

NBER WORKING PAPER SERIES

THE DISTURBING “RISE” OF GLOBAL INCOME INEQUALITY

Xavier Sala-i-Martin

Working Paper 8904
<http://www.nber.org/papers/w8904>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2002

This paper was partly written when I was visiting Universitat Pompeu Fabra in Barcelona. I thank Sanket Mohapatra for extraordinary research assistance and for comments, suggestions and short speeches related to this paper. I also benefitted from the comments of Elsa V. Artadi, Tony Atkinson, François Bourguignon, Gernot Doppelhofer, Laila Haider, and Casey B. Mulligan. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2002 by Xavier Sala-i-Martin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Disturbing “Rise” of Global Income Inequality
Xavier Sala-i-Martin
NBER Working Paper No. 8904
April 2002
JEL No. D31, F0, I30, I32, O00

ABSTRACT

We use aggregate GDP data and within-country income shares for the period 1970-1998 to assign a level of income to each person in the world. We then estimate the gaussian kernel density function for the worldwide distribution of income. We compute world poverty rates by integrating the density function below the poverty lines. The \$1/day poverty rate has fallen from 20% to 5% over the last twenty five years. The \$2/day rate has fallen from 44% to 18%. There are between 300 and 500 million less poor people in 1998 than there were in the 70s.

We estimate global income inequality using seven different popular indexes: the Gini coefficient, the variance of log-income, two of Atkinson's indexes, the Mean Logarithmic Deviation, the Theil index and the coefficient of variation. All indexes show a reduction in global income inequality between 1980 and 1998. We also find that most global disparities can be accounted for by across-country, not within-country, inequalities. Within-country disparities have increased slightly during the sample period, but not nearly enough to offset the substantial reduction in across-country disparities. The across-country reductions in inequality are driven mainly, but not fully, by the large growth rate of the incomes of the 1.2 billion Chinese citizens.

Unless Africa starts growing in the near future, we project that income inequalities will start rising again. If Africa does not start growing, then China, India, the OECD and the rest of middle-income and rich countries diverge away from it, and global inequality will rise. Thus, the aggregate GDP growth of the African continent should be the priority of anyone concerned with increasing global income inequality.

Xavier Sala-i-Martin
Department of Economics
Columbia University
420 West 118th St., 1005
New York, NY 10027
and NBER
xs23@columbia.edu

1. Introduction	1
2. Methodology and Data	11
(A) Measuring Individual Income from Aggregate GDP and Income Shares ...	11
(B) Estimating Income Shares	12
3. The World Income Distribution Function	15
(A) The Kernel Density Function	15
(B) Poverty Rates	16
(C) Poverty Headcounts	18
4. Measuring Global Income Inequality	20
(A) Income Inequality Indexes	20
(i) Ad Hoc Indexes	20
(ii) Social Welfare Function Indexes	21
(iii) Axiomatic Indexes	22
(iv) Across-Country Inequality	24
(v) Decomposing Inequality: Within and Across-Country Measures ...	26
(B) Seven Estimates of Global Income Inequality	28
(i) The Gini Coefficient in a Historical Perspective	28
(ii) Non-Decomposable Measures	30
(iii) Decomposable Measures	31
5. The Assumption of “Equal Income Within Quintiles”	33
6. The Role of China	35
7. Converging-to-the-Rich, Diverging-from-the-Poor : African Growth	36
8. Conclusions	38
References	40
Figures	43
Tables	60

*“To the detached observer, noting the contrast between the presumed benefits of globalization and developments in the real world, the international economy displays a number of worrying trends. Most obviously, **poverty and inequality have grown alongside the expansion of globalization.** In a world of **disturbing contrasts, the gap between rich and poor countries, and between rich and poor people, continues to widen**”*

Kevin Wadkins

(Background Paper for the 1999 UNPD’s “Human Development Report”)

*“**The evidence strongly suggests that global income inequality has risen in the last twenty years.** The standards of measuring this change, and the reasons for it, are contested – but the trend is clear. “*

Robert Wade

(“Inequality of World Incomes: What Should be Done?”, 2001)

*“With regard to incomes, **inequality is soaring through the globalization period – within countries and across countries. And that’s expected to continue**”.*

Noam Chomsky

(“September 11th and Its Aftermath: Where is the World Heading?

