

Global Income Inequality by the Numbers: in History and Now

—An Overview—

Branko Milanovic

The World Bank
Development Research Group
Poverty and Inequality Team
November 2012



Abstract

The paper presents an overview of calculations of global inequality, recently and over the long-run as well as main controversies and political and philosophical implications of the findings. It focuses in particular on the winners and losers of the most recent episode of globalization, from 1988 to 2008. It suggests that the period might have witnessed the first decline in global inequality between world citizens since the Industrial Revolution.

The decline however can be sustained only if countries' mean incomes continue to converge (as they have been doing during the past ten years) and if internal (within-country) inequalities, which are already high, are kept in check. Mean-income convergence would also reduce the huge "citizenship premium" that is enjoyed today by the citizens of rich countries.

This paper is a product of the Poverty and Inequality Team, Development Research Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at bmilanovic@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

GLOBAL INCOME INEQUALITY BY THE NUMBERS: IN HISTORY AND NOW

--AN OVERVIEW--

Branko Milanovic¹

JEL classification: D31.

Key words: globalization, global inequality, citizenship premium.

Number of words: about 8,700.

Sector board: social protection.

¹ Branko Milanovic is Lead economist, Development Research Group, Poverty and Inequality Unit (Email: bmilanovic@worldbank.org). The paper was first presented at the Conference on Global Justice held at the Libera Università Internazionale degli Studi Sociali (LUISS) in Rome on June 6-9, 2012. It is written as part of the KCP-funded project “Changeable inequalities: facts, perceptions and policies” TF012968. The interpretations and conclusions of this paper are entirely mine. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the governments they represent.

When we think of income inequality, our first reaction is to think of it within the borders of a country. This is quite understandable for a world where the nation-state is very important in determining one's income level, access to a number of benefits, from pensions to free health care, and where by far the dominant way in which political life is organized is at the level of a country. However, in the era of globalization another way to look at inequality between individuals is to go beyond the confines of a nation-state, and to look at inequality between all individuals in the world. Once we do so, many of the things about inequalities in general that we believe or that we think we know change; it is like going from a two-dimensional flat world to a three-dimensional one.


As the world becomes more integrated the global dimension of inequality is likely to become increasingly relevant. This is for at least two reasons: because of much greater movement of factors of production across borders, and because of greater influence of other people's (foreigners') standard of living and way of life on one's perceived income position and aspirations. Greater movement of capital, goods, technology and ideas from one end of the globe to another implies greater connectivity with people who are not one's compatriots, and greater dependence on other nations for generation of one's income. Movements of labor which illustrate this interdependence in a most obvious fashion are still less important than movements of capital, but they are increasing. The knowledge of how other people live and how much money they make influences strongly our perception of own income and position in the income pyramid. An imaginary community of world citizens is thus gradually built. And once this is done, comparisons of actual incomes and welfare between different members of that imaginary community acquire importance. This is why global inequality, even if not as relevant and important for an average individual as inequality within her political community (nation state) will gain in importance. Once we compare ourselves with people from other parts of the world, we are indeed interested in global income distribution. Global inequality begins to matter.

1. Three concepts of inequality and how they evolved over the past sixty years

When we talk about inequality that transcends national borders, we really often have in mind not one but three different concepts—even when we are not fully aware of it. I am going to articulate these three concepts.

The first concept of inequality (let's call it *Inequality 1*) is focused on inequality between nations of the world. It is an inequality statistic calculated across GDPs or mean incomes obtained from household surveys of all countries in the world, without population-weighting.

Figure 1 Three concepts of inequality defined



To show how this is done, consider the three individuals in the top row of Figure 1: the height of each person represents the GDP or mean income of his or her country. Somebody from a poor country would be represented as a short person, somebody from a middle-income country as a person of medium height, and somebody from a rich country as a very tall person. When we calculate this concept of inequality, we take all countries with their mean incomes –we have some

150 countries in the world with such data- and calculate the Gini coefficient.² China and Luxembourg have the same importance, because we do not take population sizes into account. Every country counts the same, somewhat like in the UN General Assembly.

Consider now the second row of the figure which would help us define Concept 2 inequality or *Inequality 2*. There, individuals from poor countries are all equally short as before and those from rich countries all equally tall, but the difference lies in the fact that countries' population sizes are now taken into account. We do exactly the same as we did in *Inequality 1*, but now China and Luxemburg (or any other country) enter the calculation with their populations. In Figure 1, the poor country is the most populous (5 individuals out of total of 10 displayed there), and the middle-income country, the least populous (2 individuals). Introducing population is very important. As we shall see in the next section, during the past 25 years, the movements in Concept 1 and Concept 2 inequalities were very different. Recall, however, that in both cases the calculation takes into account not actual incomes of individuals, but country averages.

