

Forum essay

A poverty-inequality trade off?

MARTIN RAVALLION*

Development Research Group, World Bank

(Accepted: 17 June 2005)

Abstract. The idea that developing countries face a trade off between poverty and inequality has had considerable influence on thinking about development policy. The experience of developing countries in the 1990s does not, however, reveal any sign of a systematic trade off between measures of absolute poverty and relative inequality. Indeed, falling inequality tends to come with falling poverty incidence. And rising inequality appears more likely to be putting a brake on poverty reduction than to be facilitating it. However, there is evidence of a trade off for absolute inequality, suggesting that those who want a lower absolute gap between the rich and the poor must in general be willing to see lower absolute levels of living for poor people.

Key words: inequality, poverty, growth.

1. Introduction

It is often said that a development strategy that emphasizes pro-poor redistribution will stifle growth and (hence) lead to higher absolute poverty.¹ Similarly, it is often said that policy makers in developing countries should not worry about rising inequality, on the grounds that it is a more-or-less unavoidable by-product of rapid economic growth, which serves to bring down the incidence of absolute poverty. In other words, there is a view that developing countries face a poverty-inequality trade off.

What evidence is there to support this view? If such a trade off existed as a general rule, then one would expect to find signs of its presence in the diverse experiences of developing and transition economies. One would tend to see falling poverty with rising inequality, and rising poverty with falling inequality.

Before we can test this implication we need to be more specific about what we mean by “poverty” and “inequality.” On the former, this paper will confine attention to absolute poverty in the space of consumption or income, meaning that poverty is assessed relative to poverty lines in local currencies that attempt to have fixed purchasing power over commodities across countries. This seems closest to the

* These are the views of the author, and need not reflect those of the World Bank or any affiliated organization.

concept of poverty underlying the view that there is a trade off between poverty and inequality and it is the most widely used concept of poverty in development policy discussions more generally. However, it is certainly not the only concept. Other definitions have been proposed, including “relative poverty” in which the poverty line rises with average income, and poverty measures that allow for aspects of welfare not fully captured by household command over commodities, such as measures that better reflect inequalities within households or access to public goods.

What do we mean by “inequality”? Here there are two possible interpretations of what those who are concerned about a trade off between poverty and inequality mean by “inequality.” The most common concept of inequality found in applied economics is what is often termed “relative inequality” in the theoretical literature.² Relative inequality is a function of the ratios of individual incomes to the mean; this property is implied by the “scale independence axiom” in inequality measurement, which says that when all incomes are multiplied by a constant inequality is unchanged. The alternative concept is *absolute inequality*, which depends on the absolute differences in levels of living, rather than relative differences [13]. A measure of absolute inequality is unchanged if all incomes increase by the same amount. To see the difference between these two concepts, consider an economy with just two households with incomes: \$1,000 and \$10,000. If both incomes double in size then relative inequality will remain the same; the richer household is still 10 times richer. But the absolute difference in their incomes has doubled, from \$9,000 to \$18,000. Relative inequality is unchanged but absolute inequality has risen sharply.³ There is survey-based evidence suggesting that many people (though not the majority) think about “inequality” in absolute terms [1]. In popular writings and policy discussions it is often extremely unclear which of these two concepts is deemed relevant.⁴

The paper looks for evidence of the trade off using estimates of poverty and inequality measures over time for 70 developing countries in the 1990s. The principal finding is that whether one thinks about inequality in relative or absolute terms is crucial to the position one takes on the trade off between poverty and inequality; there is no sign in the available data of such a trade off for relative inequality, but that there is evidence of it for absolute inequality. The following section reviews the theoretical and empirical arguments that have been made in past work. Section 3 then presents new evidence using both relative and absolute inequality measures. Section 4 discusses the results, while Section 5 concludes.

