

Inequality is Bad for the Poor

Martin Ravallion*

*Development Research Group, World Bank
1818 H Street NW, Washington DC*

World Bank Policy Research Working Paper 3677, August 2005

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the view of the World Bank, its Executive Directors, or the countries they represent. Policy Research Working Papers are available online at <http://econ.worldbank.org>.

* This is a Background Paper to the 2006 *World Development Report on Equity and Development*.

1. Introduction

It has been argued that inequality should be of little concern in poor countries on the grounds that: (i) absolute poverty in terms of consumption (or income) is the overriding issue in poor countries, and (ii) the only thing that really matters to reducing absolute income poverty is the rate of economic growth. This article takes (i) as given but questions (ii). It is argued that there are a number of ways in which the extent of inequality in a society, and how it evolves over time, influences the extent of poverty today and the prospects for rapid poverty reduction in the future.

The following section looks at the empirical relationship across countries between inequality and growth, while section 3 turns to the relationship between inequality and poverty reduction. Section 4 examines whether the evidence from the experiences of developing countries supports the view that there is an aggregate trade-off between growth and reducing inequality. Section 5 returns to the issues of the preceding sections in the context of recent research for the two largest countries, China and India. Finally, section 6 tries to draw out some lessons for policy and for policy-relevant research.

2. Inequality and growth revisited

A number of papers in the literature have found that changes in inequality at the country level have virtually zero correlation with rates of economic growth; see, for example, Ravallion and Chen (1997), Ravallion (2001) and Dollar and Kraay (2002). Among growing economies, inequality tends to fall about as often as it rises, i.e., growth tends to be distribution neutral on average. If all levels of income grow at roughly the same rate then of course absolute poverty must fall. This makes it unsurprising that the literature has also found that absolute poverty measures tend to fall with growth — that “growth is good for the poor” (to quote the title of an influential paper by Dollar and Kraay, 2002). Supportive evidence for the view that absolute

poverty tends to fall with economic growth can be found in World Bank (1990, 2000), Ravallion (1995, 2001), Ravallion and Chen (1997), Fields (2001) and Kraay (2005).

There are a number of reasons for caution in interpreting this finding of a lack of correlation between changes in inequality and growth. Firstly, finding that there is no change in overall inequality can be consistent with considerable “churning” under the surface, with gainers and losers at all levels of living. This cannot be seen in cross-sectional surveys. The (more limited) panel data sets available point to churning.¹ Simulations of the impacts of specific policy changes intended to promote economic growth also point to such heterogeneity, or “horizontal inequality,” in the impacts of reform. In the context of trade reform, Ravallion (2005a) reviews evidence on the extent of horizontal inequality, as indicated by the dispersion in welfare impacts of reform at any given level of pre-intervention income. This dispersion reflects differences in variables such as household demographics and location that influence the net trading positions in relevant markets and (hence) the welfare impacts of trade reform.

Secondly, the measures of “inequality” in this literature are typically measures of relative inequality, whereby multiplying all incomes by a constant leaves the measure of inequality unchanged. Finding that a relative inequality measure is unchanged during an aggregate economic expansion is perfectly consistent with large increases in absolute income disparities. Growth in average income tends to come with higher absolute disparities between the “rich” and the “poor” (Ravallion, 2003). Arguably, it is the absolute changes that are more obvious to people living in a growing developing economy than the proportionate changes.² So it may well

¹ A useful compilation of studies using panel data can be found in the August 2000 special issue of the *Journal of Development Studies*; see the introduction by Baulch and Hoddinott (2000). The churning also stems in part from time-varying measurement errors, though plausible covariates have been evident in the studies that tested for this (see, for example, Jalan and Ravallion, 2000).

² For further discussion see Amiel and Cowell (1999), Atkinson and Brandolini (2004) and Ravallion (2004).

be the case that much of the debate about what is happening to inequality in the world is actually a debate about the meaning of “inequality” (Ravallion, 2004).

Thirdly, there are also signs that the growth processes seen in many reforming economies in the 1990s have been putting more systematic upward pressure on inequality. Lopez (2005) reports evidence to support this view (though based on a smaller, selected, sample of countries than will be studied in this paper). To re-examine the relationship between growth and changes in inequality, I created 290 “spells” defined by two household surveys for a given country with more than one observation for most countries; there are about 80 countries represented, spanning 1980-2000.³ I then compared the changes in the Gini index with the changes in the survey mean (in real terms, using local CPIs). Figure 1 gives a scatter plot of changes in the log Gini index against changes in the log real survey mean between successive household surveys. The correlation coefficient is -0.13 and is not statistically significant (at the 10% level). Among growing economies, inequality increased about as often as it fell, and similarly among contracting economies. Figure 2 focuses on the period after 1992, dividing the sample in two. There is now a mild positive correlation coefficient of 0.26 , which is significant at the 5% level.⁴

Fourthly, it must be acknowledged that there is likely to be considerable measurement error in the changes in inequality and the survey means. The errors can come from a variety of sources, including sampling errors (probably a minor concern in most cases for the surveys used here), errors arising from selective compliance (whereby certain types of households participate in surveys with lower probability than others), under-reporting of incomes and comparability

³ The data are drawn from *PovcalNet* and the *World Development Indicators*. *PovcalNet* is a new interactive tool for poverty analysis that provides the primary distributional data for about 500 surveys for 100 developing countries, drawing on the World Bank’s data base; see <http://iresearch.worldbank.org/povcalnet>. The primary background paper is Chen and Ravallion (2004).

⁴ All significance tests in this paper are based on White standard errors (corrected for heteroscedasticity, which is clearly present).

problems between surveys arising from differences in questionnaires, interviewing procedures or processing methods. These errors can greatly weaken the power of the tests found in the literature using cross-country and inter-temporal comparisons for detecting the true relationship.

There are a couple of things we can do to test robustness to time varying measurement errors. One is to use data over longer periods. Figures 1 and 2 use whatever time periods are available between successive surveys. If instead one uses changes over three surveys (taking the log difference between the survey for date t and $t-2$) the correlation over the whole period becomes significantly negative ($r=0.24$, $n=206$), and that remains true for the data points after 1992. Alternatively, one can use the longest spell for each country; again there is no significant correlation ($r=0.10$, $n=80$).

A second test is to use growth rates in consumption from national accounts (NAS) as the instrumental variable for the growth rates based on the survey means. This assumes that the measurement errors in the two data sets are uncorrelated. While in practice there are sometimes overlaps in the underlying data sources used (such as when specific consumption items in the national accounts are benchmarked from household survey data), by-and-large the assumption is probably defensible for the purpose of testing robustness. Using this test, one finds no significant correlation (in either direction) between changes in inequality and (instrumented) growth in survey means for either the 1990s, or the period as a whole since the early 1980s. When estimated over all available observations, the IV estimate of the regression coefficient of the change in log Gini index on change in log survey mean using the change in log private consumption per capita from the NAS as the instrument is 0.04 with a standard error of 0.26. Confining the estimation to the post 1992 period, the IV regression coefficient rises substantially to 0.15, but this is only significantly different from zero at the 15% level (White standard error of

0.11). So the claim that growth has been inequality increasing in the 1990s is not robust to allowing for time-varying measurement errors.

While acknowledging these data issues and caveats, the lack of correlation between changes in relative inequality and growth does not imply that policy makers aiming to fight poverty in any given country can safely focus on growth alone. All this empirical finding tells us is that, on average, there was little effective redistribution in favor of the poor. It does not tell us that re-distribution rarely happens or that distribution is unimportant to the outcomes for poor people from economic growth. The rest of this article takes up these issues.

3. Inequality and the pace of poverty reduction

While it may be readily agreed that economic growth tends to lead to lower measures of absolute poverty, there is nonetheless a wide variation in the impact of a given rate of growth on poverty. Ravallion (2001) estimates that the 95% confidence interval implies that a 2% annual growth rate in average household income will bring anything from a modest drop in the poverty rate of 1% to a more dramatic 7% annual decline. (So for a country with a headcount index of 40%, we have 95% confidence that the index will fall by somewhere between 0.4 percentage points and 2.8 points in the first year.)

Why do we find that the same rate of growth can bring such different rates of poverty reduction? In answering this question it is convenient to start with the identity that the proportionate rate of poverty reduction is the product of the “growth elasticity of poverty reduction” and the rate of growth. Note that this is not the same as the elasticity of poverty with respect to the mean holding distribution constant (Kakwani, 1993). The latter can be thought of as the partial elasticity, as distinct from the total elasticity given by the proportionate rate of poverty reduction divided by the rate of growth. Of course, if growth is distribution-neutral on

average then the two elasticities will be similar on average, though they may differ greatly in specific countries and time periods.

Two factors can be identified as the main proximate causes of the differing total elasticities of poverty reduction: the initial level of inequality and inequality changes over time.

Initial inequality. It is intuitive that the higher the initial inequality in a country, the less the poor will share in the gains from growth; unless there is sufficient change in distribution, a larger (smaller) initial share of the pie will tend to come with a larger (smaller) share in the pie's expansions. While this intuition is compelling, it is theoretically ambiguous as to how differences in initial inequality will affect the growth elasticity of poverty reduction. Consider two countries, one with a Lorenz curve that unambiguously dominates the other, i.e., inequality is higher in one country for all possible inequality measures (Atkinson, 1970). Suppose first that the Lorenz curves remain unchanged over time. It can be readily shown that the proportion of the population below any given level of income will then be homogeneous of degree zero in the mean and the level of income considered.⁵ Then it is plain that the growth elasticity of poverty reduction for the headcount index (H) is (minus one times) the elasticity of the cumulative distribution function evaluated at the poverty line.⁶ Next note that there can be no presumption that the country with higher inequality will have a higher H ; depending on the specific properties of the Lorenz curve at H , the higher inequality country could have either a higher or lower headcount index.⁷ The implications for the growth elasticity are then also ambiguous. Non-neutralities in the growth process add a further source of ambiguity in the implications of initial

⁵ This follows from the fact that $L'(p) = y / \mu$ where $L(p)$ is the Lorenz curve and $p = F(y)$ is the cumulative distribution function (Gastwirth, 1971).

⁶ In other words, the growth elasticity is $-zf(z)/H$ where $H=F(z)$ is the headcount index at the poverty line z and $f(\cdot)$ is the density function.

⁷ This ambiguity stems from the fact that H is found at the tangency of the Lorenz curve at z / μ where μ is the mean (i.e., $L'(H) = z / \mu$).

differences in inequality for the elasticity of the headcount index to the mean (allowing the Lorenz curve to change). Even when the initial share held by the poor is low, their gains from growth can be sizeable if growth is accompanied by sufficient pro-poor redistribution.

Some special cases yield unambiguous results, which are achieved by collapsing the potential differences in initial distribution into just one parameter. Analytic predictions obtained under the assumption that household income or consumption is log-normally distributed predict that the partial growth elasticity of poverty reduction holding distribution constant will fall (in absolute value) as inequality rises (Bourguignon, 2003). Son and Kakwani (2004) invoke the Kakwani (1993) assumption that the Lorenz curves across countries only differ in a special way, namely that the entire curve shifts by a constant proportion of the difference between the actual value on the Lorenz curve and the line of equality. They also assume that the growth process is distribution-neutral and that the poverty line is less than the mean. Under these assumptions, Son and Kakwani show that the growth elasticity of poverty reduction for the Foster-Greer-Thorbecke class of poverty measures is monotonically decreasing in the initial value of the Gini index, which essentially becomes the sole parameter locating the Lorenz curve.

