that money raised from a French citizen will end

- 27 Note that in a world of large between-country income differences, and very small within-country inequalities, there would be no overlap at all, and 100 per cent of global inequality would have been 'caused' by between-country differences.
- 28 This is, of course, true at the level of ventiles. It is quite possible, even likely, that there are some individuals in rural Indonesia who are richer than some individuals in France. If we conducted the analysis in terms of national percentiles, rather than ventiles, there would be some overlap. But it would be clearly minimal.
- 29 Even if at each given ventile, the income of the French is higher than the income of the Brazilians. This means that French income distribution is first-order dominant over the Brazilian distribution (as is, for example, the Sri Lankan over the rural Indonesian), even though the French and Brazilian distributions do overlap (unlike the French and the rural Indonesian).

up in the pockets of somebody who is richer than he, is higher if money is transferred from France to Brazil than if it is transferred from France to rural Indonesia. But we shall return to this topic below.

Is there a link between globalization and global inequality?

It is often implicitly assumed that the changes in global inequality can be interpreted as telling us whether globalization leads to widening or shrinking income differences among individuals in the world.³⁰ However, the causal link between globalization and global inequality is very difficult to make. To see this, consider several ways in which globalization affects inequality among individuals in the world. The first channel goes through globalization's effects on within-country distributions. As we would expect from economic theory, the effect varies between rich and poor countries. In the simplest Heckscher-Ohlin world, globalization would increase demand for, and the wages of, low-skilled labour in poor countries and the wages of high-skilled workers in the rich world. Consequently, we would expect income distribution in poor countries to become 'better' and income distribution in rich countries to get 'worse'.

This is not, however, consistent with what has been observed over the last twenty years when distributions in poor, middle-income and rich countries have tended to grow more unequal (Cornia and Kiiski, 2001). This is an issue which has recently been studied a lot and is still the subject of intense debate: Is openness to blame for increasing wage and income differences in the US? Is openness associated with rising income inequality in poor countries? For example, Milanovic (2005) and Ravallion (2001) find that openness is associated with increased inequality in poor countries, and lower inequality in rich countries, while Dollar and Kraay (2002) argue that there is no systematic effect of openness on inequality.³¹

Then, and this is the second channel, globalization may differently affect mean incomes in poor and rich countries: in other words, it might lead to divergence or convergence in country incomes. There is no unanimity on this point either. Most authors agree that openness is positively associated with mean income growth, but some of them (Sachs and Warner, 1997; World Bank, 2002) find the effect stronger for poor countries, while others (DeLong and Dowrick, 2003; Dowrick and Golley, 2004) argue that the openness premium has been larger for rich than for poor countries during the last twenty years.³² The first group of authors would expect openness to lead to shrinking differences in national mean incomes. Therefore, they have to explain away the observed divergence in mean country incomes by the lack of openness among the laggards. According to the second group of authors, the divergence is an indication that the effects of openness might change over time, and that openness, even if positive for all on balance, may exacerbate inter-country inequality.

Third, the effects of globalization may vary between populous and small countries. This area has not been much explored except in the context of the rather limited (in scope and number) studies of small island

30 We define globalization in terms of an outcome variable: increased share of trade and direct foreign investments in GDI. This is quite acceptable when we have income inequality as the left-hand side variable since inequality moves in response to outcomes (higher or lower trade). But one could also define globalization in terms of policies (e.g. lower trade barriers).


31 For a review of the literature, see Winters, McCulloch and McKay (2004). The role of trade in increasing wage differentials in rich countries is the subject of a voluminous controversy (for some examples, see Freeman, 1995; Slaughter, 1999).

32 According to them, the openness premium was larger for poor countries in the 1960-1980 period, but then changed in the last two decades. For some speculation on what might have triggered that change, see Dowrick and Golley (2004: 53).

economies. Yet, one can imagine that globalization may play out differently in populous countries with large domestic markets, or in small niche economies like Hong Kong, Singapore or Luxembourg, than in middle-size countries.

