

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).

We assign to each fifth of the population, $\frac{N_{it}}{5}$, the income level $5 s_{ikt} y_{it}$. Within each

quintile, therefore, each individual is assumed to have the same level of income.¹⁶ Once we have this estimate of the income of each individual in the planet, we can estimate the world distribution of income.

(B) Estimating Income Shares

The procedure we just described combines within-country income share numbers with aggregate income numbers. Unfortunately, income shares are not available for all years and all countries. Berry et al. (1983) solve this problem by assuming that the income shares do not change over time so the same shares to all the years of their sample. They also ignore countries that do not have any estimates of income shares. Dowrick and Akmal (2001), on the other hand, analyze the income distribution in just two years (1980 and 1993) and they assign the income shares of the closest year to both 1980 and 1993. They also ignore countries that do not have data on income shares.

We use all the countries with available PPP-adjusted GDP data as computed by Heston, Summers, and Aten (2001). We restrict our analysis to the post 1970 period so that we can extend the number of countries to 125. Together, these countries have close to 90% of the world's population. The last year with available data is 1998 (although the data for 1997 and 1998 are preliminary and should be taken with caution).¹⁷

¹⁶ Obviously, the assumption that all individuals within a quintile have the same level of income will lead to underestimation of the true global inequality. It is unclear, however, how this assumption biases the “trend” of inequality over time. In Section 5 we report results from Sala-i-Martin (2002) where individual income for each country is computed using a gaussian kernel density function. We show that the assumption of equal income within quintiles does, indeed, lead to a slight underestimation of the level of global income inequality, but the main results about the trend of inequality over time are identical.

¹⁷ Earlier drafts of this paper used the original Penn World Tables Mark 6 which go up to 1996 only. For the common years, the main results of the paper do not change at all.

The within-country income shares are from the Deininger and Squire (DS) data set, extended with the World Development Indicators (WDI) of the World Bank. These data are based on national-level income and expenditure surveys.¹⁸

Using these data we have three broad groups of countries (listed in Appendix Table 1):

Group A.- Countries for which we have a time series of income shares by quintiles (by time series we mean that we have a number of observations over time, although we may not have observations for every year between 1970 and 1998).

Group B.- Countries for which we have only one observation between 1970 and 1998.

Group C.- Countries for which we have NO observations of income shares.

We have 68 countries in group A. Together, they have 4.7 billion inhabitants, which account for 88% of our sample population. For them, we plot the income shares over time and we observe that they tend to follow very smooth trends (see the Appendix Figures). In other words, although the income shares estimated by Deininger and Squire and the World Bank are not constant, they do not seem to experience large movements. If anything, they seem to have small time trends.¹⁹ Using this information, we regress income shares on time to get a linear trend for each economy. This was done using two methods – first, with income shares estimated for all five quintiles, and second, with income shares estimated for the top two and the bottom two

¹⁸ These survey data have been criticized by Atkinson and Brandolini (2001) on various grounds (in fact, Atkinson and Brandolini warn all researchers that use any “secondary data-sets” that, many times, these data are incomplete, not comparable across countries or over time and hold a whole array of additional problems.) Future research should make use of the best data and comparable data only. The evolution of the big picture for the world is unlikely to change much, however. The main reason is that, as we will see, most of the *movement* in global inequality comes from cross-country disparities rather than of within-country disparities. In fact, we will also see that the main determinant of the *level* global inequality is the across-country component. It is possible, of course, that within-country inequality is vastly mismeasured and that, when we eventually do measure it correctly, differences in incomes within countries end up being larger than across countries. With the data we have available, however, there is nothing that indicates that this is likely to be true. The low quality income shares, therefore, will certainly introduce errors in our measures of world inequality. But they are not likely to change the main conclusions and trends.

¹⁹ Obviously, these trends can only be temporary since income shares are bounded between 0 and 1.

quintiles with income share of the middle quintile calculated as a residual. Both methods gave identical results.²⁰ The income shares , s_{ikt} , used to compute global income inequality are the projections of these regression lines.²¹

There are 29 nations in group B, with 316 million people (or 6% of total population). The income shares were assumed to be constant for the period 1970-98. Hence, for group B, we measure the level of within-country inequality, but we do not let it change over time. That is, we assume $s_{ikt} = s_{ik}$ for all t. Berry et al. (1983) make the assumption that ALL countries'

inequality is constant over time. To the extent that income inequality within these countries has changed, our assumption introduces a measurement error in global inequality. However, given that we do not know the direction in which this within-inequality has changed, the direction of the error is unknown. An alternative would have been to restrict our analysis to the countries that have time-series data (countries in Group A). Excluding them from the analysis, as is done by other researchers (see for example Dowrick and Akmal (2001)), may introduce substantial bias which turn out to change some of the important results. The biases arise from two sources. First, and as we just mentioned, the within-country inequality may have changed. Second, and more

²⁰ It can be persuasively argued that India experienced a large increase in inequality after the liberalization policies enacted after 1991. Imposing the time trend estimated before liberalization to the post-liberalization shares may seriously underestimate within-country inequality in India. This, of course, may be true for other countries also. But since our global measure of within-country inequality is weighted by the size of the country, this problem might be particularly important for India. We repeat our inequality calculations under the assumption that the Indian income shares follow a different trend before and after 1991. The results remain largely unchanged. The reason is that, as we will see, most within-country measures of inequality remain fairly constant over time whereas the across-country measures decline substantially. Thus, even if the change in inequality in India were to be underestimated, the correction would have to be enormous for the within-country index to grow by so much as to overturn the large decline in overall inequality.

²¹ We also used the actual income shares for the years when they are available and the projected shares for the missing years, but the aggregate indexes that are reported in the next section do not change substantially.

importantly, the countries that are excluded tend to be countries that are poor and have “diverged”. Their exclusion from our analysis, therefore, is subject to an important sample selection bias, which tends to bias the results towards finding an excessive reduction in income inequality. Thus, out of the 125 countries in the Summers-Heston data set, income inequality based on quintile income shares could be calculated for 97 countries, which cover 95 % of the sample population.

The nations of group C (for which there are no within-country inequality data) are treated as if all citizens in the country had the income per capita of the country. In other words, for Group C, we assume $s_{ikt} = 1/5$. Again, we could exclude these from the analysis, but we prefer not to do so because, as we already stated, their exclusion may lead to important biases in the results. There are 28 nations in this category, which brings the total in our analysis to 125, with a combined 1998 population of 5.23 billion (or 88% of the world’s 5.9 billion inhabitants in 1998).²²

3. The World Income Distribution Function

(A) The Kernel Density Function

We can now use the individual income numbers constructed in the previous section to estimate the world distribution of income. We estimate a gaussian kernel density function. Using the convention in the literature, we use a bandwidth $w=0.9*sd*(n^{-1/5})=0.34$, where sd is the standard deviation of (log) income ($sd=1.3$ in our data set) and n is the number of observations

²² The largest countries excluded from our sample are those from the former Soviet Union. There is little we can do to incorporate those countries because they did not exist until the early 1990s. It is unclear how these countries would affect our total inequality measures. On the one hand, it seems clear that disparities within these countries have increased. On the other hand, these countries were “rich” relative to most of the citizens of the world, and have experienced negative aggregate growth rates. Thus, the individual incomes for these countries has “converged” towards those of the 1.2 billion Chinese, 1 billion Indians and 700 million Africans. This second effect has probably reduced across-country disparities and, as a result, global inequality. The overall effect of excluding the former Soviet Union, therefore, is unclear.

(n=513).²³ Previous researchers have used kernel densities to estimate income distributions. For example, Quah (1996, 1997), Jones (1997), and Kremer, Onatski and Stock (2001) estimate the world distribution by assuming that each country is one data point (and the concept of income is the per capita GDP). Instead, we use the individual income estimated in the previous section so our unit of analysis is not a country but a person. We depart from the conventional literature in that we do not normalize our data. Quah (1996, 1997) divide each country's GDP per capita by the world average GDP per capita. Jones (1997) divides by the GDP per capita of the richest country. Kremer et al. experiment with various normalisations. The reason for normalizing the data is to prevent the distribution from “shifting to the right” as the world becomes richer. This allows for an easy comparison of the various distributions over time. We do not normalize for two reasons. First, because we think it is interesting to visualize if the world income distribution shifts to the right and how fast. Second, and most importantly, we want to compute poverty rates by integrating the kernel density function below a poverty threshold, as we discuss in the next subsection.

Figure 4 reports the estimates of the density functions for 1970, 1980, 1990 and 1998. The first thing we notice is that the distribution shifts to the right: individual incomes across the world increase over time. The second point worth emphasizing is that the distribution in 1970 had a large mode close to \$900 and a second smaller mode at around \$9,000. Over time, the “tall” mode of the poor gets smaller. In other words, the fraction of the world population with very low levels of income gets smaller over time. Both the rich and the poor modes shift to the right. By 1998, the poor mode is at \$1,900 and the rich mode is at \$16,200. The interesting thing is that a substantial fraction of the poor have now shifted into a “middle class”: by 1998, a third mode has emerged at around \$4,200.

(B) Poverty Rates

Absolute poverty rates can be inferred from our estimated kernel density functions. Poverty rates are defined as the fraction of the world's population that live below the absolute

²³ We also calculated the Epanechnikov kernel with the same bandwidth and the results were virtually identical so we do not report them here.

poverty line. We use the conventional definitions of absolute poverty: less than one dollar per day. The original definition is due to Ravallion et. al (1991). These researchers used “perceptions of poverty” in the poorest countries to place the poverty line at \$31 per month. Later , the definition was changed to \$30.42, and it then was rounded off to \$1 per day. The \$1/day line was later adopted by the World Bank as the “official” definition of “absolute poverty”. For some reason, another poverty line mysteriously appeared in the literature that doubled the original figure to two dollars per day. The United Nations some times uses four dollars per day. Of course, if one is allowed to raise the poverty line arbitrarily, then one is bound to find that all persons in the world are poor. In this paper we will stick with the original definition of poverty and we will also analyze the \$2/day line.^{24,25}

It should be noted that all these definitions are expressed in 1985 values. Since the Summers and Heston data are reported in 1996 dollars, the annual incomes that define \$1/day and \$2/day poverty in our data set are \$532 and \$1064 respectively.²⁶

²⁴ Ravallion et al. (1991) one-dollar line refers to individual consumption. We will compute income (not consumption) poverty. This means that our level poverty rates may differ from those computed by Ravallion et al. (1991) or Chen and Ravallion (1997, 1999). Since, for poor economies, more than 50% of GDP is consumed (savings and investment are quite small and so is government spending), this means that the consumption poverty line should be located somewhere between our one-dollar and two-dollar lines. Since absolute poverty is an arbitrary line, the important question is how many people “jump over” it during a particular period. As we will see later in this section, our answer is: many.

²⁵ Some analysts (including the UNDP’s HDR) use a different definition of poverty for rich nations than they use for poor countries (the reason appears to be that, if the 1\$/day line is used in these developed nations, almost nobody is found to be poor). The definition for developing countries is one half of the median income. This (arbitrary) definition has obviously more to do with inequality than with poverty because, if all the incomes were to double (or increase by a factor of one hundred), the poverty rate would be the same, a property of poverty that is clearly undesirable.

²⁶ Notice that by assigning the same level of income to each person within a quintile, we label all the citizens of the quintile as poor if the mean income of the quintile is below the poverty line. Similarly, we label all of them as non-poor if the average income of the quintile is above the line. For the countries in group C, we assign the country’s per capita income to all individuals. If per capita income is below the poverty line, then the whole country is labeled as poor. There is one relatively large country that fell below the \$1/day poverty line in the 1990s:

The \$1/day poverty rate can be computed by dividing the integral of the density function below the value \$532 by the integral of the entire density. The \$2/day poverty rate can be computed by integrating below \$1064. The poverty rate is represented by the fraction of the overall area under the density function that lies to the left of the poverty line.

Since the underlying income data used to compute the distribution is in constant 1996 dollars, we can compare all the estimated density functions to the same the poverty lines. Thus, we plot two vertical lines representing the \$1/day and \$2/day thresholds in Figure 4. The evolution of poverty over time can be analyzed by checking the size of the area below each year's density and to the left of the poverty line. In 1970, the area below the two-dollar line seems to be about half the overall area. This suggests that the poverty rate was close to 50%. The area below the one-dollar line is smaller, but still appears to be quite large. It is interesting to note that the mode of the 1970 distribution was slightly above the one-dollar line and below the two-dollar line. The fraction of population that lived below or around the poverty lines in 1970 was substantial. If we look now at the 1998 density function, we see that the areas to the left of the two poverty lines have shrank dramatically. Thus, poverty rates declined substantially between 1970 and 1998.

The exact poverty rates estimates are displayed graphically in Figure 5. For the \$1/day definition, we observe that, after increasing during the early 1970s and peaking at about 20 percent in 1974, poverty rates have been declining almost monotonically throughout the period. The smallest poverty rate (5.4 percent) occurred in the last year of our sample, 1998. In other words, the poverty rate in 1998 was cut to almost one-fourth from its peak value in 1974 and by about a third since 1970. The \$2/day poverty rate declines monotonically throughout the period. It falls from 44.5% of the people in 1970 to 18.7% in 1998, a reduction of close to 60%.

(C) Poverty Headcounts

Congo (former Zaire). Obviously, the fact that we do not look at within-quintile variation for countries in groups A and B or within-country variations for countries in group C means that our estimates of poverty ratios are subject to errors.

Some have argued that the poverty rates are irrelevant and that the really important information is the number of people in the world that live in poverty (some times this is called “poverty headcount”). The distinction is important because, even though the poverty rates have declined, since world population has increased, it could very well be the case that the total number of poor citizens has been rising. A veil of ignorance argument, however, suggests that the world improves if poverty rates decline. We could ask ourselves whether, with the veil of ignorance, we would prefer to be born in a country of a million people with half a million poor (poverty rate of 50%) or in a country of two million people and 600,000 poor (a poverty rate of 30%). Since our chance of being poor is much smaller in the country with a larger headcount, we would prefer the country with a smaller poverty rate.²⁷ Thus, we should say that the world is improving if the poverty rates, not the headcounts, decrease. Of course the best of the worlds would be one in which both the poverty rates and poverty headcounts decline over time. But this would be too much to ask in this era of neoliberal globalization. Wouldn’t it?

We can estimate poverty headcounts by simply multiplying our poverty rates by the overall population each year. The results are displayed in Figure 8. We see that, using the one-dollar-a-day definition, the overall number of poor declined by over 400 million people: from close to 700 million citizens in the peak year of 1974, to less than 300 million in 1998. Using the two-dollar definition the number of poor declined by about 500 million: from 1.48 billion in 1976 to 980 million in 1998. The last year of the sample, 1998, is the year with the smallest number of poor according to both definitions.

In sum, world poverty has declined substantially over the last thirty years. This is true if we use the one-dollar-a-day or the two-dollar-a-day definition and whether we use poverty ratios or poverty counts.

Having said this, the fact that the world is improving does not mean that all is good. Actually, one disturbing fact (which cannot be seen from the picture) is that more than 95% of the “one-dollar poor” live in Africa. What’s even more disturbing is that the whole Democratic

²⁷ Those people who might prefer the country with smaller headcount should ask themselves if they would also prefer a country of half a million people with 499,999 poor.

Republic of Congo belongs in this category,²⁸ four quintiles (80% of the people) of Tanzania and three quintiles of the Central African Republic, Ethiopia, Guinea-Bissau and Sierra Leone live on less than one dollar a day. The massive concentration of poverty in the African continent suggests that the lack of economic growth of Africa is the most serious economic problem we face today.

4. Measuring Global Income Inequality

Poverty rates are one interesting aspect of the world distribution of income as they represent the fraction of this distribution that lies below a commonly accepted poverty line. There are other aspects of the distribution that might be worth analyzing. A particularly important one is its dispersion. The reason why the dispersion is the focus of attention of many researchers is that it reflects the inequality across individual incomes. The vast theoretical and empirical literatures on inequalities has produced a substantial number of measures.²⁹ Moreover, different indexes tend to give different answers. Hence, instead of trying to find the single best index, we will next review some of the most widely accepted and we will report estimates for seven of them.

(A) Income Inequality Indexes

(i) Ad Hoc Indexes

Some indexes of income inequality discussed in the literature are intuitive and attractive at first sight, but ad hoc and unattractive upon analysis. Of these ad hoc measures, the most popular are the **Gini coefficient**, the **variance of incomes** or the **variance of the logarithm of income**. The simple variance is universally known, but it is unattractive because if we double every person's income, this measure of inequality quadruples (in principle, it would seem reasonable that our measures of inequality do not change when all incomes in the world double).

²⁸ The Democratic Republic of Congo (with close to 50 million inhabitants) being below the one-dollar poverty line explains the small blip that can be observed at the leftmost end of the density function in 1998.

²⁹ See Cowell (1995) for an excellent survey of measures and their potential drawbacks.