These theoretical results are instructive, and consistent with intuition. In practice, however, distributions vary by more than one parameter and growth processes are only (roughly) distribution-neutral on average. Growth in specific countries and time periods is rarely distribution-neutral, so that assumption can be quite deceptive in predicting outcomes of specific growth episodes. For example, consider the growth process in Brazil in the 1980s. Datt and Ravallion (1992) show that if one had assumed at the outset of the decade that growth would be distribution neutral then one would have predicted a 4.5% point decline in the headcount index

of poverty. In fact, there was no change over the decade, with the headcount index staying at 26.5%. Distributional shifts working against the poor exactly offset the gains from growth.

What does the empirical evidence suggest about the relationship between initial inequality and the growth elasticity of poverty reduction? Support for the intuition that higher inequality countries tend to have lower (absolute) elasticities was first presented in Ravallion (1997) and subsequently verified by Ravallion (2001) and Kraay (2005). These papers have used regression-based methods (in which rates of change in poverty are regressed on rates of growth both on its own and interacted with initial inequality). We will return to this approach shortly, but first it is instructive to look at the empirical relationship seen in a more flexible way.

A simpler non-parametric method is to calculate the elasticity as the log difference in the headcount index divided by the log difference in the mean, all based on successive household surveys. There is clearly a lot of noise in such a measure. To help reduce the noise, I smoothed the period-specific elasticities by taking the simple average of two-period elasticities (across three surveys). I also trimmed 15 extreme elasticities (below -20 or above 20). Figure 3 gives the results for the “\$1 a day” poverty rate. The elasticity is negative in 80% of cases. We see a rather weak tendency for the elasticity to rise (become less negative) as inequality rises, from an average of about -4 at the lowest Gini index to roughly zero at the highest. The correlation coefficient is 0.26, which is significant at the 1% level. The two high positive elasticities in Figure 3 are almost certainly measured with large errors, and this is exaggerating the slope of the line of best fit. Dropping these two observations, the correlation is still significant at the 1% level, and the line of best fit passes through an elasticity of zero at Gini index of about 60%.

In modeling the relationship between poverty reduction and growth, Ravallion (1997) postulated that the rate of poverty reduction (measured as the difference in the log of the measure

of poverty) is directly proportional to the “distribution corrected rate of growth” where the latter is given by the ordinary rate of growth (log difference in mean consumption or income) times a distributional term. In Ravallion (1997) the distributional correction used is one minus the initial Gini index. This model can be improved (in terms of fit with data on actual spells of changes in poverty matched with growth) by using instead an adjustment for nonlinearity in the relationship between the growth elasticity of poverty and the initial inequality, giving a simple model of the expected rate of poverty reduction over any period:

$$\text{Rate of poverty reduction} = [\text{Constant} \times (1 - \text{Inequality index})^\theta] \times \text{Ordinary growth rate}$$

The constant term is negative and θ is a parameter not less than one. The total growth elasticity of poverty reduction is the term in square brackets. At high levels of inequality the poor will gain little or nothing from growth; at the extreme in which the inequality index is one, the richest person has all the income and so all the gains from growth will go to that person; the elasticity will be zero. For values of θ strictly greater than one, higher levels of initial inequality will have progressively smaller impacts on the elasticity as inequality rises. The above model can be augmented by adding one or more terms for changes in distribution, to isolate the partial elasticity. This raises the R^2 but does not affect the results of interest here, given that (as we have already seen) changes in distribution tend to be uncorrelated with growth rates.

Quite a good fit with data on actual rates of poverty reduction across developing countries can be obtained using the initial Gini index as the measure of inequality and using $\theta = 3$.⁸ By this simple model, the rate of poverty reduction in a given time period is directly

⁸ The nonlinear least squares estimate of θ on a sample of estimates of the changes in the log of the “\$1/day” poverty rates for the longest available spells between surveys for 62 countries gave 3.031 with a standard error of 0.491. Using the full sample of all the spells gave a lower estimate, of 2.056 with a standard error of 0.493. However, the noise in the data is probably attenuating the coefficient.

proportional to $(1-G)^3$ times the rate of growth over that period, where G is the Gini index at the beginning of the time period. Using a sample of estimates of the changes in the log of the “\$1/day” poverty rates for the longest available spells between successive surveys for 62 countries the constant of proportionality is -9.33 , with a standard error of 0.75 and $R^2 = 0.71$.⁹ Figure 4 plots the implied growth elasticity of poverty against the Gini index; the elasticity ranges from -4.3 to -0.6 .

To help interpret this model, consider the rate of poverty reduction with a 2% rate of growth in per capita income (roughly the mean rate for the developing world in 1980-2000) with no change in distribution and a headcount index of 40% (the mean poverty rate for the developing world around 1980). In a low-inequality country, with a Gini index of 0.30, say, the headcount index will fall by 6.4% per year, or 2.6 percentage points in the first year; the headcount index will be halved in 10.5 years. By contrast, in a high inequality country, with a Gini index of 0.60 growing at the same rate and with the same initial headcount index, the latter will fall at an annual rate of 1.2%, representing a decline of only 0.7 percentage points in the first year; it will then take 57 years to halve the initial poverty rate. Poverty responds slowly to growth in high inequality countries; or (to put the same point slightly differently) high inequality countries will need unusually high growth rates to achieve rapid poverty reduction.

Two further observations can be made. Firstly, the argument works in reverse too; high inequality will help protect the poor from the adverse impact of aggregate economic contraction

⁹ If one simply regresses the rate of poverty reduction on the rate of growth (both as log differences) then one obtains $R^2 = 0.56$. Thus incorporating the nonlinear interaction effect with initial inequality adds 15 percentage points to the variance in rates of poverty reduction that can be explained by rates of growth. The “long spells” series was possible for 70 countries, but eight were dropped on the grounds that the measured rates of poverty reduction relative to rates of growth were either far too large or far too small to be believed (elasticities less than -10 or greater than 0.5). On the full sample of 70 countries, I obtained $R^2 = 0.46$ using the ordinary rate of growth versus 0.58 using the distribution corrected rate of growth, as above.

(Ravallion, 1997). Low inequality can thus be a mixed blessing for poor people living in an unstable macroeconomic environment; it helps them share in the benefits of growth, but it also exposes them to the costs of contraction. There is evidence that this also happens at the local level during an economy-wide crisis; high inequality districts of Indonesia experienced less dramatic rates of increase in poverty during the 1998 financial crisis than did low inequality districts (Ravallion and Lokshin, 2004).

Secondly, I find very little robust evidence of a significant correlation between the growth elasticity of poverty reduction and the initial mean (either on its own, or controlling for initial inequality). The theoretical relationship between the elasticity of the headcount index with respect to distribution-neutral growth and the mean is known to be ambiguous, though for a log-normal distribution of income, the partial elasticity is strictly decreasing in the mean (Bourguignon, 2003) and this also holds for the poverty gap index and other “higher order” poverty measures in the Foster-Greer-Thorbecke class under quite general conditions (Son and Kawkani, 2004). However, the empirical evidence does not offer much support for this theoretical prediction.

None of this is inconsistent with the findings in the literature indicating that a large share of the variance in rates of poverty reduction can be attributed to differences in ordinary rates of growth (Ravallion, 1995; Ravallion and Chen, 1997; Fields, 2001; Kraay, 2005). In a recent contribution, Kraay (2005) presents Datt-Ravallion decompositions of changes in “\$1/day” poverty measures into growth and redistribution components for as many countries as possible. Kraay’s growth component is the product of the growth rate and the partial elasticity.¹⁰ Kraay

¹⁰ Recall that it is a partial elasticity because it holds distribution constant; by contrast the “total elasticity” lets distribution vary consistently with the data; the elasticity in square brackets in the above equation is a total elasticity. The analytic elasticities of poverty measures discussed in Kakwani (1993) and Bourguignon (2001) are partial elasticities.

finds that the variance in the growth component is largely attributable to the growth rate, rather than the partial elasticity or its covariance with growth. For example, he attributes 81% of the variance in the log absolute value of the growth component of changes in the headcount index to the variance in the log absolute growth rate.

All this is perfectly consistent with finding that poverty is relatively unresponsive to growth in specific countries. Kraay's results are based on averages formed from cross-country comparisons. (A variance is an average too, namely the mean of the squared deviations from the ordinary mean.) For a developing country with average inequality and for which inequality does increase with growth, Kraay's results offer some support for his policy conclusion that for reducing poverty the main thing to worry about is achieving a higher rate of growth. However, that does not mean that growth is sufficient even when inequality is low. If growth in a low inequality country comes with a sufficient increase in inequality then it will by-pass the poor. And, as already noted, the empirical finding that growth is roughly distribution neutral on average is consistent with the fact that it increases roughly half the time during spells of growth (Ravallion, 2001). Policy effort to keep inequality low may then be crucial to pro-poor growth in many low-inequality countries.

Furthermore, as we have seen, for high inequality countries, growth will be quite a blunt instrument against poverty unless that growth comes with falling inequality. The heterogeneity in country circumstances is key here. Averages formed across countries can be quite uninformative about how best to achieve pro-poor growth in specific countries.

While initial inequality is an important proximate determinant of differing rates of poverty reduction at a given rate of growth, to help inform policy we need to probe more deeply into the relevant sources of inequality. There are inequalities in a number of dimensions that are

likely to matter, including access to both private (human and physical) assets and public goods. Inequalities in access to infrastructure and social services naturally make it harder for poor people to take up the opportunities afforded by aggregate economic growth.

Changing income distribution. A second factor influencing the rate of poverty reduction at a given rate of growth is changing income distribution. Finding that growth tends to be distribution neutral on average does not, of course, mean that distribution is unchanging. Whether inequality is rising or not can make a big difference to the rate of poverty reduction. Among growing economies, the median rate of decline in the “\$1/day” headcount index is 10% per year among countries that combined growth with falling inequality, while it is only 1% per year for those countries for which growth came with rising inequality (Ravallion, 2001). Either way poverty tends to fall, but at very different rates. (And similarly among contracting economies; poverty rises on average, but much more rapidly when inequality is rising than falling.) As one would expect, changes in distribution matter even more for higher-order poverty measures, which can respond quite elastically to even small changes in overall inequality.

What underlies the changes in distribution, as they affect poverty? There are a great many country-specific idiosyncratic factors (such as shocks to agricultural incomes, changes in trade regime, shifts in relative prices, tax reforms, welfare-policy reforms and changes in demographics). Generalizations across country experience are never easy, but one factor that is likely to matter in many developing countries is the geographic and sectoral pattern of growth. The marked concentrations of poor people in specific regions and/or sectors that one finds in many countries point to the importance of the pattern of growth to overall poverty reduction. The extent to which growth favors the rural sector is often key to its impact on aggregate poverty. The geographic incidence of both rural and urban economic growth is often important

as well. However, the extent to which the pattern of growth (rather than simply the overall growth rate) matters to the rate of poverty reduction is likely to vary from country to country, depending on (*inter alia*) how unbalanced the growth process has been in the past and (hence) how much difference one currently finds between sectors or regions in levels of poverty.

While it still appears to be the case that (relative) inequality falls about as often as it increases during spells of aggregate economic expansion, there are also signs that higher growth in a number developing countries has come with widening regional disparities and often little or no growth in lagging poor areas. China and India are examples, to which we return.

4. A growth-equity trade-off?

Making growth more pro-poor requires a combination of more growth, a more pro-poor pattern of growth and success in reducing the antecedent inequalities that limit the prospects for poor people to share in the opportunities unleashed in a growing economy. The ideal combination will naturally vary with country circumstances. In some countries, attention can safely focus on the overall rate of growth to assure rapid poverty reduction; elsewhere, a broader approach will be called for. This begs the question as to whether there might be a trade-off between interventions to make growth more pro-poor and the rate of growth.

