
Inequality as We Know It

Inequality is a subject of inquiry for every brand of social scientist—philosophers, economists, sociologists, psychiatrists, and even politicians. It is unclear what the net value added of all the opinion is. On one related issue, however, there is agreement: What is most important, both in moral and in practical terms, is that *equality of opportunity* is a goal worthy of pursuit.

The dichotomy, both conceptually and for policy, between concern with inequality and concern with equality of opportunity is stark and obvious. There is very little one can, or should do, with either inequality or trends in inequality.¹ The reason one should do little is partly on practical grounds; moralists have devised ingenious ways to tax the rich so that inequality is altered, and ingeniously have the wealthy avoided taxation (including bribery, of course). The other objection is philosophical: Reducing inequality implies knowledge of a preferred inequality *level* and a preferred dynamic path. This may seem laughable now, but the popularity of central planning and *loco parentis* (the state knows better) in the 1960s, 1970s, and 1980s was due to the presumed power of the state to achieve whatever goal it desired. This brings us back to the original problem: if one wanted to change inequality levels, and inequality trends, what policies would one implement? What confidence would one have in the goals? And do the unequal even care?

The concern with inequality is not new; indeed, it has commanded the attention of social scientists for several centuries. Philosopher Immanuel

1. Although facts are stressed throughout this book, this is much more in the nature of opinion.

Kant (as quoted in Dahrendorf 1968) called “inequality among men a rich source of much that is evil, but also of everything that is good.” But the sociologist Dahrendorf admits to failure in making any progress in doing anything about inequality:

So far, however, as far as the problem of inequality itself is concerned, this history [of sociological research] has achieved little more than given it a different name: what was called in the eighteenth century the origin of inequality and in the nineteenth the formation of classes, we describe today as the theory of social stratification—all this even though the original problem has not changed and no satisfactory solution to it has been found. (1968, 152)

Over the years, several prominent economists (e.g., Adam Smith, Friedrich von Hayek) have agreed with the observation that there is no easy solution to be found; indeed, given its confusing complexity, perhaps no solution need be found. Marx thought he had found the solution: perfect equality. History has judged that notion to be extremely problematical, and thrown out the sage, along with his followers. Nowadays, economists generally tend to be more concerned with the growth in the pie than in its distribution; for many of them, the economics of inequality per se is not much different from the “economics of envy.”

Rawls (1958), in his classic essay “Justice as Fairness” (later published in an expanded form as *A Theory of Justice*, 1971), outlined the *maximin* principle of justice (and equality). This principle argued that the concept of justice dictated that policies that improved the position of the least advantaged member of society should be preferred over all others. (Obviously, it would follow that among competing policies, those that improved the lot of the least advantaged more would be preferred.) In 1971, I criticized this principle on the grounds that it ignored the role of inequality, and inequality trends, in a person’s utility function, something that was likely to be prominent among the citizens of industrialized economies (Bhalla 1971).

At a minimum, therefore, it is a moot question whether inequality per se, at the world or country level, is important. No doubt it is important for politicians because populism is almost exclusively based on the economics of envy. But is it important for policy? Do poor people really care about inequality or do they care more, much more, about absolute standards of living? Very likely, concern for inequality is a highly income-elastic superior good—the more one becomes richer, the more one is (naturally) concerned with the Joneses and the Veblins who have more.

The fact remains that inequality, and its study, remains an intellectually fascinating and emotive issue. One of the most profound insights into what actually happens to inequality within a society was provided by Simon Kuznets in the mid-1950s. In one of the most justifiably celebrated articles, Kuznets postulated that inequality charts out an inverted U-curve with development; that is, as economies develop, inequality first increases

(the left leg of the inverted U) and then declines. Research over the years, which is documented extensively in this chapter, supports this result. However, it should be emphasized that the Kuznets inverted U-curve applies to what happens *within* countries over time; what happens to inequality among the collection of individuals in the world, over time, is quite another story. Indeed, it is theoretically possible (and empirically true!) that each country can experience a worsening of inequality and yet overall improving inequality—a possibility, and finding, missed by most research to date.

Measuring Inequality

This chapter discusses issues relating to the measurement of inequality, and it introduces such concepts as the Lorenz curve and the inequality Gini index. After measurement is out of the way, the chapter goes on to discuss perhaps the most important development in the study of the economics of inequality during the past 50 years—the Kuznets inverted U-curve, which relates the path of inequality to the path of development.

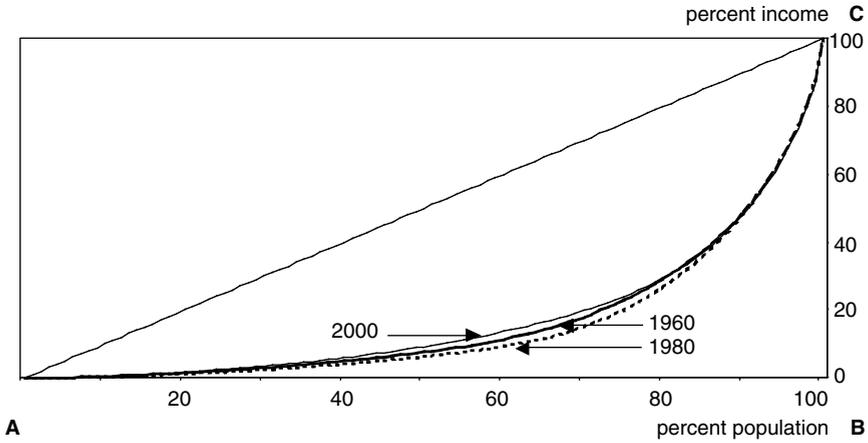
This inverted U-curve can be tested in a variety of ways. The chapter reviews the results of the various tests and reaches a threefold conclusion: During short periods of time (a decade or so), inequality does have a tendency to worsen. During long periods, it appears that the level of inequality on the right side of the inverted U is approximately equal to that on the left. And data for the past 20 years indicate that, on average, country inequality has definitely worsened.

The study of income distribution is a study of shares in the pie—who has what amount. The best-known of the descriptions of the pattern of income distribution is the Lorenz curve: If the population is ordered according to income, the Lorenz curve plots the cumulative share in income (*y*-axis) versus the cumulative share in the population (*x*-axis).