And finally, and possibly, most importantly, the effect of globalization on global inequality will depend on history, that is on whether populous countries happen to be poor or rich at a given point in time. To see this, assume for a moment that globalization has a positive impact on the growth rates of populous and poor countries, and has no effect on within-country income distributions. This will, in the current constellation of world incomes (see Figure 2), mean that India and China would be expected to catch up with the rich world, while their national distributions will not change, and global inequality will tend to decrease.³³ There would be both mean-income convergence and reduction of global inequality. But let us decouple the poor and populous countries. Suppose that India and China are rich (and still populous) and let most poor countries be relatively small. Now, mean income convergence will continue, but the effect on global inequality will be ambiguous. China and India will benefit from the pro-big bias of globalization, but since they would be rich, globalization will be less beneficial to them than to poor countries. These two effects will pull in opposite directions, and global inequality may go down or up. Moreover, if populous countries are generally poor, the convergence effect is nil, globalization on average favours small countries and leads to the widening of national income distributions, then the overall effect must be to increase global inequality.

Figure 2:
Distribution of people according to GDI
per capita of the country where they live, 2000



Source: Milanovic (2005: 94).

Key:
horizontal axis: GDP per capita
in 1995 international prices;
vertical axis: share of world
population.

This illustrates a key point: even if the effects of globalization on within-country inequality, mean income convergence, and populous vs. small countries, are unambiguous and do not change over time, globalization's impact on global inequality will vary depending on where along the international income distribution, countries with different attributes happen to lie at a given point in time. The implication is, of course, that all statements about the relationship between globalization and global inequality are highly time-specific, contingent on past income history, and not general.

33 We are concerned with effects at one point in time only. Independent changes in population may, by affecting the weights in the inequality statistics, influence changes in global inequality on their own. For example, China's impact on global inequality can be decomposed between income per capita growth effects and population growth effects. Jiang (2006) finds that one-third of China's contribution to reducing global inequality is due to its population growth.

Does global inequality matter?

There are two views on this matter (as on pretty much everything else discussed here). One group of people believe that global inequality is irrelevant. There are two reasons why it may be so. According to Bhagwati (2004), even calculation of global inequality is ‘lunacy’ as it is a mere number. There is no ‘addressee’ to whom this mere number matters because there is no global government and there is no global civil society. According to this view, national inequalities matter because they become the stuff of political discourse; they are used to form political parties or platforms, and to organize interest groups. But at the global level, none of that exists because there is no global polity.

Another reason adduced for the irrelevance of global (or for that matter, all) inequality is that only changes in absolute income matters to the poor and the rich alike (Krueger, 2002; Feldstein, 1999). In the words of Anne Krueger (2002), “Poor people are desperate to improve their material conditions...rather than to march up the income distribution [ladder]”. Thus, even if the absolute income gap between an average American and an average African increases, these authors are unconcerned. After all, they argue, the average African would be a bit less poor. This, of course, assumes that our income relative to the incomes of others does not matter. Yet, this conclusion is at odds with psychological studies that invariably show that people do not care only about their absolute income, but also about where they stand in the social pyramid, and also whether they think this position to be fair (Graham and Felton, 2005; Frank, 2005).

Or—differently—global inequality may matter. On this side of the issue, there are also different approaches. For Pogge and Reddy (2002) and Singer (2002), global poverty and global inequality are ethical issues. Hence, the rich world cannot disown all interest in global poverty and inequality: to some extent, the fate of every individual in the world affects us. Distributional justice within a nation, and in the world as a whole, is -- from an ethical perspective -- the same thing (see Singer, 2002: Chapter 5).

There are also more pragmatic reasons why global inequality may matter. Kuznets (1965 [first published in 1954]: 173-174) formulated the following half a century ago: “Since it is only through contact that recognition and tension are created, one could argue that the reduction of physical misery associated with low income and consumption levels...permit[s] an increase rather than a diminution of political tensions [because] the political misery of the poor, the tension created by the observation of the much greater wealth of other communities...may have only increased.”

When people observe each other and interact, it is no longer simply a national yardstick that they have in mind when they compare their incomes with the incomes of others, but an international or global one. What globalization does is to increase awareness of other people’s incomes, and therefore, the perception (knowledge) of inequalities among both the poor and the rich. If it does so among the poor, then their aspirations change: they may no longer be satisfied with small increases in their own real income, if they know that other people are gaining much more. Therefore, the process of globalization by itself changes the perception of one’s position, and even if globalization may raise everybody’s real income, it could exacerbate, rather than moderate, feelings of despondency and deprivation among the poor.