While poverty is more often seen as a consequence of low average income, there are reasons for thinking that there is a feedback effect whereby high inequality also impedes future growth.¹¹ In many developing countries, a plausible way this can happen stems from credit market failures, which mean that some people are unable to exploit growth-promoting opportunities for investment. And it will tend to be the poor for whom these constraints are most

¹¹ There is now a sizeable theoretical literature on the various ways in which inequality can impede growth. Contributions include Galor and Zeira (1993), Banerjee and Newman (1993), Benabou (1996), Aghion et al., (1999) and Bardhan et al., (1999).

likely to be binding. With declining marginal products of capital, the output loss from the market failure will be greater for the poor. So the higher the proportion of poor people there are in the economy the lower the rate of growth. Then poverty is self-perpetuating.

There are other ways in which initial distribution matters to growth prospects. In the presence of capital market failures due to moral hazard, high inequality can dull incentives for wealth accumulation. It has also been argued that high inequality can foster macroeconomic instability and impede efficiency-promoting reforms that require cooperation and trust.¹²

There is supportive evidence for the view that inequality is bad for growth from cross-country comparisons of growth rates, suggesting that countries with higher initial inequality experienced lower rates of growth controlling for other factors such as initial average income, openness to trade and the rate of inflation.¹³ At the same time, there are also a number of concerns about the data and methods underlying these findings based on cross-country comparative analysis (Ravallion, 2001). Future research will hopefully throw more light on the magnitude of the efficiency costs of inequality.

5. China and India

The world's two largest countries differ in a great many ways, but their overall development paths since the 1980s have shared some common features: more-or-less sustained economic growth (since the early 1980s for China and since the early 1990s for India), falling absolute poverty, and signs of rising overall inequality (though more persistently so in China's

¹² Aghion et al (1998) and Bardhan et al. (1999) review these and related arguments as to why high inequality can reduce aggregate output.

¹³ See Persson and Tabellini (1994), Alesina and Rodrik (1994), Clarke (1995), Birdsall et al., (1995), Perotti (1996), Deininger and Squire (1998) and Easterly (2002).

case), reflecting in part geographic and sectoral “imbalances” in the growth process, which have dulled the impact of growth on poverty.

China

China since around 1980 is often cited as an example in which rising inequality allowed rapid growth and (hence) rapid poverty reduction. There can be no doubt that absolute poverty in China has fallen greatly since around 1980. While China’s poverty rate today is probably slightly lower than the average for the world as a whole,¹⁴ it was a very different story around 1980, when the incidence of extreme poverty in China was one of the highest in the world.¹⁵ That is huge progress. However, there were some significant setbacks for China’s poor. Poverty reduction stalled in the late 1980s and early 1990s, recovered pace in the mid-1990s, but stagnated again in the late 1990s. About half of the decline in poverty came in the first half of the 1980s.

Income inequality has also been rising, though not continuously and more in some periods and provinces. Figure 5 gives the estimates of the Gini index, which rose from 28% in 1981 to 39% in 2001.¹⁶ The Gini index is only one possible measure of inequality, and may not reflect well how we would weight gains at different level of living (Atkinson, 1970). A more flexible way of representing the distributional impacts of China’s growth is the growth incidence curve (GIC) given in Figure 6 (following Ravallion and Chen, 2003). This gives the rate of

¹⁴ See Chen and Ravallion (2004) who estimate that in 2001, 17% of China’s population live below \$1 a day at 1993 Purchasing Power Parity; the corresponding figure for the world as a whole is 18% (21% for developing countries alone).

¹⁵ The proportion of China’s population living below \$1 a day in 1981 is estimated to have been 64%. Based on the “\$1 a day” poverty rates for 1981 from <http://iresearch.worldbank.org/povcalnet>, only four countries (Cambodia, Burkina Faso, Mali and Uganda) had a higher poverty rate than this in 1981.

¹⁶ Note that the latter figure is somewhat lower than past estimates for China; this is because corrections have been made for urban-rural cost-of-living differences, which have tended to rise over time because of higher inflation in urban areas. Without these corrections the Gini index for 2001 rises to 45%.

growth over the relevant time period at each percentile of the distribution (ranked by income or consumption per person). We see that growth rates in China in the 1990s tend to rise as we move up the distribution; the annual rate of growth in the 1990s varies from about 3% for the poorest percentile to nine percent for the richest. While the growth rate in the overall mean was 6.2%, the mean growth rate for the poorest 20% (roughly according with China's "\$1 a day" poverty rate in 1990) was 4.0%.¹⁷

What lies behind these distributional shifts in China? Like many developing countries, living standards tend to be lower in rural areas than urban areas. In the case of China, mean income is about 70% higher in urban areas (adjusted for cost-of-living differences). However, one does not find that inequality between urban and rural areas has shown a trend increase since reforms began, though there have been sub-periods (such as the late 1980s to the early 1990s) in which the urban-rural disparity rose over a number of years.¹⁸

To understand the rise in overall inequality one must understand what has happened within urban and rural areas, and particularly the latter, which naturally carries larger weight given that 60% of the population still live in rural areas, and that 80% did so at the outset of the reform period. Indeed, around 1980, a staggering 98% of China's poor lived in rural areas (Ravallion and Chen, 2004). While there are various ways (trade, migration, transfers) that non-farm economic growth will spill over to the farm economy, the sheer weight of the rural sector in absolute poverty at the outset of China's reform period means that agricultural and rural

¹⁷ This is the Ravallion-Chen (2003) "rate of pro-poor growth," namely the mean growth rate of the poor. This gives the change in the Watts index per unit time divided by the initial headcount index. Notice that the mean growth rate of the poor is not the same thing as the growth rate in the mean for the poor, which will not in general be consistent with even the direction of change in any sensible measure of the level of poverty.

¹⁸ An important difference with past results on this point is that Ravallion and Chen allowed for the fact that the rate of increase in the urban cost-of-living exceeded that for rural areas.

economic growth would have been crucial. This also carries an important lesson for other developing countries hoping to emulate China's success, which we return to in the final section.

Table 1 gives a regression decomposition of the rate of change in poverty over time (difference in the log headcount index) on the share-weighted growth rates of rural and urban mean incomes and the populations shift effect; the Appendix derives this decomposition in more formal terms. (The table gives results for both China and India; we turn to India shortly.) It can be seen that only rural economic growth is statistically significant. An alternative decomposition exploits the analytic (additivity) properties of the headcount index, whereby the national index is the population-weighted mean of the urban and rural indices. This decomposition makes somewhat different assumptions to the regression decomposition. However, it confirms the quantitative importance of rural economic growth; about 72% of the reduction in the headcount index is attributable to rural poverty reduction, versus 5% due to urban and 23% due to the population shift from rural to urban areas (Ravallion and Chen, 2004).¹⁹

Table 2 gives an alternative decomposition by source of GDP. (Again the Appendix gives the decomposition in more formal terms.) The overall elasticity of the headcount index to GDP growth is -2.6 . However, when one decomposes growth into "primary" (mainly agriculture), "secondary" (manufacturing and construction) and "tertiary" (services and trade) it becomes clear that the sectoral composition of growth matters greatly to the rate of poverty reduction. The primary sector has far higher impact (by a factor of about four) than either the secondary or tertiary sectors (Table 2). The impacts of the latter two sectors are similar (and we cannot reject the null that they have the same impact).

¹⁹ Corresponding results for "higher-order" poverty measures can be found in Ravallion and Chen (2004). These are similar to the results reported here for the headcount index.

These aggregate results do not tell us about the source of the poverty-reducing impact of primary sector growth. With a relatively equitable distribution of access to agricultural land and higher incidence and depth of poverty in rural areas it is plausible that agricultural growth will bring large gains to the poor. There is evidence for China that this may also involve external effects at the farm-household level. One important source of externalities in rural development is the composition of economic activity locally. In poor areas of southwest China, Ravallion (2005b) finds that the composition of local economic activity has non-negligible impacts on consumption growth at the household level. There are significant positive effects of local economic activity in a given sector on income growth from that sector. And there are a number of significant cross-effects, notably from farming to certain nonfarm activities. The sector that matters most as a generator of positive externalities turns out to be agriculture (Ravallion, 2005b).

A natural counterfactual for measuring the contribution of the sectoral composition of growth is the rate of poverty reduction if all three sectors had grown at the same rate. We call this “balanced growth.” Then the sector shares of GDP in 1981 would have remained constant over time. For the same GDP growth rate, the mean rate of poverty reduction would then have been 16.3% per year, rather than 9.5% (Ravallion and Chen, 2004). Instead of 20 years to bring the headcount index down from 53% to 8% it would have taken about 10 years.

This calculation would be deceptive if the same overall growth rate would not have been possible with balanced growth. There may well be a trade-off, arising from limited substitution possibilities in production and rigidities in some aggregate factor supplies; or the trade-off could stem from aggregate fiscal constraints facing the government in supplying key public infrastructure inputs to private production. It is suggestive in this respect that there is a

correlation of -0.414 between the two growth components identified from Table 2. However, this correlation is only significant at the 6% level, and it is clear that there were sub-periods (1983-84, 1987-88 and 1994-96) in which both primary sector growth and combined growth in the secondary and tertiary sectors were both above average. So these data do not offer strong support for the view that more balanced growth would have meant lower growth.

Economic growth is rarely balanced across regions or sectors of a developing economy, and China is no exception. However, it is clear that for China, the pattern of growth has mattered to the evolution of both poverty and inequality measures. The research findings reviewed above suggest that the sectoral and geographic pattern of growth has not been particularly pro-poor. Migration to urban areas helped reduce poverty nationally. However, growth in the primary sector (primarily agriculture) did more to reduce poverty and inequality than growth in either the secondary or tertiary sectors.

The geographic composition of growth also mattered. Progress was geographically uneven with some provinces seeing far more rapid reduction in poverty than others. In particular, the coastal areas fared better than inland areas. The trend rate of decline in the poverty rate was 8% per year for inland provinces, versus 17% for the coastal provinces. However, while provinces with higher rural income growth tended to have higher poverty reduction, by-and-large growth was not higher in the provinces where it would have had the most impact on poverty nationally. This pattern of growth naturally also influenced the evolution of inequality. Rural and (in particular) agricultural growth tended to bring inequality down (Ravallion and Chen, 2004). Rural economic growth reduced inequality within both urban and rural areas, as well as between them.

Has China faced a growth-equity trade-off? One of the most striking aspects of China's success against poverty to emerge from recent research is that there is very little evidence of an aggregate growth-equity trade-off (Ravallion and Chen, 2004). Inequality in China has clearly shown a tendency to rise over time (Figure 5). The regression coefficient of the Gini index on GDP per capita has a t-ratio of 9.22 (a correlation coefficient of 0.90). But this correlation could well be spurious; the Durbin-Watson statistic is 0.45, indicating strong residual auto-correlation. This is not surprising since both inequality and mean income have strong trends, though possibly associated with different causative factors.

A better test is to compare the growth rates with changes in inequality over time.²⁰ Then it becomes far less clear that higher inequality has been the price of China's growth. The correlation between the growth rate of GDP and log difference in the Gini index is -0.05. Now the regression coefficient has a t-ratio of only 0.22 (and a Durbin-Watson of 1.75). This test does not suggest that higher growth *per se* meant a steeper rise in inequality.

The periods of more rapid growth did not bring more rapid increases in inequality; indeed, the periods of falling inequality (1981-85 and 1995-98) had the highest growth in average household income (Table 3). Also, the sub-periods of highest growth in the primary sector (1983-84, 1987-88 and 1994-96) did not come with lower growth in other sectors (Ravallion and Chen, 2004). Nor does one find that the provinces with more rapid rural income growth experienced a steeper increase in inequality; if anything it was the opposite.