2. Arguments and evidence of a poverty-inequality trade off

The classic theoretical argument for believing in the existence of a poverty-inequality trade off (PIT) in developing countries is based on the Kuznets Hypothesis (KH). Under the KH relative inequality increases in the early stages of growth in a developing country but begins to fall after some point, i.e., the relationship between inequality (on the vertical axis) and average income (horizontal) is predicted to

trace out an inverted U [14]. As typically formalized in the literature, the KH assumes that the economy comprises a low-inequality and low-mean rural sector, and a richer urban sector with higher inequality.⁵ Growth occurs by rural labor shifting to the urban sector. This happens in a rather special way such that a representative slice of the rural distribution is transformed into a representative slice of the urban distribution. Thus (by assumption) distribution is unchanged *within* each sector. Starting with all the population in the rural sector, when the first worker moves to the urban sector inequality must increase, even though the incidence of poverty has fallen. And when the last rural worker leaves, inequality must clearly fall again. Between these extremes, the relationship between inequality and average income will follow an inverted U. The PIT will be found in countries with income less than the turning point of this inverted U.

There are a number of ways in which the KH might fail to hold in practice. One possibility is that the rural–urban migration process may not be distribution-neutral *within* sectors. Selective migration processes may well cloud the overall impacts of rural–urban migration on inequality. Additionally, the Kuznets process of poverty reduction through population urbanization may also account for rather little of the overall change in inequality, given the importance of within-sector growth processes. For example, in tests using decomposition methods for both India and China, the Kuznets process of migration could account for only a small proportion of the evolution of aggregate poverty measures, which were much more heavily influenced by the pattern of growth across sectors and regions of the economy. The “Kuznets component” of the decomposition of changes in poverty accounted for about 10% of total poverty reduction in China in the 1980s and 1990s and even less in India over the period from about 1970 to the mid 1990s.⁶

Though less well formalized than the KH, another argument for believing that poor countries may face a PIT rests on the assumption that governmental controls on economic activity in poor countries are used to keep inequality lower than it would be without those controls. With economic reform, inequality rises as these controls are removed, and with that there is both growth and poverty reduction. This is a widely heard characterization of the process of economic transition from a command economy to a market economy.

This argument is suggestive of the existence of a PIT in a period of economic transition, though not in steady state. There are a number of concerns about the argument even in a transition economy. The assumption that the control regime keeps inequality low is questionable; one can imagine other political-economy stories in which controls restricting market activity are used to protect the living standards of the rich rather than the poor, such as when controls on trade and financial flows are used to maintain high returns to domestic suppliers of scarce factors, which typically do not include unskilled labor, which is the main asset of the poor. Presumably there must also be considerable heterogeneity in the specifics of the reform process and how it impacts on inequality, leading one to question whether inequality would respond similarly in different settings.

What does the existing evidence suggest about the existence of a PIT? There is a large literature testing the KH. Supportive evidence has often been found in cross-sectional comparisons of the level of inequality with mean income. In the early days of testing the KH this was the only type of data available. However, as time series data has accumulated over time it has been possible to see whether inequality tended to rise over time in growing but initially poor countries. A number of studies based on time series evidence have found very little support for the KH.⁷ Indeed, there appear to be very few cases of a trend increase in inequality in the early stages of development [4].

Another strand of the empirical literature has tested for a conditional relationship between poverty and inequality, controlling for mean income. The simple mathematics of how poverty is typically measured tells us that the measure obtained will depend on both the mean of the distribution on which the measure is based and the properties of the Lorenz curve, reflecting the distribution around that mean. The mathematical properties of poverty measures do not provide an unambiguous prediction about what relationship one would expect to see between a standard measure of overall relative inequality and the measure of poverty. For example, an unambiguous outward shift in the Lorenz curve (indicating that relative inequality has risen for all standard measures) while holding the mean constant could be consistent with either an increase or decrease in the most widely used measure of poverty, namely the headcount index (H) given by the proportion of the population living in households with consumption per person less than the poverty line.⁸ If one adds to this ambiguity the possibility that the mean is also changing endogenously with the change in inequality, then the relationship becomes even harder to predict. The idea of a PIT rests on the claim that higher inequality will on balance yield lower poverty. This effect is seen to be transmitted through economic growth, which is assumed to be poverty reducing.