Inequality 3 is the global inequality, which is the most important concept for those interested in the world as composed of individuals, not nations. Unlike the first two concepts, this one is individual-based: each person, regardless of her country, enters in the calculation with her actual income. In Figure 1, this is represented by the different heights of individuals who belong to the same country. Not all Americans have the average income of the United States, nor do all Chinese have the average income of China. And indeed in Figure 1, the poorest person is from the middle-income country, while her compatriot is the second richest (the second tallest) in our group of ten individuals.

But moving from Concept 2 to Concept 3 inequality is not easy. The chief difficulty comes from the fact that to calculate the latter we need access to household surveys with data on individual

² Gini coefficient is a statistical measure of inequality which takes its name from the Italian statistician and economist Corrado Gini. The Gini index is the most frequently used measure of inequality, ranging from 0 – when everybody has the same income - to 1, or 100 (expressed as a percentage or an index), when one person gets the entire income of a city (province, nation, world)—whatever is the relevant population over which we calculate inequality.

incomes or consumption. They have to be measured using the same or similar methodology, and need to include as many countries of the world as possible. Perhaps at least 120-130 surveys are needed in order to cover more than 90 percent of the world population and account for 95 or more percent of world income.³ Ideally of course we would like to have surveys from all the countries in the world. This is a very hard requirement. There are still quite a few countries, mostly in Africa, where household surveys are not regularly conducted and where methodologies change, rather brusquely, from one survey to another, thus rendering comparisons difficult.


Because the calculation of global inequality relies on household surveys, we cannot calculate Inequality 3 with much precision for the period before the mid- or late-1980s. There are simply no household surveys available for too many parts of the world. The first available Chinese household surveys are from 1982, the first usable surveys from the former Soviet Union are from 1988, and for many Sub-Saharan African countries, the earliest household surveys date from the mid-1980s. Thus, for the past, we have to rely on much more tentative data, where countries' income distributions are only approximated, using various more or less reliable methods. This is particularly so if we wish to study global inequality in the long-run, covering the 19th century as well—a topic which I will discuss in Section 3 below.

Figure 2 displays the movements of the three types of inequalities after the Second World War. The Gini coefficient is on the vertical axis. Inequality 1 was stable from 1960 to 1980. This means that there was no systematically faster or slower growth of poor or rich countries. Neither were poor catching on with the rich, nor were the two groups growing further apart. Divergence started only at the beginning of globalization, around 1980, and went on until the turn of the century. These two decades were very bad as far as convergence, or catching up by poor countries, is concerned: rich countries grew, on average, faster than poor countries. However China and India,

³ The coverage is always greater for total world income than population because countries that do not have household surveys are generally poor countries whose importance in global output is small.

which are the huge success cases of that period and the two most populous countries in the world, do not enter into calculation of Inequality 1 with greater weights than any other country.

Figure 2. International and global inequality, 1952-2011:
 “The mother of all inequality disputes”



Let us now consider further Figure 2. Why is it called “the mother of all inequality disputes”? To see what the dispute is about, consider the difference in the movements of Inequality 1 and Inequality 2. While the first, as we just saw, rose during the globalization era, the second declined, at times even dramatically. Measured by Inequality 2, the world has certainly become a much better (“more convergent” or more equal) place precisely during the same period. Thus, those who desire to emphasize the unevenness of globalization tend to focus on growing inter-country gaps, not taking into account sizes of population, and prefer Inequality 1. Those who, on the contrary, wish to focus on positive aspects of globalization tend to favor Concept 2, and to point to the indubitable successes of China and India. In effect, to grasp intuitively why and how Concept 2 inequality declined, we need just to recall that in these calculations, China counts a lot because of its large population size. And China, starting in the 1980s from an extremely low level of income,

has during the past three decades grown very fast, converging on the rich world. Until recently, it was China alone that had been preventing a rise in global inequality as measured by Concept 2. But now it has “support” from India which is also registering high rates of growth, and is also starting from a very low baseline. High rates of growth of these two countries are thus the major factor underlying the downward trend of Inequality 2.

Inequality 3 can be calculated, as mentioned before, only from the mid-1980s because we do not have household surveys going further back in time. Figure 2 shows that Inequality 3 is higher than Inequality 2. This is true by definition because in Inequality 3 people enter the calculations with their actual incomes, not with country averages. A quick glance at Figure 1 shows that the variability of heights is greater in the third row than in the second. Averaging-out reduces measured inequality.

To calculate “true” global inequality or Concept 3 inequality, we have to adjust people’s incomes with the price levels they face and which, of course, differ between countries. We are interested in real welfare of people and those living in “cheaper” countries will get a boost in their incomes compared to what they make in nominal dollar terms. The currency we use is international (or “PPP” for purchasing power parity) dollar with which, in principle, one can buy the same amount of goods and services in any country of the world. Indeed, if we were not to adjust for the differences in price levels, and were to use nominal dollars, global inequality would have been even higher. This is because price levels tend to be lower in poorer countries, and income of people living in poorer countries thus gets a significant “boost” when we use PPP dollars.