To consider one of these periods more closely, Figure 7 gives the GIC for China in 1993-96, which took on an inverted U shape, with highest growth rates observed at around the 25th percentile. The growth rate for the poorest quintile for this sub-period was 10.1% per annum —

²⁰ There is still positive first-order serial correlation of 0.48 in the first difference of log GDP though there is no sign of serial correlation in the residuals from the regression of the first difference of log Gini on log GDP. So the (first-order) differenced specification is appropriate.

above the ordinary growth rate of 8.2%, indicating the extent to which the distributional shift in this sub-period favored the poor. (Note also that the overall rate of growth was higher in this sub-period than for the 1990s as a whole.) Ravallion and Chen (2004) argue that the main reason for this change in the mid-1990s was a sharp reduction in the taxation of farmers, associated with a rise in the government's procurement price of foodgrains. (China had a long-term policy of taxing farmers this way to provide cheap food to urban areas; naturally this was inequality increasing.)

This lack of any evident aggregate trade-off has important implications. On the one hand, it means that growth will tend to reduce absolute poverty. Naturally, with the same growth rate and no rise in inequality, the number of poor in China would be lower; indeed, it would be less than one-quarter of its actual value (a poverty rate in 2001 of less than 1.5% rather than 8%). This calculation would clearly be deceptive if inequality rises with economic growth, as the "price" of that growth. However, the evidence does not support that view. On the other hand, the absence of such a trade-off also means that rising inequality put a serious brake on China's pace of poverty reduction. That is also borne out by the finding of Ravallion and Chen (2004) that the provinces that saw a more rapid rise in rural inequality saw less progress against poverty, not more.

As China's policy makers now realize, it will be harder for China to maintain its past rate of progress against poverty without addressing the problem of rising inequality. To the extent that recent history is any guide to the future, we can expect that the historically high levels of inequality found in many provinces today will inhibit future prospects of poverty reduction — just as we find that the provinces that started the reform period with relatively high inequality

faced a double handicap in future poverty reduction: they had lower subsequent growth and the poor shared less in the gains from that growth.

Other factors point to the same conclusion. It appears that aggregate economic growth in China is increasingly coming from sources that bring more limited gains to the poorest. The low-lying fruit of efficiency-enhancing pro-poor reforms are possibly getting scarce. Inequality is continuing to rise and poverty is becoming more responsive to rising inequality. At the outset of China's current transition period to a market economy, levels of poverty were so high that inequality was not an important concern. That has changed.

It also appears that perceptions of what "poverty" means are evolving in China. It can hardly be surprising to find that the standards that defined poverty 20 years ago have lost relevance to an economy that quadrupled its mean income over that period. China could well be entering a stage of its development in which relative poverty emerges as a more important concern than in the past. Economic growth will then be a blunter instrument for fighting poverty in the future.

India

As in the case of China, it is clear that economic growth has tended to reduce poverty in India. And, as in China, the poverty impact of accelerated growth in the 1990s has been dulled by rising inequality.

Assessing what has been happening to inequality in India has been clouded by a comparability problem between the two main surveys available for the 1990s (Deaton, 2001; Datt and Ravallion, 2002). Figure 8 gives three estimates of the GIC for the 1990s. The "unadjusted" estimate is based on the actual surveys with no attempt to correct for the comparability problem. One of the other two is based on comparable distributions of

consumption per person based on a common “mixed reference period” for categories of consumption as obtained by Sundaram and Tendulkar (2003). The other uses the alternative method of estimating a “common reference period” made by Deaton (2001). In all three cases, the rural and urban distributions are aggregated assuming urban-rural cost-of-living differentials of 33% and 38% for 1993/94 and 1999/00 respectively; these are based on updated poverty lines as used in Ravallion and Datt (2002). Using either adjustment method, the GIC for the 1990s tends to show a U shape, with lowest growth rates for people around the 20th percentile. Overall, the Gini index rose in the 1990s, by either adjustment method. However, it is too early to say if this is going to be the similar to the trend increase in inequality that China has experienced. Looking back over time, rising inequality in India is a recent phenomenon (Figure 9). (Longer term comparisons are only possible using the Deaton method of correcting for the comparability problem in the 1999/00 data.)

As in China, India’s recent rise in inequality appears to have stemmed in part from geographic and sectoral imbalances in the growth process, evident as regional divergence and a lagging rural economy. As in the case of China, one finds that growth in mean rural incomes has been far more effective against poverty in India (Table 1) and that the sectoral composition of growth has been important (Table 4), though tertiary sector growth was relatively more important in India than we find for China. This could well reflect the difference between the two countries in the distribution of agricultural land, which is clearly more unequal in India, which naturally attenuates the impact of agricultural growth on poverty relative to that found in China.

By one estimate, if not for the sectoral and geographic imbalance of growth, the national rate of growth since reforms began in full force in the early 1990s would have generated a rate of poverty reduction that was double India’s historical trend rate (Datt and Ravallion, 2002). States

with relatively low levels of initial rural development and human capital development were not as well-suited to reducing poverty through economic growth.

Higher average farm yields, higher public spending on development, higher (urban and rural) non-farm output and lower inflation were all poverty reducing (Ravallion and Datt, 2002). However, the response of poverty to non-farm output growth in India has varied significantly between states. The differences reflect systematic differences in initial conditions. Low farm productivity, low rural living standards relative to urban areas and poor basic education all inhibited the prospects of the poor participating in growth of the non-farm sector (Ravallion and Datt, 2002). Rural and human resource development appear to be strongly synergistic with poverty reduction through an expanding non-farm economy.

For example, non-farm economic growth in India has not occurred in the states where it would have the most impact on poverty nationally (Datt and Ravallion, 2002). This is clear from Figure 5, which plots the non-farm growth rates by states of India against the (share-weighted) elasticity of poverty reduction with respect to non-farm economic growth at the beginning of the period. It is clear that the non-agricultural growth has not been concentrated in the states where it would have had the greatest impact on poverty nationally. A more pro-poor geographic pattern of growth in India's non-agricultural economy would have required higher growth in states such as Bihar, Madhya Pradesh, Orissa and Uttar Pradesh. As a result, the overall non-farm growth process in India has tended to become less pro-poor over time. This is evident from Figure 10, which plots the elasticity of the national headcount index of poverty to non-farm economic growth over time.²¹

²¹ These are weighted sums of the state-specific elasticities from Ravallion and Datt (2002); the elasticity of the national headcount index is: $\eta = \sum s_j \eta_j (d \ln y_j / d \ln y)$ where η_j is the elasticity of the headcount index w.r.t. non-ag output per capita in state j , s_j is the share of state j in national poverty

Nor has the geographic pattern of agricultural growth in India been particularly pro-poor. The states with higher growth in agricultural yields were not the key states with higher shares of India's poverty; indeed, there is a mild negative correlation, although not statistically significant (Datt and Ravallion, 2002). Agricultural growth as a whole has also lagged relative to India's (primarily urban) non-farm economy.

Which country has had more pro-poor growth?

In addressing this question we must first confront a semantic point that has been a source of some confusion in recent development policy discussions. "Pro-poor growth" has sometimes been taken to mean that poverty falls more than it would have if all incomes had grown at the same rate (Baulch and McCulloch, 2000; Kakwani and Pernia, 2000). This definition focuses on the distributional shifts during the growth process; roughly speaking, for growth to be deemed "pro-poor" the incomes of the poor should grow at a higher rate than those of the nonpoor. By this definition, growth has not been pro-poor in either China or India (after adjusting for the comparability problem in the survey data noted above). In both countries, a distribution-neutral growth process would have had more impact on poverty than actually observed in the 1990s.

However, it is surely problematic to identify a growth process as not being "pro-poor" when the poor benefit as much as they have in India and (especially) China. As the experience of both countries has exemplified, rising inequality during a period of overall economic expansion can come with large absolute gains to the poor. Similarly, a recession will be deemed pro-poor if poor people lose proportionately less than others, even though they are in fact worse off. An alternative definition of "pro-poor growth" proposed by Ravallion and Chen (2003) avoids this problem by focusing instead on what happens to poverty. By this definition "pro-

and the term in () is the empirical elasticity of non-farm output per capita in state i to national non-farm output per capita. I am grateful to Gaurav Datt for suggesting this calculation.

poor growth” is growth that reduces poverty. The focus then shifts to the extent to which growth is pro-poor, i.e., the speed at which poverty falls. Naturally this will depend in part on what happens to distribution, but only in part — it will naturally also depend on what happens to average living standards. By this alternative definition, it is clear that growth has been more pro-poor in China. Figure 12 compares the headcount index in both countries over 1981-2001 on as comparable a basis as is currently feasible with the data available (Chen and Ravallion, 2004). The poverty line is about \$33 per month at 1993 Purchasing Power Parity. It can be seen that China started this period with the higher poverty rate, but soon overtook India.

6. Lessons for development policy and future research

If we accept that inequality is bad for the poor, what should policy makers do about it? First we must be clear on the objective. If we agree that poverty reduction is a far more important overall goal for development policy than reducing inequality *per se* then we should not accept redistributive policies that come at the expense of lower longer-term living standards for poor people. Accepting that there is no aggregate trade-off between mean income and inequality does not mean that there are no trade-offs at the level of specific policies. Reducing inequality by adding further distortions to an economy may well have ambiguous effects on growth and poverty reduction. But nor should it be presumed that there will be such a trade-off with all redistributive policies. The potential for “win-win” policies stems from the fact that some of the factors that impede growth also entail that the poor share less in the opportunities unleashed by growth.

We have learnt that more rapid poverty reduction requires a combination of more growth, a more pro-poor pattern of growth and success in reducing the antecedent inequalities that limit the prospects for poor people to share in the opportunities unleashed by a growth economy.

Even a distribution-neutral growth process — which hardly seems a high standard for “equitable growth” in high-inequality countries — can leave many poor people behind. The challenge for future research is to better understand the specific factors that constrain some poor people from participating in the benefits of a growing economy, and to draw out the lessons for the types of policies that are needed for rapid poverty reduction in addition to promoting economic growth.

A majority of the world’s poor still live in rural areas and this is likely to remain true for some time to come (Ravallion, 2002). It can be expected that agriculture and non-farm rural development will remain a high priority for sectoral policies. However, past interventions have had a mixed record. New approaches based on community-driven development have held promise but need careful monitoring and evaluation, recognizing the likely heterogeneity in performance across different institutional settings depending on how successful local elites are in capturing the gains (Mansuri and Rao, 2004; Galasso and Ravallion, 2005).

The continuing existence of marked regional disparities in living standards has prompted renewed interest in explicit geographic dimensions in policy making, such as “poor-area programs” and attempts to set up “growth poles.” However, many questions remain. Are infrastructure investments in poor areas (often with poor natural resources) effective in reducing poverty? Does it make more sense to move jobs to people, or people to jobs? Is there a trade-off between achieving greater regional equity — such as by focusing on areas with high poverty *rates* but low poverty *densities* — and poverty reduction in the aggregate?

Recognizing that it is typically the poor rather than the rich who are locked out of profitable opportunities for self-advancement by the failures of markets and governments, interventions that make these institutions work better for poor people today can also help promote pro-poor growth in the future. Successful policies can focus on either correcting the

underlying market and governmental failure or on directly intervening to redress the asset inequalities, by fostering accumulation of (physical and human) assets by poor people. One can point to the potential importance of a wide range of policies including sound public investments in rural infrastructure, better policies for delivering quality health and education services to poor people, and policies that allow key product and factor markets (for land, labor and credit) to work better from the point of view of poor people. The combination of interventions needed will naturally depend on country and regional circumstances. There is still much we do not know about the most appropriate policy combinations in specific circumstances, although some pointers have emerged from research. Making the provision of health and education services more responsive to the needs of poor people is likely to be crucial to achieving pro-poor growth in most settings (World Bank, 2004a). In rural economies, security of access to land through tenancy reform and titling programs is arguably no less important (World Bank, 2004b). In some circumstances, rural infrastructure development can also play a decisive role; for example, research has revealed the importance of rural roads to achieving more pro-poor growth processes in rural China (Jalan and Ravallion, 2002) and that quite reasonable rates of return are possible from well-designed poor-area development programs (Ravallion and Chen, 2005). Better instruments for credit and insurance can also help, both in smoothing consumption and underpinning otherwise risky growth-promoting strategies. Removing biases against the poor in taxation, spending and regulatory (including migration) policies can also play an important role. Again taking an example from China, reducing the government's taxation of farmers through foodgrain procurement quotas has been a powerful instrument against poverty (Ravallion and Chen, 2004). China's recent policy to give tax breaks to farmers in poor regions is surely welcome.