By definition, therefore, the Lorenz curve starts at zero and ends at unity (figure 3.1 plots such a curve for three different years—1960, 1980, and 2000). If incomes are equally distributed, then the Lorenz curve is described by the diagonal (i.e., the share of income at each point is equal to the share of population). If all the income in society accrued to just one person, then the Lorenz curve would be the zero line (A to B) and then at the end the share would “spike” up to 1 (B to C). The Gini coefficient² is equal to 1 in the latter case (perfect inequality) and 0 in the former (perfect equality). For the curve shown in the figure (world distribution in 2000), the Gini coefficient is 0.651 (or 65.1).

2. The Gini coefficient is the “area between the Lorenz curve and the diagonal, relative to the whole triangle below the diagonal, or half the mean difference relative to the mean” (Atkinson and Brandolini 1999, 6).

Figure 3.1 World Lorenz curve, 1960, 1980, and 2000



Note: World income inequality is computed by the simple accounting procedure (SAP) method of aggregating individual country data on distributions of incomes into a world distribution of individuals; see chapter 11 for the details.

Sources: Deininger and Squire (1996); World Income Inequality Database, available at <http://www.wider.unu.edu/wiid>; Asian Development Bank (2002); Milanovic and Yitzhaki (2001).

Measuring Intertemporal Inequality

Along with the Gini, researchers have used other definitions of inequality—the share in incomes of a particular quintile (e.g., the bottom 20 percent) or the ratio in incomes of different groups (e.g., the ratio in incomes of the top vs. the bottom decile). This logic has even been extended to the study of incomes of the richest country relative to the poorest country, as we saw in the previous chapter. Decile or quintile shares are descriptions or summary statistics for different *parts* of the Lorenz curve.

The trends in these components can be used to reveal whether the incomes of a particular quintile are growing at a faster or slower rate than the average. A faster rate implies that the *share* of that group in total income is increasing, a slower rate that the *share* is decreasing. This is definitional and studies of inequality that look at changes in shares are *equivalent* to studies of inequality that look at relative rates of growth in incomes. In economic terms, the much talked about *unitary* elasticity of incomes of poor people with respect to average income (“elasticity of connection,” according to Timmer 1997) is the same as a *zero* elasticity of the share of income with respect to income growth.

Simple Inequality Mathematics

Country data are often presented in the form of quintile and mean incomes. Knowledge of the mean income at each point in time allows

one to compute absolute levels, and changes. The *absolute* income at any point in time is revealed by a simple formula:

$$(U*N)/(Y*100) = X \tag{3.1}$$

where N represents the number of percentiles, U the mean income of these percentiles, Y the mean income of the population, and X the share in income of U percentiles. (A population has 100 percentiles, and each quintile has 20.)

For an N with a value of 20, or the bottom 20 percent, the equation becomes

$$\begin{aligned} U*20 &= X*Y*100, \text{ or} \\ U &= (100/N)*X*Y \\ U &= k*X*Y \end{aligned} \tag{3.2}$$

In log terms, the equation becomes

$$\log(U) = \log(k) + \log(X) + \log(Y)$$

In (log) rate-of-growth terms, the above equation reduces to

$$u = y + x \tag{3.3}$$

where lowercase u represents the rate of growth in incomes for a particular quintile (or decile or percentile), y is the rate of growth in average income of the population, and x is the growth in share in income of this particular classification (decile, percentile, etc.). Equation 3.3 is important to appreciate because it underpins a considerable part of the simple accounting that is needed to talk about inequality changes. The equation is an identity—it must hold, and no assumptions are required.

The simple computations involved in equation 3.3 allow one to interpret, at a glance, whether the incomes of the poorest quintile have grown at a faster (increase in share) or slower (decrease in share) rate than the richest quintile or the middle quintile, or whatever. Note that the coefficients of both y and x are unity. In other words, *if there is no change in inequality (shares are constant, or x is zero) then the rate of growth of incomes in the particular quintile is equal to the growth in average incomes.*

Data for China can be used to illustrate the ease, and power, of equation 3.3. In 1985, the share of the bottom quintile was 7.7 percent, which declined to 5.6 percent in 1995. This decline is a (log) change of -16.4 percent. The average per capita income during the same period rose from purchasing power parity (PPP) \$1.95 (at 1993 prices) to PPP \$3.62, or a 61.8 percent (log) increase. According to equation 3.3, the change in income of the first quintile, u , is equal to the change in average income (variable y),

plus the change in inequality (variable x). The value of y is 61.8; that of x is -16.4 . Hence, the increase in income of the first quintile is equal to $61.8 + (-16.4)$, or 45.4 percent.

Armed with this simple inequality mathematics, one can begin to analyze various relationships. The decline in share of income of the first quintile is analogous (equivalent) to lower relative growth (45.4 vs. 61.8 percent); this decline is also analogous (equivalent) to the elasticity of incomes of the group (defined as poor people or the first quintile) being lower than 1; in this case, the elasticity is equal to 0.73 (relative growth rates, or 45.4 divided by 61.8). Because the growth in incomes of poor people was lower than average growth, we can state with some basis that growth was antipoor; in other words, the term “propoor growth” is nothing more than a restatement of the fact that the incomes of poor people rose at a lower rate than average income.³

Inequality: Kuznets Curve and Data Requirements

On the basis of the data available, Kuznets observed that rich economies displayed lower inequality than poor countries. If this observation is combined with assumptions about rural-to-urban transformation—with the smaller urban sector growing faster (due to industrialization) than the larger rural sector⁴—one would logically observe inequality increasing, and upon “maturation” of the economy, observe decreasing inequality. Hence, the inverted U-shaped pattern.

When Kuznets posited his hypothesis about income distribution worsening with development, he had very little data on income inequality in developing countries—countries for which his hypothesis had relevance. As a part of its drive to study inequality and poverty in detail, the World Bank launched its Redistribution with Growth project in the mid-1970s. This project consisted of compiling existing data, analyzing these data (among the earliest efforts in this regard are the pioneering papers by Ahluwalia and Chenery and their associates: Ahluwalia 1974, 1976; Ahluwalia, Carter, and Chenery 1979), and encouraging, financially and otherwise, developing countries to conduct household surveys. The dataset used by Ahluwalia (1976) was published in a book by Jain (1975). This

3. Actually, as we will see in chapter 10, the correct derivation of whether growth has been propoor or not is more, much more, than observing this simple relationship. Indeed, the elegance and the “obviousness” of the simple relationship has prevented researchers from correctly identifying whether a particular process of growth was propoor or not.