Often, a key issue of concern regarding global inequality is not only its level, but its trend: has it been going up or down during the globalization era? Global inequality is calculated at approximately five-year intervals, from 1988 (the first dot on the left) to 2008 (the dot on the right). If we compare this last dot with a couple of dots for the earlier years, we see something that may be historically important: perhaps for the first time since the Industrial Revolution, there may be a

decline in global inequality.⁴ Between 2002 and 2008, global Gini decreased by 1.4 points. We must not rush to conclude that what we see in the most recent years represents a real or irreversible decline, or a new trend, since we do not know if the decline of global inequality will continue in the next decades. It is so far just a tiny drop, a kink in the trend, but is indeed a hopeful sign. For the first time in almost two hundred years—after a long period during which global inequality rose and then reached a very high plateau—it may be setting onto a downward path.

The main reason for this break in the previous trend is what also underlies the decrease in Concept 2 inequality: fast growth of relatively poor and very populous countries, most notably China and India. Their growth, reflected in the rising real incomes of their populations, has not only curbed the rise in global inequality, but pushed it slightly down. China's and India's roles stand in marked contrast to the two other factors that influence global inequality and which have both been clearly pro-inequality. The first is the divergence of countries' mean incomes which lasted from around 1980 to 2000; the second were rising within-national inequalities in many countries. The catching-up of poor and large countries has been the sole factor offsetting these upward pressures. But it has been such a strong factor that it has either kept global inequality from rising or, more recently with the acceleration of Indian growth, reduced it.


What can we say about the level of global inequality? What does the Gini of about 70, which is the value of global inequality (see Figure 2), mean? One way to look at it is to take the whole income of the world and divide it into two halves: the richest 8% will take one-half and the other 92% of the population will take another half. So, it is a 92-8 world. Applying the same type of division to the US income, the numbers are 78 and 22. Or using Germany, the numbers are 71 and 29. Another way to look at it is to compare what percentage of world population, ranked from the poorest to the richest, is needed to get to the cumulative one-fifths of global income. Three-

⁴ Our knowledge of the long-run evolution of global inequality is indeed very tentative, as far as its exact levels are concerned, but very clear as far as broad tendencies since the mid-19th century: the Industrial Revolution, by creating a massive divergence between the rich Western countries and the rest of the world, has pushed global inequality up (see also Section 3 below).

quarters of (the poorer) world population are needed to get to the first 1/5th of total income, but only 1.7% of those at the top suffice to get to the last one-fifth.

Global inequality is much greater than inequality within any individual country. In Figure 3, global Gini of 70 is shown together with the Ginis for several countries. Global inequality is substantially greater than inequality in Brazil, a country that is often held, despite the recent improvements under the Lula presidency, as an exemplar of excessive inequality. And it is almost twice as great as inequality in the United States.⁵

Figure 3. Global Gini coefficient compared to the Ginis of selected countries



How confident are we that these numbers truly reflect what is happening to inequality among world citizens? The global inequality numbers come from calculations done across representative national surveys which monitor incomes or consumption of households. About 120 such surveys stand behind each of the six dots shown in Figures 2 and 3. These 120 surveys include actual incomes or consumption levels for about 10 million people in the world. This is about 1.5% of the current world population, not a negligible number, and in principle, sufficiently

⁵ The vertical axis in Figure 3 shows Gini coefficient in its “natural” values, i.e., not in percentages. Thus a Gini of 0.7 displayed there is the same as a Gini of 70. For simplicity, we use the second approach throughout the paper.

representative for the world as a whole even if recently there has been an apparent greater reluctance of the rich to participate in national household surveys. This in turn likely imparts a downward bias to national and perhaps to global inequality estimates.⁶

Table 1 shows the coverage of world population by household surveys. The second row from the bottom shows the overall coverage which was in all years but one greater than 90%. This is quite good, but should not make us forget that the countries that are omitted because they do not conduct household surveys are not a random draw from among all countries in the world but are all poor countries such as Afghanistan, Sudan, Congo, Somalia, Eritrea, etc. This is reflected in substantially lower population and income coverage of Africa. While the population coverage of other continents never falls below 92 percent, African coverage at its peak is 78 percent (see Table 1). And, in a worrisome development, the number and availability of household surveys in Africa is currently less than five or ten years ago. If we could include all of the omitted countries, global inequality would increase. In other words, what we calculate here, the Gini of about 70, is a lower bound to global inequality, simply because we do not have data from many of the poorest countries. Thus both the decreasing participation of rich individuals in national surveys, and the fact that countries that do not have surveys are overwhelmingly poor, bias the global inequality numbers down.

⁶ The effects of non-participation in surveys on measured inequality is by definition difficult to estimate since income of people who refuse to participate is not known. It is only indirectly (e.g., by looking at the geographical distributions of refusals as in Korinek, Mistiaen and Ravallion, 2005) that we can conclude that it is the rich who comply less. The difficulty of figuring out the effects of rich's non-participation exists despite the intuition that it must underestimate actual inequality. In a model proposed by Angus Deaton (2005), where compliance decreases with income following a Pareto-like function, standard deviation of income distribution does not change and inequality, by most measures, is unaffected. However, with different non-compliance functions, inequality may indeed be underestimated.