This is why researchers prefer to estimate the variance of log-incomes.³⁰ If we have a set of m countries, the variance of the log of individual income is given by

$$\text{var}(\ln y_{ijt}) = \frac{\sum_{i=1}^m \sum_{j=1}^{N_{it}} (\ln y_{ijt} - \mu_t)^2}{N_t}, \quad (1)$$

where y_{ijt} is the level of income for a person j in country i at time t , and N_t is the world population (equal to the sum of country populations, N_{it}).³¹ Using the income share notation from Section 2, the world variance of individual log-income can be written as:

$$\text{var}(\ln y_{ijt}) = \sum_{i=1}^m \frac{N_{it}}{5 N_t} \sum_{k=1}^5 \left[\ln(5 s_{ikt} y_{it}) \right]^2 - \left[\sum_{i=1}^m \frac{N_{it}}{5 N_t} \sum_{k=1}^5 \ln(5 s_{ikt} y_{it}) \right]^2 \quad (2)$$

(ii) Social Welfare Function Indexes

Some researchers have proposed a derivation of indexes of income inequality from Social Welfare Functions (SWF). For example, Atkinson (1970) shows that a good measure of inequality would relate the “equally distributed equivalent level of income” to the actual mean

³⁰ The variance of log-income corresponds to the variance of the distributions reported in Figure 4.

³¹ See Cowell (1995) for the definitions of the other indexes.

income.³² The Atkinson measure, therefore, depends on society's degree of "inequality aversion". The Atkinson index is given by

$$A_{\epsilon} = 1 - \left[\frac{1}{N_t} \sum_{i=1}^m \sum_{j=1}^{N_{it}} \left(\frac{y_{ijt}}{\bar{y}_t} \right)^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}} = 1 - \left[\frac{1}{N_t} \sum_{i=1}^m \frac{N_{it}}{5} \sum_{k=1}^5 \left(\frac{5 s_{ikt} y_{it}}{\bar{y}_t} \right)^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}}, \quad (3)$$

where, $\bar{y}_t = \frac{1}{N_t} \sum_{i=1}^m \sum_{j=1}^{N_{it}} y_{ijt} = \sum_{i=1}^m \frac{N_{it}}{N_t} y_{it}$ is the world average per capita income and ϵ is

the coefficient of relative inequality aversion.³³

(iii) Axiomatic Indexes

A third set of measures is derived from some axiomatic principles which are thought to be desirable for a good index of inequality. Cowell (1995) highlights four principles:

The *Pigou-Dalton Transfer Principle* (see Dalton (1920) and Pigou (1952)).

This principle states that a good inequality measure should rise in response to a mean-preserving redistribution from a poor to a rich person. Most measures used in the literature, including Gini, Generalized, Entropy Class and the Atkinson class satisfy this principle. The main exception is the variance of logarithms.

The *Income Scale Independence Principle*.

This principle states that a desirable measure should be homogeneous of degree zero. That is, if we scale all the incomes in the world by the same number, our measure of inequality should not

³² The "equally distributed equivalent level of income" is the level of income that each person should receive if income were to be equally distributed and we were to achieve the same level of social welfare as we get with the actual unequal distribution.

0. The SWF that delivers this index has a constant relative inequality aversion:

$$W(.) = \sum_{i=1}^m \sum_{j=1}^{N_{it}} \left(\frac{y_{ijt}^{1-\epsilon} - 1}{1-\epsilon} \right).$$

change. Again, most measures satisfy this principle. The main exception is the variance of income (if we double all incomes, the variance quadruples).

The Principle of Population.

This postulates that the distribution of the cake should not depend on the number of the cake receivers. That is, if we measure inequality in an economy with N people and then merge it with another identical economy, inequality in the larger economy should be the same (Dalton (1920).

The Principle of Decomposability. We would like to be able to write down total inequality as the sum of inequality *within* groups and inequality *across* groups. Moreover, the inequality within groups should be expressed as something like the average of the inequality of each individual group. This is an important property for the purposes of this paper because we hope to find out how the evolution of global or world inequality depends on the convergence across countries (or across-country inequality) and on the evolution of the within-country differences (or within-country inequality.)

Bourguignon (1979), Shorrocks (1980) and Cowell (1995) show that the only inequality indexes that simultaneously satisfy all the above principles are the Generalized Entropy Indexes,

$GEI(\theta)_t$:

$$\begin{aligned}
 GEI(\theta)_t &= \frac{1}{\theta^2 - \theta} \left[\frac{1}{N_t} \sum_{i=1}^m \sum_{j=1}^{N_{it}} \left(\frac{y_{ijt}}{\bar{y}_t} \right)^\theta - 1 \right] = \\
 &= \frac{1}{\theta^2 - \theta} \left[\frac{1}{N_t} \sum_{i=1}^m \sum_{k=1}^5 \left(\frac{s_{ikt} y_{it}}{\bar{y}_t} \right)^\theta - 1 \right]
 \end{aligned} \tag{4}$$

where θ is a constant parameter. One particular case of $GEI(\theta)_t$ is the limit as θ tends to zero. Using l'Hôpital's rule, we get:

$$\begin{aligned}
GEI(0)_t &\equiv MLD_t = \frac{1}{N_t} \sum_{i=1}^m \sum_{j=1}^{N_{it}} \ln \left(\frac{\bar{y}_t}{y_{ijt}} \right) = \\
&= \ln(\bar{y}_t) - \frac{1}{N_t} \sum_{i=1}^m \frac{N_{it}}{5} \sum_{k=1}^5 \ln(5 s_{ikt} y_{it}),
\end{aligned} \tag{5}$$

which is the Mean Logarithmic Deviation (MLD). Notice that the MLD is the difference between the log of average income and the average of the logarithms of incomes.

A second particular case of GEI can be found by taking the limit as $\theta \Rightarrow 1$. Again, using l'Hôpital's rule, we get:

$$\begin{aligned}
GEI(1)_t &\equiv T_t = \frac{1}{N_t} \sum_{i=1}^m \sum_{j=1}^{N_{it}} \frac{y_{ijt}}{\bar{y}_t} \ln \left(\frac{y_{ijt}}{\bar{y}_t} \right) \\
&= \sum_{i=1}^m \sum_{k=1}^5 \frac{N_{it}}{N_t} \frac{s_{ikt} y_{it}}{\bar{y}_t} \ln(5 s_{ikt} y_{it}) - \ln(\bar{y}_t),
\end{aligned} \tag{6}$$

which is Theil's index (see Theil (1967)).

By setting $\theta = 2$, we get a third particular case:

$$\begin{aligned}
GEI(2)_t &\equiv \frac{1}{2} (CV_t)^2 = \frac{1}{2} \frac{1}{\bar{y}_t^2} \left(\frac{1}{N_t} \sum_{i=1}^m \sum_{j=1}^{N_{it}} y_{ijt}^2 - \bar{y}_t^2 \right) = \\
&= \frac{1}{2} \frac{1}{\bar{y}_t^2} \left(\frac{1}{N_t} \sum_{i=1}^m \frac{N_{it}}{5} \sum_{k=1}^5 (5 s_{ikt} y_{it})^2 - \bar{y}_t^2 \right),
\end{aligned} \tag{7}$$

which is equal to one-half the square of the coefficient of variation.

(iv) Across-Country Inequality

It will be convenient to compare our measures of global income inequality with the measures one would obtain if all individuals within each country had the same level of income. That is, if all inequalities in the world came from differences in per capita incomes across countries. These will be called “across-country indexes” of income inequality and correspond to the same indexes we just described, but under the assumption that the income shares are all equal to $s_{ikt} = 1/5$. For example, the across-country variance of log income is given by

$$\text{var}(\text{across}) = \sum_{i=1}^m \frac{N_{it}}{N_t} \left[\ln(y_{it}) \right]^2 - \left[\sum_{i=1}^m \frac{N_{it}}{N_t} \ln(y_{it}) \right]^2. \quad (8)$$

Notice that this index is the population-weighted variance and differs from the variance of (log) income that one would get if one takes each country as a single data point.

An important point about across-country indexes is that they measure world inequality across individuals when we assume that all individuals within a country have the same level of income. This is very different from assuming that each country is one data point. For example, the cross-country variance of log-income is given by:

$$\hat{\sigma}_t^2 = \frac{1}{m} \sum_{i=1}^m \ln(y_{it})^2 - \left[\frac{1}{m} \sum_{i=1}^m \ln(y_{it}) \right]^2. \quad (9)$$

When we compare Equations (2), (8) and (9), we see that the latter two assume that the income shares for all countries are $s_{ikt} = 1/5$ (that is, there is no within-country inequality).

In addition, Eq. (9) assumes that all countries have the same population size $N_{it} = N_t / 125$. In

other words, it assumes that a 1% growth in the per capita income in Lesotho has the same effect on world income inequality as a 1% growth in the per capita income in China, although China has 600 times as many people. As we saw when we compared Figures 2 and 3, ignoring

population weights turns out to be a mistake that delivers a very misleading picture and one is led to conclude (wrongly) that there has been “divergence: big time”.

The across-country versions of Gini, Atkinson indexes, and General Entropy Indexes discussed above are computed also by setting $s_{ikt} = 1/5$ in the corresponding formulae (we do not reproduce all the formulae here because of space constraints).

(v) Decomposing Inequality: Within and Across-Country Measures

The discussion of Section 1 suggests that it might be important to decompose the global index inequality into an “across-country” component (which shows the degree to which the levels of income converge to each other) and a “within-country” component that emphasizes the disparities within countries. For example, it has been argued by Schultz (1998), Dowrick and Akmal (2001), and Melchior et al. (2000) among others, that world income inequality might have been falling since 1978 because China and its 1.2 billion inhabitants have been converging to the levels of income of the OECD. Critics of these papers argue that, along with the growth of the average per capita income, China has experienced a radical increase in inequality among its citizens. The question is whether the increase in inequality within China has more than offset its process of aggregate convergence. To answer these questions, we want to analyze indexes of income inequality that can be decomposed into within-country and across-country inequality. The problem is that the difference between the global index and the across-country index is not always a measure of within-country inequality. That is, not all measures can be “decomposed”.

As previously stated, Bourguignon (1979), Shorrocks (1980), and Cowell (1995) show that the only indexes that can be decomposed are those in the class of the Generalized Entropy, GEI. For example, the MLD index (which is the GEI with coefficient 0) can be decomposed by

simply adding and subtracting the term $\frac{1}{N_t} \sum_{i=1}^m N_{it} \ln (y_{it})$ to Eq. (5) to get

$$\begin{aligned}
MLD_t = & \left(\ln(\bar{y}_t) - \sum_{i=1}^m \frac{N_{it}}{N_t} \ln(y_{it}) \right) + \\
& + \sum_{i=1}^m \frac{N_{it}}{N_t} \left(\ln(y_{it}) - \frac{1}{N_{it}} \sum_{j=1}^{N_{it}} \ln(y_{ijt}) \right).
\end{aligned} \tag{10}$$

The first term is the MLD measure of “across-country inequality” and the second term is a weighted average of the MLDs within each country, where the weights are the size of the population. This weighted average is what we call “within-country inequality”. In practice, with our within-country income decomposition using quintile income shares, the within country

inequality index becomes $\sum_{i=1}^m \frac{N_{it}}{N_t} \left(\ln(y_{it}) - \frac{1}{5} \sum_{k=1}^5 \ln(5 s_{ikt} y_{it}) \right)$.

Similarly, the Theil inequality index can be decomposed by subtracting and adding the term $\sum_{i=1}^m \frac{N_{it} y_{it}}{N_t \bar{y}_t} \ln(y_{it})$ to Eq. (6) to get

$$T = \sum_{i=1}^m \frac{N_{it} y_{ijt}}{N_t \bar{y}_t} \ln\left(\frac{y_{it}}{\bar{y}_t}\right) + \sum_{i=1}^m \frac{N_{it} y_{it}}{N_t \bar{y}_t} T_{it} \tag{11}$$

where $T_{it} = \frac{1}{N_{it}} \sum_{j=1}^{N_{it}} \frac{y_{ijt}}{y_{it}} \ln(y_{ijt}) - \ln(y_{it})$ is the Theil inequality index for country i at

time t . Using our income share notation, the within-country Theil index becomes

$$T_{it} = \sum_{k=1}^5 s_{ikt} \ln(5 s_{ikt} y_{it}) - \ln(y_{it}) . \text{ Notice that, again, the global inequality index}$$

decomposes into a between or across-country index and a weighted average of within-country inequalities. It is interesting to note that the weights are the aggregate incomes rather than the population sizes.³⁴

(B) Seven Estimates of Global Income Inequality

(i) The Gini Coefficient in a Historical Perspective

We can now use the individual income data constructed in Section 2 to estimate global income inequality. To put things in a historical perspective, in Figure 7 we plot the estimates of the global Gini coefficient from Bourguignon and Morrisson (2002). In their ambitious study, Bourguignon and Morrisson estimate global inequality for selected years going back to 1820. Their estimates (which are displayed in Figure 7 with a small hollow squares) show that the global Gini continuously increased between 1820 and the 1980s. Inequality then stabilized at more or less constant level between 1980 and 1992.

Although, once the income shares are computed, our methodology is similar to theirs, our estimates of global income inequality are not necessarily comparable to those of Bourguignon and Morrisson for a number of reasons. First, the sample of countries is different. Second, instead of using individual country data, they put countries in groups according to their presumed similarities. Third, the source of within-country shares that are needed to compute within country inequality are different. Despite these differences, the level of our Gini for 1970 is remarkably similar to theirs. The key difference between our estimates and Bourguignon and Morrisson's is that, for the 1980s and 1990s, we get a substantial decline in the world Gini coefficient whereas their estimates remain constant through 1992, the last year of their study.³⁵ Rather than

³⁴ The calculation of the between-country Theil index takes each country's per capita income as a point in the distribution, but it gives it a weight proportional to the country's population. In other words, notice that the between-country index is the weighted average of each

$$\text{country's } \frac{y_{jit}}{\bar{y}_t} \ln \left(\frac{y_{it}}{\bar{y}_t} \right) .$$

³⁵ When they use the Mean Logarithmic deviation, Bourguignon and Morrisson estimate a reduction in inequality for 1980 and 1992. On the other hand, their estimate of Theil suggests a

experiencing an “explosion” or a “disturbing increase”, we find that world income inequality has declined by more than 5% over the last two decades.

The reduction in world income disparities is largely explained by the decline in differences across countries. To see this, we use the “across-country” Gini coefficient which, as discussed above, is the Gini that arises when all the income shares are assumed to be 1/5. The across-country Gini is also reported in Figure 7.³⁶ As expected, the across-country Gini is smaller than the global Gini. This is reasonable because all within-country inequalities have been eliminated by construction. It is interesting that the pattern is remarkably similar to that of the global Gini: the cross-country Gini increases during the 1970s, and falls during the following two decades. Thus, the decline in global income inequality is accounted for by the substantial convergence in per capita incomes across countries.

The reason for the decline in global income inequality after 1978 is that the most populated country in the world, China, experienced substantial growth rates. Hence, the incomes of a big fraction of the world’s population (approximately 20%) started converging towards the rich economies after 1978. The process was later reinforced by the positive growth performance of India (which is another very highly populated poor country: it hosts about 15% of the world’s population). Notice that the unweighted inequality measures discussed in Section 1 treat China and India as simply two data points and give them the same importance as Lesotho (with 2 million inhabitants) or Luxembourg (with half a million). Thus, the unweighted measures give a Chinese citizen 1/600th of the weight that it gives to a citizen of Lesotho and 1/3000th the weight it gives a citizen of Luxembourg. The global Gini estimated in this section treats each person in the world equally. Thus, the economic progress of China since 1978 and India after 1991 has resulted in large improvements in the incomes of more than a third of the poorest citizens of the world, and this has implied a substantial reduction in global income inequality.

small increase between 1980 and 1992. They interpret their evidence as showing that inequality stopped growing in the 1980s. Instead, we will show that inequality has declined substantially according to the three decomposable measures.

³⁶ We only report our estimated across-country Gini because Bourguignon and Morrisson (2002) do not report theirs.

It is true that inequalities within China and India have increased in the process. However, the decline of the global Gini over the last two decades clearly indicates that the growth in within-country inequality has not been nearly large enough to offset the substantial decline in across country disparities.³⁷

(ii) Non-Decomposable Measures

Figures 8 through 13 report the estimates of six other popular indexes of income inequality. The exact numbers used to construct these figures are also reported in Tables 1 and 2. Figure 8 reports the variance of the logarithm of income (varlog). The global variance increased during the 1970s but enjoyed a sharp decline during the early 1980s. During the last couple of years of the decade, it increased slightly. During the 1990s, this measure of global inequality remained constant. Overall, the variance declined by more than 7% between 1978 and 1998. The across-country variance follows a similar pattern during the 1980s and it also declines during the 1990s. The fact that the decline of the global index is not as sharp as the across-country index suggests that “within country” inequality has increased during this period. However, the calculation cannot be made by subtracting the across-country measure from the overall measure because, as we showed in the previous section, the variance of log income cannot be decomposed.

Figure 9 displays Atkinson’s $A(0.5)$ index, that is, the index with a coefficient of inequality aversion of 0.5. The pattern of global inequality is remarkably similar to that of the Gini coefficient: after increasing during the 1970s, it has experienced a substantial decline (close to 12%) over the next 20 years, with a small blip at the end of the 1980s. The across-country $A(0.5)$ also mimics the behavior of the global index, with a sharp decline over the last twenty years. Figure 10 reports the behavior of the Atkinson index with a coefficient of 1. The decline in

³⁷ Later in this section we estimate the evolution of within-country inequality. In Section 5 we analyze the role played by China in the process of world income convergence. And finally, in Section 6 we conjecture what will happen to global inequality measures if the income levels of the two billion inhabitants of China and India keep growing and the incomes of the 700 million Africans remain stagnated.

world income inequality over the last two decades is substantial. Again, the reduction in $A(1)$ can be explained by the large drop in across-country inequality, which is also reported in Figure 10.