4. This pattern of growth has also been argued by Lewis (1954); an elegant mathematical formulation, and simulation, of the inverted U-curve hypothesis is contained in Robinson (1976).

compilation of data, warts and all, was to be a major input for the analysis, and debate, over poverty, inequality, and growth. This dataset had expanded the earlier collection by Adelman and Morris (1973) and Paukert (1973).

Ahluwalia put the Jain inequality data through several tests and concluded thus:

There is strong support for the proposition that relative inequality increases substantially in the early stages of development, with a reversal of this tendency in the later stages. This proposition holds whether we restrict the sample to developing countries or expand it to include developed and socialist countries. Furthermore, it appears that the process is most prolonged for the poorest group . . . the cross section results do not support the stronger hypothesis that the deterioration in relative inequality reflects a prolonged absolute impoverishment of large sections of the population in the course of development. The cross country pattern shows average absolute incomes of the lower percentile groups rising as per capita GNP rises, although slower than for upper income groups. (1976, 338)

Ahluwalia's paper was the first extensive study of inequality and development. This seminal effort set off a flurry of research, in both data collection and replication. The conclusions also appear to have stood the test of time, though some not as strongly as first identified. The "prolonged" experience, when it has occurred, has not necessarily been particular to the poorest group; and the lower percentile groups have not necessarily had a lower elasticity of growth, as Ahluwalia's own concurrent research (1974) had shown.

It is important to note what this important cross-sectional result (the inverted U-curve) does and does not indicate. It does not indicate that worsening inequality will happen. It does indicate that developing countries (at lower levels of income or development) were *observed* to have lower shares in income of the first quintile. If it is now assumed that developing countries would follow, *over time*, the pattern of income shares observed at a *point in time*, then and only then can one conclude that the Kuznets curve is predictive.⁵

The Ahluwalia data, and results, were subjected to intense scrutiny and analysis by Anand and Kanbur (1993a, 1993b), who demonstrated that data inaccuracies and misspecifications (e.g., mixing household data with individual data, and consumption surveys with income surveys) may have had a considerable effect on the finding of the inverted-U-curve result. Further, if these data were corrected for such measurement mistakes, instead of an inverted U-curve, a normally shaped U-curve was found; that is, developing countries tended first to show a decline in

5. The brilliance of Kuznets's insight is underlined by the fact that his observation about the inverted U was based on very limited data. Starting with Ahluwalia (1976), there have been scores of important studies on the Kuznets curve. The conclusion 40 years later and with tons of new data is that it is one of the most remarkably accurate "forecasts" in economics.

inequality before a subsequent increase. Anand and Kanbur were silent on what could be generating this anomalous finding, because the theory about inequality first worsening due to the rural urban transformation seemed to be intuitive and plausible.

The increased availability of distribution data also led to an analysis of “spells” of experience (i.e., noting what happened to inequality and levels of development for the *same* country over time).⁶ The conclusions of these studies were mixed, but they broadly supported the proposition that inequality stayed relatively constant, within a country, over time (i.e., they rejected the Kuznets curve).

Thus, by the early 1990s, the received wisdom on country inequality and changes therein was that inequality levels in developing countries had stayed relatively constant; some countries had improved, others had deteriorated, but there was no major story either way. Anand and Kanbur’s conclusions, while technically correct, were more in the nature of a statistical, not economic, result. The reality was of a moderately Kuznet-sian variety. Further, whereas data availability and documentation had improved considerably since the mid-1970s, there was still a relative paucity of data for some of the poorest countries in the world, especially those in sub-Saharan Africa.⁷

There are three other stylized facts pertaining to income distribution. First, consumption distributions tend to be more equally distributed than do income distributions, often by about 10 to 15 percent. This result can be verified by observing the data for the same country for the same point in time. Both Berry, Bourguignon, and Morrisson (1981) and Deininger and Squire (1996) suggest that the consumption Gini is about 6 points lower than the income Gini—a result also corroborated by our data. Second, after-tax distributions tend to be more equal than before-tax distributions. Third, inequality appears to be region specific. The industrialized world and Asia show low levels of inequality, and sub-Saharan Africa and Latin America show very high (almost “double” the Gini level of Asia) levels of inequality.

Recent Evidence on Country Inequality

For the study of trends in inequality, researchers have used essentially two different methods—either a study of the change in shares of income of different quintiles (and/or changes in the Gini), or the analysis of the elasticity of the mean income of a particular income group (e.g., the first

6. See Fields (1980); World Bank, *World Development Report 1990: Poverty and Development*.

7. E.g., by the end of the 1970s, results for only three sub-Saharan African economies were available, compared with more than 16 surveys for the five major countries in South Asia.

quintile) with respect to average income in the economy. As was discussed in the context of the equations above, the two yield identical results: a zero change in shares is equivalent to a unitary elasticity of quintile income.

Underlying all of this analysis is the Kuznets framework (i.e., how is inequality related to development). The variable of interest is an *index* of inequality, for example, the Gini, or the share of income of a particular quintile, or the ratio of incomes of two deciles (the income of the top 10 percent relative to the income of the bottom 10 percent, etc.), and the relationship of the chosen economic variable to economic growth.

The first important post-Ahluwalia and post-Anand and Kanbur result about inequality trends was provided by Deininger and Squire (1996). They assembled the second large inequality dataset (the first was that of Ahluwalia and Jain), subjected it to various tests, and found that inequality does not change much *within* countries over time. To be sure, there are problems with the data,⁸ and there are exceptions to the rule; but the fundamental trend was one of little change.

This new dataset was analyzed further by Li, Squire, and Zou (1998), who ran several regressions with the Gini coefficient as the dependent variable. Their conclusion was the same: income inequality was relatively stable within countries over time. *In other words, there was no support for the Kuznets hypothesis.*

Cornia and Kiiski (2001), using a third compilation of data (the World Institute for Development Economics Research, or WIDER, dataset⁹) on inequality, conducted detailed tests of inequality changes within countries, and particularly the analysis of changes during the globalization period. Their conclusion: "Over the last twenty years, this trend towards equality was halted in parallel with the emergence, consolidation and diffusion of a new economic paradigm often referred to as the 'Washington Consensus'."