Table 1. Population coverage by household surveys, 1988-2008 (in percent)

	1988	1993	1998	2002	2005	2008
Africa	48	76	67	77	78	75
Asia	93	95	94	96	94	98
Latin America and the Caribbean	87	92	93	96	96	95
Post-communist countries	99	95	100	97	93	92
Rich world (Western Europe, North America and Oceania)	92	95	97	99	99	97
<i>World</i>	87	92	92	94	93	94
<i>Number of countries with household surveys</i>	103	122	124	122	122	116

Note: Post-communist countries include Eastern European countries (many of which are members of the EU), and former Soviet republics. This is not an ideal classification, and in the future it may have to be changed.


2. From the fall of the Berlin Wall to the global financial crisis: Who won and who lost

It is generally thought that there are two groups who are the big winners of the past two decades of globalization: first, the very rich, those at the top of national and global income distributions, and second, the middle classes of emerging market economies, in particular China, India, Indonesia and Brazil. Is this true? Figure 4 provides an answer by showing the change in real income (measured in constant international or PPP dollars) between 1988 and 2008 at various percentiles of the global income distribution.

What parts of the global income distribution registered the largest gains between 1988 and 2008? As the figure shows, it is indeed among the very top of the global income distribution and among the “emerging global middle class”, which includes more than a third of world population, that we find most significant increases in per capita income. The top 1% has seen its real income rise by more than 60% over those two decades. The largest increases however were registered around the median: 80% real increase at the median itself and some 70% around it. It is there, between the 50th and 60th percentile of the global income distribution that we find some 200 million Chinese, 90 million Indians, and about 30 million people each from Indonesia, Brazil and Egypt. These two groups—the global top 1% and the middle classes of the emerging market economies—are indeed the main winners of globalization.

The surprise is that those at the bottom third of the global income distribution have also made significant gains, with real incomes rising between more than 40% and almost 70%. The only exception is the poorest 5% of the population whose real incomes have remained the same. It is this income increase at the bottom of the global pyramid that has allowed the proportion of what the World Bank calls the absolute poor (people whose per capita income is less than 1.25 PPP dollars per day) to go down from 44% to 23% over approximately the same 20 years.

Figure 4. Change in real income between 1988 and 2008 at various percentiles of global income distribution (calculated in 2005 international dollars)



Note: The vertical axis shows the percentage change in real income, measured in constant international dollars. The horizontal axis shows the percentile position in the global income distribution. The percentile positions run from 5 to 95, in increments of five, while the top 5% are divided into two groups: the top 1%, and those between 95th and 99th percentiles.

But the biggest losers (other than the very poorest 5%), or at least the “non-winners,” of globalization were those between the 75th and 90th percentiles of the global income distribution whose real income gains were essentially nil. These people, who may be called a global upper-middle class, include many from former Communist countries and Latin America, as well as those citizens of rich countries whose incomes stagnated.

Global income distribution has thus changed in a remarkable way. It was probably the profoundest global reshuffle of people’s economic positions since the Industrial revolution. Broadly speaking, the bottom third, with the exception of the very poorest, became significantly better-off, and many of the people there escaped absolute poverty. The middle third or more became much richer, seeing their real incomes rise by approximately 3% per capita annually.

The most interesting developments, though, happened among the top quartile: the top 1%, and somewhat less so the top 5%, gained significantly, while the next 20% either gained very little or faced stagnant real incomes. This created polarization among the richest quartile of world population, allowing the top 1% to pull ahead of the other rich and to reaffirm in fact -- and even more so in public perception -- its preponderant role as winners of globalization.


Who are the people in the global top 1%? Despite its name, it is a less “exclusive” club than the US top 1 percent: the global top 1% consists of more than 60 million people, the US top 1% of only 3 million. Thus, among the global top percent, we find the richest 12 percent of Americans (more than 30 million people) and between 3 and 6 percent of the richest Britons, Japanese, Germans, and French. It is a “club” still overwhelmingly composed of the “old rich” world of western Europe, northern America and Japan. The richest 1% of the embattled Euro countries of Italy, Spain, Portugal and Greece are all part of the global top 1 percentile. However, the richest 1% of Brazilians, Russians and South Africans belong there, too.

To which countries and income groups do the winners and losers belong? Consider the people in the median of their national income distributions in 1988 and 2008. In 1988, a person with a median income in China was richer than only 10% of world population. Twenty years later, a person at that same position within Chinese income distribution, was richer than more than one-half of world population. Thus, he or she leapfrogged over more than 40% of people in the world. For India, the improvement was more modest, but still remarkable. A person with a median income went from being at the 10th percentile globally to the 27th. A person at the same income position in Indonesia went from the 25th to 39th global percentile. A person with the median income in Brazil gained as well. She went from being around the 40th percentile of the global income distribution to about the 66th percentile. Meanwhile, the position of large European countries and the United States remained about the same, with median income recipients there in the 80s and 90s of global percentiles. But if the economic crisis that currently affects these countries persists, we should not be surprised to find the median individual in the “rich world” becoming globally somewhat poorer.