Conference presented in Chennai (Madras), India: November 10, 2001)

*“The dramatic advance of globalization and neoliberalism ... has been accompanied by an **explosive growth in inequality** and a return of mass poverty and unemployment. The very opposite of everything which the modern state and modern citizenship is supposed to stand for.”*

Ignacio Ramonet

(in Le Monde Diplomatique, May 1998)

1. Introduction

We live in a world that displays disturbing trends. Among them, nothing is more disturbing than the exploding rises of poverty and income inequality. The quotes above are only a very small sample of what is rapidly becoming the general consensus among journalists, antiglobalization activists and social scientists concerned about the state of the world in which we live. This pessimistic view is fed by one of the most widely cited publications on the issue: the United Nation's Human Development Report (HDR). In the 1999 issue, the HDR says¹: *"Poverty is everywhere. Gaps between the poorest and the richest people and countries have continued to widen. In 1960, the 20% of the world's people in the richest countries had 30 times the income of the poorest 20%. In 1977, 74 times as much. This continues the trend of nearly two centuries. Some have predicted convergence, but the past decade has shown increasing concentration of income among people, corporations and countries."*²

Although poverty and income inequality are usually mentioned in the same statement, they are very different concepts. While it is unanimously agreed that poverty is bad, it is less clear that income inequalities are undesirable. After all, increases in inequality can arise from the worsening of the poor (a situation that is clearly bad) or the improvement of the rich (a situation

¹ <http://www.undp.org/hdro/99.htm> pages 29 and 36.

² It is also widely agreed that these disturbing trends are to be blamed on globalization and the neo-liberal policies imposed by the World Bank, the IMF and the WTO. The HDR says: *"National and international economic policies shifted sharply in the 1970s and 1980s towards more reliance on the market -diminishing role of the state. By 1997, India had reduced its tariffs from an average of 82% in 1990 to 30%, Brazil from 25% to 12% and China from 43% to 18%. Driven by technocrats, the changes were supported by the IMF and the World Bank as part of comprehensive economic reform and liberalization packages. Conditions for membership to the WTO were important incentives. Country after country undertook deep unilateral liberalization, not just in trade but in foreign investment. These changes have sped the pace of globalization and deepened the interactions among people. The new rules of globalization focus on integrating global markets, neglecting the needs of people that markets cannot meet. The process is concentrating power and marginalizing the poor, both countries and people."*

that is clearly not bad). It can be argued that excessive income equality is not good for the economy as it tends to kill the incentives to invest in both physical and human capital. In this sense, income inequalities are seen as the rate of return to investment. On the other hand, it can also be argued that excessive inequalities create social tensions and political instability. In this case, inequalities are seen by the poor as the rate of return to social and economic disruption. In this paper, we will take an agnostic view. Our main goal is to simply estimate the evolution of the world distribution of income, poverty and income inequality over the last three decades without really and we will leave it to others to judge the welfare implications of our results.

Before presenting our estimates, however, we would like to analyze the methodology used and conclusions reached by the HDR. The evidence put forth by the HDR is important because, as argued above, the HDR is widely cited and reproduced in hundreds of popular publications all over the world.

The HDR shows that inequalities have “exploded” over the last two decades (which are the so-called “globalization decades”) by following three steps:

Step 1. Show that within-country inequality is rising.

The first step is to show that income disparities are widening within each country or at least within a large number of countries. In page 36 of the report we read: *“The transition economies of Eastern Europe have experienced the fastest rise in inequality ever. Income inequalities also grew markedly in China, Indonesia, Thailand and other East and South-East Asian countries. Recent studies show inequality rising in most OECD countries during the 1980s and into the early 1990s. Of 19 countries, only one showed a slight improvement.”*

The statement is very clear. However, it is interesting to see a claim that the transition economies of Eastern Europe have experienced the fastest rise in inequality ever, because in the Table they report on page 39 they show that, for Ukraine, Russia and Lithuania, inequality has, indeed, risen. However, for Hungary and Poland it has remained virtually unchanged. Equally interesting (or mysterious) is the statement about income inequalities in Indonesia and other East and South-East Asian countries also growing markedly because the Gini coefficient for Indonesia went down by 16.6% between 1970 and 1990, the one for South Korea went down by 15.2%

between 1980 and 1988, and the one for Hong-Kong went down by 10.2% between 1981 and 1991.³

The most curious of all is the statement about the 19 OECD countries. To document the claim that only one OECD country has experienced reductions in inequality, the HDR reports a table on page 39 with a list of 19 countries. If we focus on the column for disposable income, we see that nine of the 19 countries do not experience any change in inequality and only one experiences a decline (Italy). Only the UK, the US and Sweden display a “large deterioration” (defined as more than 16%).