Cornia and Kiiski's analysis was in the spirit of "spell analysis," that is, looking at individual countries' changes in inequality. Using the larger WIDER dataset (which includes 73 countries with spells, as compared with the 49 countries with spells used by Li et al.), Cornia and Kiiski

8. The Deininger and Squire (1996) dataset, along with a corresponding assembling of data by the United Nations University's World Institute for Development Economics Research (the World Income Inequality Database, available at <http://www.wider.unu.edu/wiid>), and the data collected by the Asian Development Bank (RETA-5917, Research Project on Developing a Poverty Database), forms the core of the income distribution data used for this study. There are definite problems with the non-Asian Development Bank distribution data, and better documentation and "cleaning" would be warranted. E.g., for some observations, the quintile shares are inconsistent (i.e., a poorer quintile being shown to have a larger share of income than a richer quintile, a definitionally impossible situation). Atkinson-Brandolini (1999) point out several problems with these datasets.

9. This dataset expands on the observations contained in the Deininger and Squire (1996) dataset.

reveal that the conclusion of Li and his colleagues of “no change in inequality” was somewhat inaccurate; that there are trends not caught by the simple model of Li and his colleagues, but captured by their non-linear techniques. These techniques suggest that within-country inequality declined between 1950 and 1970, and increased thereafter. However, during a much longer period (1960-2000), there was little change in intracountry inequality, thus reaffirming the conclusion of Deininger and Squire and of Li and his colleagues.

The results are indicative of one simple fact: When the hurly-burly is done, little change in inequality is observed in most countries—the ones that started off unequal stayed unequal, and the ones that started equal stayed equal.¹⁰ Of course, some striking changes are observed (e.g., the large increase in inequality in China), and after a century of constancy, a large increase in inequality in the United States in the late 1980s. But these are exceptions. The norm is small long-term changes in country inequality.

The Evidence Once Again

The method developed to analyze inequality at an individual level, through construction of Lorenz curves, is outlined in appendix A. Construction of the dataset required the use of data for more than 1,000 household surveys and is possibly the most exhaustive dataset of its kind. This dataset covers, for most countries, all the years between 1950 and 2000 for four major variables: population, real PPP income, real PPP consumption, and income (and consumption) distribution.

Does this new dataset suggest any “new” results on inequality levels and trends? Let us reestimate the relationship between inequality and development. Table 3.1 reports the results for income distribution data for selected countries and regions. Three years are selected—1960, 1980, and 2000.¹¹ The results, according to the Gini inequality index, is that for a large set of countries inequality *has* worsened over time. The worsening inequality is particularly noticeable for the United Kingdom, United States, China, Nigeria, and Russia. The initially highly unequal economies like Brazil and Mexico buck the trend, and show significant improvement in inequality. This “equality convergence” pattern has not gone unnoticed, and has been commented upon by Benabou (1996), Roemer and Gugerty (1997), and Ravallion (2001b).

10. Or as T.S. Eliot might have said: “We arrive where we started, and know inequality for the first time.”

11. If a survey was not conducted in a particular year, data for the nearest year prior to the year in question were used. If there was no survey earlier, then the result of the first such survey subsequent to the year in question was used.

Table 3.1 Country and regional inequality; Gini and quintile income shares

Region or country	Income shares (percent) ^b								
	Gini ^a			Share of 1st quintile			Share of 5th quintile		
	1960	1980	2000	1960	1980	2000	1960	1980	2000
Industrialized world									
United States	36.3	39.7	45.7	4.8	4.3	3.6	41.3	43.6	49.4
United Kingdom	27.7	27.4	36.2	9.4	10.2	6.6	36.9	37.6	42.9
Australia	32.0	39.3	35.4	6.6	4.6	5.9	38.8	43.4	41.3
Germany	27.3	32.1	30.1	10.5	6.6	8.2	37.7	39.0	38.4
Sweden	40.0	32.7	25.1	4.4	7.0	9.6	44.0	39.5	34.6
Developing world									
Bangladesh	37.4	35.3	45.1	6.9	7.4	6.2	44.5	42.9	52.1
Brazil	49.7	57.9	52.1	3.8	2.9	3.5	54.0	61.6	56.2
China	29.5	29.5	40.5	7.9	7.9	5.9	36.7	36.7	46.6
South Korea	35.5	40.0	31.7	5.8	5.1	7.5	41.9	45.4	39.3
Mexico	56.0	56.2	51.8	4.4	3.1	4.0	61.4	60.2	56.6
Malaysia	37.4	51.3	49.3	6.5	3.7	4.4	44.0	55.8	54.3
Indonesia	40.2	40.5	37.6	6.5	5.9	7.3	47.0	46.8	45.5
Russia	24.7	24.7	35.9	9.5	9.5	6.7	34.0	34.0	42.4
Nigeria	44.7	44.7	56.9	4.7	4.7	3.3	49.6	49.6	61.2
Industrialized world				6.1	6.4	6.8	42.1	40.5	42.2
Developing world				6.7	6.5	5.7	45.2	45.8	49.2
World				6.8	6.7	5.9	43.7	44.1	47.7

a. Gini is computed by the simple accounting procedure method and not “published” estimates because there are problems with some of these Gini estimates; see chapter 8.

b. In many cases, income distribution data are not available for 1960, 1980, or 2000. In such cases, the table presents either the closest earlier year for which data are available, or, where earlier data are not available, data for the earliest later year. E.g., if the latest survey took place in 1995, the 2000 figures reflect these values; if the first survey took place in 1975, the 1960 figures reflect those values.

Sources: Deininger and Squire (1996); World Income Inequality Database, available at <http://www.wider.unu.edu/wiid>; Asian Development Bank (2002); Milanovic and Yitzhaki (2001).

Not reported in table 3.1 is the pattern of distribution for consumption. The pattern is the same, except that four economies with relatively large populations, all in Asia, grew at an accelerated pace during the globalization period and yet did not show any trend in consumption inequality: India, Indonesia, South Korea, and Vietnam. On a probability basis, all these economies are due for a worsening; when and whether it will happen, and why inequality has not worsened for 20 years, is a subject for research and perhaps another book!

The result for trends in regional inequality as measured by quintile shares is consistent with the result observed for individual countries. Among the regions, the industrialized countries show a decrease in

inequality (see the figures for the first quintile) from 1960 to 2000, despite the fact that such countries as the United States and United Kingdom significantly worsen.¹² However, the fifth decile also gains, suggesting that the Lorenz curves are “crossing.” The developing world shows a similar pattern—little change in inequality from 1960 to 1980, and a large increase in inequality from 1980 to 2000. Note that the above results hold regardless of which measure of inequality is used—share of the first quintile or share of the fifth quintile.