(iii) Decomposable Measures

The next three figures display three popular General Entropy Indexes (GEI): the Mean Logarithmic Deviation (or $GEI(0)$), the Theil Index (or $GEI(1)$) and the square of the coefficient of variation (or $GEI(2)$). As shown in Section 2, these indexes can be decomposed into an “across-country” and a “within-country” inequality component (the within-country index is a weighted average of within-country measures of income inequality, where the weights are proportional to population or aggregate income). Thus, Figures 11, 12 and 13 display three curves each: the global, the across-country and the within-country inequality.

To put, again, things in historical perspective, Figure 11 merges our results with those of Bourguignon and Morrisson (2002) for the Mean Logarithmic Deviation (or $GEI(0)$). Their estimates are depicted with hollow symbols whereas ours are depicted with filled symbols. There are a few interesting things to observe in Figure 11. First, global inequality increased, as measured by MLD, continuously between 1820 and 1980 and it declined slightly between 1980 and 1992, the last year of their study. In 1970, our estimates and theirs were almost identical. Second, according to our estimates, inequality declined after 1980. Our estimated decline is much larger than reported by Bourguignon and Morrisson. Third, across-country inequality increased dramatically between 1820 and 1970. Fourth, within-country inequality was virtually unchanged between 1820 and 1910, it declined between then and 1950, and then it increased slightly over the following four decades. Fifth, our within and across-country measures of inequality differ substantially from those of Bourguignon and Morrisson. As explained in Section 1, the reason is that they group countries according to economic, social and historical similarities. Thus, what they call “within-inequality” is a mixture of inequality within countries and inequality across countries within the group. We assign the latter to the across-country component. This explains why our within-country inequality is smaller than theirs whereas our across-country measure is larger.

If we focus our attention on our own measures of global, within and across-country inequality, the first thing we note is that the bulk of global income disparities in the world are accounted for by differences across countries rather than differences within countries (in other words, the across-country inequality line always lies above the within-country line). In 1970, only 21% of the world's differences in income could be accounted for by within-country disparities.³⁸ Although the fraction increased over time, the fraction of worldwide inequality accounted for by within-country disparities remains at 30%. Thus, most of the global inequality in the world comes from the fact that the average American income is so much larger than the average Senegalese rather than from the fact that, within the United States or within Senegal, there are very rich people and very poor people.

A second interesting lesson drawn from Figure 11 is that the pattern of global inequality over time is remarkably similar to the evolution of the global Gini or the two Atkinson indexes: an increase during the 1970s and a substantial decline during the following 20 years. Overall, the MLD declined by close to 14% over this period.

A third phenomenon observed in Figure 11 is that this substantial decline is accounted for by the reduction in across-country disparities: we see that the across-country line follows very closely the global inequality line.

A fourth worthy observation is that within-country inequalities have increased monotonically over the whole sample. Obviously, some countries have experienced declining and some others have experienced increasing income disparities. Figure 11 shows that the population-weighted average of all these movements delivers an overall increase over the last three decades.³⁹ The important point, however, we see that this increase is very small compared

³⁸ Since we use quintile data instead of individual income data, the true within-country inequality is actually larger than we estimate. In the next section we show that, if the worldwide distribution of income is computed using individual-country distributions rather than by assuming that all citizens within a quintile have the same level of income, then within-country inequalities are, indeed, larger. However, it is still true that the majority of worldwide inequalities remains explained by across-country disparities. See Sala-i-Martin (2002) for details.

³⁹ An important reason for this rise in within-country inequality is the increased disparities within China and India. Remember that the within-country index is the weighted

to the substantial decline in across-country inequality: between 1978 and 1998, the within-country index increases from 0.18 to 0.23 whereas the across-country declines from 0.67 to 0.51. Naturally, the small increase in within-country inequality is not nearly enough to offset the decline in across-country disparities. Global inequality, therefore, has fallen substantially over the last twenty years.

The overall pattern of global, across and within-country inequality is confirmed by the use of the Theil Index (or GEI(1)), which is reported in Figure 12. The only difference is that the fraction of total inequality accounted for by within-country disparities is now 23% in 1970 and 28% in 1998. Thus, the increase in within-country inequality is slightly smaller according to the Theil index.

Finally, Figure 13 displays the evolution of inequality according to the (squared of the) coefficient of variation (which corresponds to GEI(2)). The two key differences are as follows. Firstly, the small increase in overall inequality that occurred for a couple of years in the late 1980s (an increase that was captured by all the previous measures) is now a bit larger. And secondly, the decrease in global inequality during the 1990s is larger according to GEI(2) than it was according to the other measures. Within-country inequality remains flat over the entire period so the overall decline observed during the last two decades mimics the decline in across-country disparities.

5. The Assumption of “Equal Income Within Quintiles”

The world distribution of income and the various inequality measures reported up to this point were constructed under the assumption that all individuals within a given quintile (for all countries and years), have the same level of income. Obviously, this assumption tends to hide some individual inequality because, in reality, there are differences in income within quintiles. To solve this problem, Sala-i-Martin (2002) uses the five individual quintiles for each country and year to estimate a gaussian kernel density function for every country and every year. He then

average of the inequalities within countries, where the weights are proportional to the size of the country. China and India have more than one-third the world's population and, therefore, the evolution of inequality within these two countries tends to dominate the aggregate index.

integrates the individual income density functions of all countries to find the world income distribution. This global density function can then be used to estimate the various measures of income inequality presented in the previous section. The results are reported in Table 13.

As expected, the global inequality measures are larger when the assumption of equal-income-within-quintile is dropped. For example, the Gini coefficient (displayed in Figure 14) is about 3.5% larger when it is computed from kernels than when it is computed with quintiles. For 1970 the coefficient is equal to 0.63 when individuals are assumed to have the same income within quintiles and it becomes 0.65 when their income is approximated with a gaussian kernel density function. For 1998, the numbers are 0.61 and 0.63 respectively. The increase in the level of global income inequality that one gets by estimating the individual kernels, therefore, is positive but modest. Obviously, the estimates of global disparities a bit larger only because the within-country disparities are estimated to be larger (by construction, the across-country index assumes that all individuals within the country have the same level of income, so these measures are independent of the assumptions we make about the shape of the within-country distribution). In the previous section we showed that, when the MLD is used, the within-country disparities accounted for between 21% and 30% of total inequalities. If we allow for differences of income within-quintiles, the fraction of inequality accounted for by the within-country measure increases by about 5% points and it moves between 26% and 35%. Hence, most of the global income disparities are still accounted by across-country, not within-country, inequalities. The same is true if we use the Theil index or the square of the coefficient of variation, the other two decomposable measures used in the previous section.

The most interesting result is that the behavior of the Gini coefficient over time seems to be exactly the same: the curves representing the ginis computed using the two methods are exactly parallel. The Gini coefficient declines by 4.9% between 1978 and 1998 if all citizens are assumed to have the same income within quintiles and it declines by 5.1% if we allow for within-quintile disparities. The results are identical for all other indexes of income inequality reported in Table 3. Figure 15 displays similar results for the MLD index.

The main lesson is that, although neglecting income inequalities within quintiles might lead to a slight underestimation of global income inequality, it does not change its evolution over

time. Thus, the main conclusions of this paper, namely, that world income inequality has declined during the last two decades, does not depend on the assumption that all individuals within a quintile have the same level of income.

6. The Role of China

In Section 4 we showed that the decline in global income inequality was largely explained by the decline in the across-country component. We also argued that one of the reasons, perhaps the main reason, for this reduction in across-country disparities was the large growth rate enjoyed by the incomes of the people of China. Many observers have pointed out that the “winds of falsification” have apparently influenced Chinese statistical reporting during the last few years Ren (1997), Maddison (1998), Meng and Wang (2000), and Rawski (2001). The bulk of the complains pertain to the period starting in 1996 and especially after 1998 (see Rawski (2001)) so they occur at the very end or after of our sample period.⁴⁰ It should be pointed out that we do not use the official statistics of Net Material Product supplied by Chinese officials but the numbers estimated by Heston, Summers and Aten (2001), who introduce some of the anomalies following the work of Maddison (1998). The result is that the growth rate of Chinese GDP per capita in our data set is 4.8% per year, a full two percentage points lower than the official estimates (the growth rate for the period 1978-1998 is 6.1% in our data set as opposed to the 8.0% reported by the Chinese Statistical Office).

Although it is clear that the success of China is part of the story, the question is whether it is the entire story. To check this, we do two things. First, we recompute all the inequality indexes reported in Section 3 excluding the 1.2 billion Chinese incomes from the sample. Since all the measures deliver similar results, we simply report the decomposition of the MLD in Figure 15. The overall index is basically flat during the whole sample period. It increases slightly during the 1970s, it falls during the early 1980s, it increases again slightly during the late 1980s and it remains constant during the 1990s. The figure also shows that the across-country index mimics

⁴⁰ Some of the inconsistencies found relate to declines in energy consumption or air travel between 1998 and 2001, a period during which the official estimates of GDP growth continued to deliver staggering numbers.

the global index. The reason is that within-country inequality remains virtually unchanged during the whole sample period. Hence, China seems to be responsible for both the reduction in across-country inequality and the small increase in within-country inequality documented in Section 5. However, even if we were to eliminate China from our analysis, we would fail to see an “explosion” in global income inequality.

The elimination of one-fifth of the people in the sample (which is what we do when we simply exclude China from our analysis) may not be the best way to evaluate what would have happened had China failed to grow at the rapid rates it experienced during the last twenty years. Throwing away 20% of one’s sample point is never a good idea. An alternative would be to assume that China grew at the average growth rate of the world per capita income. We recomputed all the indexes for this counterfactual. The main result is that the global inequality index declines slightly during the last 20 years. The across-country inequality also experiences a negligible decline whereas the within-country index is almost unchanged. Thus, the large growth rates in China help us explain the overall decline in income inequality during the last two decades, but even if China did not grow at the extraordinary rates we would have failed to see exploding income inequalities across the globe.

7. Converging-to-the-Rich, Diverging-from-the-Poor : African Growth

We have argued throughout the paper that, as far as world income inequality is concerned, the big story of the last twenty years has been the extraordinary growth rate of the incomes of 1.2 billion Chinese citizens (and to a lesser extent, the growth of the income of one billion Indians). Despite the increasing disparities within these two large countries, the growth of their average income has triggered a process of convergence that led to substantial reductions in world income inequality. And this is true independent of the exact measure of inequality used.

As China and India grew, the incomes of their numerous citizens tended to converge to those of the citizens of the rich world. At the same time, their incomes diverged from those of the people living in countries that were left behind. Initially, when China and India were very poor, the convergence-to-the-rich effect dominated. As they got richer, the divergence-from-the-poor effect got stronger and stronger. If China and India continue to grow over the next 50 years and

the incomes of the 700 million Africans remain stagnant, world income inequalities will start rising again. To see this, we project incomes of all the citizens of the world 50 years into the future and we estimate the evolution of world inequality. We assume that the growth rates of all countries over the next half a century are the same rates they experienced between 1980 and 1998 (thus, we assume that China and India will continue to converge towards the OECD). We also assume that poor (mainly African) countries that have experienced negative growth rates will not grow at all between now and 1950.⁴¹ Finally, we assume that population and income shares remain frozen at their 1998 levels.⁴² With these projected incomes, we compute various measures of income inequality for the period 1998-2050. The results for two of these indexes are reported in Figure 17. The main lesson is that world income inequality is expected to fall for a few more years and then, when the Chinese (and Indian) convergence-to-the-rich effect is over, inequality is expected to rise again as the divergence-from-the-poor effects begins to dominate.

Perhaps one lesson we can take from all this is that we worrying about income inequality is not a good idea. One would like to think that it is unambiguously good that more than a third of the poorest citizens see their incomes grow and converge to the levels enjoyed by the richest people in the world. And if our indexes say that inequality rises, then rising inequality must be good and we should not worry about it!⁴³

Another lesson is that, those who worry about income disparities should focus on the lack of aggregate growth of some countries in Latin America and Asia and especially, the lack of

⁴¹ Quah (1997), Jones (1996) and Kremer et al. (2001) forecast the world distribution of income but they assume that each African country has a small but positive probability of engaging in a rapid growth process. Their goal is to predict the steady-state distribution of cross-country per capita income. Our goal here is simply to show that if Africa does not grow, inequalities will increase in the near future.

⁴² We cannot assume that shares will change over the last twenty years will continue over the next fifty because, if so, many of these shares would larger than one, which is impossible. We do not believe this assumption is crucial because within-country inequalities have not changed much over the last two decades so most of the action has come from across-country inequality.

⁴³ This goes back to the question discussed in Section 1 of whether we should worry about inequality or poverty.

growth of almost all the African continent. In 1998, over 80% of the people in the world's lowest quintile outside China and India lived in Africa. The entire populations of Benin, Chad, Congo, Malawi and Togo (all of them belong to group C) and Tanzania are in the lowest quintile. 80% of the population of Burkina Faso, Burundi, Central African Republic, Ethiopia, Gambia, Guinea Bissau, Kenya, Madagascar, Mali, Mozambique, Niger, Nigeria, Rwanda, Sierra Leone, Uganda and Zambia, and 60% of people of Ghana, Mauritania, and Senegal, also belong to the world's lowest quintile. Only Botswana, Gabon, Guinea and South Africa have less than 20% of their citizens in the world's lowest quintile. The growth rates of almost all these countries over the last three decades have been either negative or negligible.⁴⁴ Thus, these people's incomes have been diverging from those of the rich countries and from those who live in the successful economies of China and India. If Africa's growth rate remains stagnant, income inequalities across citizens of the globe will start diverging once again.

8. Conclusions

The main lessons of this paper are as follows:

Firstly, we estimated kernel density functions for the world distribution of income and found that (a) it shifts to the right (so the world as a whole is becoming richer), (b) in 1970, it had a large mode at a level very close to the poverty line and (c) the mode has been shrinking and giving rise to a large middle class.

Secondly, we used the estimated density functions to compute poverty rates and poverty headcounts. Both have declined dramatically. The one-dollar-a-day poverty rate fell from 20% in 1970 to 5% in 1998. The two-dollar rate fell from 44% to 8%. Poverty headcounts have also declined substantially: there were close to 400 million less poor in 1998 than there were in the 1970s.

⁴⁴ Botswana is the most notable exception, as it experienced the largest growth rates in the world during the last 30 years. The rapid aggregate growth led this country to go from having two quintiles below the poverty line in 1970 to none in 1998, and from having 60% of its people in the lowest world quintile in 1970 to 20% in 1998.

Thirdly, the “dramatic” and “disturbing rise” in income inequality during the globalization period is nowhere to be seen. On the contrary, income disparities during the last two decades have declined substantially. This result is confirmed by the seven most popular measures of income inequality: the Gini coefficient, the variance of log-income, two Atkinson’s Indexes, three Generalized Entropy Indexes (which include the Mean Logarithmic Deviation, the Theil Index and the squared coefficient of variation).

Fourthly, the reduction in income inequality can be fully accounted for by the decline in across-income inequalities. A substantial part of the story (although not the entire story) has been the important growth rates experienced by the incomes of 1.2 billion Chinese individuals.

Fifthly, most income inequalities in the world are explained by differences in per capita incomes across countries rather than differences within countries. If inequalities within countries miraculously ceased to exist, about 70% of the world inequalities would remain. Thus, the best strategy to reduce world income inequalities is to induce aggregate economic growth in poor countries and, in particular, in those countries in the African continent which account for over 95% of the world’s poor.

Sixthly, within-country inequalities have increased slightly over the last thirty years. However, this increase has been so small that it does not offset the substantial reduction in across-country disparities.

And finally, we predict that if Africa remains stagnant and all other countries, including China and India, keep growing at rates similar to the ones they experienced during the last two decades, then world income inequalities will resume their long-term upward trend some time during the next twenty years.

Perhaps the main message of this paper is that any economist concerned with poverty or with income inequality should focus his attention on the aggregate growth rate of Africa.

References

- Atkinson A. B. (1970), "On the Measurement of Inequality", *Journal of Economic Theory*, 2, 244-263.
- Atkinson A. B. Ad A. Brandolini, (2001), "Promise and Pitfalls in the use of 'Secondary' Data-Sets: Income Inequality in OECD Countries as a Case Study", *Journal of Economic Literature*, vol XXXIX, number 3, pp.771-800, September.
- Balassa, B., (1964), "The Purchasing Power Parity Doctrine: A Reappraisal", *Journal of Political Economy*, 72: 584:596.
- Barro, R. J. and X. Sala-i-Martin, (1992), "Convergence", *Journal of Political Economy*, 100, 21, p.223-251, April
- Barro R. J. and X. Sala-i-Martin, (1998) "*Economic Growth*", MIT Press.
- Berry, A., F. Bourguignon and C. Morrisson, (1983), "Changes in the World Distribution of Income Between 1950 and 1977", *The Economic Journal*, June, pp. 331-350.
- Bhagwati, J., (1984), "Why are Services Cheaper in Poor Countries?", *Economic Journal*, 93: 331-350.
- Bourguignon, F., (1979) "Decomposable Income Inequality Measures", *Econometrica* 47, pp. 901-920.
- Bourguignon, F. and C. Morrisson, C., (2002) , "Inequality Among World Citizens: 1820-1992", forthcoming *American Economic Review*.
- Chen, S., and M. Ravallion, (2001), "How did the World's Poorest Fare in the 1990s?", *Review of Income and Wealth*, September.
- Cowell, F.A., (1995), "*Measuring Income Inequality*", 2nd Edition, Harvester Wheatsheaf, Hemel Hempstead.
- Dalton, H., (1920), "Measurment of the Inequality of Incomes", *Economic Journal*, 30(9), 348-361.
- Deininger, K. and L. Squire, (1996), "A New Data Set Measuring Income Inequality," *World Bank Economic Review*, Vol. 10, pp. 565-91.
- Dowrick, S. and M. Akmal, (2001), "Contradictory Trends in Global Income Inequality: A Tale of Two Biases", mimeo Australian National University, March.
- Firebaugh, G., (1999), "Empirics of World Income Inequality", *American Journal of Sociology*, 104: 1597-1630.
- Grosh, M. and E.W. Nafziger, (1986), "The Computation of World Income Distribution", *Economic Development and Cultural Change*, pp 347-359.
- Heston, A. R. Summers, and B. Aten, (2001), Penn World Table Version 6.0, Center for International Comparisons at the University of Pennsylvania (CICUP), December.
- Jones, C. I., (1997) "On the Evolution of the World Income Distribution", *Journal of Economic Perspectives*, Summer.