So who lost between 1988 and 2008? Mostly people in Africa, some in Latin America and post-Communist countries. The average Kenyan went down from the 22nd to the 12th percentile globally, the average Nigerian from the 16th to 13th percentile. A different way to see this is to look at how far behind the global median was an average African in 1988 and twenty years later. In 1988, an African with the median income of the continent had an income equal to two-thirds of the global median. In 2008, that proportion had declined to less than one-half. The position of a median-income person in post-Communist countries slid from around the 75th global percentile to the 73rd. The relative declines of Africa, and Eastern Europe and the former Soviet Union confirm the failure of these two parts of the world to adjust well to globalization, at least up to the early years of the 21st century. Their improved more recent performance is still too fragile to have been reflected in the data.

The Lorenz curves, which plot the percentage of cumulative income (running from 1 to 100) on the vertical axis against the percentage of cumulative population (running also from 1 to 100) on the horizontal axis, for 1988 and 2008 intersect in an almost textbook-like fashion (see Figure 5). Neither distribution is Lorenz-dominant. The gains at below and around the median make the Lorenz curve for 2008 lie above the one for 1988 all the way up to the 80th percentile. For example, the bottom two-thirds of world population received 12.7 percent of world income in 2008 as against 9.3 percent in 1998. But the stagnation or decline in real income of the global upper middle class, and big gains realized by the top 1%, reverse the position of the Lorenz curves for the last one-fifth of the distribution. Here, the top 1% in 2008 receives almost 15 percent of global income vs. 11½ twenty years earlier.

Figure 5. Lorenz curves for global income distributions in 1988 and 2008



Note: The Lorenz curve shows on the horizontal axis, the cumulative percentage of population, ranked from the poorest to the richest percentile, and on the vertical axis, the cumulative percentage of total income received by such population percentiles. If, for a given value of x, y is greater, it means that the bottom x percent of population receives a greater share of total income.

The bottom line is that these results show a remarkable change in the underlying global income distribution. We now live in a world with a bulge around the median with significantly rising incomes for the entire second third (or more) of the global income distribution. That is the new aspiring global middle class. We also see growing wealth and probably power of those at the very top and, remarkably, stagnant incomes for both the people just below the “enchanted” richest 1 or 5 percent, and those poorest in the world.

3. Global inequality over the long-run: From proletarians to migrants


I will now look at global inequality over the long sweep of history. It is here that we can establish an important finding, which goes, I think, into some core issues of political philosophy and economics.

Let us try to do for the entire period since the Industrial Revolution the same type of global inequality calculations which we have just shown for the last 20 years. We ask, “what was global inequality then—say, around the mid-19th century?” It is a question impossible to answer with any precision, because we do not have household surveys or any other reliable sources of income data for these times. Nonetheless, some important attempts to estimate it have been made before, notably by François Bourguignon and Christian Morrisson (2002), who were the first to take such a long-run view. They used income levels (GDP per capita) from Angus Maddison’s database (2004; 2007), and some of their own (necessarily often shaky) estimates of income distributions for different parts of the world, to create global income distributions for eleven benchmark years spanning the period 1820-1992. They did the best one could do with the available data—and their results have been corroborated, to the extent that it is possible to corroborate something as tentative, by several later authors (Van Zanden, Baten, Foldvari and van Leeuwen 2010; Milanovic, 2011). The basic story that emerges from these calculations of income inequality in far-away times is that since the Industrial Revolution, which launched a score of European countries and their overseas off-shoots onto a path of faster growth, global inequality kept on rising until the mid-20th century. There was a period of more than a century of steady increase in global inequality, followed by perhaps fifty years (between the end of the Second World War and the turn of the 21st century) when global inequality remained on a high plateau, changing very little. We saw this in Figure 2 where the six dots are all within several Gini points of each other, that is, within one-standard error of the calculated Gini coefficients. It is only in the early 21st century that global inequality might have commenced its downward course. If indeed this happens to pass, global

inequality would have charted a gigantic inverted U-shaped curve and perhaps in some fifty years—if the emerging market economies continue to grow faster than the rich world—we might be back to the state of affairs that existed around the time of the Industrial Revolution.

But, for now, we are still very far from it. And perhaps nothing shows it better than Figure 6. There the height of the bar represents the Theil coefficient of global inequality in two baseline years: 1870 and 2000.⁷ The height of the bar is much greater now, meaning that global inequality today is greater than in 1870, which of course is not a surprise.⁸

Figure 6. A Non-Marxian world:
Level and composition of global inequality in the 19th century and around year 2000
(measured by the Theil index)



Note: I use Theil mean log deviation because it is exactly decomposable (as between “class” and “location”) and because the importance of each component does not depend on the rest of the decomposition. Anand and Segal (2008) in their review of global inequality studies suggest that it is the most appropriate inequality index for this kind of decomposition.