The challenge for policy is to combine growth-promoting policies with the right policies to assure that the poor can participate fully in the opportunities unleashed, and so contribute to that growth. If a country gets the combination of policies right then both growth and poverty reduction can be rapid. Get it wrong, both may well be stalled. Future research can help meet this challenge by:

- throwing light on the country-specific and sub-national factors that influence the distributional nature of aggregate growth;
- identifying to what extent those factors are amenable to policy intervention; and
- quantifying the trade-offs between alternative policies for promoting pro-poor growth, embracing both redistributive social policies and alternative growth strategies.

Appendix: Regression decompositions for rates of poverty reduction

Consider first the urban-rural decomposition for the survey mean. The overall mean at date t is $\mu_t = n_t^r \mu_t^r + n_t^u \mu_t^u$ where μ_t^i is the mean for sector $i=r,u$ for rural and urban areas. It is readily verified that the growth rate in the overall mean can be written as:

$$\Delta \ln \mu_t = s_t^r \Delta \ln \mu_t^r + s_t^u \Delta \ln \mu_t^u + [s_t^r - s_t^u (n_t^r / n_t^u)] \Delta \ln n_t^r$$

where $s_t^i = n_t^i \mu_t^i / \mu_t$ (for $i=r,u$) is the income share. We can write down the following regression for testing whether the composition of growth matters:

$$(A1) \quad \Delta \ln P_t = \eta_0 + \eta^r s_t^r \Delta \ln \mu_t^r + \eta^u s_t^u \Delta \ln \mu_t^u + \eta^n (s_t^r - s_t^u \cdot \frac{n_t^r}{n_t^u}) \Delta \ln n_t^r + \varepsilon_t$$

where ε_t is a white-noise error term. The motivation for writing the regression this way is evident when one notes that if the η^i ($i=r,u,n$) parameters are the same then equation (A1) collapses to a simple regression of the rate of poverty reduction on the rate of growth ($\Delta \ln \mu_t$). Thus testing $H_0: \eta^i = \eta$ for all i tells us whether the urban-rural composition of growth matters.

A second decomposition is possible for GDP per capita which we can divide into n sources to estimate a test equation of the following form:

$$(A2) \quad \Delta \ln P_t = \pi_0 + \sum_{i=1}^n \pi_i s_{it} \Delta \ln Y_{it} + \varepsilon_t$$

where Y_{it} is GDP per capita from source i , $s_{it} = Y_{it} / Y_t$ is the source's share, and ε_t is a white-noise error term. In the special case in which $\pi_i = \pi$ for $i=1, \dots, n$, equation (A2) collapses to a simple regression of the rate of poverty reduction on the rate of GDP growth ($\Delta \ln Y_t$).

References

- Aghion, Philippe, Caroli, Eva and Garcia-Penalosa, Cecilia (1999), 'Inequality and Economic Growth: The Perspectives of the New Growth Theories', *Journal of Economic Literature*, 37(4), 1615-1660.
- Alesina, Alberto and Rodrik, Dani (1994), 'Distributive Politics and Economic Growth', *Quarterly Journal of Economics*, 108, 465-90.
- Amiel, Yoram and Frank Cowell, 1999, *Thinking about Inequality: Personal Judgment and Income Distributions*, Cambridge: Cambridge University Press.
- Atkinson, Anthony B. (1970) 'On the Measurement of Inequality', *Journal of Economic Theory* 2:244-263.
- _____, (1987), 'On the Measurement of Poverty,' *Econometrica* 55: 749-764.
- Atkinson, Anthony B. and Andrea Brandolini (2004), "Global World Inequality: Absolute, Relative or Intermediate," mimeo, Nuffield College, Oxford University.
- Banerjee, Abhijit V., and Andrew F. Newman, (1993), 'Occupational Choice and the Process of Development,' *Journal of Political Economy* 101(2): 274-298.
- Bardhan, Pranab, Samuel Bowles and Herbert Gintis (1999), 'Wealth Inequality, Wealth Constraints and Economic Performance,' in A.B. Atkinson and F. Bourguignon (eds) *Handbook of Income Distribution*, vol. I., Amsterdam: North-Holland.
- Baulch, Robert and John Hoddinott (2000), "Economic Mobility and Poverty Dynamics in Developing Countries," *Journal of Development Studies* 36(6): 1-24.
- Baulch, Robert and Neil McCulloch (2000), 'Tracking pro-poor growth.' ID21 insights No. 31. Sussex: Institute of Development Studies.
- Benabou, Roland (1996), 'Inequality and Growth', in Ben Bernanke and Julio Rotemberg (eds), *National Bureau of Economic Research Macroeconomics Annual*, Cambridge: MIT Press, pp.11-74.
- Birdsall, Nancy, Ross, D. and Sabot, R. (1995), 'Inequality and Growth Reconsidered: Lessons from East Asia', *World Bank Economic Review*, 9(3), 477-508.
- Bourguignon, Francois (2001), 'The Pace of Economic Growth and Poverty Reduction.' Paper presented at LACEA 2001 Conference.
- _____, (2003), 'The Growth Elasticity of Poverty Reduction: Explaining Heterogeneity across Countries and Time-periods.' In T. Eichler and S. Turnovsky, eds.,

- Growth and Inequality*. Cambridge, Mass.: MIT Press.
- Chen, Shaohua and Ravallion, Martin (2001), 'How Did the World's Poorest Fare in the 1990s?', *Review of Income and Wealth*, 47, 283-300.
- _____ and _____ (2004), 'How Have the World's Poorest Fared Since the Early 1980s?', *World Bank Research Observer*, 19(2): 141-170.
- Clarke, George R.G., (1995), 'More Evidence on Income Distribution and Growth', *Journal of Development Economics*, 47, 403-428.
- Datt, Gaurav and Ravallion, Martin (1992), 'Growth and Redistribution Components of Changes in Poverty Measures: A Decomposition with Applications to Brazil and India in the 1980s', *Journal of Development Economics*, 38, 275-295.
- _____ and _____ (2002), 'Has India's Post-Reform Economic Growth Left the Poor Behind', *Journal of Economic Perspectives*, 16(3), 89-108.
- Deaton, Angus. 2001. "Adjusted Indian Poverty Estimates for 1999-00," mimeo, Research Program in Development Studies, Princeton University.
- Deininger, Klaus and Squire, Lyn (1998), 'New Ways of Looking at Old Issues: Inequality and Growth', *Journal of Development Economics*, 57(2), 259-87.
- Dollar, David and Aart Kraay, (2002), 'Growth is Good for the Poor,' *Journal of Economic Growth*, 7(3): 195-225.
- Easterly, William (2002), 'Inequality does Cause Underdevelopment: New Evidence,' Working Paper 1, Center for Global Development, Washington DC.
- Fields, Gary S., (2001), *Distribution and Development*. New York: Russell Sage Foundation.
- Foster, James., J. Greer, and E. Thorbecke, (1984), 'A Class of Decomposable Poverty Measures,' *Econometrica* 52: 761-765.
- Galasso, Emanuela and Martin Ravallion, (2005), 'Decentralized Targeting of an Anti-Poverty Program,' *Journal of Public Economics*, 85: 705-727.
- Galor, Oded and Joseph Zeira (1993), 'Income Distribution and Macroeconomics,' *Review of Economic Studies* 60(1), 35-52.
- Gastwirth, J.L., 1971, 'A General Definition of the Lorenz Curve,' *Econometrica* 39: 1037-39.
- Jalan, Jyotsna and Martin Ravallion, 2000, 'Is Transient Poverty Different? Evidence for Rural China,' *Journal of Development Studies*, 36(6): 82-99.

- _____ and _____, 2002, 'Geographic Poverty Traps? A Micro Model of Consumption Growth in Rural China?' *Journal of Applied Econometrics*, 17:329-46.
- Kakwani, Nanak (1993), 'Poverty and Economic Growth with Application to Côte D'Ivoire,' *Review of Income and Wealth*, 39, 121-139.
- Kakwani, Nanak and E. Pernia, (2000), 'What Is Pro-Poor Growth?' *Asian Development Review*. 18(1): 1-16.
- Kraay, Aart (2005), 'When is Growth Pro-Poor? Evidence from a Panel of Countries,' *Journal of Development Economics*, in press.
- Lopez, Humberto (2005), 'Growth and Inequality: Are they Connected?' mimeo, World Bank.
- Mansuri, Ghazala and Vijayendra Rao (2004), 'Community-Based and -Driven Development,' *World Bank Research Observer* 19(1): 1-40.
- Perotti, Roberto (1996), 'Growth, Income Distribution and Democracy: What the Data Say,' *Journal of Economic Growth*, 1(2), 149-87.
- Persson, Torsten and Tabellini, Guido (1994), 'Is Inequality Harmful for Growth?,' *American Economic Review*, 84, 600-621.
- Ravallion, Martin (1995), 'Growth and Poverty: Evidence for Developing Countries in the 1980s', *Economics Letters*, 48, 411-417.
- _____ (1997), 'Can High Inequality Developing Countries Escape Absolute Poverty?,' *Economics Letters*, 56, 51-57.
- _____ (2001), 'Growth, Inequality and Poverty: Looking Beyond Averages', *World Development*, 29(11), 1803-1815.
- _____ (2002), 'On the Urbanization of Poverty,' *Journal of Development Economics* 68, 435-442.
- _____ (2003), 'The Debate on Globalization, Poverty and Inequality: Why Measurement Matters', *International Affairs*, 79(4), 739-754.
- _____ (2004), 'Competing Concepts of Inequality in the Globalization Debate,' *Brookings Trade Forum 2004*, pp.1-38.
- _____ (2005a), 'Looking Beyond Averages in the Trade and Poverty Debate,' Policy Research Working Paper 3461, World Bank, Washington DC.
- _____ (2005b), 'Externalities in Rural Development: Evidence for China,' in Ravi Kanbur and Tony Venables (eds) *Spatial Inequality and Development*, Oxford University

- Press.
- Ravallion, Martin and Monika Huppi (1991), 'Measuring Changes in Poverty: A Methodological Case Study of Indonesia During and Adjustment Period', *World Bank Economic Review*, 5, 57-82.
- Ravallion, Martin and Shaohua Chen (1997), 'What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty?', *World Bank Economic Review*, 11(2), 357-82.
- _____ and _____ (2003), 'Measuring Pro-Poor Growth,' *Economics Letters*, 78(1), 93-99.
- _____ and _____ (2004), 'China's (Uneven) progress Against Poverty,' Policy Research Working Paper 3408, World Bank.
- _____ and _____ (2005), 'Hidden Impact: Household Saving in Response to a Poor-Area Development Project,' *Journal of Public Economics*, forthcoming.
- Ravallion, Martin and Gaurav Datt (1996), "How Important to India's Poor is the Sectoral Composition of Economic Growth?," *World Bank Economic Review*, 10: 1-26.
- _____ and _____ (2002), 'Why Has Economic Growth Been More Pro-Poor in Some States of India than Others?', *Journal of Development Economics*, 68, 381-400.
- Ravallion, Martin and Michael Lokshin (2004), 'Lasting Local Level Impacts of a Crisis,' Policy Research Working Paper 3503, World Bank Washington DC.
- Son, Hyun and Nanak Kakwani (2004), 'Economic Growth and Poverty Reduction: Initial Conditions Matter,' Working Paper 2, International Poverty Center, UNDP.
- Sundaram, K., and Suresh D. Tendulkar (2003), 'Poverty in India in the 1990s: Revised Results for All-India and 15 Major States for 1993-94,' *Economic and Political Weekly*, 38 (November 15), 4865-73.
- World Bank (1990), *World Development Report: Poverty*, New York: Oxford University Press.
- _____ (2000), *World Development Report: Attacking Poverty*, New York: Oxford University Press.
- _____ (2004a), *World Development Report: Making Services Work for Poor People*, New York: Oxford University Press.
- _____ (2004b), *Land Policies for Growth and Poverty Reduction*.