- Korzeniewicz, P. and R. Moran, (1997), "Inequality Decomposition Analysis and the GINI coefficient Revisited" *American Journal of Sociology*, vol 102:1000-39.
- Kremer, M., A. Onaski, and J. Stock, (2001), "Searching for Prosperity", Carnegie-Rochester Conference Series on Public Policy, vol 55, December.
- Maddison, A., (1998), "*Chinese Economic Performance in the Long Run*" Paris: OECD
- Melchior, A, K. Telle and H. Wiig, (2000), *Globalisation and Inequality – World Income Distribution and Living Standards, 1960-1998*. Oslo: Royal Norwegian Ministry of Foreign Affairs, Report 6b/2000.
- Meng, L. and X. Wang, (2000), "An Estimate of the Reliability of Statistical Data on China's Economic Growth", *Jingji yanji (economic Research)*, 10: 3-13.
- Milanovic, B., "True World Inequality 1988 and 1993: First Calculations Based on Household Surveys Alone", World Bank Working Paper, 2000.
- Pigou, A. C., (1920), "*The Economics of Welfare*", 4th Edition, McMillan, London.
- Pritchett, L., (1997), "Divergence Big Time", *Journal of Economic Perspectives*, vol 11, number 3, Summer, pp. 3-17.
- Quah, D., (1997), "Empirics for Growth and Distribution: Polarization, Stratification, and Convergence Clubs", *Journal of Economic Growth*, 2.
- Quah, D., (1996), "Twin Peaks: Growth and Convergence in Models of Distribution Dynamics", *Economic Journal*, July.
- Ravallion, M. and S. Chen, (1997), "What Can New Survey Data Tell Us About Recent Changes in Distribution and Poverty?" *World Bank Economic Review*. 11: 357-82.
- Ravallion, M., Datt, G., and van de Walle, D. "Qualifying Absolute Poverty in the Developing World", *Review of Income and Wealth*, 37, 345-361, 1991.
- Rawski, T. G., "What's Happening to Chinese Statistics?", *China Economic Review*, December.
- Ren, R., (1997), "China's Economic Performance in International Perspective", *OECD*, Paris.
- Sala-i-Martin, X., (2002), "Estimating Poverty Rates, Inequality, and the World Distribution of Income from Individual Country Distributions: 1970-1998", mimeograph Columbia University, March.
- Samuelson, P.A., (1964), "Theoretical Notes on Trade Problems", *Review of Economics and Statistics*, 46(2),: 145-154.
- Schultz, T.P., (1998), "Inequality and the Distribution of Personal Income in the World: How it is Changing and Why", *Journal of Population Economics*, pp. 307-344.
- Shorrocks, A. F., (1980), "The Class of Additively Decomposable Inequality Measures", *Econometrica* 48, pp 613-625.
- Summers, R. and A. Heston, (1988), "A new set of international comparisons of real product and price levels estimates for 130 countries, 1950-1985", *Review of Income and Wealth*, 34(1), March, 1-25.

- Summers, Robert and Heston, Alan., (1991), The Penn World Table (Mark 5): an expanded set of international comparisons, 1950-1988, *Quarterly Journal of Economics*, 106(2), May, 327-68.
- Theil, H., (1967), "*Economics and Information Theory*", North-Holland, Amsterdam,
- Theil, H., (1979), "World Income Inequality", *Economics Letters*, 2: 99-102.
- Theil, H., (1996), "*Studies in Global Econometrics*", Kluwert Academic Publishers, Amsterdam.
- Theil, H. And J. L. Seale, (1994), "The Geographic Distribution of World Income, 1950-1990", *De Economist*, 4.
- United Nations Development Program, (1999), Human Development Report.
- Wade, R., (2001), "The Rising Inequality of World Income Distribution", *Finance and Development*, International Monetary Fund, December.

Figure 1. Per Capita Income Ratios:
20% Highest to 20% Lowest




Figure 2. Variance of Log- Per Capita Income: 125 Countries

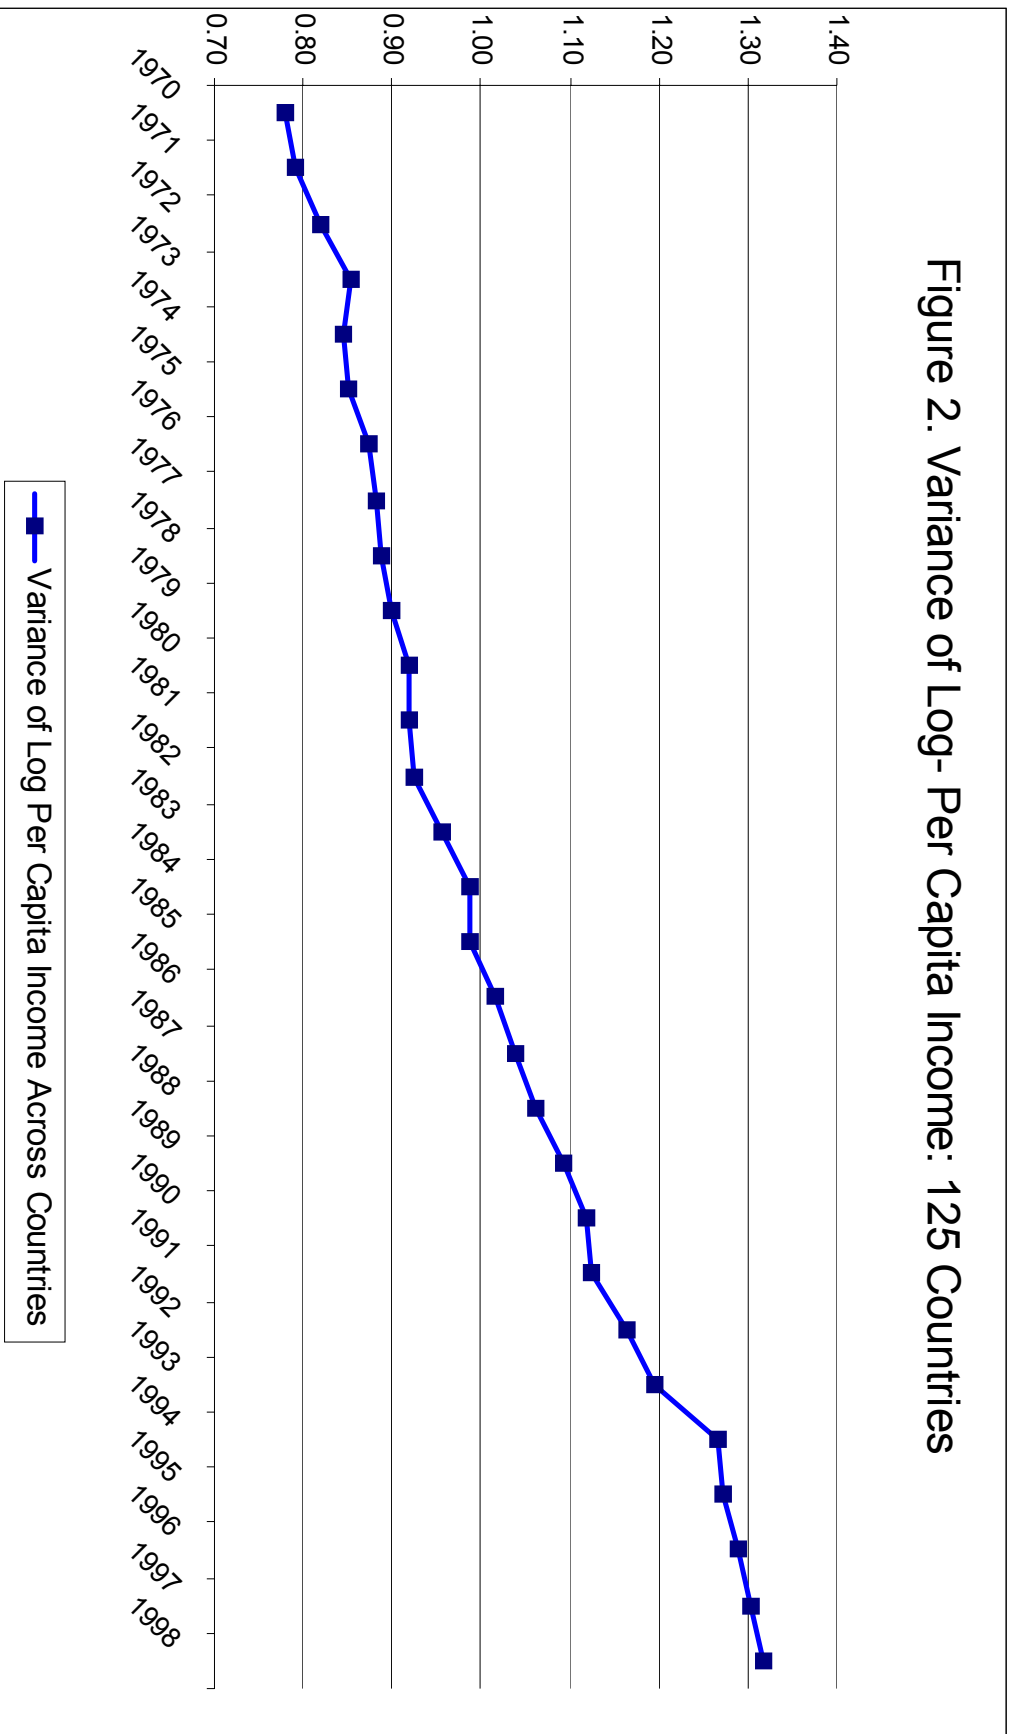


Figure 3. Population-Weighted Variance of Log Per Capita Income:
125 Countries

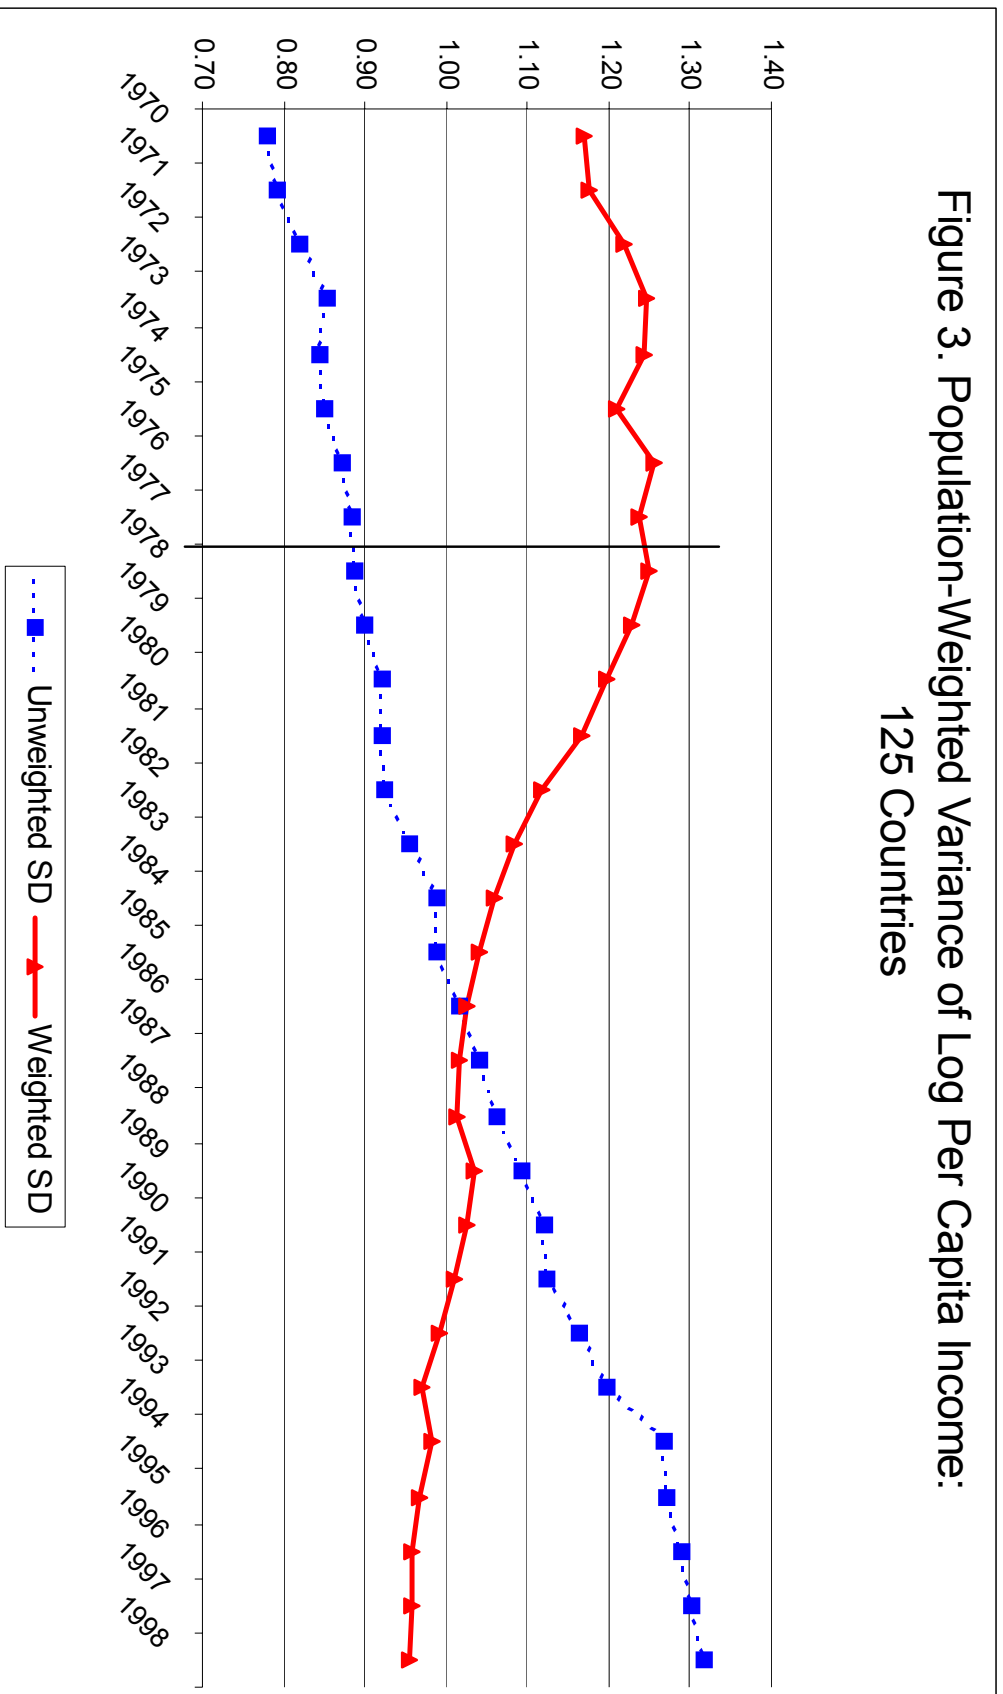


Figure 4. Estimated World Income Distributions (Various Years)




Figure 5. Poverty Rates




Figure 6. Poverty Headcount (in millions of people)