The final “screening” calculations on the data are reported in figure 3.2. The averages for the annual (log) change in inequality income and consumption distribution changes are reported for three classifications of data and two time periods: all changes experienced during the very short term (less than or equal to 3 years), medium-term changes (between 4 and 8 years), and long-term changes (9 to 15 years). For the 1960-80 period, there seems to be a definite trend toward equality, and the change is larger, and more positive, the larger the time frame of observation. For the next 20 years, there is a complete reversal. Now there is a trend toward greater inequality, and the magnitude is again larger the longer the period of observation.

Thus, the conclusion that *country* inequality worsened during the past 20 years appears to be fairly robust. These results directly contradict the findings of Deininger and Squire and Li and his colleagues and are more along the lines of the results reached by Cornia and Kiiski.

Tests of Inequality Change

The large number of countries for which data on distributions exist for more than 1 year allows the *direct* estimation of the time profile of inequality. Individual quintile income levels (derived from a multiplication of quintile shares with mean income) can be regressed on mean income and the “elasticity of connection” observed. If the elasticity is equal to 1, it follows that inequality did not change for a typical country.

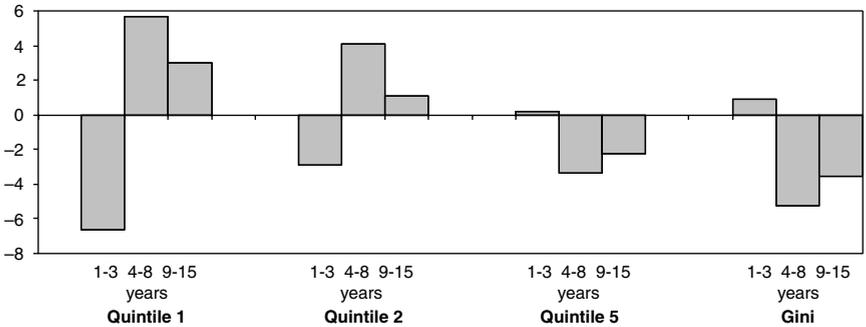
This 45-degree method (so named because the result obtained is a 45-degree line between average income of the first quintile and average income in the economy), which is quite the rage today, was first used by Ahluwalia in 1974. He regressed the growth rate of income of the lowest 40 percent on the growth rate of GNP for 13 developing countries, and he failed to find any strong pattern; that is, he found a unitary elasticity between the growth in incomes of poor people (defined as the bottom 40 percent of the developing world) and the growth in average incomes of

12. Regional figures for the first and fifth quintiles are obtained by weighting the quintile shares of the countries by their population in the region. This cannot be done with the Gini; for what happened to regional Ginis, see chapter 11.

Figure 3.2 The time dimension of world inequality change, 1960-2000

1960-80

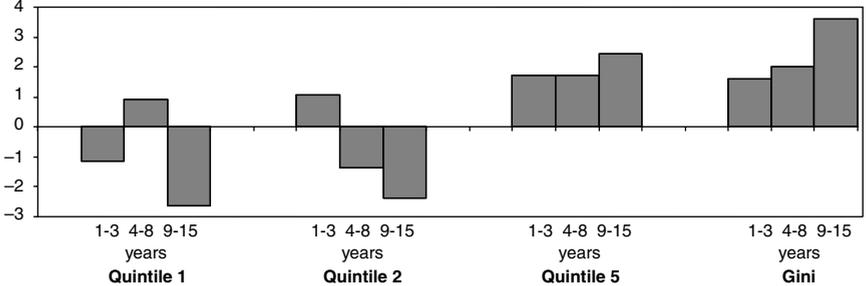
percent (log) / change^a



measure and time frame

1980-2000

percent (log) / change^b



measure and time frame

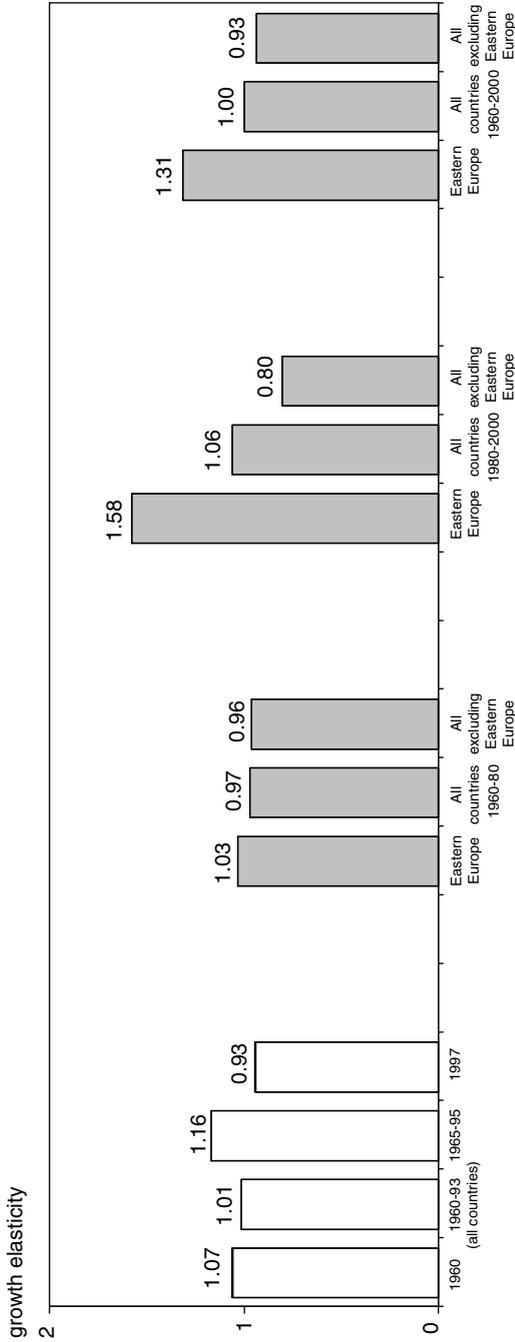
a. Percent changes, all in log terms, reflect the *total* changes over each period, rather than annual changes.

Note: For each measure of inequality—income/consumption shares of quintile 1 (Q1), Q2, and Q5, and income/consumption Gini—(log) changes are calculated over various time frames with reference to each year—(t, t-1), (t, t-2), (t, t-3) . . . (t, t-15). The mean is then calculated, for each year, for sets of changes, i.e., the 1-year, 2-year, and 3-year changes are one set; 4-year to 8-year changes are another set; and 9-year to 15-year changes are a third set. (The shorter (1-3 years) and longer (9-15 years) sets give, respectively, the short- and long-term trends in inequality.)