Figure 1: Changes in inequality and growth in the mean between successive surveys, 1980-2000

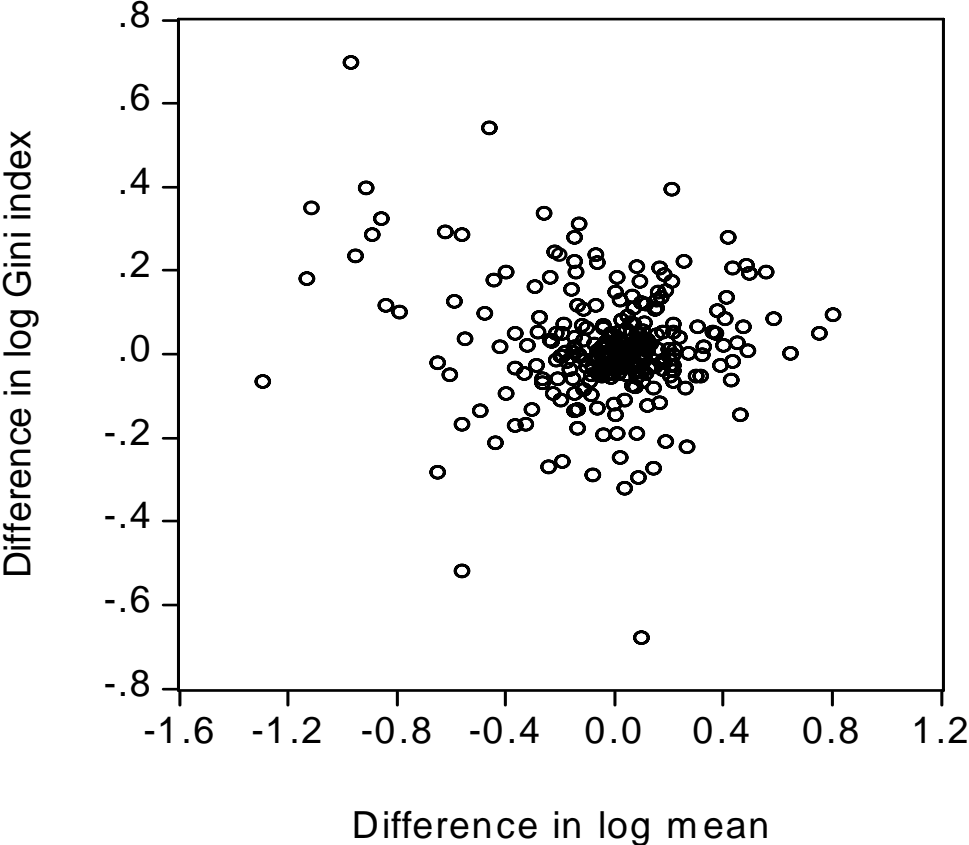


Figure 2: Changes in inequality and growth in the mean between successive surveys, post-1992

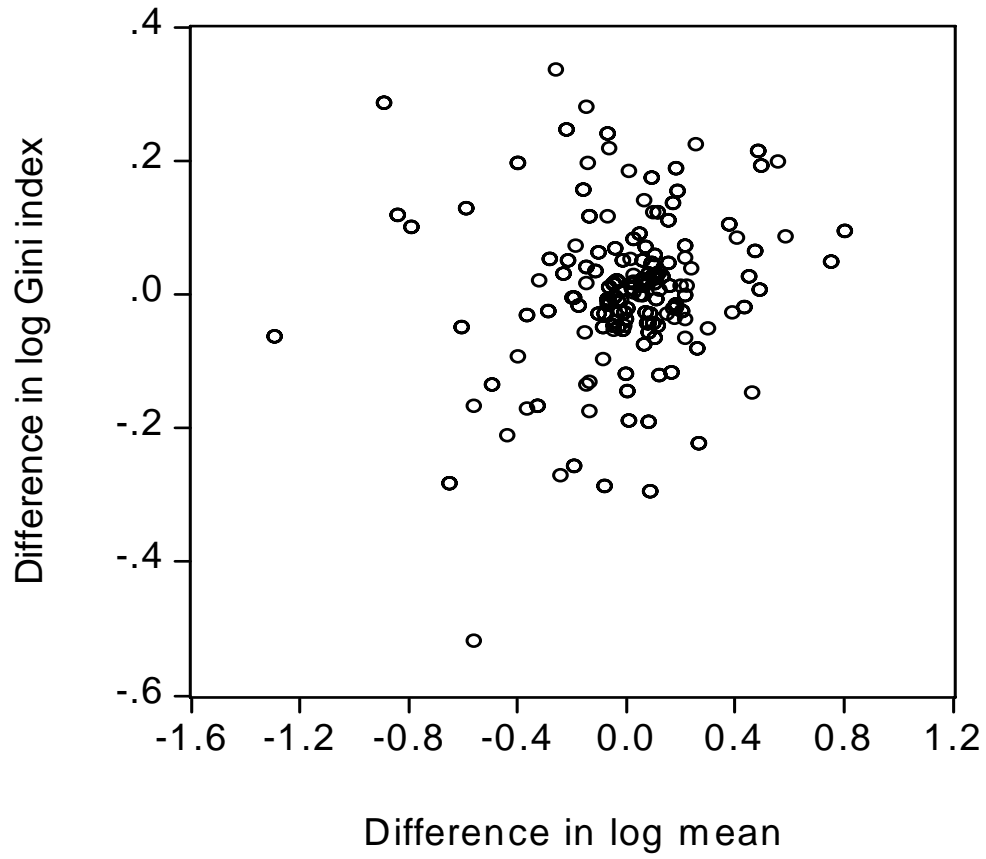
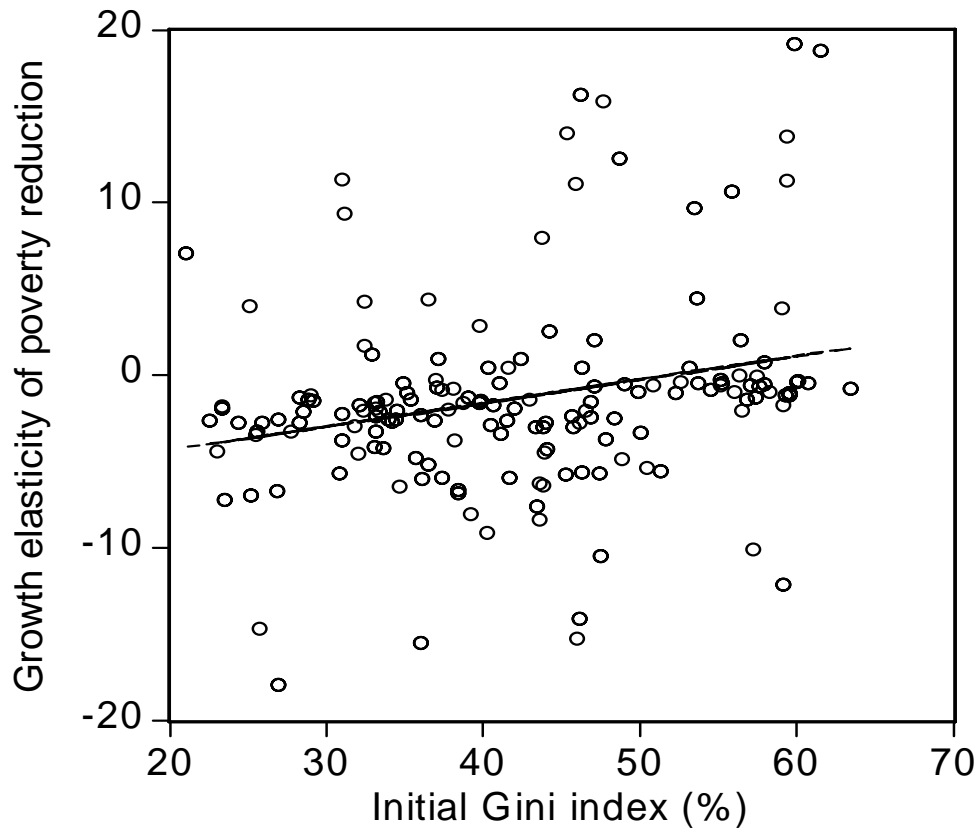
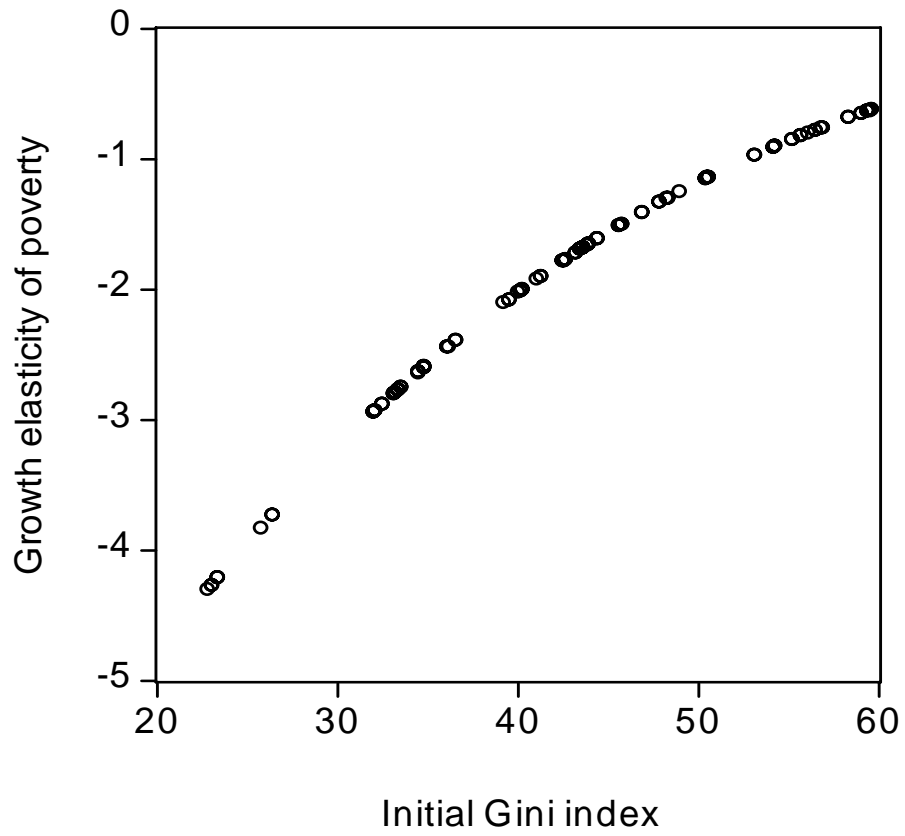


Figure 3: Empirical growth elasticities of poverty reduction against initial Gini index