What is curious is the number of countries reported, 19. This is curious because in 1999 (the year the report was written) the OECD had 29 members.⁴ Why were only 19 reported? A quick look at the U.N.’s own website reveals that, of the 10 missing countries, only Iceland and the Czech Republic do not have data on Gini coefficients for some extended period during the 1980s. However, Ginis are reported for the remaining countries: Turkey’s Gini experienced a decline of 15% between 1973 and 1987, Mexico’s declined by 7% between 1977 to 1989, Korea’s went down by 15% between 1980 and 1988, and Switzerland’s Gini declined by 9% between 1982 and 1992. According to the definitions on page 39 of the HDR, Turkey and Korea experienced “large declines” in inequality, whereas Mexico and Switzerland enjoyed “small declines”. Of the remaining countries, Luxembourg, Austria and Greece experienced “zero” changes in inequality and Hungary and Poland are reported in the table for Eastern Europe. In sum, the HDR does not report data for seven OECD countries for which data is available in their own website. All of them experience a “large decline”, a “small decline” or “no change” in

³ The numbers are taken from the United Nations, the same institution that sponsors the HDR (see www.wider.unu.edu/wiid/wiid.htm) and the percentage changes are computed exactly like the ones reported in the Tables on page 39. The sample periods are chosen so that they are comparable with the ones reported by the HDR.

⁴ It now has 30 members since the Slovak Republic joined in 2000.

inequality.⁵ The conclusion of the HDR text that “only ONE OECD country shows a slight improvement” is, therefore, rather peculiar.

Despite these sample selection problems, it is probably still true that within-country inequality has risen. The increase has certainly not been uniform across countries, but the increases probably have more than offset the declines. In order to estimate the overall changes, however, one would have to average out the individual country measures, maybe giving more weight to larger countries. Section 2 of this paper shows that the proper measure is a weighted average of individual indexes, where the weights are either the fraction of the world’s population that live in each country or the fraction of world’s income enjoyed by the country.

Step 2. Show that across-country inequality is also rising.

The second step is to show that inequalities across countries are also increasing. Interestingly, to show that across-country inequalities have increased, the HDR does not use the same statistical tools used in its within-country analysis (that is, the estimation of the Gini coefficients). Instead, it reports the ratio of the per capita income of the richest countries to the poorest countries in the world: “*in 1960, the 20% of the world’s people in the richest countries had 30 times the income of the poorest 20%. The ratios increased to 60 to 1 in 1990 and 74 to 1 in 1997*”.⁶ Curiously, these ratios are computed using current-exchange rates and without adjusting for purchasing power parity. Since the HDR is concerned with human welfare (at least the part of human welfare that relates to income), one would imagine that they should worry about the amount of goods and services that each person’s income is able to purchase. Because citizens tend to buy goods where they live, the only reasonable way to compare incomes across

⁵ If you add up the number of countries we mentioned to the 19 reported by the HDR, you will see that you get 30. How can that be, given that OECD had only 29 members? The answer is that the HDR mysteriously includes Israel (which, of course, shows an increase in inequality) even though it is not a member of OECD.

The list of members and the year each joined the OECD can be found in www.oecd.org/oecd/pages/document/displaywithoutnav/0,3376,EN-document-notheme-1-no-no-9464-0,00.html. Israel is not on the list.

⁶ HDR, pages 3 and 36.

people who live in different countries is to adjust for purchasing power differences. Whether a person in Ethiopia can afford a hotel in Manhattan is not as relevant.⁷ We recompute the ratio by dividing the average per capita income of the countries in which the richest 20% of the people live by the bottom 20% using Summers and Heston PPP-adjusted data (Summers and Heston (1998, 1991).) The picture looks quite different: rather than rising from 20 to 60 to 74, the ratio increases from 11.3 in 1960 to 15.9 in 1980, but then declines slowly to 15.09 in 1998. The explanation for these very different results is that poor countries tend to be cheaper. Thus, when one compares incomes across countries using nominal exchange rates, one tends to find that the purchasing power of the people who live in poor countries is smaller than it actually is. Once one corrects for differences in purchasing power, the disparities become a lot smaller.⁸ Figure 1 displays the difference visually: by using nominal exchange rate data the HDR gets an enormous