Once obtained, these figures are averaged for the two time periods, 1960-80 and 1980-2000, to see whether there has been a shift in the *direction* of inequality change between the preglobalization (1960-80) and globalization (1980-2000) periods. In most instances, the results are dramatic and conclusive.

Sources: Deininger and Squire (1996); World Income Inequality Database, available at <http://www.wider.unu.edu/wiid>; Asian Development Bank (2002); Milanovic and Yitzhaki (2001).

Figure 3.3 Elasticity of incomes of poor people (bottom 20 percent) with respect to average incomes



Note: Income elasticities are computed by regressing the (log) average income of the 1st quintile as a function of the (log) average income, at 1993 purchasing power parity prices. Differences in these estimates are due to differences in the choice of countries and time periods, as well as to different methods of estimation. The simple accounting procedure results are for the countries in the nonindustrialized world. The estimation method incorporates fixed effects with robust standard errors.

Sources: *Unshaded bars*: All countries, 1960-93; Roemer and Gugerty (1997); 1965-95: Gallup, Radelet, and Warner (1998); 1997: Timmer (1997); 1960: Dollar and Kraay (2000). *Shaded bars*: Simple accounting procedure method.

the entire economy: the predicted line was one of 45 degrees (Ahluwalia 1974).

Three studies undertaken at the Harvard Institute of International Development (HIID) in 1997 and 1998 (Roemer and Gugerty 1997; Timmer 1997; Gallup, Radelet, and Warner 1998) used the newly available Deininger and Squire dataset to update, and methodologically improve upon, Ahluwalia's regressions. Except for one regression by Timmer, which used dummy variables for each decade of analysis, all the regressions reported in the HIID papers show the first quintile elasticity to be not significantly different from unity. The period covered by the HIID studies was mostly the years after 1960, and the studies used data for approximately 30 to 60 countries (depending on the author and analysis).

A few years later, Dollar and Kraay (2000) subjected the same distribution data to additional tests and reached the same conclusion as the HIID studies. They found that "growth was good for the poor" because poor people shared equally in the fruits of average growth (i.e., the 45-degree line).

Thus, a strong result emerging from these log-log studies is that the average country has not witnessed any change in inequality during the past 40 years—a "flat" U-curve. The bottom 20 percent of individuals in a typical country experienced the same growth, or fall, in incomes as the average person. In other words, the aggregate average experience is one of *rejection* of the Kuznets hypothesis—a contrast to Cornia and Kiiski's result and our preliminary result reported in table 3.1. For example, the share of the first quintile for the world had declined from 6.7 to 5.9 percent for the developing world, or an implied elasticity of 0.87.

This important unitary elasticity or no-Kuznets-curve result has been interpreted differently by different authors—some legitimate interpretations, some not so. The correct explanation of a "flat" U-curve is that the data reveal that a "typical" country will experience no change in inequality. The unit of observation is a country-year; and if it is assumed that all the observations are drawn from the same "structural universe," there is no reason to worry that the above regression gives equal weights to a small country showing no change in inequality and to a large country showing a large change in inequality.

The *incorrect* interpretations start with the assumption that a regression of the bottom 20 percent containing all the world's countries says anything about absolute poverty. The titles of all four major papers on the subject (Roemer and Gugerty 1997; Timmer 1997; Gallup, Radelet, and Warner 1998; Dollar and Kraay 2000) contain the words "growth" and "the poor," even though the papers' texts contain the necessary disclaimers and statements about the differences between relative poverty (incomes of the first quintile) and absolute poverty (incomes below an absolute level of income).

It would be correct to assume that the regressions reflect changes in absolute poverty if the sample of countries chosen contained *only* countries with large populations of poor people in the base year. The later papers (Roemer and Gugerty 1997; Dollar and Kraay 2000) use *both* industrialized and developing-country data to rerun Ahluwalia's regressions. Though it must be admitted that it is good to have more data, their regressions no longer retain the connection with the incomes of poor people. The reason Ahluwalia chose the bottom 40 percent was because the poverty line was *defined* as the bottom 40 percent of developing economies. It is a case of the methodological trees being copied without paying due attention to the definitional forest.

Elasticity of Connection, 1960-2000

Figure 3.3 reports on the estimates of the "elasticity of connection" for the bottom quintile found by various studies. The model that had been developed by Ahluwalia and the HIID researchers was reestimated using data for 617 household surveys in nonindustrialized countries, and the unitary elasticity result appears to be grossly in error. For the period 1980-2000, the elasticity of connection for the first quintile is not 1; it is 0.80, with a robust standard error of 0.1.¹³

It is not the newly available data that lead to a rejection of the original robust conclusion of a unitary elasticity. Nor did the results reached before the mid-1980s contain an error (e.g., Ahluwalia 1976; Fields 1980; World Bank, *World Development Report 1990: Poverty and Development*). But studies after them do have an error. And the error is a simple one: the nonrecognition of the reality, and the structurally different experience, of Eastern European countries and countries that had been part of the USSR. Though Eastwood and Lipton (2001) (an excellent survey of the issues related to poverty, inequality, and growth) do not reestimate the Dollar-Kraay regressions, they broadly hint at the possibility that the presence of the Eastern European economies was biasing Dollar and Kraay's results strongly in favor of a unitary elasticity.¹⁴

In an unusual development (in the sense of historical patterns), the Eastern European economies simultaneously experienced *declining* equal-

13. The log-log model was estimated with fixed effects; incorporation of industrialized-country (nonpoor) data raises the elasticity to 0.83 with a standard error of 0.098.

14. Dollar and Kraay run several regressions and report one set of results, which incorporate regional dummies (2001, table 4, 39). Their basic regional dummy specification yields an average elasticity of 0.91 with a standard error of 0.09. Reproducing their results with their data but our specification and estimation method (fixed-effects model with robust standard errors, and instead of regional dummies all Eastern European countries excluded from the sample for the period 1980-2000), an elasticity of 0.84 with a standard error of 0.116 is obtained.

ity and *declining* incomes; this joint occurrence results in the elasticity of connection being close to 1.6 for the 1980-2000 period. During this period, a 1 percent decline in average incomes led to a 1.6 percent decline in the incomes of the bottom 20 percent!