Globalization, in that sense, is no different from the process which led to the creation of modern nation states out of isolated, and often mutually estranged, hamlets. National income distribution was similarly an abstraction for the people who did not interact with each other, and almost ignored each others’ existence and way of life. However, once nation-states came into existence, national inequality became an issue—sim-

ply because people were able to compare their own standards of living and to make judgments as to whether these income differences were deserved or not. If one believes that the process of globalization would slowly lead to the formation of a global polity, then global inequality will indeed become a relevant issue. For it is difficult to envisage that a fully free exchange of goods, technology and information, transfer of capital, and some freedom in the movement of people can go on for a long time without creating a global polity of sorts and requiring decision-making processes at the global level.

If so, then we need to develop some rules for global redistribution. The first rule, which may be called Progressivity 1 rule (a companion to Concept 1 inequality), is that funds should flow from a (mean-income) rich to a (mean-income) poor country. This requirement is easily satisfied. Even today, bilateral aid is given by rich to poor countries (not the other way round). But in a globalized world, this is not enough. Redistribution needs to be globally progressive—that is, to satisfy the same criteria that we require from redistribution within a nation-state. This means that the tax-payer ought to be richer than the beneficiary of the transfer. But both Progressivity 1 and global Progressivity may be satisfied (as shown in Figure 3 by points B_1 and T_1) while the beneficiary is a relatively rich individual in a poor country and the tax payer a relatively poor individual in a rich country. And it is precisely the perception that many transfers end up in the pockets of the rich elite in poor countries which is fuelling the current discontent with multilateral and bilateral aid. Thus, the third requirement ought to be that transfers be such that inequality decreases in both donor and recipient countries. This will happen only if the tax payer is relatively rich, even within his/her own country, and the beneficiary is relatively poor in his/her country. This situation is illustrated by points such as T_2 and B_2 .

Now, these three requirements regarding global transfers will be more easily satisfied when the income distributions of rich and poor nations do not overlap. This is, for example, the case of France and rural Indonesia (illustrated in Figure 1). Even if the distribution of aid money among Indonesian beneficiaries is random, global progressivity will be satisfied since there are practically no people in rural Indonesia who are better off than even the poorest Frenchmen. But this is not the case if we look at a transfer between France and Brazil. There, assuming that the tax payer belongs to the French middle class (say, around the median of

Figure 3:
Globally progressive transfer



Key: T: tax payer. B: beneficiary.

French income distribution), a purely random allocation of aid to Brazil will still yield a non-negligible probability of 10 per cent for a globally regressive transfer.³⁴ This means that in the design of global transfers, one needs to take into account national income distributions to determine the actual beneficiaries of aid. This is likely to give preference to poor and egalitarian countries since transfers to them are unlikely to be globally regressive.