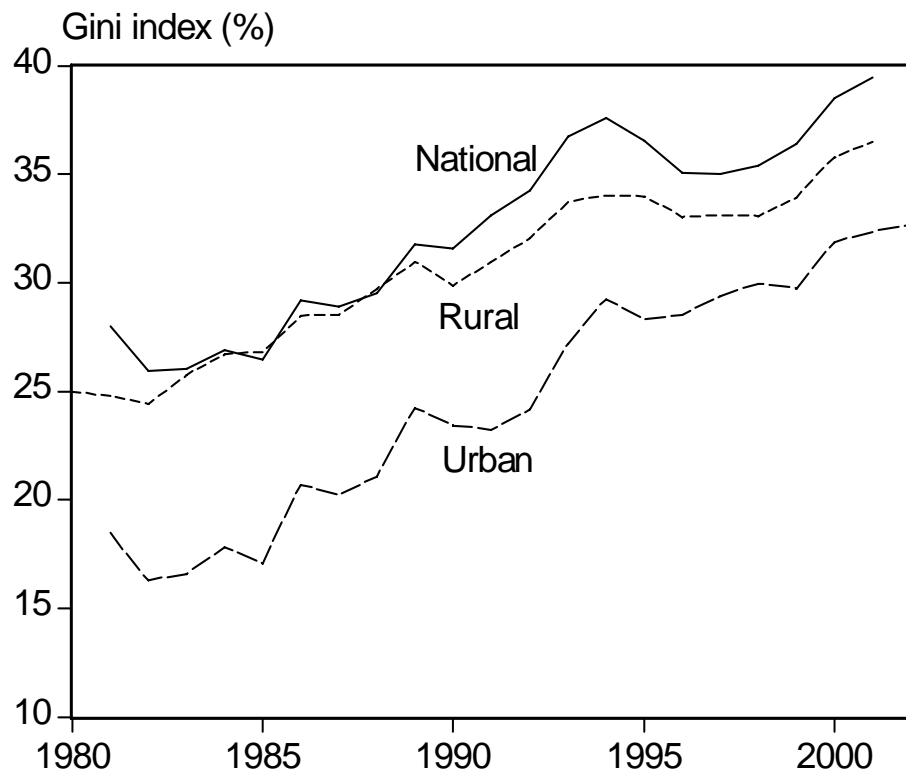
Source: Author's calculations (see text)

Figure 4: Growth elasticity of poverty as a function of the initial Gini index



Source: Authors calculations (see text)

Figure 5: China: Income inequality in rural and urban areas and nationally



Source: Ravallion and Chen (2004).