This simple result for Eastern Europe demonstrates, without the use of fancy econometrics,¹⁵ that Dollar and Kraay's conclusion—"incomes of the poor do not fall more than proportionately during economic crises"—is invalid. Note the elasticity for Eastern Europe: 1.6. Ordinarily, one would say that this economy was becoming most equal—a 1 percent increase in average incomes is resulting in a 1.6 percent increase in the incomes of poor people. But what is actually happening is a fast jump toward inequality—a 1 percent *decline* in *average* incomes is leading to a 1.6 percent *decline* in the mean incomes of poor people.

This large 1.6 elasticity has just the opposite consequence of the elasticity that would have been observed if the average income had risen, not fallen. In this instance, the incomes of poor people are falling considerably faster than the incomes of nonpoor people. Negative growth has been worse for poor than nonpoor people. Now, when this negative-negative elasticity of 1.6 (two negatives make a positive) is combined with a world excluding Eastern Europe (positive-positive) elasticity of 0.8, one obtains the controversial,¹⁶ and wrong, finding of a unitary elasticity for the average country (a finding reached by Roemer and Gugerty 1997; Timmer 1997; Gallup, Radelet, and Warner 1998; and Dollar and Kraay 2000).

Using virtually identical data to those used by the HIID researchers and Dollar and Kraay, an elasticity that is significantly lower than unity is obtained. This "true" elasticity for all the world's countries, excluding those in Eastern Europe, is about 0.8, with a standard error of 0.10 (i.e., significantly less than 1). For the previous 20-year period (1960-80), the elasticity was 0.96, with a standard error of 0.12. The conventional wisdom—that inequality has essentially stayed constant for long periods—is incorrect. It was constant during the 1960-80 period, but it has clearly not been constant since then.

Country-Level Inequality

Five conclusions follow from this detailed analysis of country inequality. First, over long periods of time, inequality at the country level does not

15. Dollar and Kraay run several models, and use sophisticated econometrics, to justify their conclusions, something approvingly noted by others, e.g., Srinivasan (2001), who concludes in a review of several such studies that "in many ways, the most careful (economically) study is that of Dollar and Kraay" (20).

16. There was a lot of objection to Dollar and Kraay's finding that the poor shared equally in the growth process. If this objection occurred due to the conceptual error of not accounting for the fact that Eastern Europe contained the experience of declining incomes and worsening

display any significant change either way. Second, during the past few decades, some countries have registered sharp changes toward inequality; the notable instances are the United States, China, the United Kingdom, and those in Eastern Europe. (The change in the case of Eastern Europe is unprecedented.)

Third, there have been equally sharp movements toward equality in some countries (e.g., Sweden). Fourth, many countries have remained at the same level of inequality during the past decades, like India and Indonesia. Fifth, *all possible formulations of intercountry inequality suggest that inequality worsened during the 1980s and 1990s*. This last result contradicts most of the “no-change” results reported to date, but it is in accordance with Cornia and Kiiski’s findings. Whether this increase in *intercountry* inequality leads to an increase in world *individual* inequality is the subject of chapter 11.

Toward Individual Inequality Estimates

If country data are aggregated (assume each country is an individual), and population weights are used, then a crude, first-cut profile of a world distribution could be obtained. With such a distribution, one would just be able to read off what happened to quintile shares, Ginis, or any other index of inequality in a given year. One would also be able to interpret the trends.

Korzeniewicz and Moran (1997) constructed a distribution for each country that uses a quintile as a unit of observation (i.e., the average income of each quintile is attributed to all residents in the quintile). In the case of a country with a population of 10 million, this method assumes that each segment of 2 million has the same income. For India, it means each 200-million segment; for China, each 250-million one.

Though useful, such an approach may yield more wrong than right results. Why? Because a lot is going on within each country, especially large countries. If a detailed distribution for each country is not available (as is most often the case), the next-best procedure is to parameterize the distribution. Schultz (1998) converts each country’s distribution into a log-variance model; this yields a mean and a standard deviation, parameters that can yield an estimate of mean income at *any* level of aggregation.

The problem with this approach is that it is prone to large errors in the approximation. Schultz finds that PPP-measured inequality has been relatively stable throughout the past 20 years; indeed, such inequality has had a slight tendency to decline, a result echoed by Melchior, Telle, and

inequality, and therefore two negatives making a positive, then the objection was for the right reason. If the objection occurred because of a prior belief that this result “cannot” occur, then it was an objection for the wrong reason.

Wiig (2000).¹⁷ There is no basis, therefore, for concluding that individual inequality has worsened during the globalization period.

Individual Inequality

The John Lennon song “Imagine” contains the fundamental idea behind the generation of these *world individual income inequality* estimates, or W3i for short. What happens if there is no country, only one world? What if we have world distributions, distributions that imply that individuals in the world are lined up according to their income and their calculated share of world income—in other words, a distribution that has no regard for sex, color, national origin, or individual prejudices?

The generation of both intercountry inequality (dubbed “international inequality” in Bourguignon and Morrisson 2001) and individual inequality require essentially the same data, that is, the population of each country, its per capita income in a common currency,¹⁸ and its distribution of income. A summary of research on international inequality computations is as follows. The first level of aggregation is to assume nothing about the distribution and allocate the per capita income to each resident.

The next lower level of aggregation is to allocate each individual their mean quintile income (as was done by Korzeniewicz and Moran). The next lower level of aggregation is to allocate incomes according to each decile of the population and, depending on data availability, even finer gradations. This method would yield a reasonably close approximation to the desired goal of obtaining an index for global individual inequality; this was the procedure outlined by Berry, Bourguignon, and Morrisson (1981, 1983) in the early 1980s.

There can be some controversy over the selection and estimation of *average* income (it can be obtained either from surveys or from national

17. In two recent papers, Sala-i-Martin (2002a) constructs world distributions like those of Korzeniewicz and Moran (1997) (using country quintile income) and in a later paper (2002b) extends the methodology to map individual country quintile data onto a world distribution. Sala-i-Martin’s results are not analyzed in any detail here because they came just after the draft was submitted to referees.