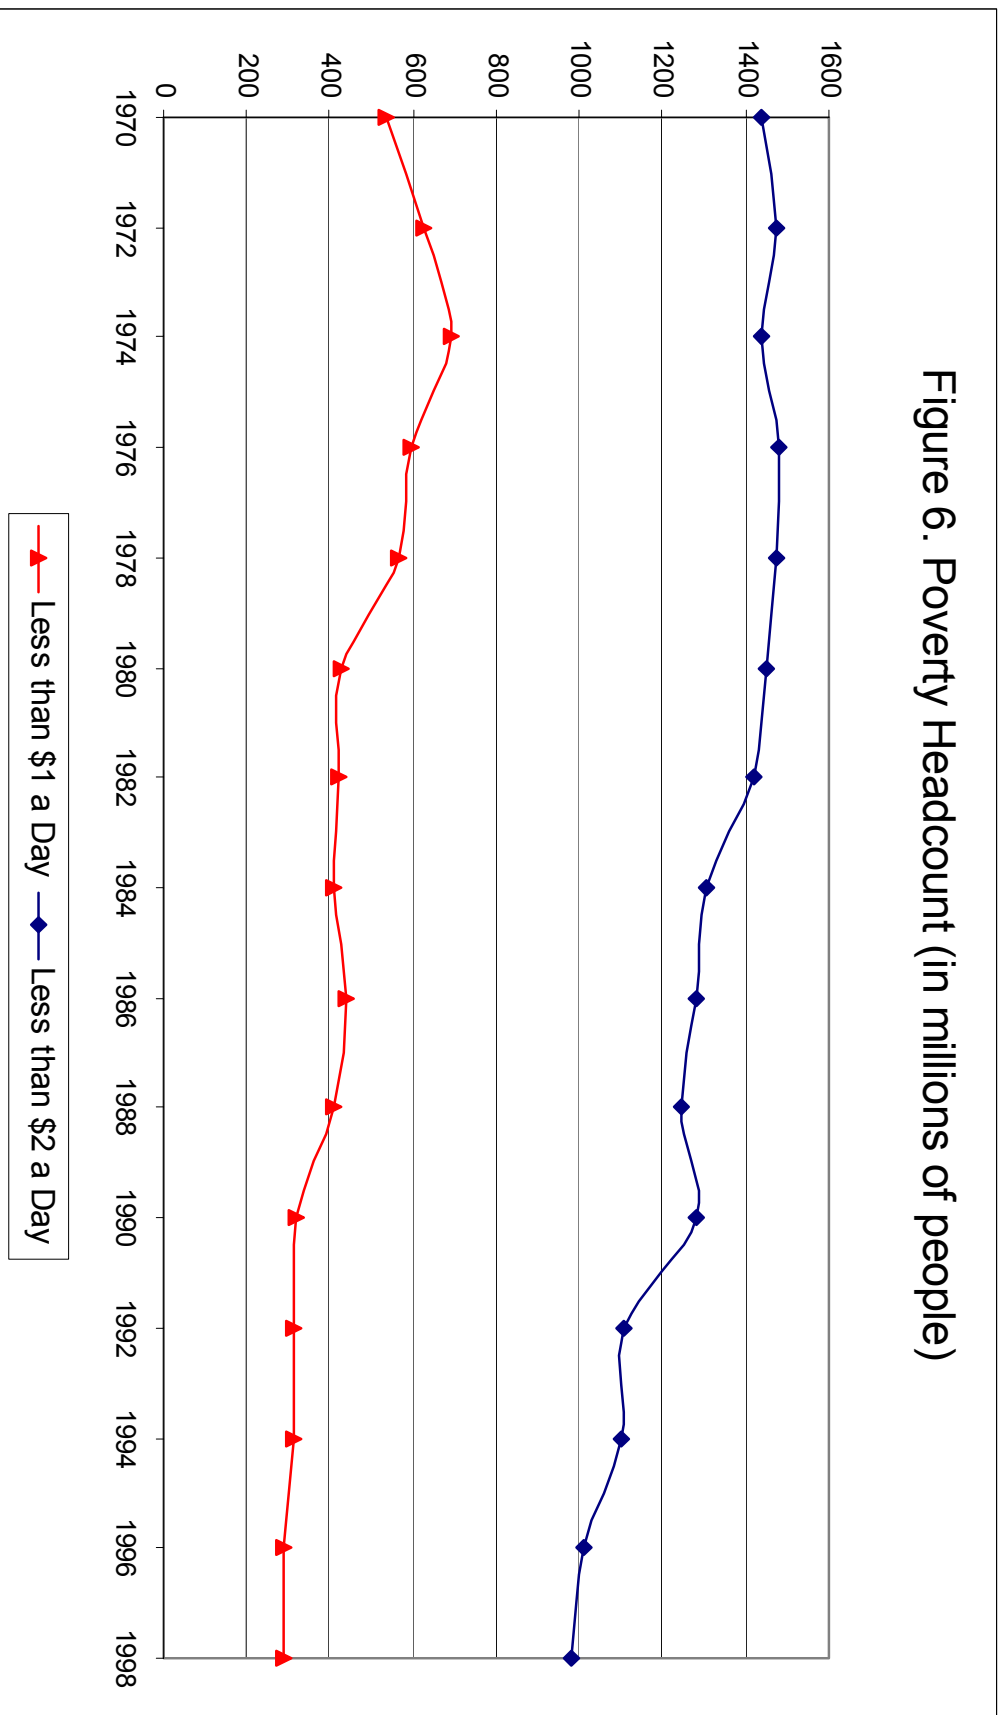


Figure 7. Bourguignon-Morrisson and Sala-i-Martin: Global and Across-Country Gini




Figure 8. Global Inequality: Variance of log-Income




Figure 9. Global Inequality: Atkinson (0.5)




Figure 10: Global Inequality: Atkinson (1)




Figure 11. Bourguignon-Morrison and Sala-i-Martin Decomposition
MLD (or GEI(0))

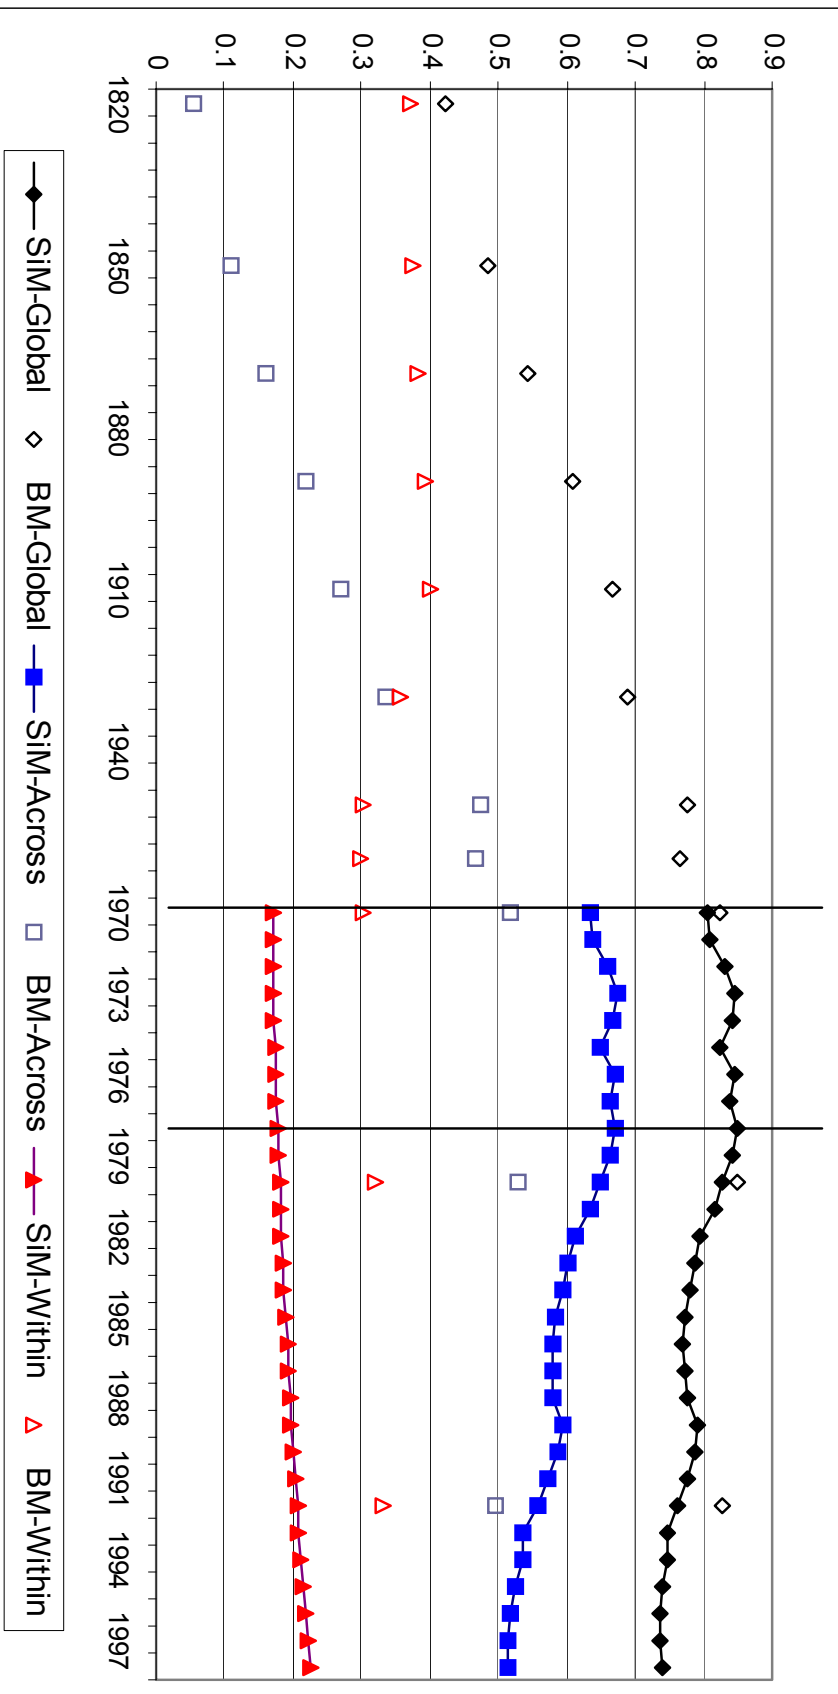


Figure 12. Inequality Decomposition: Theil (or GEI(1))

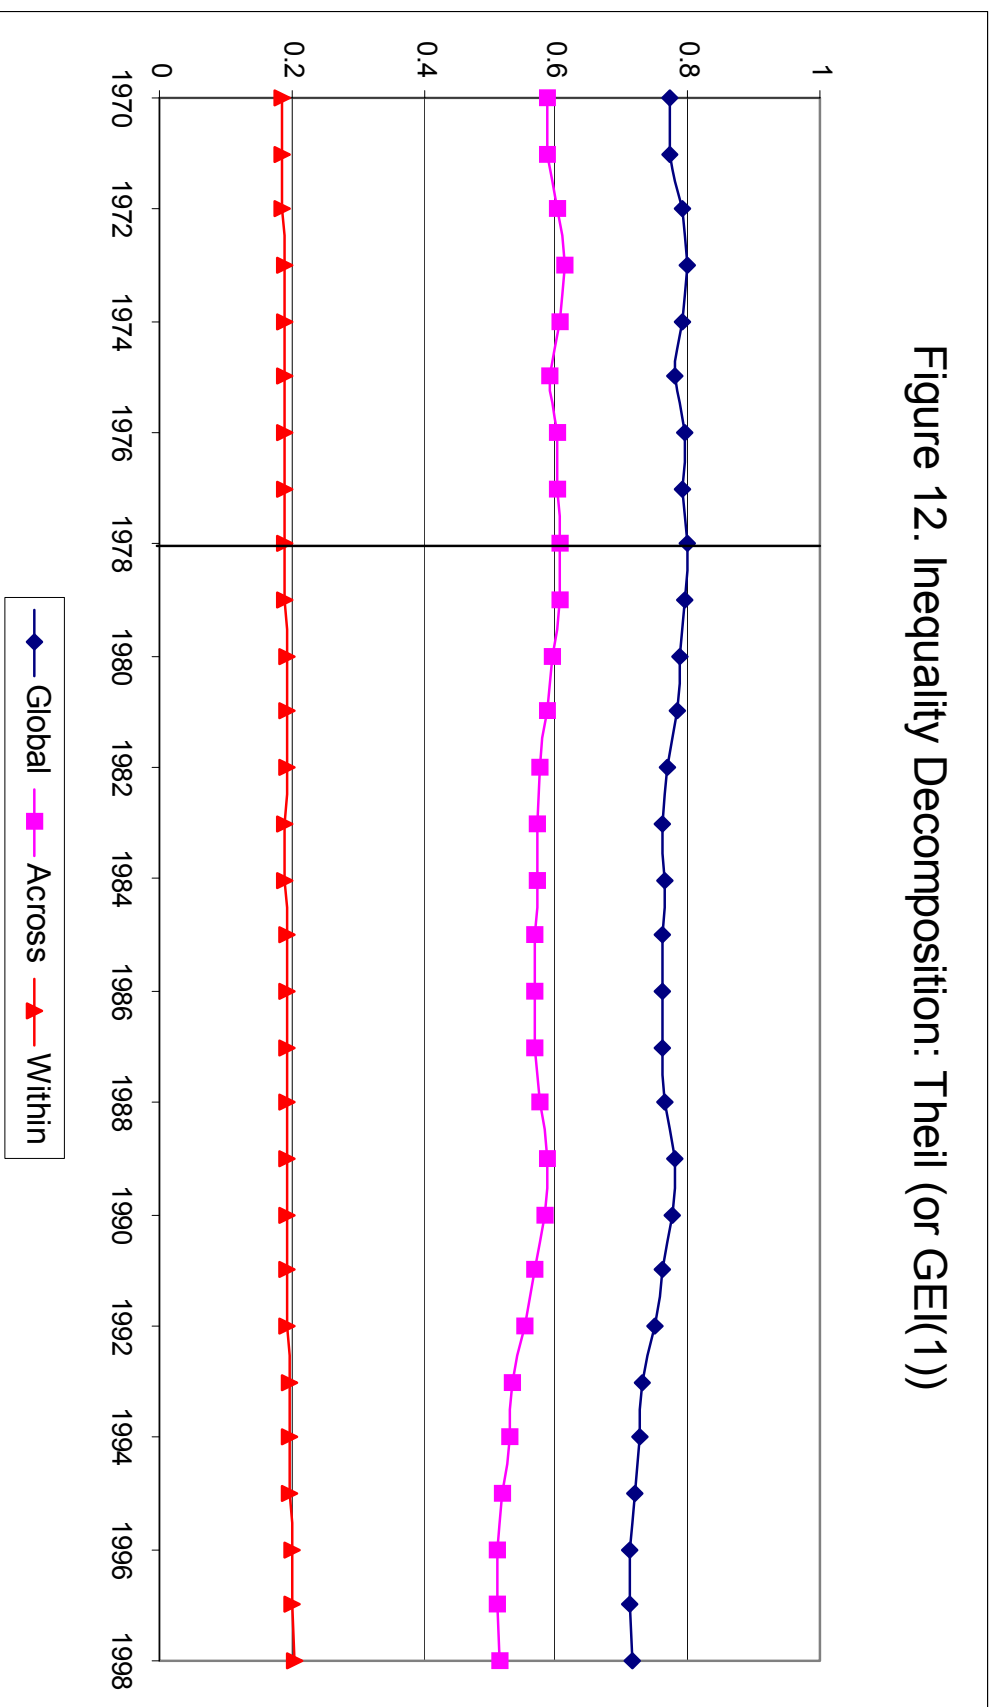


Figure 13. Inequality Decomposition: Squared of CV (or GEI(2))