⁷ Not everyone agrees with this statement. Robert Wade (2001) makes a very creative argument for not adjusting for PPP: “Indeed, for most of the issues that concern the world at large—such as... the economic and geopolitical impact of a country (or region) on the rest of the world—we should use market exchange rates to convert incomes in different countries into a common numeraire. After all, the reason why many poor countries are hardly represented in international negotiations whose outcomes profoundly affect them is that the cost of hotels, offices, and salaries in places like New York, Washington, and Geneva must be paid in U.S. dollars, not in purchasing power parity-adjusted dollars.”

Even if it is true that poor countries do not have international political influence because they cannot afford hotels, offices and salaries in New York, if one wants to make claims about the welfare people get from their incomes, these must be adjusted for purchasing power because people get welfare from consumption, and consumption is purchased at the prices paid in their country of residence. Thus, reasonable numeraires must be adjusted for PPP.

⁸ It is well known that international currency markets with mobility of capital and goods tend to undervalue the domestic purchasing power of the currencies of poor countries (see Balassa (1964), Samuelson (1964) and Bhagwati (1984).) The intuition is that real wages are low in countries with low productivity, so that non-traded (labor-intensive) goods tend to be relatively cheap in poor countries. Since current exchange rates will tend to underestimate the incomes of poor countries, inequalities will tend to be over estimated. Dowrick and Akmal (2001) point out that adjusting for PPP may introduce another bias arising from the fact that households in poor countries substitute their consumption towards goods that are locally cheap, even though they may look relatively expensive when evaluated at international prices. They attempt to correct this using Afriat true indexes.

increase in disparities, an increase that disappears if one adjusts, as one should, for PPP differences across countries.

If, rather than the ratio of the top to the bottom 20%, the HDR had used another popular estimate of income disparities, namely the variance of the (log) GDP per capita, it would have actually concluded that cross-country inequalities have risen over the last three decades.⁹ We have done the calculations for a sample of 125 countries in the Heston, Summers and Aten (2001)¹⁰ data set for the period 1970 to 1998.¹¹ The results, reported in Figure 2, are well known: the variance increased markedly during this period. As Lant Pritchett (1997) famously put it: *there has been “Divergence Big Time”*.

In sum, if personal income disparities are computed using PPP-adjusted data, then the ratio of the top 20% to the bottom 20% not only did not increase as much as the HDR claims, but it even declined during the 1990s. However, if we use the cross-country variance of the log per capita GDP, then it is true that countries' GDP per capita over the last 30 years has diverged.

Step 3. Conclude: therefore, global income inequality is rising

If one observes that income inequality has risen within countries and has also risen across countries, the natural inference must be that overall inequality has also increased. Although the conclusion seems to follow from the premises, it actually does not. The reason is that the within and the across-country inequalities discussed above are not strictly comparable: the within-country measure refers to “individuals” whereas the across-country measure refers to “countries”. To see why this confusion may indeed lead the researcher to reach the wrong answer, consider the following example. Imagine that differences in per capita GDP across countries are large.

⁹ The recent economic growth literature has dealt very often with the question of cross-country “income convergence” (see Barro and Sala-i-Martin (1998), Chapters 10, 11 and 12 for a survey of results). This literature uses the evidence on convergence (or the lack thereof) as a “test” of economic growth theories and does not make claims about personal income inequality.

¹⁰ The list of countries in the sample are reported in Appendix Table 1.

¹¹ This is what the convergence literature, following Barro and Sala-i-Martin (1992) calls σ -convergence analysis.

Suppose that five-sixths of the world population live in very poor economies which do not grow and one-sixth lives in rich economies that grow. The variance across countries, therefore, increases over time as the income of the rich countries grows whereas the income of the poor remains stagnant. In other words, we observe “divergence, big time”. Compared to this large and growing variance, imagine that the distributions across individuals within each country are extremely compact. Imagine now that one very poor but populated economy (with, say, one-sixth of the world’s population) enjoys large aggregate growth rates so that the incomes of its people tends to converge to those of the rich countries. Simultaneously, the distribution within country increases a little bit so that within-country inequality rises. Since the poor country that grows very fast is only one data point, the cross-country variance will still be growing: poor and rich economies still “diverge big time”. However, the incomes of one billion people (one-sixth of the world) are getting closer to the incomes of the rich. Under these circumstances, it may very well be the case that the reduction of the gap between these billion people and the rich, more than offsets the assumed increase in within-country inequality. Hence, we may observe increasing inequalities in the per capita incomes across countries and increasing within-country inequalities, but overall reductions across world individual income disparities! Thus, the three-step procedure followed by the United Nations and described in this section says nothing about the true evolution of world income inequality. Although this may seem like a farfetched example, we will show that this is exactly what happened in the world during the last two decades (where the large, poor and growing economy is China).