18. There is (almost) universal agreement that the currency of choice is purchasing power parity incomes; for reasons that are not clear, the United Nations *Human Development Report 1999* used constant US dollars as the “numeraire” and not surprisingly found that rich countries were getting considerably richer while poor countries were found entering the gates of oblivion. This study was unique in other respects as well—it forced the United Nations to set up a committee to investigate the statistical methods used by its own staff; see Report of the Friends of the Chair of the Statistical Commission (2000). The investigation into the UN statistical methods seems to have had little impact. The latest UN *Human Development Report* (2002) follows the earlier tradition. It reaches the conclusion that world inequality has worsened, thus completely ignoring all available evidence suggesting the opposite.

accounts), so it is not so simple. And distribution data may not be available for every country for every year. So it is not so straightforward. But assumption, thy name is economist—and Berry and his colleagues made reasonable assumptions to arrive at an income distribution estimate for the world (124 countries) in 1970. Their assumptions involved both the generation of income distribution for countries for which such data are not available, as well as the growth in incomes of countries like China.

Berry and his colleagues eschewed the generation of a Lorenz curve for each country, a major intermediate step for the simple accounting procedure (see chapter 8). Instead, they aggregated the available distributional data—not by deciles or quintiles (as was done by Milanovic) but rather by “income brackets, so that none of them include a significant percentage of the world’s population . . . and no income bracket was to include more than 2% of the world population in the low income range and 1% in the high income range” (Berry, Bourguignon, and Morrisson 1983, 333).

Using PPP income growth rates, Berry and his colleagues reported on the evolution of world inequality between 1950 and 1977. The assumption that individual country inequality did not change for 27 years (they use the same distribution for each country as was generated by them for 1970) may be a bit extreme but is nevertheless understandable, given the paucity of country distribution data at that time. Their result: a highly unequal Gini index of 66.2 for the world income distribution, and only a slightly lower inequality level of 65.2 for the distribution of consumption.¹⁹

What Berry and his colleagues did was exactly what one would do if confronted with the task of estimating the distribution for an individual economy. Collect the data for each economy for only three variables: population, average income, and the distribution of this income. Once these data are assembled, it is then a simple *accounting* procedure to estimate the Lorenz curve and such associated inequality measures as the Gini, Theil, log-variance, ratio of the incomes of the 10th to the 1st decile, and so on.

Much like the 1970s, the availability of newer survey data in the 1980s and 1990s, and the compilation of the World Bank and WIDER datasets, have allowed for a resurgence of individual inequality studies. The first to undertake such a study were Bourguignon and Morrisson (2001),²⁰ two of the authors of the original first study on W3i by Berry, Bourguignon, and Morrisson. They not only improved and updated the original study to 1992, but also provided the first set of estimates for several years going

19. The authors present estimates under alternative assumptions in the two papers, but the figures reported here can be considered “representative.”

20. This is the date of the final draft; the first draft was dated June 1999.

back to 1820. Soon after, three other studies appeared (Milanovic 1999; Bhalla 2000d; Dikhanov and Ward 2000).²¹

Bourguignon and Morrisson's improved method still does not involve the estimation of a Lorenz curve; instead, they attempt to collect as much disaggregated data as possible—in particular, the average income of each decile of the population, as well as the mean income of the top two ventiles (90th-95th and 95th-100th percentiles). Because most of the published data exist only for quintiles, and at most deciles, they must have used some interpolation (i.e., some attempt must have been made to generate portions of the Lorenz curve).²² They use data for 33 selected countries to approximate regional distributions. Clearly, much ambition was involved in constructing estimates back to 1820, but both the assumptions made, and the results obtained, seem compelling.

Milanovic (1999) assembled the survey data gathered at the World Bank and estimated inequality for 1988 and 1993. Sometimes these data were available in decile form, sometimes only in quintile form. For such large countries as India and China, Milanovic had data on both rural and urban areas, so the level of disaggregation achieved for these countries is twice as large as for others (i.e., if original data are available in decile form, then there are 20 units of observation for each country, or 20 units of data to estimate average income). Milanovic's estimate is a "pure" estimate of inequality; *no* interpolation is made of survey data to arrive at any portion of the Lorenz curve. On *average*, Milanovic probably achieved the same level of disaggregation as Bourguignon and Morrisson; that is, about 12 inequality data points per country-year.

To sum up the results of these studies: Bourguignon and Morrisson's results suggest that *W3i* worsened until 1980 and since then has remained relatively flat, at a Gini of 65.7. Milanovic shows a *huge* increase in inequality in the space of just 5 years, 1988 to 1993—from 62.5 to 65.9. Dikhanov and Ward reach the same conclusion as Milanovic; Bhalla reaches the conclusion that not only had *W3i* declined in the past 40 years, but it was at its most equal in the late 1990s (the latest estimates of this method are presented in chapter 11).

21. An earlier version of the SAP was used to generate world inequality and world poverty estimates in the paper presented in June 2000 at a seminar at the International Monetary Fund. The SAP method has undergone improvements, especially in the generation of a Lorenz curve from limited data. The present version can be considered the final version! The June 2000 method, however, had a slight tendency to overestimate world individual inequality, but this made little difference to the generation, and estimation, of world poverty levels.

22. Bourguignon and Morrisson (2001) do not use the expanded distribution set made available by Deininger and Squire (1996) and the World Institute for Development Economics Research (the database referred to above in note 9), but instead use an updated version of the data in Berry, Bourguignon, and Morrisson (1983). Of course, data from before 1950 are unique to the Bourguignon and Morrisson study.

Summing Up: The Facts as We Know Them

The trends in *country* inequalities documented above highlight one major conclusion: World inequality (as measured by the inequality of a “representative” average economy), after having stayed relatively constant and even improved until the late 1970s, significantly deteriorated in the post-1980s globalization period. If the world is disaggregated by regions, one obtains the same conclusion. The observation of trends in individual large countries yields the same conclusion. Observations over long periods of time lead to the same conclusion—though some dampening in the worsening-inequality result is obtained if the time period used is more than 8 years. The use of popular log-log regressions (regressing the log of quintile income on the log of average income) yields the same result—the average elasticity is 0.8, suggesting that the first quintile’s share of the pie has declined over time (or inequality has increased).

Moving toward a lower level of aggregation—the individual rather than a country as a unit of analysis—yields a mixed result. Milanovic obtains the same result as that yielded by country inequalities, that is, a significant worsening in the globalization period. His results are an extension of the individual country results. Bourguignon and Morrisson report a mild increase in inequality (their last date of observation is 1992). In an earlier version of the findings reported here, Bhalla (2000d) suggests that the result for individual inequality is completely contrary to that for country inequality: Not only did individual inequality not worsen in the 1980s and 1990s, it actually improved, and did so significantly! The veracity of these findings is explored in detail in chapter 11.