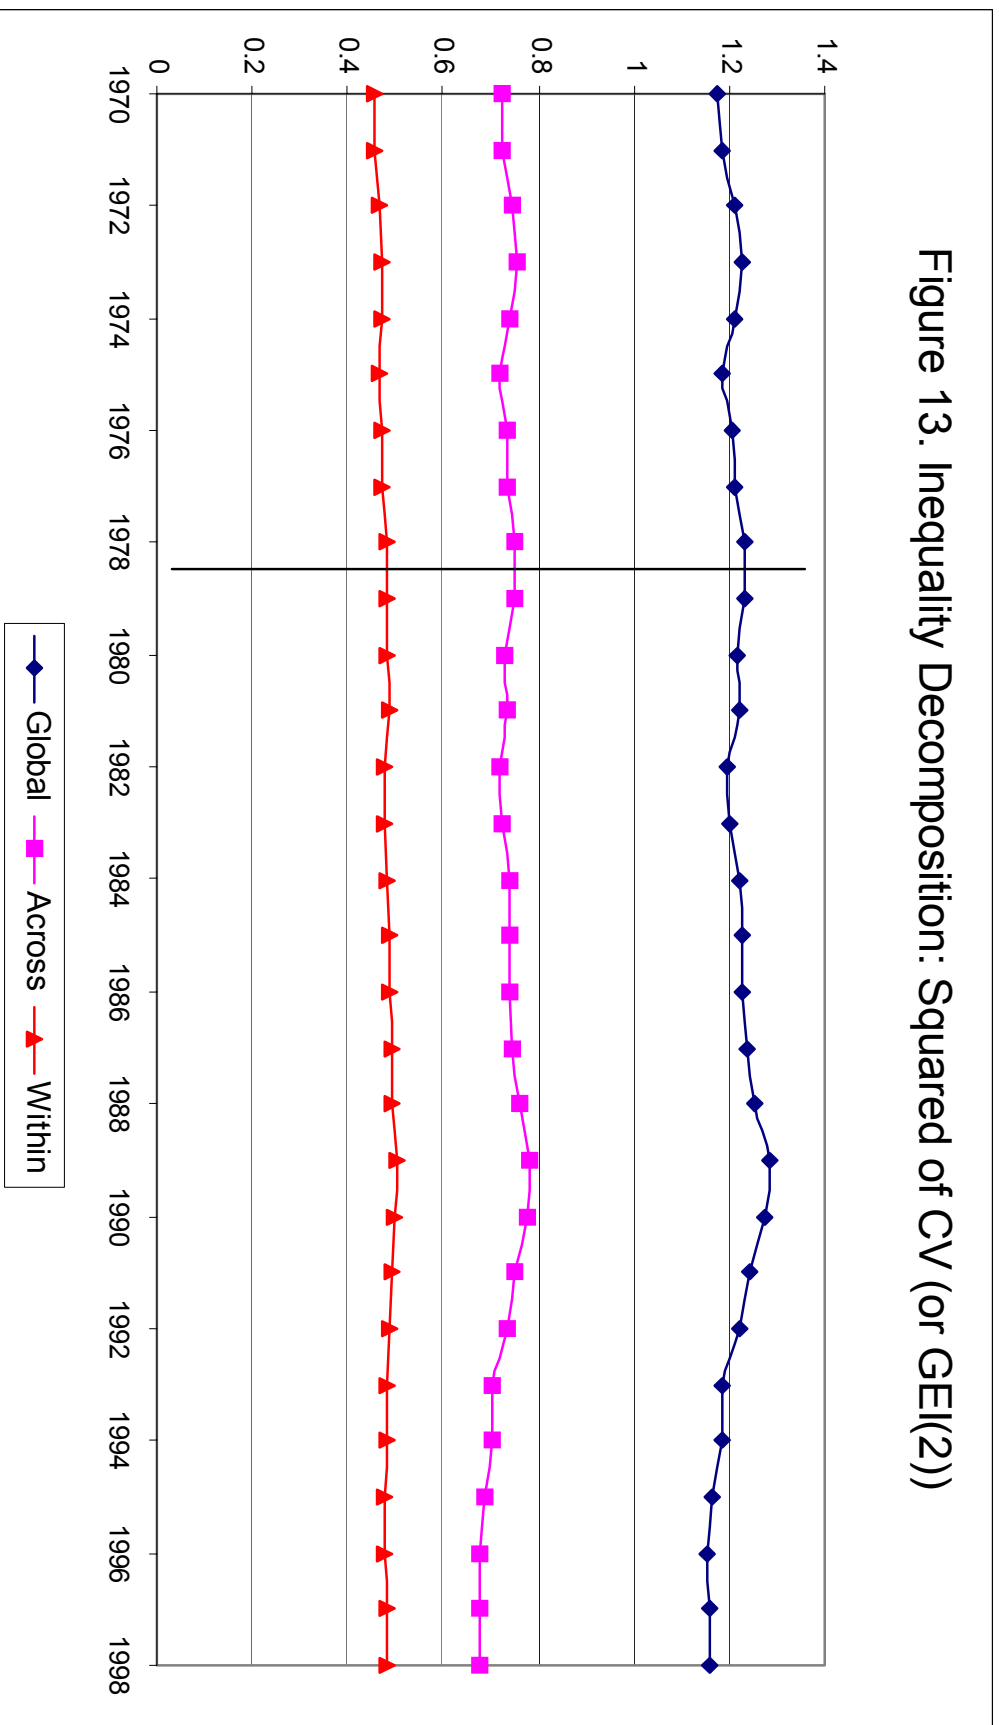


Figure 14. Quintiles vs. Kernels: Gini Coefficient

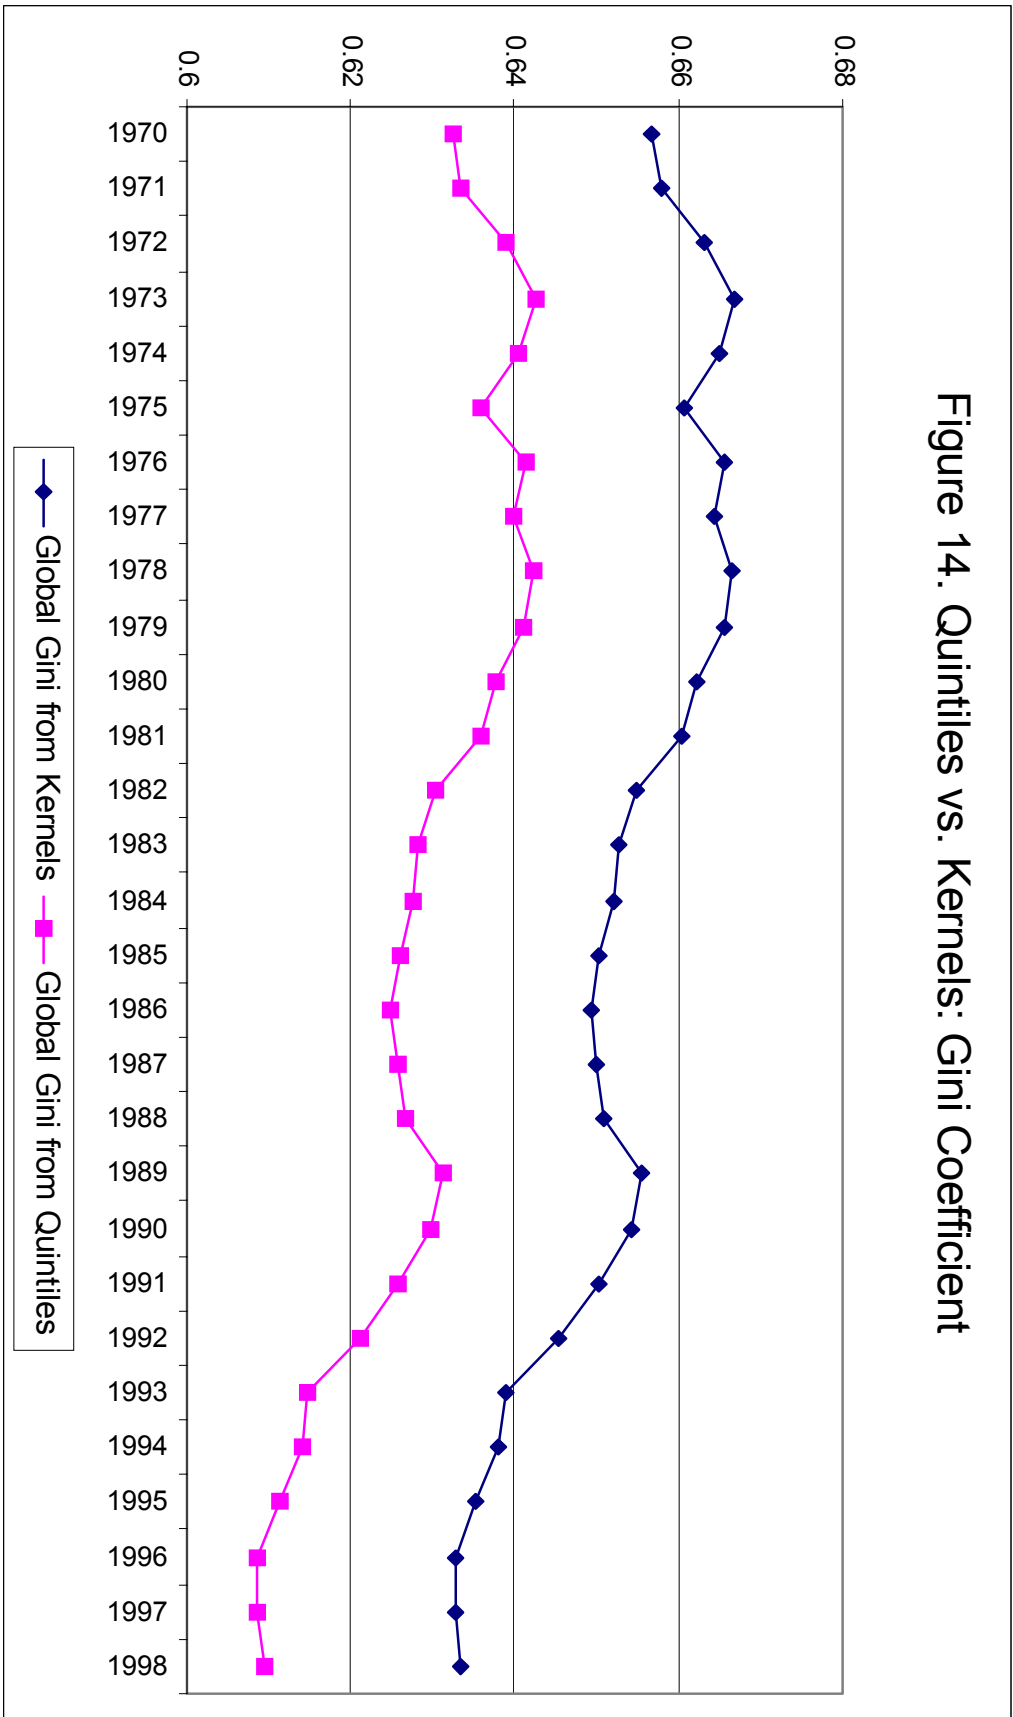


Figure 15. Quintiles vs. Kernels: MLD (GE(0))