References

- Atkinson, Anthony B., and Andrea Brandolini (2004). Global world inequality: Absolute, relative or intermediate. Paper prepared for the 28th Conference of the International Association for Research in Income and Wealth, Cork, Ireland, 22-28 August.
- Bauer, John, and Andrew Mason (1992). The distribution of income and wealth in Japan. *Review of Income and Wealth* 38 (4): 403-428.
- Baumol, William (1986). Productivity growth, convergence, and welfare: What the long-run data show. *American Economic Review* 76, December: 1072-1116.
- Berry, Albert, Francois Bourguignon and Christian Morrisson (1983). Changes in the world distribution of income between 1950 and 1977. *Economic Journal* 93, June: 331-350.
- Bhagwati, Jagdish (2004). *In Defense of Globalization*. Oxford University Press, New York.
- Bhalla, Surjit (2002). *Imagine There is No Country*. Institute for International Economics, Washington, DC.
- Boltho, Andrea, and Gianni Toniolo (1999). The assessment: the twentieth century-achievements, failures, lessons. *Oxford Review of Economic Policy* 15 (4), Winter: 1-17.
- Bourguignon, Francois, and Christian Morrisson (2002). The size distribution of income among world citizens, 1820-1990. *American Economic Review* September: 727-744.
- Brandt, Loren, and Carsten Holz (2006). Spatial price differences in China: Estimates and implications. *Economic Development and Cultural Change*, forthcoming. <http://ihome.ust.hk/~socholz/SpatialDeflators.html>.
- Capeau, Bart, and Andre Decoster (2004). The rise or fall of world inequality: A spurious controversy. WIDER Discussion Paper No. 2004/02, UNU-WIDER, Helsinki. <http://www.wider.unu.edu/publications/publications.htm>.
- 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: 533-546.
- Concialdi, Pierre (1997). Income distribution in France: The mid-1980s turning point. In Peter Gottschalk, Bjorn Gustafson and Edward Palmer (eds). *Changing Patterns in the Distribution of Economic Welfare: An International Perspective*. Cambridge University Press, Cambridge: 239-264.
- Cornia, Giovanni Andrea, and Sampsa Kiiski (2001). Trends in income distribution in the post WWII period: Evidence and interpretation. WIDER Discussion Paper No. 2001/89, UNU-WIDER, Helsinki.: <http://www.wider.unu.edu/research/1998-1999-3.1.publications.htm>. Revised version published in Giovanni Andrea Cornia (ed.). *Inequality, Growth and Poverty in an Era of Liberalization and Globalization*. Oxford University Press, Oxford, for UNU-WIDER, Helsinki, 2004: 26-55.
- Deaton, Angus (2003). Measuring poverty in a growing world (or measuring growth in a poor world). NBER Working Paper No. 9822, June, National Bureau of Economic Research, Cambridge, MA. <http://www.wws.princeton.edu/~deaton/working.htm>. Revised version published in *Review of Economics and Statistics* 87, February 2005: 1-19.
- Deaton, Angus, and Jean Drèze (2002). Poverty and inequality in India: A re-examination. *Economic and Political Weekly* September 7: 3729-3748.
- DeLong, Bradford, and Steve Dowrick (2003). Globalization and convergence. In M. Bordo, A.M. Taylor and J. Williamson (eds). *Globalization in Historical Perspective*. Chicago University Press, Chicago.
- Dikhanov, Yuri, and Michael Ward (2001). Evolution of the global distribution of income, 1970-99. Processed, August, World Bank, Washington, DC.
- Dollar, David, and Aart Kraay (2002). Growth is good for the poor. *Journal of Economic Growth* 7: 195-225.