⁷ Theil coefficient, named after the Dutch econometrician Henri Theil, is another way to measure inequality. It is not as popular, nor is its meaning as intuitive, as Gini, but in this case, when we have to decompose inequality into two components, Theil coefficient is preferable to Gini whose decomposition is not “exact”. That is, with the Gini there is a residual terms whose interpretation is not always clear.

⁸ The results would have been the same with the Gini.

What is less obvious and less well known is that the shares of the two factors determining global inequality have changed in a remarkable fashion. Global inequality can be decomposed into two parts. The first part is due to differences in incomes within nations, which means that that part of total inequality is due to income differences between rich and poor Americans, rich and poor Chinese, rich and poor Egyptians and so on for all countries in the world. If one adds up all of these within-national inequalities, one gets their aggregate contribution to global inequality. This is what I call the “class” component to global inequality because it accounts for (the sum) of income inequalities between different “income classes” within countries. The second component, which I call the “location” component, refers to the differences between mean incomes of all the countries in the world. So there, one actually asks “how much are the gaps in average incomes between England and China, between the Netherlands and India, between the United States and Mexico and so on influencing global inequality?” It is the sum of inter-country differences in mean incomes. In technical terms the first part - “class” - is also called “within inequality”, the second part - “location” - is called “between inequality”.

Figure 6 plots these two parts, class and location, for the years 1870 and 2000. Around 1870, class explained more than 2/3 of global inequality. And now? The proportions have exactly flipped: more than 2/3 of total inequality is due to location. The implication of this overwhelming importance of location, or which is the same, citizenship (*i.e.*, being a member of a rich or poor country), for our lifetime incomes can be also very well captured by another exercise. We divide the population of each country into 100 income percentiles, ranked from the lowest to the richest. Now, if we run a regression with income levels of these percentiles (for 120 countries, this gives 12,000 observations) as the dependent variable, and on the other side of the regression, use as the only explanatory variable the mean income of the country where each percentile comes from, we explain between more than one-half of variability in individual incomes. This is a remarkable achievement for a single explanatory variable. Differently put, more than fifty percent of one’s income depends on the average income of the country where a person lives or was born (the two

things being, for 97% of world population, the same). This gives the importance of the location element today. There are of course other factors that matter for one's income, from gender and parental education which are, from an individual point of view externally given circumstances, to factors like own education, effort and luck that are not. They all influence our income level. But the remarkable thing is that a very large chunk of our income will be determined by only one variable, citizenship, that we, generally, acquire at birth. It is almost the same as saying, that if I know nothing about any given individual in the world, I can, with a reasonably good confidence, predict her income just from the knowledge of her citizenship.

As stated in the title of Figure 6, we live today in a non-Marxian world. Karl Marx could indeed eloquently write in 1867 in "Das Kapital", or earlier in "The Manifesto" about proletarians in different parts of the world—peasants in India, workers in England, France or Germany—sharing the same political interests. They were invariably poor and, what is important, they were all about equally poor, eking out a barely above-subsistence existence, regardless of the country in which they lived. There was not much of a difference in their material positions. One could imagine and promote proletarian solidarity, and consequently—because equally poor people of different nations faced equally rich people each in their own nation—a generalized class conflict. This was the idea behind Trotsky's "permanent revolution". There were no national contradictions, just a worldwide class contradiction.


But if the world's actual situation is such that the greatest disparities are due to the income gaps between nations, then proletarian solidarity does not make much sense. Indeed income levels of poor individuals in poor countries are much lower than income levels of poor people in rich countries. Those who are considered nationally poor in the United States or the European Union have incomes which are many times greater than the incomes of the poor people in poor countries and moreover often greater than the incomes of the middle class in poor countries. And if that gap is so wide, then one cannot expect any kind of coalition between these income-heterogeneous groups of nationally poor people, or at least not any coalition based on the similarity of their material

positions and near-identity of their economic interests. Proletarian solidarity is then simply dead because there is no longer such a thing as the global proletariat. This is why ours is a distinctly non-Marxian world. But what kind of the world is it? To this question I turn next.

4. Gaps between country incomes today

In Milanovic (2012), I have argued that a proper analysis of global inequality today requires an empirical and mental shift from concerns with class to concerns with location, in other words, a movement “from proletarians to migrants”. This was meant to summarize a macro-development that has taken place over the last two centuries. If the main determinant of one’s income is now location, who are the underdogs? People who live in poor countries. And what do underdogs want to do? They want to become richer at home, or failing that, to migrate to richer places.

Figure 7. Different countries and income classes in global income distribution, 2005



Note: The line drawn at $y=60$ shows the global position of the poorest 5% of the US population.