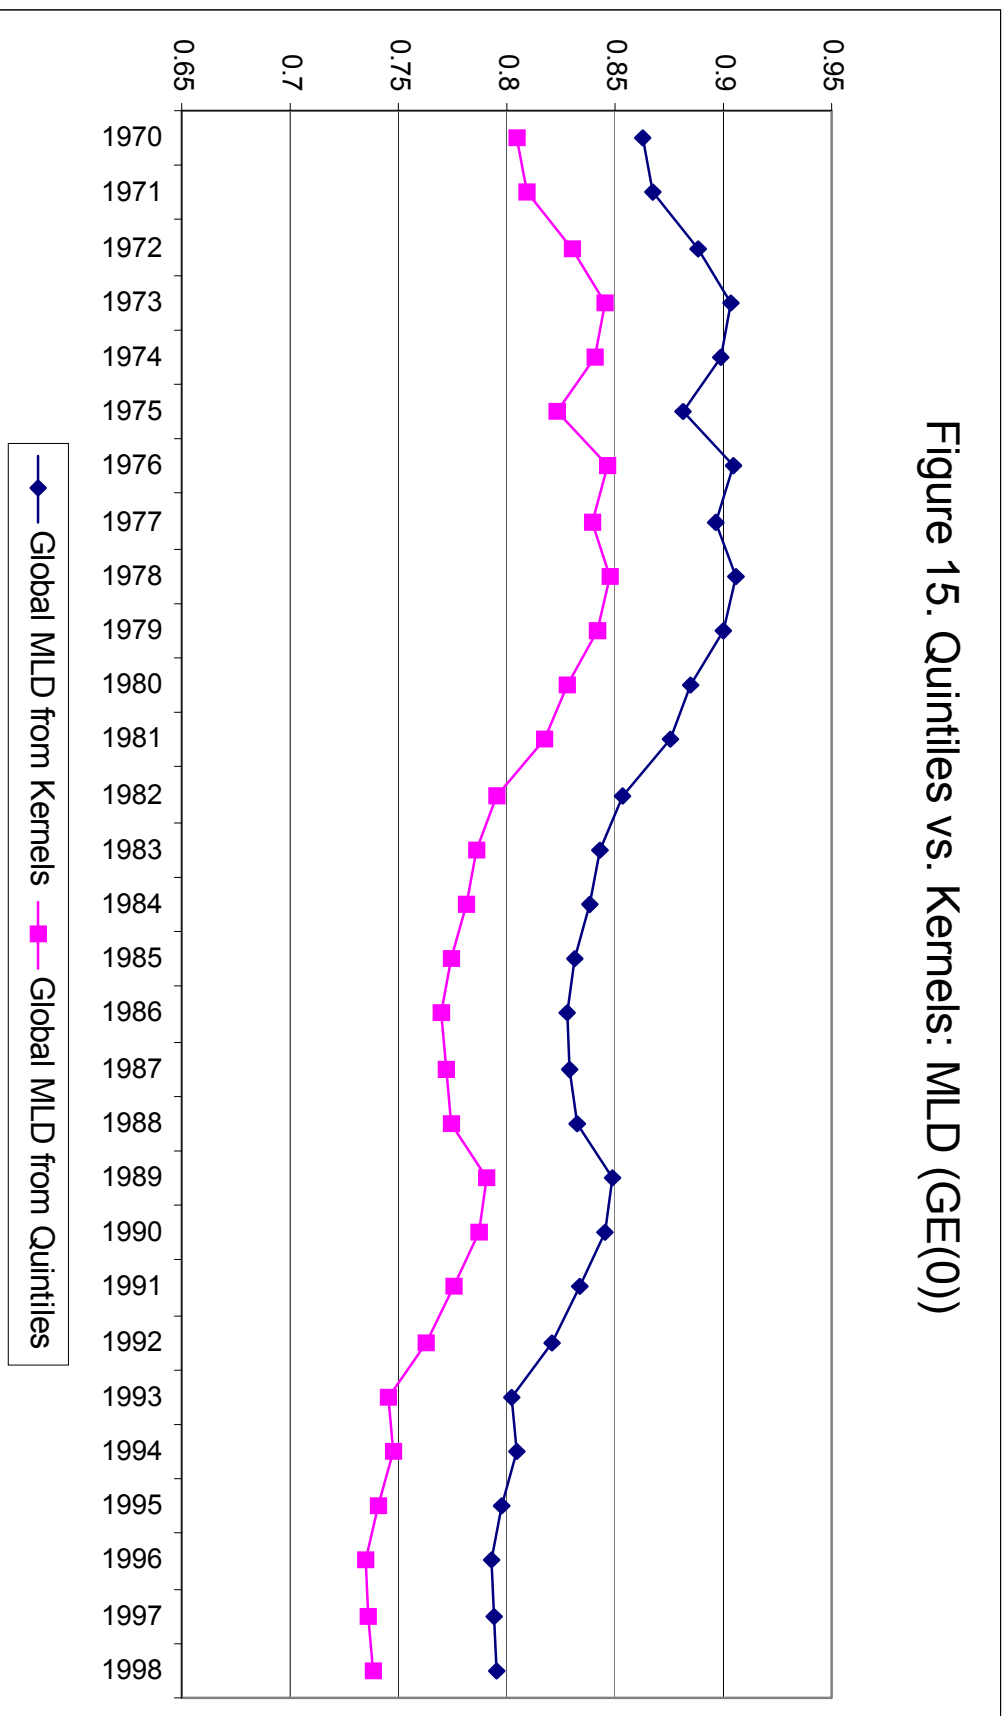


Figure 16. MLD Excluding China

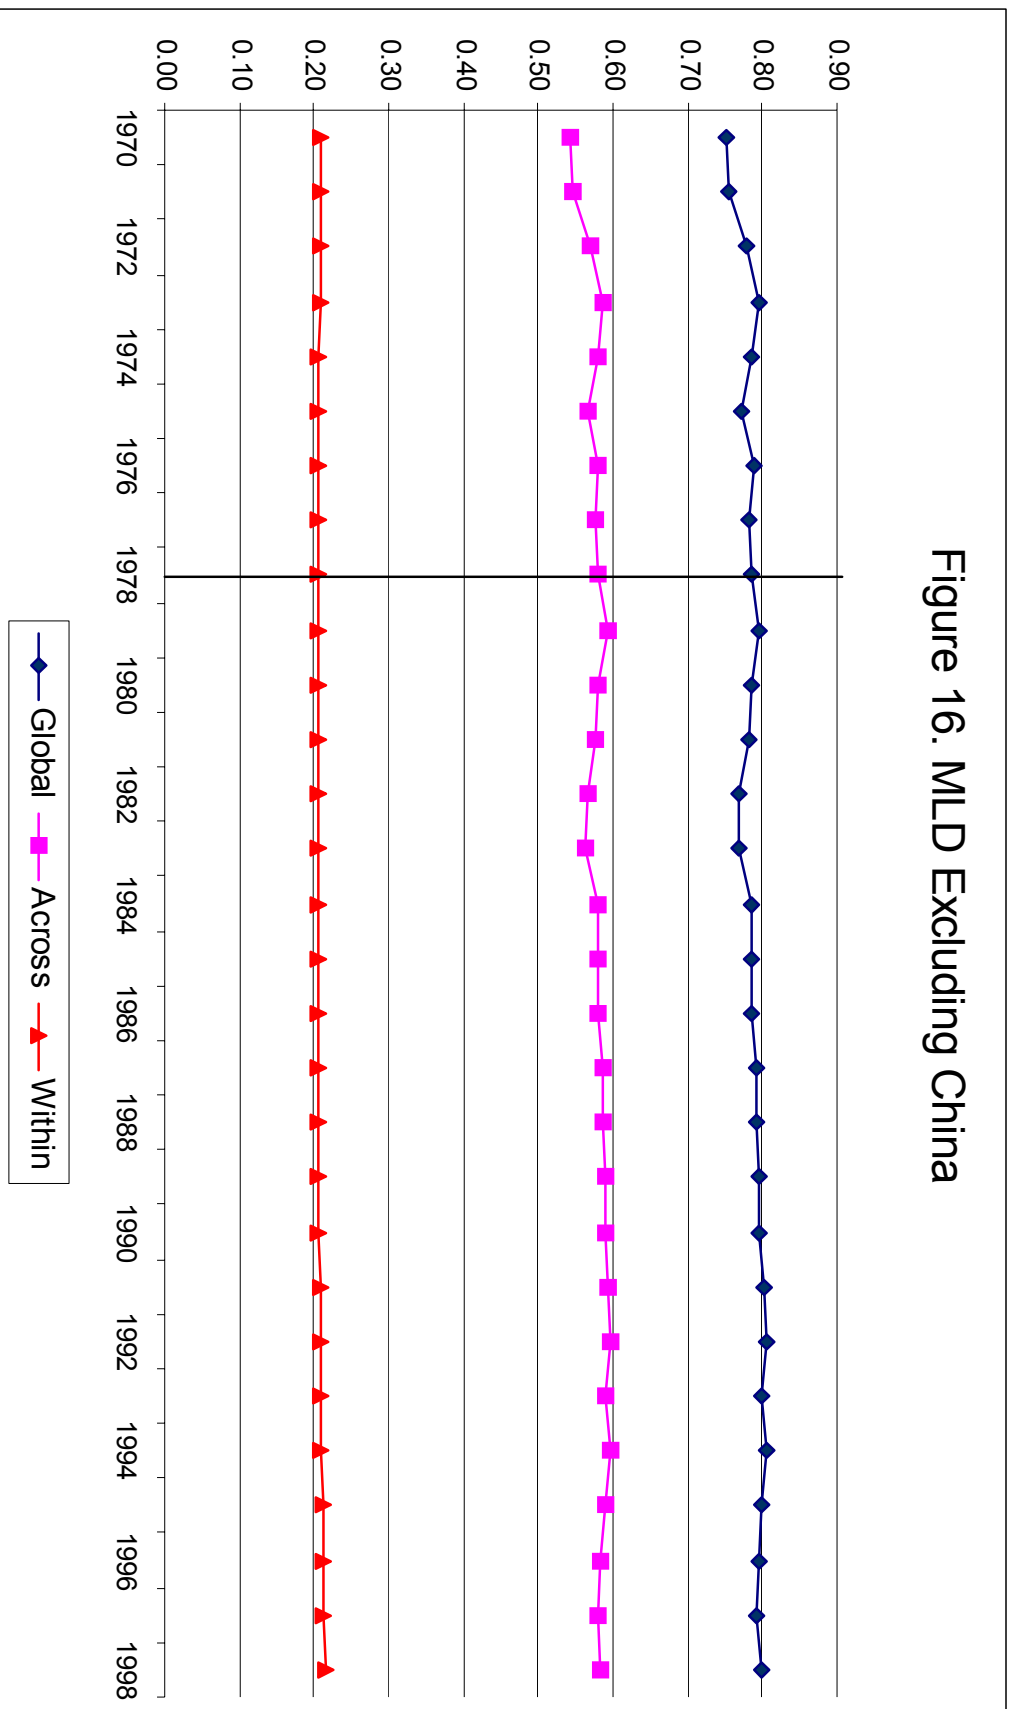


Figure 17. Global Projections if Same Growth as 1980-98

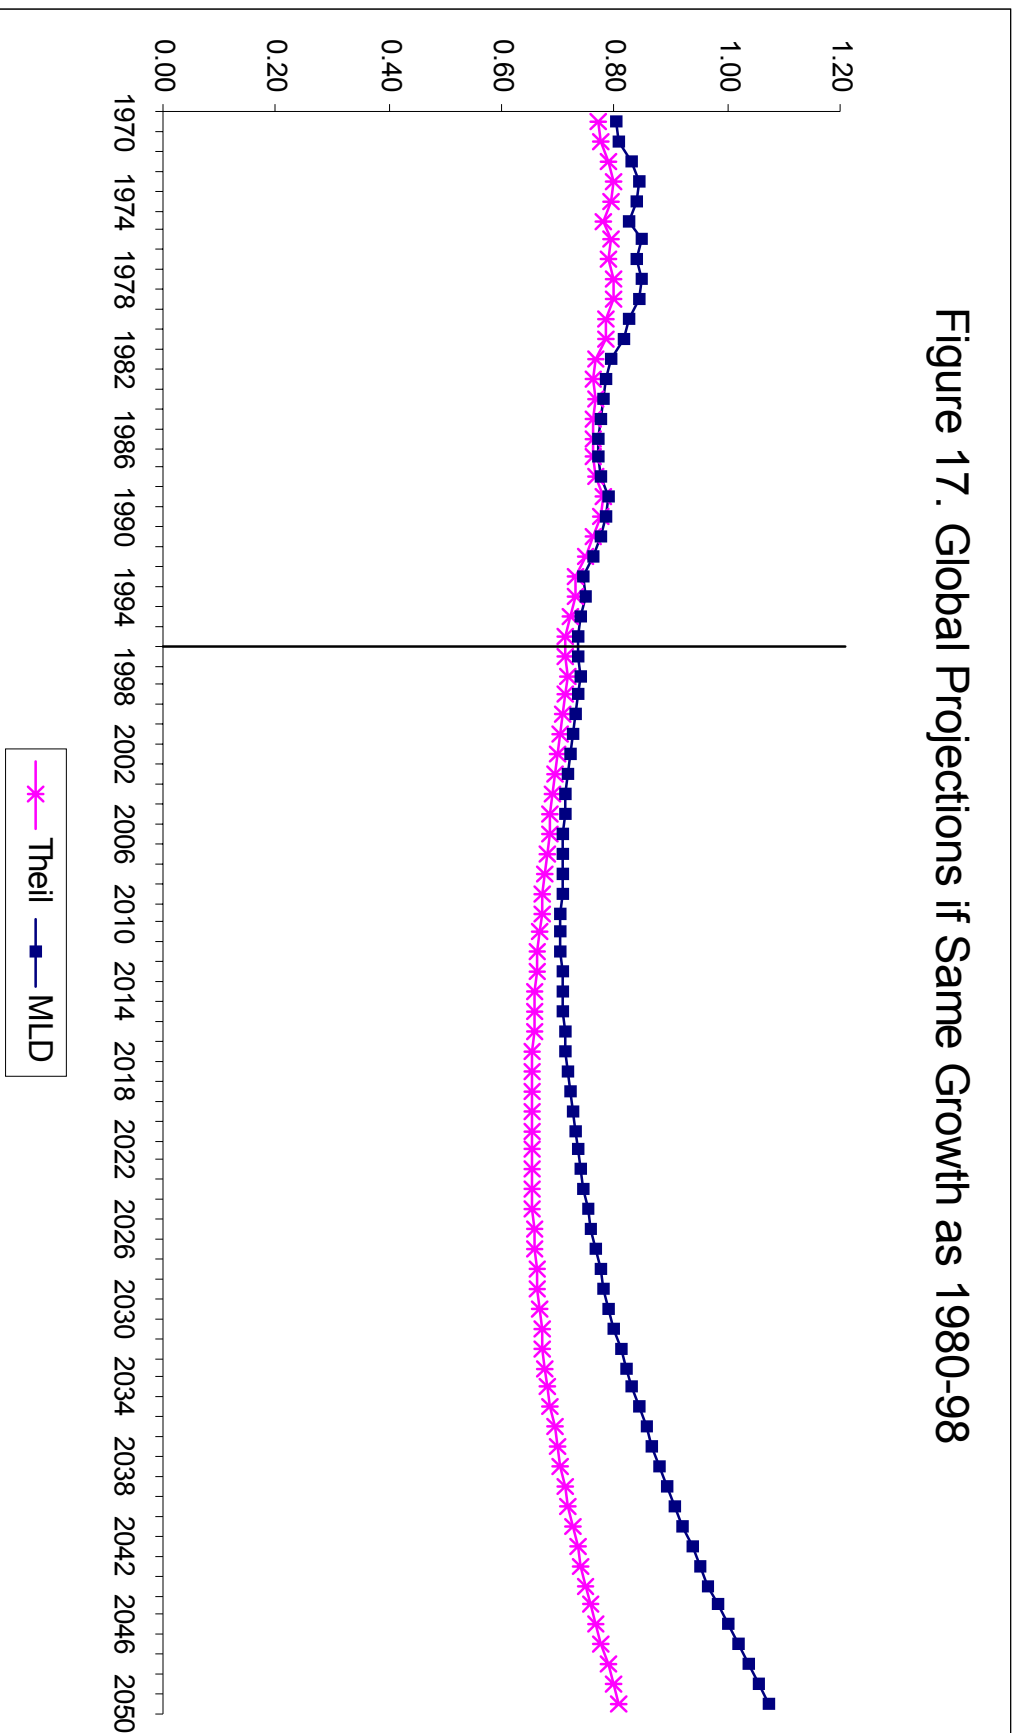


TABLE 1: Non-Decomposable Inequality Indexes

YEAR	Gini Coefficient Global Across	Variance of Log Global Across	Atkinson A(0.5) Global Across	Atkinson A(1) Global Across
1970	0.633	1.419	0.337	0.553
1971	0.634	1.426	0.338	0.555
1972	0.639	1.467	0.345	0.564
1973	0.643	1.497	0.349	0.571
1974	0.640	1.495	0.347	0.569
1975	0.636	1.465	0.342	0.561
1976	0.641	1.514	0.348	0.571
1977	0.640	1.500	0.346	0.568
1978	0.642	1.515	0.349	0.572
1979	0.641	1.502	0.347	0.569
1980	0.638	1.477	0.343	0.563
1981	0.636	1.454	0.340	0.558
1982	0.630	1.414	0.334	0.549
1983	0.628	1.394	0.331	0.544
1984	0.628	1.381	0.330	0.542
1985	0.626	1.370	0.328	0.539
1986	0.625	1.362	0.327	0.537
1987	0.626	1.363	0.328	0.538
1988	0.627	1.368	0.329	0.539
1989	0.631	1.395	0.334	0.546
1990	0.630	1.396	0.333	0.545
1991	0.626	1.390	0.328	0.540
1992	0.621	1.384	0.323	0.534
1993	0.615	1.371	0.315	0.525
1994	0.614	1.393	0.315	0.526
1995	0.611	1.390	0.312	0.523
1996	0.609	1.394	0.309	0.521
1997	0.609	1.404	0.309	0.521
1998	0.609	1.411	0.310	0.522

Percent Change
Since 1978 -0.049 -0.064

-0.061 -0.263

-0.109 -0.194

-0.084 -0.189

TABLE 2: Decomposable Inequality Indexes


YEAR	MLD or GEI(0)			Theil or GEI(1)			Squared CV or GEI(2)		
	Global	Across	Within	Global	Across	Within	Global	Across	Within
1970	0.805	0.634	0.170	0.771	0.586	0.186	1.175	0.722	0.454
1971	0.809	0.638	0.171	0.775	0.588	0.186	1.183	0.726	0.457
1972	0.830	0.659	0.171	0.790	0.603	0.187	1.210	0.743	0.466
1973	0.846	0.674	0.172	0.800	0.613	0.188	1.228	0.755	0.473
1974	0.841	0.668	0.173	0.793	0.605	0.189	1.209	0.739	0.470
1975	0.823	0.650	0.174	0.780	0.591	0.189	1.185	0.719	0.466
1976	0.846	0.671	0.175	0.794	0.604	0.190	1.208	0.735	0.474
1977	0.839	0.663	0.176	0.791	0.601	0.190	1.210	0.736	0.474
1978	0.847	0.670	0.178	0.798	0.608	0.191	1.233	0.750	0.482
1979	0.842	0.663	0.179	0.796	0.605	0.191	1.233	0.749	0.484
1980	0.828	0.647	0.181	0.786	0.593	0.193	1.214	0.731	0.483
1981	0.817	0.635	0.182	0.783	0.589	0.194	1.219	0.733	0.487
1982	0.796	0.612	0.184	0.767	0.574	0.193	1.195	0.716	0.479
1983	0.786	0.600	0.186	0.763	0.572	0.191	1.200	0.724	0.476
1984	0.781	0.593	0.188	0.765	0.573	0.191	1.222	0.737	0.484
1985	0.774	0.585	0.190	0.761	0.570	0.191	1.225	0.739	0.486
1986	0.770	0.578	0.192	0.759	0.568	0.192	1.227	0.740	0.487
1987	0.772	0.578	0.194	0.762	0.570	0.193	1.237	0.746	0.490
1988	0.775	0.579	0.196	0.767	0.574	0.193	1.252	0.758	0.494
1989	0.791	0.592	0.198	0.781	0.587	0.194	1.282	0.779	0.503
1990	0.787	0.586	0.201	0.776	0.583	0.194	1.273	0.774	0.500
1991	0.776	0.573	0.203	0.763	0.568	0.194	1.242	0.751	0.490
1992	0.763	0.557	0.206	0.749	0.554	0.195	1.221	0.734	0.487
1993	0.745	0.536	0.209	0.729	0.533	0.196	1.185	0.705	0.480
1994	0.747	0.535	0.212	0.727	0.529	0.198	1.183	0.700	0.483
1995	0.740	0.525	0.215	0.719	0.520	0.199	1.166	0.687	0.479
1996	0.735	0.517	0.219	0.712	0.512	0.200	1.155	0.678	0.477
1997	0.736	0.514	0.222	0.712	0.511	0.201	1.161	0.678	0.483
1998	0.739	0.513	0.226	0.716	0.513	0.203	1.160	0.678	0.483
Percent Change Since 1978	-0.128	-0.257	0.248	-0.099	-0.158	0.067	-0.042	-0.082	0.018

Table 3: Global Inequality Measures from Kernel Estimates


YEAR	Gini	Var(log y)	Atkinson(0.5)	Atkinson(1)	MLD	Theil	CV ²
1970	0.657	1.530	0.356	0.578	0.863	0.835	1.503
1971	0.658	1.536	0.358	0.580	0.867	0.838	1.505
1972	0.663	1.577	0.364	0.589	0.889	0.853	1.526
1973	0.667	1.607	0.368	0.595	0.904	0.863	1.536
1974	0.665	1.605	0.366	0.593	0.899	0.856	1.512
1975	0.661	1.575	0.361	0.586	0.881	0.843	1.484
1976	0.666	1.624	0.367	0.595	0.904	0.856	1.503
1977	0.664	1.610	0.365	0.592	0.897	0.853	1.499
1978	0.667	1.626	0.368	0.596	0.905	0.861	1.520
1979	0.666	1.612	0.366	0.593	0.900	0.858	1.515
1980	0.662	1.587	0.362	0.587	0.885	0.848	1.492
1981	0.660	1.564	0.360	0.583	0.875	0.844	1.496
1982	0.655	1.524	0.353	0.574	0.853	0.829	1.469
1983	0.653	1.504	0.350	0.570	0.843	0.825	1.472
1984	0.652	1.491	0.350	0.568	0.839	0.826	1.491
1985	0.650	1.480	0.348	0.565	0.832	0.823	1.491
1986	0.649	1.472	0.346	0.563	0.828	0.821	1.491
1987	0.650	1.472	0.347	0.564	0.829	0.824	1.499
1988	0.651	1.477	0.348	0.565	0.832	0.828	1.513
1989	0.656	1.504	0.354	0.572	0.848	0.842	1.545
1990	0.654	1.505	0.352	0.570	0.845	0.837	1.533
1991	0.650	1.500	0.347	0.566	0.834	0.823	1.497
1992	0.645	1.494	0.342	0.560	0.821	0.810	1.471
1993	0.639	1.480	0.335	0.552	0.803	0.790	1.428
1994	0.638	1.502	0.335	0.553	0.805	0.787	1.424
1995	0.635	1.499	0.332	0.550	0.798	0.779	1.402
1996	0.633	1.503	0.329	0.547	0.793	0.772	1.388
1997	0.633	1.513	0.329	0.548	0.794	0.772	1.383
1998	0.633	1.520	0.330	0.549	0.796	0.776	1.380
Percent Change Since 1978	-0.051	-0.067	-0.108	-0.082	-0.129	-0.103	-0.096