34 Differently put, 10 per cent of Brazilians have higher income than the median Frenchman.

- Dowrick, Steve, and Jane Golley (2004). Trade openness and growth: Who benefits. *Oxford Review of Economic Policy* 20 (1): 38-56.
- Dowrick, Steve, and Muhammed Akmal (2001). Contradictory trends in global income inequality: A tale of two biases. Processed, 29 March, Australian National University, Canberra. <http://ecocomm.anu.edu.au/economics/staff/dowrick/dowrick.html>, published in *Review of Income and Wealth* 51 (2), June 2005: 201-230.
- Feldstein, Martin (1999). Reducing poverty not inequality. *Public Interest* Fall: 33-43.
- Firebaugh, Glenn (2003). *The New Geography of Global Income Inequality*. Harvard University Press, Cambridge, MA.
- Frank, Robert H. (2005). Positional externalities cause large and preventable welfare losses. *American Economic Review* 95 (2): 137-151.
- Freeman, Richard B. (1995). Are your wages set in Beijing? *Journal of Economic Perspectives* 9 (3), Summer: 15-32.
- Graham, Carol, and A. Felton (2005). Does inequality matter to individual welfare? An exploration based on household surveys in Latin America. Center on Economic and Social Dynamics Working Paper No. 38, Brookings Institution, Washington, DC.
- Grosh, Margaret E., and E. Wayne Nafziger (1986). The computation of world income distribution. *Economic Development and Cultural Change* 34: 347-359.
- Ishizaki Tadao (1985). Is Japan's income distribution equal? An international comparison. *Japanese Economic Studies* 14 (2): 30-55.
- Jiang Zhiyong (2006). China component in international income inequality. *Population Research and Policy Review*, forthcoming.
- Kravis, Irving, Alan Heston and Robert Summers (1978). Real GNP per capita for more than one hundred countries. *Economic Journal* 88: 215-242.
- Krueger, Anne O. (2002). Supporting globalization. Remarks at the 2002 Eisenhower National Security Conference on 'National Security for the 21st Century: Anticipating Challenges, Seizing Opportunities, Building Capabilities', 26 September. <http://www.imf.org/external/np/speeches/2002/092602a.htm>.
- Kuznets, Simon (1965). *Economic Growth and Structure: Selected Essays*. Oxford University Press, New Delhi.
- Lanjouw, Jean, Peter Lanjouw, Branko Milanovic and Stefano Paternostro (2004). Relative price shifts, economies of scale and poverty during economic transition. *Economics of Transition* 12, September: 509-536.
- Maddison, Angus (2004). World population, GDP and GDP per capita, 1-2000 AD. Data set, Organization for Economic Cooperation and Development, Paris. <http://www.eco.rug.nl/~Maddison/>.
- Milanovic, Branko (2002). True world income distribution, 1988 and 1993: First calculations based on household surveys alone. *Economic Journal* 112 (476), January: 51-92.
- Milanovic, Branko (2005). *Worlds Apart: Global and International Inequality, 1950-2000*. Princeton University Press, Princeton, NJ.
- Mistiaen, Johan, and Martin Ravallion (2003). Survey compliance and the distribution of income. Working Paper No. 2956, January, World Bank, Washington, DC.
- Pogge, Thomas W., and Sanjay Reddy (2002). Unknown: The extent, distribution, and trend of global income poverty. Processed, July 26, Columbia University, New York. <http://www.columbia.edu/~sr793/povpop.pdf>.
- Quah, Danny (1999). 6×10^9 : Some dynamics of global inequality and growth. Processed, December, London School of Economics, London. <http://econ.lse.ac.uk/staff/dquah/p/9912sbn.pdf>
- Ravallion, Martin (2000). Should poverty measures be anchored to the national accounts? *Economic and Political Weekly* 34, August 26: 3245-3252.
- Ravallion, Martin (2001). Growth, inequality and poverty: Looking beyond averages, *World Development*, 29 (11), November: 1803-1815
- Ravallion, Martin (2004). Competing concepts in inequality debate. *Brookings Trade Forum, 2004*. Brookings Institution, Washington, DC: 1-23.
- Ravallion, Martin, and Shaohua Chen (2006). China's (uneven) progress against poverty. *Journal of Development Economics*, forthcoming.
- Sachs, Jeffrey, and David Warner (1997). Fundamental sources of long-run growth. *American Economic Review* 87 (2): 184-188.
- Sala-i-Martin, Xavier (2002a). The disturbing 'rise' of world income inequality. NBER Working Paper No. 8904, April, National Bureau of Economic Research, Cambridge, MA. www.nber.org.
- Sala-i-Martin, Xavier (2002b). The world distribution of income. NBER Working Paper No. 8905, May, National Bureau of Economic Research, Cambridge, MA. www.nber.org.
- 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 (3): 307-344.

- Singer, Peter (2002). *One World: The Ethics of Globalization*. Yale University Press, New Haven.
- Slaughter, Matthew (1999). Globalization and wages: A tale of two perspectives. *World Economy* 22 (5), July: 609-629.
- Sutcliffe, Bob (2003). A more or less unequal world? World income distribution in the 20th century. Working Paper No. 54, Political Economy Research Institute, University of Massachusetts, Amherst. Published as: World inequality and globalization. *Oxford Review of Economic Policy* 20 (1): 15-37.
- Svedberg, Peter (2003). World income distribution: which way. Institute for International Economic Studies, Stockholm. Published in: *Journal of Development Studies* 40 (5), June 2005: 1-32.
- Wagner, Gert, and Markus Grabka (1999). Robustness assessment report (RAR) Socio-Economic Panel, 1984-1998. Processed, German Socio-economic panel, Berlin. <http://www.lisproject.org/techdoc/ge/ge94survey.doc>.
- Winters, Alan, Neil McCulloch and Andrew McKay (2004). Trade liberalization and poverty: Evidence so far. *Journal of Economic Literature* 42, March: 72-115.
- World Bank (2002). *Globalization, Growth and Poverty: Building an Inclusive World Economy*. Policy Research Report, World Bank, Washington, DC.
- World Bank (2005). *World Development Report, 2006: Equity and Development*. Oxford University Press, New York, for World Bank, Washington, DC.