The example suggests that, rather than comparing the evolution of the variance of (log) per capita GDP to the evolution of within-country inequality, perhaps one should analyze the population-weighted variance of (log) per capita GDP.¹² The solid curve in Figure 3 is the population-weighted variance of the (log) per capita GDP for the same set of 125 countries for which Figure 2 was constructed. In fact, in order to establish comparability, Figure 3 also reports the unweighted measure that was displayed in Figure 2.¹³ The result is striking: rather than a

¹² In Section 2 we will make this point more precisely.

¹³ A similar figure, for the period 1970 to 1992 appears in Schultz (1998).

steady increase in inequality, the population-weighted variance has fallen during the last two decades. When considering individuals rather than countries, the evidence does not show “divergence, big time” but “convergence, period”.

Of course the population-weighted variance is not a good measure of personal inequality because it implicitly assumes that all individuals in a country have the same level of income (the per capita income) and, therefore, it ignores any change in the disparities that occur within the country. One of the main goals of this paper is to estimate the evolution of global inequality across individuals of the world by combining across-country inequality measures such as the population-weighted variance with the within-country measures.

Many papers in the literature have tried to estimate global inequality measures. Theil (1979, 1996), Berry, Bourguignon and Morrisson (1983), Grosh and Nafziger (1986), Theil and Seale (1994), Ravallion and Chen (1997), Schultz (1998), Firebaugh (1999), Milanovic (2000), Melchior, Telle and Wiig (2000), and Dowrick and Akmal (2001), Bourguignon and Morrisson (2002), and Sala-i-Martin (2002) are just some examples. Some of these previous studies are purely based on survey data (Ravallion and Chen (1997) and Milanovic (2000)). These studies use an income definition that is different from that of the national income accounts and ignore benefits from public spending (which are especially important when it comes to health). The advantage of survey data, on the other hand, is that it takes into account home-consumption, which may also be important in poor countries.

Theil (1994, 1996), Berry, Bourguignon and Morrisson (1983), Theil and Seale (1994), Firebaugh (1999) and Melchior, Telle and Wiig (2000) introduce population-weights to the traditional measures of PPP-adjusted GDP per capita inequality, but they explicitly ignore the evolution of intra-country inequality. Melchior, Telle and Wiig (2000) confirm Schultz’s (1998) original finding: the variance of population-weighted (log) GDP per capita has decreased, not increased, after 1978. The main reason is, of course, that China has been growing and converging to the rich at rapid rates after 1978. This means that the measures of convergence based on “each country, one data point” can show divergence, but when we give “each citizen, one data point”, the picture changes radically. The key factor is that the average Chinese person (and therefore, about 1/4 of the world population) has experienced substantial convergence of their personal

income. The main problem with these papers is that they ignore intra-country inequality. It can be persuasively argued that the growth rate of the Chinese economy has not benefitted all citizens equally so that within-China inequality has accompanied the spectacular growth process.¹⁴ The question is whether the increase in within-China inequalities more than offset the fact that one billion Chinese incomes are converging towards the levels enjoyed by OECD economies.

Schultz (1998), Dowrick and Akmal (2001) and Bourguignon and Morrisson (2002), estimate the evolution of within-country inequality. The first two studies find that, if current exchange rate measures of GDP are used, then inequality rises. But if PPP-adjusted measures are used instead, then inequality after 1978 either falls or shows little trend. They both conclude that most of the global income inequality is explained by cross-country rather than within-country inequality. Schultz uses the Deininger and Squire (1996) data set to regress the available Ginis for particular countries and periods on observed macro-magnitudes and “forecasts” the Ginis of the missing countries and periods. The problem with this approach is that the Ginis are not well explained with aggregate data so the resulting “forecasts” are very inaccurate. Dowrick and Akmal restrict their analysis to the countries that have estimates of the Gini coefficient for years “close” to 1980 and 1993 and then estimate the evolution of “world” income inequality between these two periods. The problem with this approach is that the selection of countries that do not have Gini data is not random. In particular, these are countries that are poor and that have diverged. Excluding these countries from the analysis tends to bias the results towards finding reductions in world income inequality.