To illustrate the difference in the economic positions of people from different countries, we resort to the same exercise as was sketched above: I divide the populations of all countries into groups of 5% (called ventiles, since there are 20 such groups in a population) running from the poorest to the richest. This is shown on the horizontal axis of Figure 7: the poorest ventile in any country will be at $x=1$. Consider for example the poorest 5% of people in the United States. I put them all together, and calculate their average income; I then do the same for the next 5%, then for the next 5% --all the way to the very top, richest ventile. The poorest 5% of Americans are making around \$3,000-4,000 per capita per year. How do they compare with the rest of the world? In what percentile of the global income distribution would they be? This is shown on the vertical axis. We can start with intuition: poor Americans are unlikely to be among the poorest people globally speaking, because their incomes are not that low. For example, we know that some 20 percent of global population live at less than 1 international dollar per day, while the US poverty line (below which, in principle, nobody in the United States should fall) is 13 dollars per day. Thus, intuitively and based on such very limited evidence, we can already expect the poorest Americans to be relatively high up in the global income distribution. Indeed, as shown by the graph, the poorest Americans are at the 60th percentile of world income distribution. This means that they have higher annual income than 60% of the world population. As one moves higher up, obviously each richer ventile of Americans will stand even higher in the world income distribution, with the richest 5% of Americans belonging to the global top 1 percent. (With a more detailed and finer partitioning it can be shown that the top 11 percent of Americans are all part of the highest global percentile, as we saw in Section 2.).

How does the same thing look for a country like India? The very top of the income distribution in India overlaps with the very bottom of the income distribution in the United States. Clearly, there are millionaires in India as well as other people who are quite rich, and the same graph with percentiles (rather than ventiles) would have shown the top end of India's income distribution to be a little bit higher, but even in that case it would not go past the global 80th

percentile. So these rich Indians, as a group, barely match the average income of middle-class Americans. Note that these are indeed very large groups of people and that the averages may conceal some very high individual incomes: if I use ventiles, each Indian ventile consists of some 60 million people, if I use percentile each percentile is 12 million people. The latter figure is equal to the population of the municipality of Mumbai. But the key point is that although there are in India some very rich, and even some extravagantly rich people, their numbers are not statistically significant, and the number of people who have the standard of living of the American middle class is still very limited.

Consider the same graph for China. China dominates India throughout the whole income distribution (people at a given percentile level of Chinese distribution always have higher income than people at that same percentile of India's income distribution), and the Chinese top ventile attains almost the 80th percentile of the world's income distribution. If we used percentiles, the top 1% of the Chinese would be better-off than 93% of world population.

Consider now Brazil. Not surprisingly, Brazil mimics the world. The poorest people in Brazil are at the bottom of the global income distribution, among the poorest people in the world, while its fairly large middle class enjoys income levels that place it between the 70th and 80th percentiles in the world. At the very top, the richest Brazilians are part of the top one or two global percentiles.

Figure 8. Italy and the rest of the world




Figure 8, displayed in the exactly the same fashion as Figure 7, is dedicated to Italy but could have done for any important migration-receiving country: United States, Germany, France or Spain. In the Figure, Italy's distribution is compared with that of the world, and then with the distributions of the countries providing the bulk of immigrants into Italy. First, where is Italy, compared to the rest of the world? Its poorest people are just below the 60th global percentile; it is, as we just saw, approximately the same percentile where the poorest Americans are. As we move toward the richer Italian ventiles, their global position (obviously) improves, and the richest 5% of Italians are at the top of the world, that is, among the top global percentile.

Compare now Italy with a few other countries, Germany for example. What stands out is the very high income level of the bottom ventiles for Germany. The same would be true if instead of Germany, we used Denmark, Norway or other Nordic countries. The people at the bottom of the income distributions in those countries are around the 80th percentile of the world income distribution or higher. The poorest Danes are at the 90th percentile while in the countries like

Mozambique and Uganda not even the top ventile reaches further than the 65th percentile. The poorest Danes (as a group) are richer than the richest Ugandan (as a group).

To zero in on the importance of citizenship and consequently migration, it is useful to compare the data for Italy with those for the countries where most of migrants into Italy come from. This is because the major implication of a world where location matters is that migration can significantly increase a person's income. The way to improve one's standard of living is simply to move to a richer country. In Albania, about 30% of the population have incomes that are below the poverty threshold in Italy, and obviously these people, even if they were to become the poorest people in Italy after migration, would still improve their real income. The same is true for Argentina: a very high percentage—about a quarter of the population—have incomes that are below the Italian de facto poverty threshold. And finally consider the Ivory Coast, as a representative of African countries. There, a staggering 80% of the population live below the Italian poverty threshold. So if these 80% of Ivoirians were to move to Italy, they would all become better off—even if they were just to join the poorest Italians.

5. Concluding remarks: Philosophical reflections and political implications.

I want to conclude with two points which I think can be derived from what I discussed so far.