Appendix Figures: Income Shares for Largest Countries


China Income share of Quintile 1




China Income share of Quintile 2




China Income share of Quintile 3



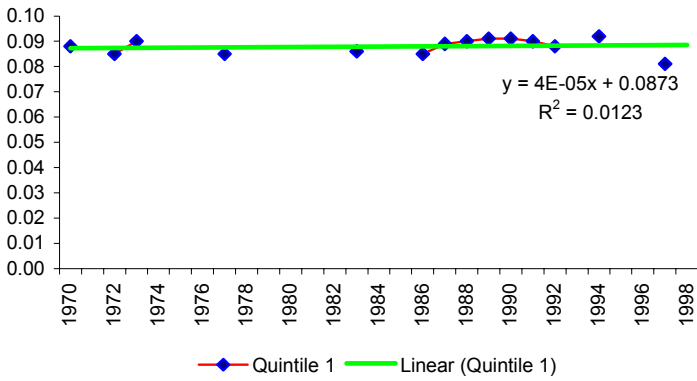
China Income share of Quintile 4



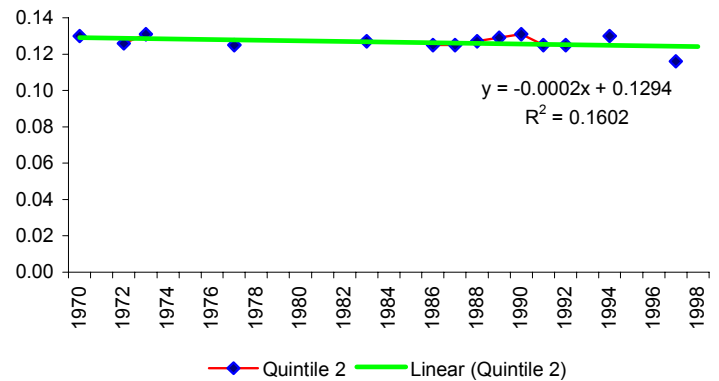
China Income share of Quintile 5



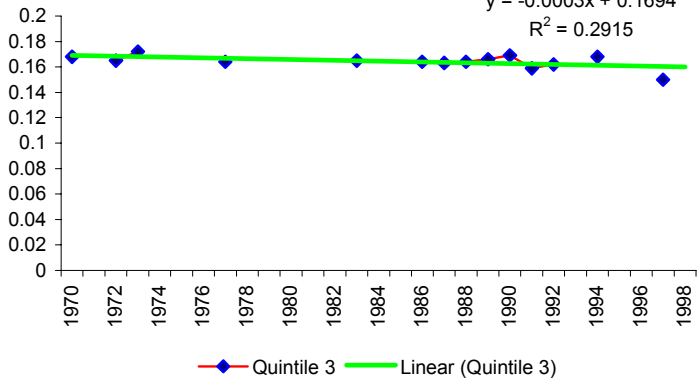
India Income share of Quintile 1



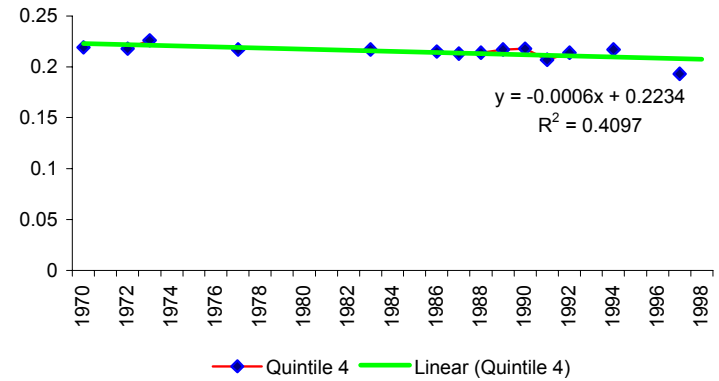
India Income share of Quintile 2



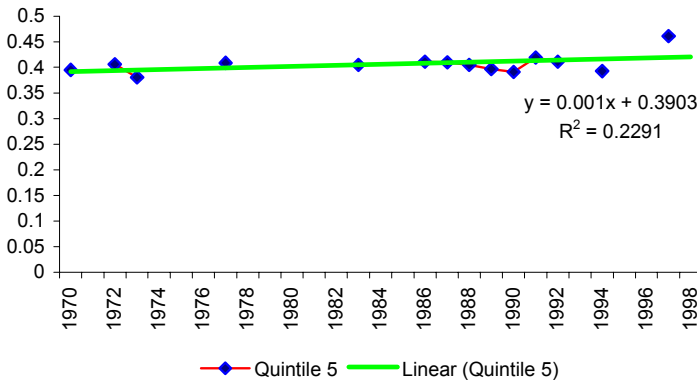
India Income share of Quintile 3



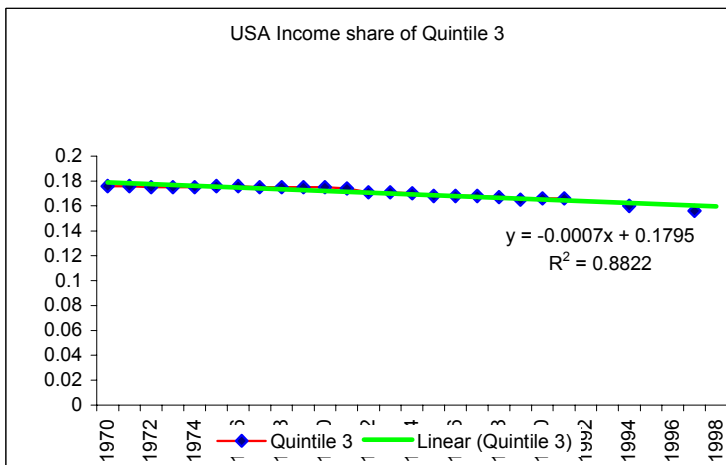
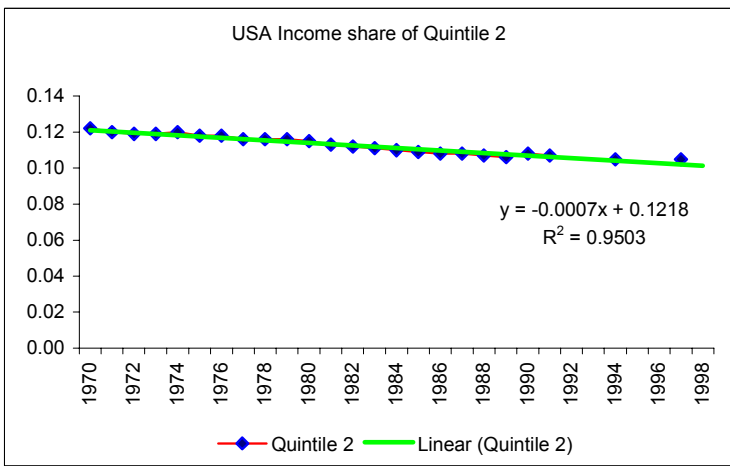
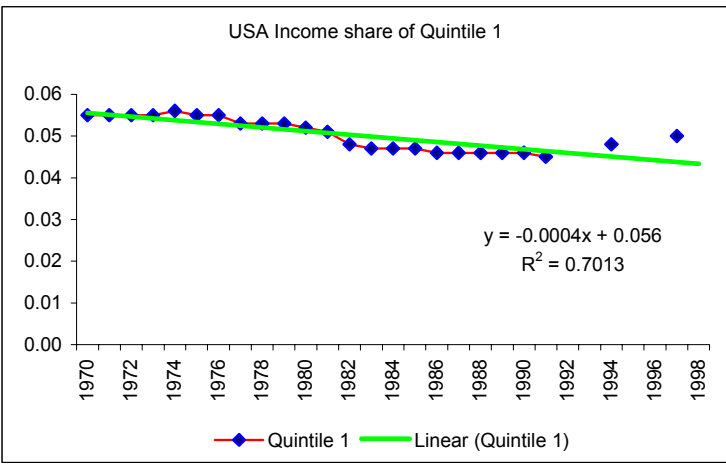


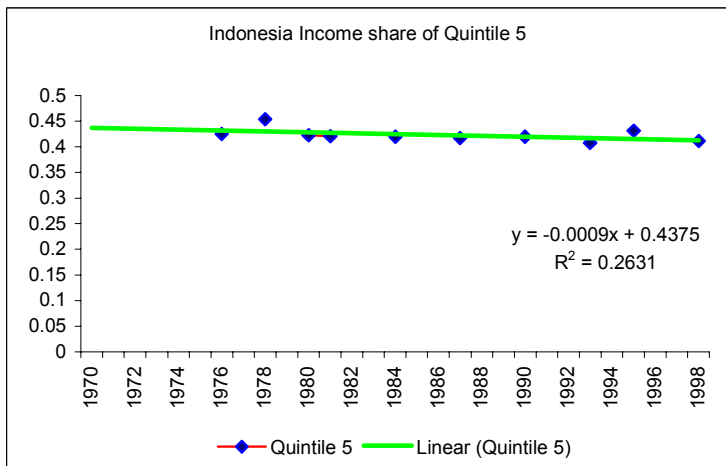
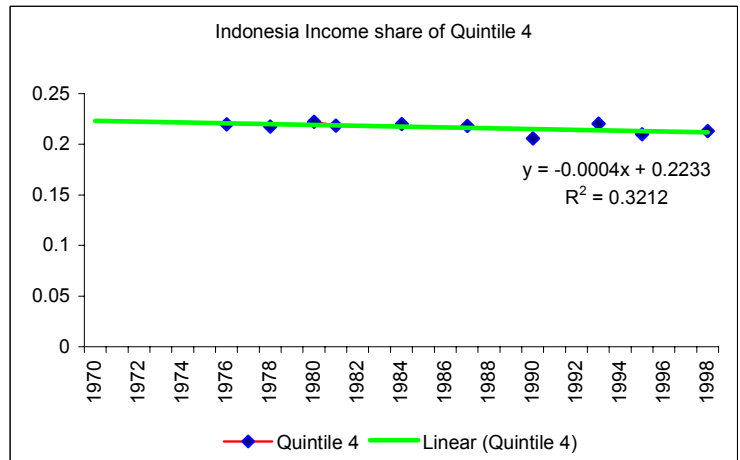
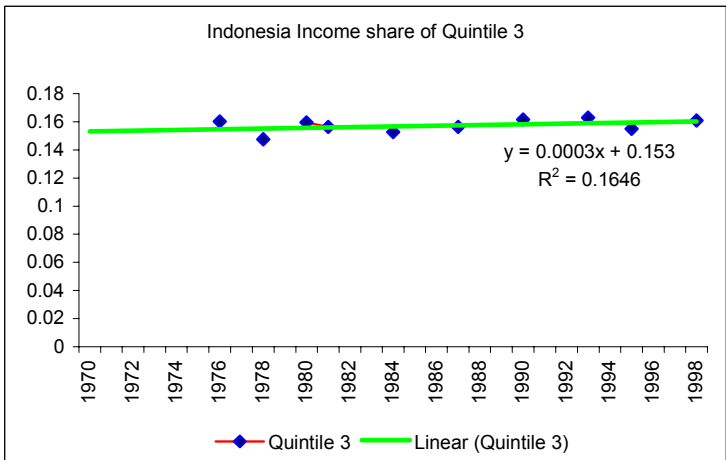
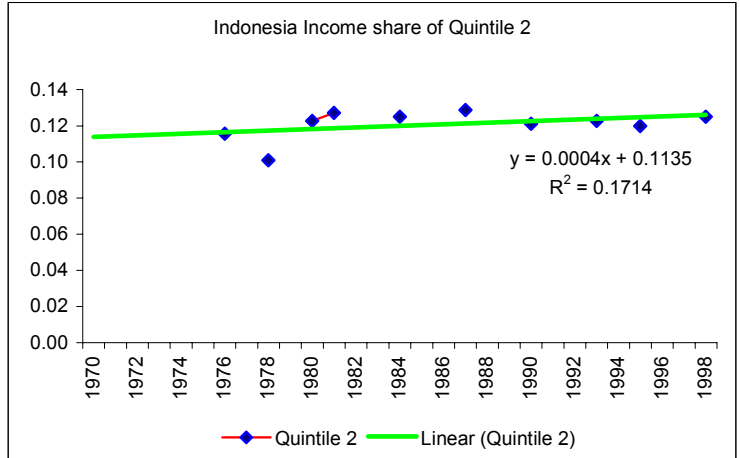
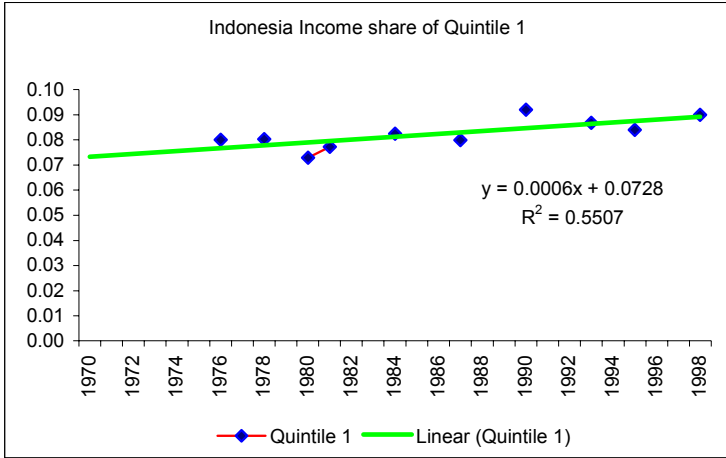
India Income share of Quintile 4




India Income share of Quintile 5








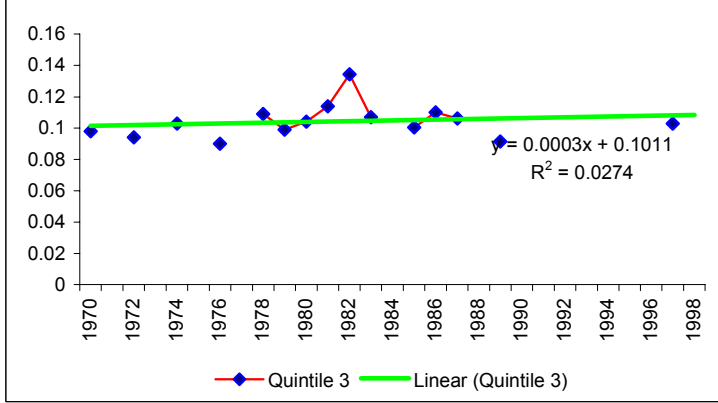
Brazil Income share of Quintile 1




Brazil Income share of Quintile 2




Brazil Income share of Quintile 3




Brazil Income share of Quintile 4




Brazil Income share of Quintile 5




Pakistan Income share of Quintile 1




Pakistan Income share of Quintile 2




Pakistan Income share of Quintile 3



Pakistan Income share of Quintile 4



Pakistan Income share of Quintile 5



Appendix Table 1: Country Listings and Groups

Country Name	Group A	Group B	Group C	1998 Population (Thousands)
Algeria	1	0	0	29922
Australia	1	0	0	18751
Bangladesh	1	0	0	125629
Belgium	1	0	0	10204
Bolivia	1	0	0	7950
Brazil	1	0	0	165874
Canada	1	0	0	30301
Chile	1	0	0	14822
China	1	0	0	1238599
Colombia	1	0	0	40804
Costa Rica	1	0	0	3526
Cote d'Ivoire	1	0	0	14492
Czechoslovakia	1	0	0	15686
Denmark	1	0	0	5301
Dominican Republic	1	0	0	8254
Egypt, Arab Rep.	1	0	0	61401
El Salvador	1	0	0	6058
Finland	1	0	0	5153
France	1	0	0	58847
Germany	1	0	0	82047
Ghana	1	0	0	18460
Greece	1	0	0	10515
Guatemala	1	0	0	10799
Honduras	1	0	0	6156
Hong Kong, China	1	0	0	6687
Hungary	1	0	0	10114
India	1	0	0	979673
Indonesia	1	0	0	203678
Ireland	1	0	0	3705
Italy	1	0	0	57589
Jamaica	1	0	0	2576
Japan	1	0	0	126410
Jordan	1	0	0	4563
Korea, Rep.	1	0	0	46430
Luxembourg	1	0	0	427
Madagascar	1	0	0	14592
Malaysia	1	0	0	22180
Mauritius	1	0	0	1160
Mexico	1	0	0	95846
Morocco	1	0	0	27775
Nepal	1	0	0	22851

Netherlands	1	0	0	15698
New Zealand	1	0	0	3812
Nicaragua	1	0	0	4794
Nigeria	1	0	0	120817
Norway	1	0	0	4432
Pakistan	1	0	0	131582
Panama	1	0	0	2764
Peru	1	0	0	24801
Philippines	1	0	0	75174
Poland	1	0	0	38666
Portugal	1	0	0	9961
Romania	1	0	0	22503
Sierra Leone	1	0	0	4862
Singapore	1	0	0	3164
Spain	1	0	0	39371
Sri Lanka	1	0	0	18778
Sweden	1	0	0	8852
Taiwan	1	0	0	21777
Thailand	1	0	0	61201
Trinidad and Tobago	1	0	0	1285
Tunisia	1	0	0	9335
Turkey	1	0	0	63451
Uganda	1	0	0	20897
United Kingdom	1	0	0	59055
United States	1	0	0	275675
Venezuela	1	0	0	23242
Zambia	1	0	0	9666
Austria	0	1	0	8078
Barbados	0	1	0	265
Botswana	0	1	0	1562
Burkina Faso	0	1	0	10730
Burundi	0	1	0	6548
Central African Republic	0	1	0	3480
Ecuador	0	1	0	12175
Ethiopia	0	1	0	61266
Gabon	0	1	0	1180
Gambia, The	0	1	0	1216
Guinea	0	1	0	7082
Guinea-Bissau	0	1	0	1161
Guyana	0	1	0	849
Israel	0	1	0	5963
Kenya	0	1	0	29295
Lesotho	0	1	0	2058
Mali	0	1	0	10596
Mauritania	0	1	0	2529
Mozambique	0	1	0	16947
Niger	0	1	0	10143
Papua New Guinea	0	1	0	4603
Paraguay	0	1	0	5219
Rwanda	0	1	0	8105

Senegal	0	1	0	9039
South Africa	0	1	0	41402
Switzerland	0	1	0	7106
Tanzania	0	1	0	32128
Uruguay	0	1	0	3289
Zimbabwe	0	1	0	11689
Angola	0	0	1	12070
Argentina	0	0	1	36125
Benin	0	0	1	5948
Cameroon	0	0	1	14303
Cape Verde	0	0	1	416
Comoros	0	0	1	531
Congo, Dem. Rep.	0	0	1	48190
Congo, Rep.	0	0	1	2783
Equatorial Guinea	0	0	1	430
Fiji	0	0	1	788
Iceland	0	0	1	274
Iran, Islamic Rep.	0	0	1	61947
Malawi	0	0	1	10534
Namibia	0	0	1	1662
Seychelles	0	0	1	79
Syrian Arab Republic	0	0	1	15277
Togo	0	0	1	4458
St. Vincent	0	0	1	113
Sao Tome e Principe	0	0	1	143
Chad	0	0	1	7283
Haiti	0	0	1	7647
St. Kitts & Nevis	0	0	1	41
Sta. Lucia	0	0	1	152
Cyprus	0	0	1	758
Grenada	0	0	1	96
Dominica	0	0	1	73
Belize	0	0	1	239
Antigua	0	0	1	68