Bourguignon and Morrisson (2002) estimate global income inequality between 1820 and 1992. They combine aggregate PPP-adjusted data and within country surveys for selected years during this very long period. Since most countries do not have survey data for most of the years to allow them to compute within country inequality, they “assign” the within-country inequalities of countries that are thought to be “similar”. They also group countries into “blocks” based on

¹⁴ Berry, Bourguignon and Morrisson (1983) allow for intra-country inequality in 1970, but they do not estimate how these inequalities change over time. On page 335 of their paper they say “*we disregard changes in intra-country inequality which may have occurred over 1950-1977. For developed countries reasonably good data indicate that those changes have been minor*”. We will show that this omission may be quantitatively important.

economic, geographic and historical similarities . Although this might be a good approximation when it comes to computing overall inequality, we should be very careful when we analyze the decomposition of global inequality into within and between inequality. The reason is that what they call “within inequality” refers to inequality “within a group” rather than “within a country”. For example, Argentina and Chile are in the same Bourguignon-Morrisson group. Suppose that their GDP per capita tends to converge and that inequality within these two countries remains constant. The convergence process will lead to a reduction in world income inequality. Bourguignon and Morrisson will assign this reduction in inequality to “within” groups although the process of convergence is clearly “across” countries. One additional problem with the Bourguignon and Morrisson study is that, since they do not have a continuous time series of income shares for each country, they cannot really compute within-country inequality for each of the years they study. To solve this problem, they assign the numbers of the closest year. This, of course, introduces potential measurement errors in the calculations of within-country inequality.

Sala-i-Martin (2002) uses income shares for each country to estimate a gaussian kernel density function for each country and each year between 1970 and 1998. He then aggregates these individual estimates across countries to estimate the worldwide income distribution function. This function is then used to estimate poverty rates and various measures of income inequality. Although, in principle, allowing for disparities across incomes within each quintile should give a better picture of the evolution of worldwide income inequalities, in Section 5 we show that the quintile approach used in this paper gives exactly the right lessons about the trends of worldwide income inequality. Since the estimation procedure based on the assumption that all persons within a quintile have the same level of income is easier and more transparent, we continue using it in this paper.

The rest of the paper is organized as follows. Section 2 discusses the methodology and the data used to infer individual incomes across the world. Section 3 estimates the world distribution of income by constructing the gaussian kernel density functions for various years, and estimates poverty rates by integrating the density functions below the poverty line. Section 4 discusses various indexes of income inequality and the theoretical decomposition of global inequality into within-country and across-country inequality. It also presents the main results on

global income inequality and its decomposition. Section 5 analyzes the assumption of equal income for all individuals within a quintile and concludes that it does not bias the results on income inequality. Section 6 comments on the quality of GDP estimates for China and analyzes the effects of eliminating this large country from the sample as well as alternative assumptions about its per capita growth rate. Section 7 forecasts the evolution of inequality over the next 50 years under the assumption that African countries continue to fail to grow. Finally, Section 8 concludes.

2. Methodology and Data

(A) Measuring Individual Income from Aggregate GDP and Income Shares

Our goal is to estimate the global distribution of individual income and to measure poverty and income inequality across the citizens of the world. Obviously, we do not have each person's income so we need to approximate it using available aggregate data. Following Theil (1994, 1996), Theil and Seale (1994), Firebaugh (1999) and Melchior et al., we could assume that all individuals in a country have the same level of income and proceed by using income per capita measures.¹⁵ This procedure, of course, ignores inequalities among individuals within a country.

Instead, we use available survey data to estimate the income shares, which we denote with the symbol s_{ikt} , for individual quintile (indexed by k), for each country (indexed by i), and each year (indexed by t).

With these income shares, we proxy the within-country distribution by dividing each country's population in five groups and assigning them a different level of income. Let N_{it} be the population in country i at time t , and let y_{it} be the income per capita for country i at time t .

¹⁵ See also Wade (2001), and some sections of Schutlz (1998) and Dowrick and Akmal (2001).