The first one is an issue for political philosophy. If most of global inequality is due to differences in location, can we treat location, and thus citizenship, as a rent or a premium (or observely, as a penalty)? Is citizenship—belonging to a given country, most often through birth—something that gives us by itself the right to greater income? Is there a difference in our view of the matter if we take a global, as opposed to national perspective? Is there a contradiction between the two?

Within a single country, society tries in principle to limit the advantages that accrue to people born in rich families. This includes having access to better education and health, to

powerful friends and private information, and of course to greater wealth. Society tries to limit these inherited advantages by either taxing wealth or by making education, health etc. available to all, regardless of their income level. But what is the case in the “global world”? The situation is, at one level, very similar. There are rich countries that have accumulated lots of wealth, and transmit that wealth, along with many other advantages, to the next generations of their citizens. This is why, for example, the poorest Americans are relatively well-off by world standards. They are lucky to have been born in the country that is rich (or has become rich; the case was different with the poorest Americans in the 17th century). And there are also people from poor countries who do not have wealth, and advantages and opportunities it confers. But—and this is in stark difference to the within-country case —this is considered unobjectionable, or rather it is not questioned whether one may keep on benefiting from something that the previous generations have created, and she has simply inherited by virtue of birth. In one case, we frown upon the transmission of family-acquired wealth to offsprings if two different individuals belong to the same nation. In the other case, we take it as normal that there is a transmission of collectively acquired wealth over generations within the same nation, and if two individuals belong to two different nations, we do not even think, much less question, such acquired differences in wealth, income and global social position.

Now, in political philosophy, there are good arguments to go on with that approach, as we implicitly do today, and there are also good arguments to disapprove of it. It is hard to decide which way is right. But what we can do is to put that argument on the table, open it for discussion.

The second implication of all of this has to do, of course, with the issue of migration. If citizenship explains 50 percent or more of variability in global incomes, then there are three ways in which global inequality can be reduced. Global inequality may be reduced by high growth rates of poor countries. This requires an acceleration of income growth of poor countries, and of course continued high rates of growth of India, China, Indonesia, etc. The second way is to introduce global redistributive schemes although it is very difficult to see how that could happen. Currently, development assistance is a little over 100 billion a year. This is just five times more than the

bonus Goldman Sachs paid itself during one crisis year. So we are not really talking about very much money that the rich countries are willing to spend to help poor countries. But the willingness to help poor countries is now, with the ongoing economic crisis in the West, probably reaching its nadir. The third way in which global inequality and poverty can be reduced is through migration. Migration is likely to become one of the key problems—or solutions, depending on one's viewpoint— of the 21st century. To give just one stark example: if you classify countries, by their GDP per capita level, into four “worlds”, going from the rich world of advanced nations, with GDPs per capita of over \$20,000 per year, to the poorest, fourth, world with incomes under \$1,000 per year, there are 7 points in the world where rich and poor countries are geographically closest to each other, whether it is because they share a border, or because the sea distance between them is minimal. You would not be surprised to find out that all these 7 points have mines, boat patrols, walls and fences to prevent free movement of people. The rich world is fencing itself in, or fencing others out. But the pressures of migration are remaining strong, despite the current crisis, simply because the differences in income levels are so huge.

I conclude with something that resembles a slogan: either poor countries will become richer, or poor people will move to rich countries. Actually, these two developments can be seen as equivalent. Development is about people: either poor people have ways to become richer where they are now, or they can become rich by moving somewhere else. Looked from above, there is no real difference between the two options. From the point of view of real politics, there is a whole world of difference though.

REFERENCES:

Anand, Sudhir and Paul Segal, "What Do We Know about Global Income Inequality?", *Journal of Economic Literature*, No. 1, 2008, p. 85.

Bourguignon, François and Christian Morrisson (2002), "The size distribution of income among world citizens, 1820-1990", *American Economic Review*, September, pp. 727-744.

Deaton, Angus (2005) "Measuring poverty in a growing world", *Review of Economics and Statistics*, vol. 87, No. 1, pp. 1-19.

Korinek, Anton, Johan A. Mistiaen, and Martin Ravallion (2005), "Survey Nonresponse and the Distribution of Income", World Bank Policy Research Working Paper 3543, March.

Maddison, Angus (2004), "World population, GDP and GDP per capita, 1-2000 AD". Available at <http://www.eco.rug.nl/~Maddison/>.

Maddison, Angus (2007), *Contours of the world economy, 1-2030: Essays in macro-economic history*, Oxford: Oxford University Press.

Milanovic, Branko (2011), "Global inequality and the global inequality extraction ratio: the story of the past two centuries", *Explorations in Economic History*, vol. 48, November, pp. 494-506.

Milanovic, Branko (2012), "Evolution of global inequality: from class to location, from proletarians to migrants", *Global Policy*, vol. 3, No. 2, May.

Van Zanden, Jan Luiten, Joerg Baten, Peter Foldvari and Bas van Leeuwen (2010), "The Changing Shape of Global Inequality, 1820-2000", typescript, December 2010.