Figure 6: Growth incidence curve for China, 1990-1999

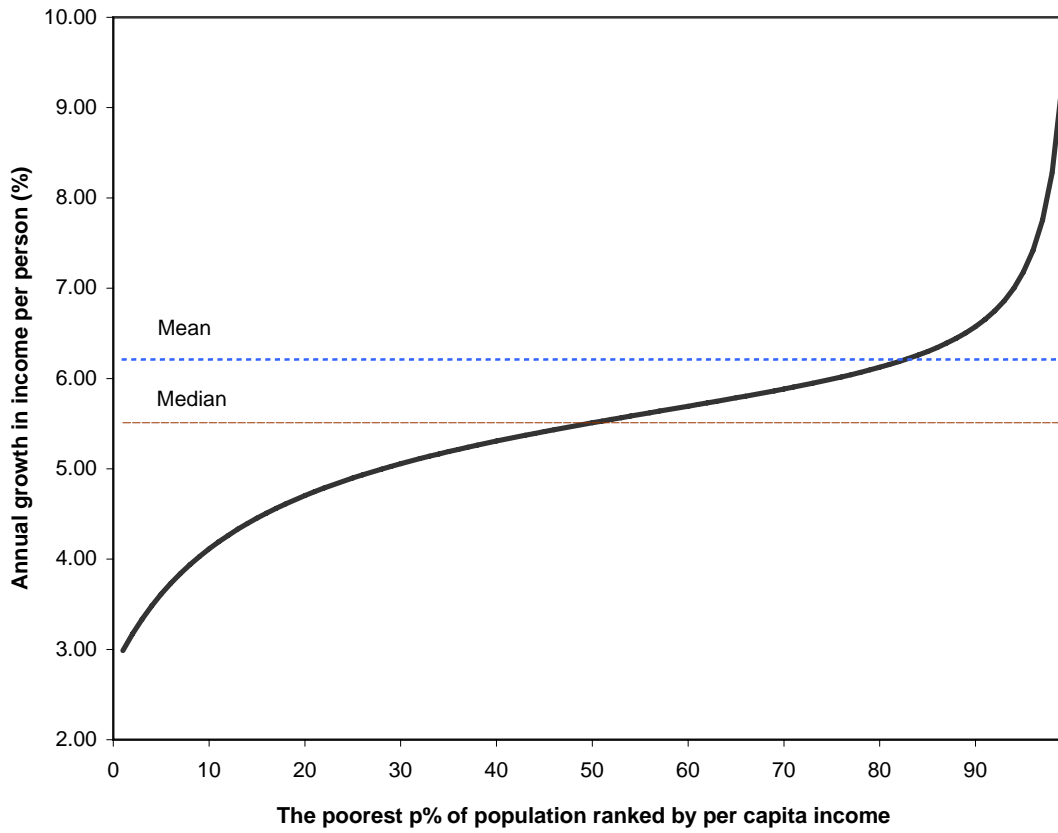


Figure 7: Growth incidence curve for China, 1993-1996

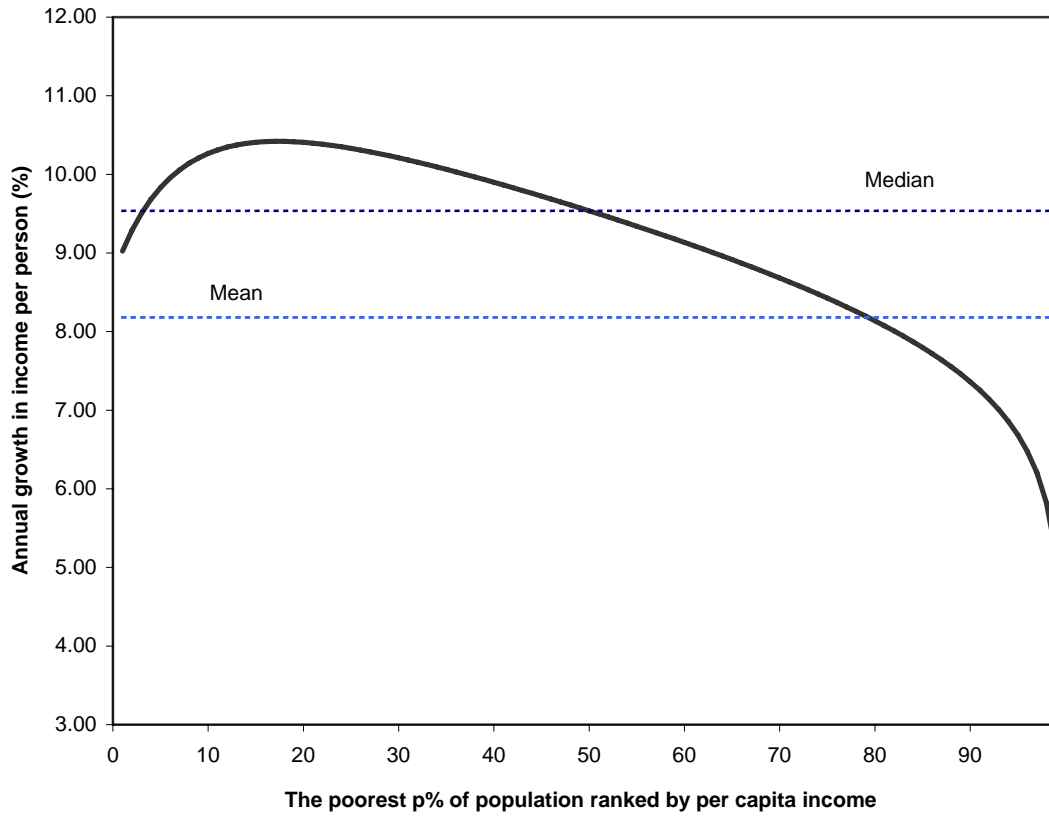


Figure 8: Growth incidence curve for India, 1993/94-1999/00

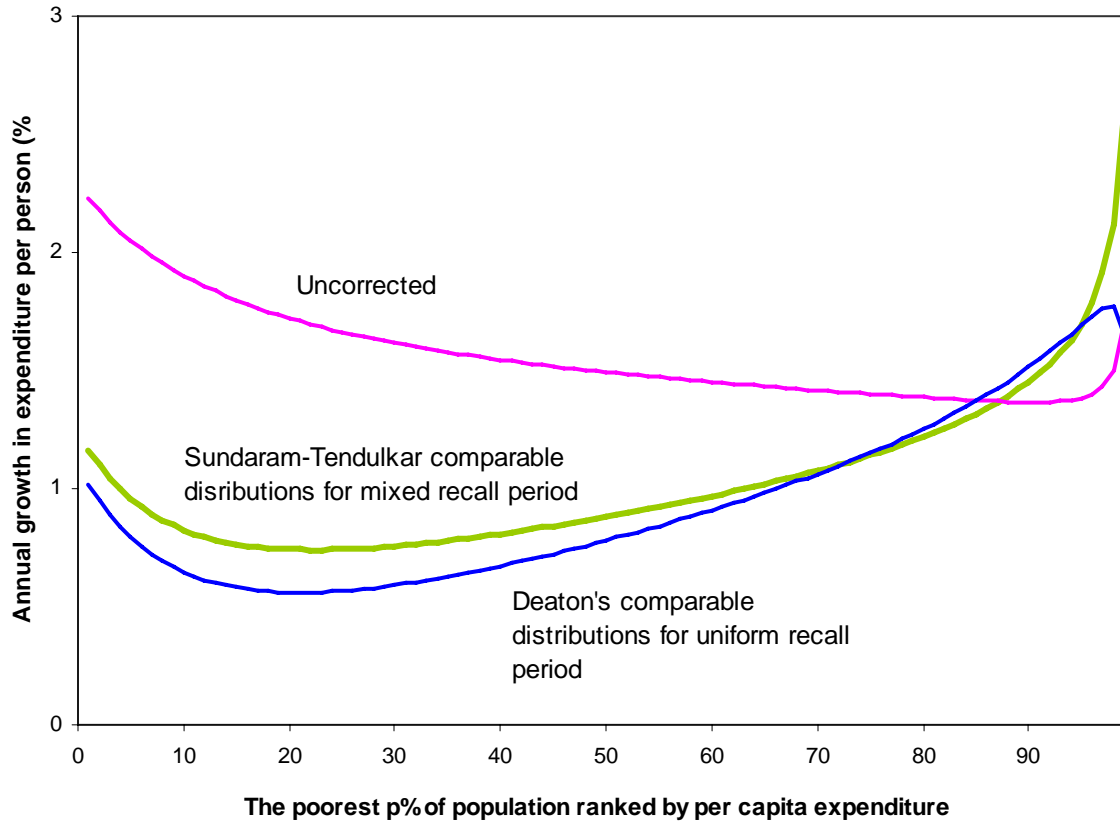


Figure 9: Inequality over time in India

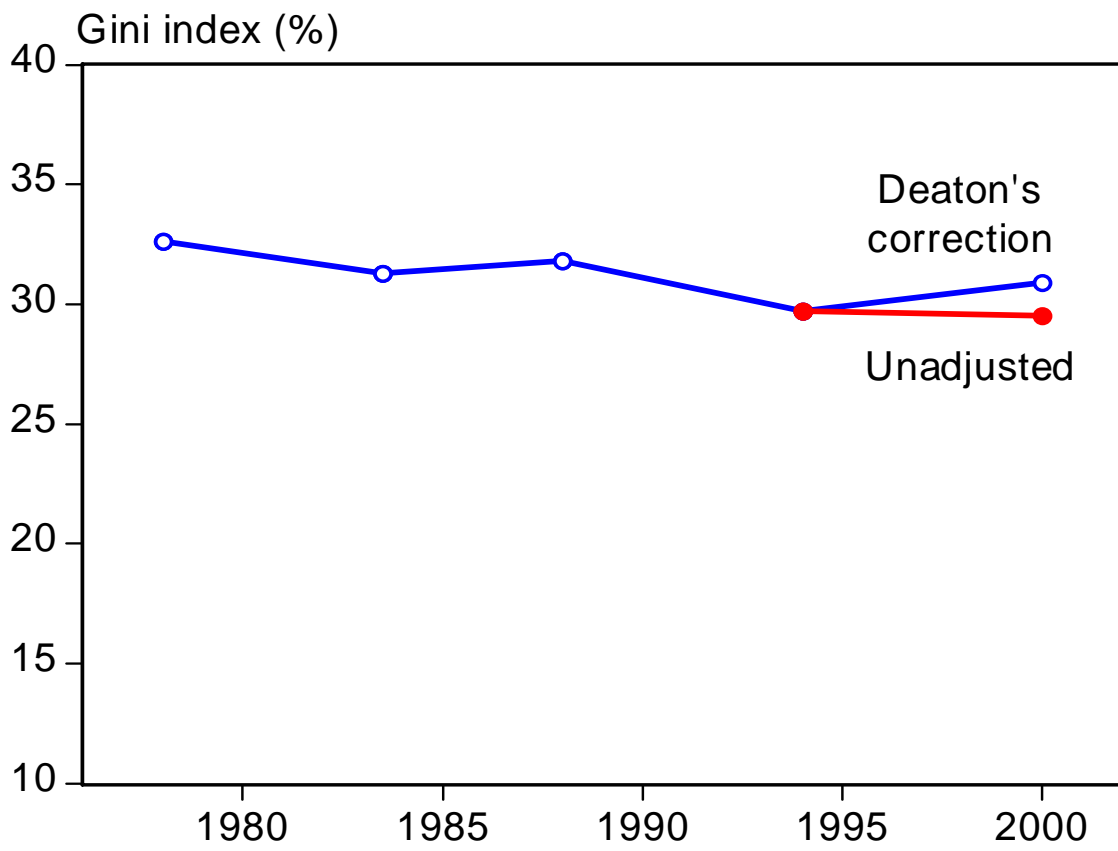
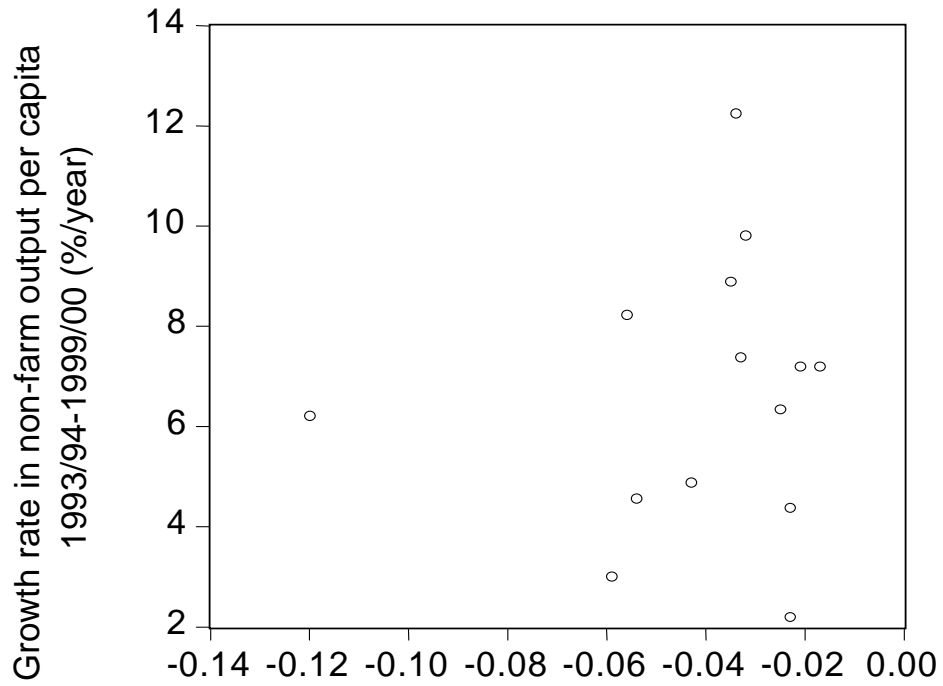


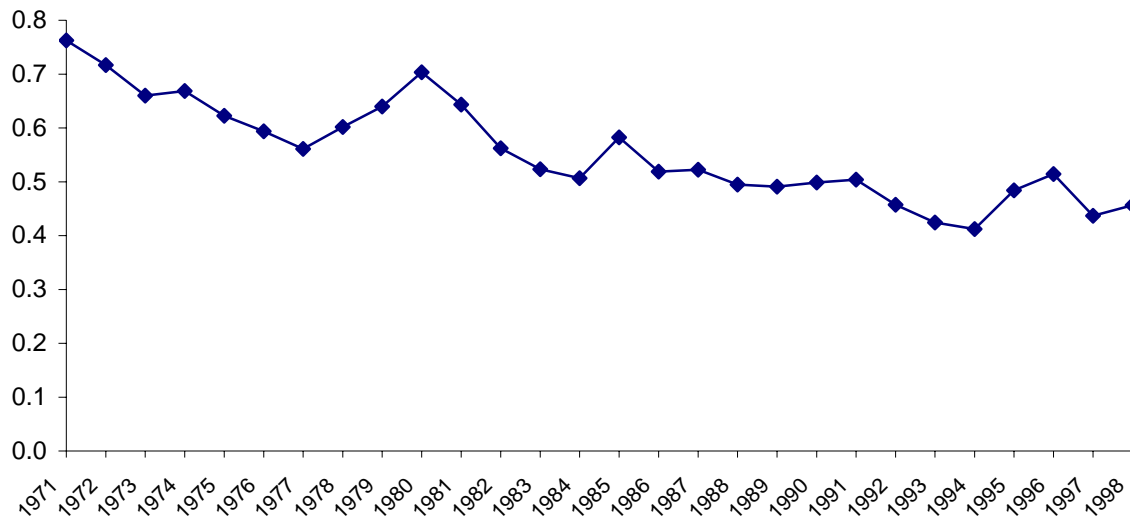
Figure 10: Non-agricultural economic growth in India in the 1990s has not been happening in the states where it would have had the most impact on poverty nationally



Impact on national poverty of non-farm output growth by state (Share-weighted elasticity for 1993/94)

Note: Based on Datt and Ravallion (2002).

Figure 11: Absolute elasticity of all-India headcount index with respect to non-farm economic growth



Source: Author's calculations from the results of Ravallion and Datt (2002)

Figure 12: Poverty incidence in China and India, 1981-2001

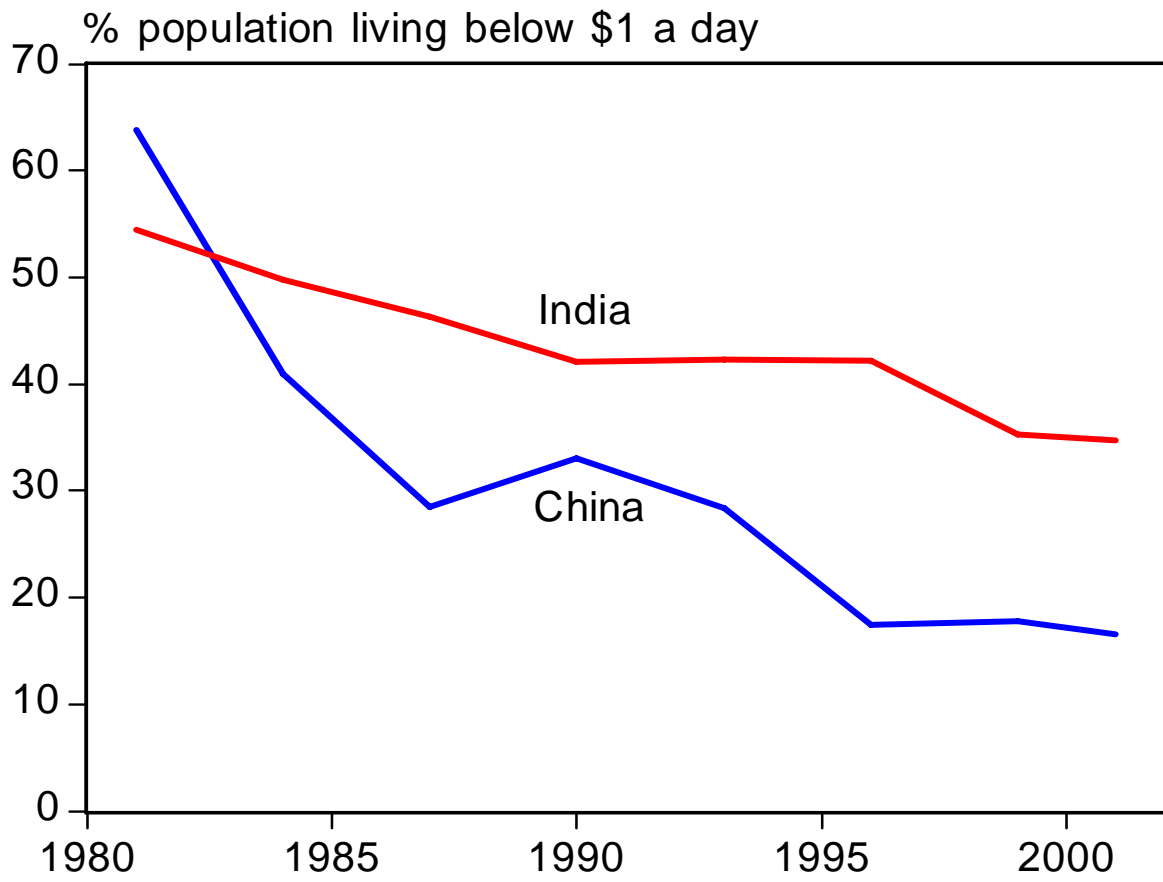


Table 1: Poverty reduction and the urban-rural composition of growth

	China	India
Growth rate of mean rural income (share-weighted)	-2.56 (-8.43)	-1.46 (12.64)
Growth rate of mean urban income (share-weighted)	0.09 (0.20)	-0.55 (-1.37)
Population shift effect	0.74 (0.16)	-4.46 (-1.31)
R ²	0.82	0.90

Source: Ravallion and Datt (1996) (for India) and Ravallion and Chen (2004) (for China)

Table 2: Poverty reduction and the sectoral composition of growth: China

Headcount index (log difference)			
Growth rate of GDP per capita	-2.60 (-2.16)		
Primary (share-weighted)		-8.07 (-3.97)	-7.85 (-4.09)
Secondary (share-weighted)		-1.75 (-1.21)	
Tertiary (share-weighted)		-3.08 (-1.24)	
Secondary+tertiary			-2.25 (-2.20)
R ²	0.21	0.43	0.42

Source: Ravallion and Chen (2004)

Table 3: Inequality and growth in China by sub-periods

	Inequality	Growth rate in household income per capita (%/year)
1. 1981-85	Falling	8.9
2. 1986-94	Rising	3.1
3. 1995-98	Falling	5.4
4. 1999-2001	Rising	4.5

Source: Ravallion and Chen (2004).

Table 4: Poverty reduction and the sectoral composition of growth: India

Headcount index (log difference)	
Growth rate of GDP per capita	-0.99 (-3.38)
Primary	-1.16 (-2.96)
Secondary	3.41 (1.84)
Tertiary	-3.42 (-2.74)
R ²	0.75

Source: Ravallion and Datt (1996)