While the effect of higher inequality on a standard measure of absolute poverty at a given mean income is ambiguous in theory, the empirical relationship is somewhat clearer. A number of empirical studies have found that higher inequality is associated with higher poverty at a given mean income.⁹ A recent example can be found in Besley and Burgess [3] who regress the (log) headcount index for the “\$1 a day” line on both the (log) mean income and a measure of inequality, namely the standard deviation of log incomes, itself predicted from the measured Gini index assuming a log-normal distribution.¹⁰ The inequality measure is found to have a significant positive coefficient. This is to be expected; the Besley–Burgess regression can be interpreted as an approximation of the underlying nonlinear formula giving H as a function of the mean and the distribution. However, this is clearly not acceptable evidence of a “reverse PIT,” whereby poverty rises with higher inequality; to test for a PIT we need to know the unconditional relationship between poverty and inequality, letting the mean vary.

There has been one empirical result reported in the literature that might be taken to be evidence against a PIT. In a comment on Besley and Burgess [3], Honohan

[10] reports that the income share of the richest 10% is significant when he adds it to a regression of the headcount index on the mean income of the poorest 90%. This prompts him to ask: “. . . what is it about societies where the rich are richer that tends to results in more people falling into poverty?” (p. 272).

However, there is no real mystery here, and Honohan’s finding cannot be interpreted as evidence against existence of a PIT. To see why, let us suppose initially that all countries have the same Lorenz curve, so that the headcount index in the data varies solely with the mean (for a fixed real poverty line). At any given mean there is only one possible value of H since (by construction) there is no distributional effect in this simple “thought experiment.” However, the overall mean (μ) is linked to the mean of the poorest $p\%$ ($m(p)$), and the income share of the richest $100-p\%$ ($s(p)$) by the identity: $\mu = p \cdot m(p)/[1 - s(p)]$. So if one follows Honohan in regressing H on $m(90)$ rather than μ then one must expect the share of the richest 10% ($s(90)$) to be a significant extra regressor even though there is no distributional effect. One would probably get an excellent statistical fit with a linear regression of $\ln H$ on both $\ln m(90)$ and $s(90)$ (given that $\ln[1 - s(90)]$ will be approximated well by $-s(90)$). Yet this has nothing to do with the distribution factors that Honohan alludes to, since the Lorenz curve is everywhere the same in this example. More realistically, on also allowing for different Lorenz curves there will be another reason for $s(90)$ to be significant in Honohan’s regression, given that $s(90)$ is very likely to be correlated with the differences in the Lorenz curves found in practice, since the data on $s(90)$ is generated by those same Lorenz curves. But again this is just another approximation to the underlying mathematical relationship generating the data. So the results reported by Honohan cannot be interpreted as evidence of a reverse PIT, as they are perfectly consistent with the expectation that the headcount index will fall with the mean and (for certain distributions at least, including log-normality) rise with inequality controlling for the mean.

In the light of the above observations, the issue to be addressed in the rest of this paper is whether there is evidence of an unconditional correlation between poverty and inequality.

3. The evidence for developing economies

Testing for the existence of such a trade off poses a number of empirical problems. The level of poverty in a country is determined by many factors, and it is unclear how well one could control for these using observable characteristics, so as to isolate the relationship with inequality. Instead, the approach taken here is to study the changes over time in both measures of poverty and measures of inequality, on the assumption that this will adequately sweep away the extraneous other factors at country level. However, this also leaves a concern that the changes over time may well contain considerable measurement error. This will cloud the true relationship, but it can also yield a systematic bias in inferences about the casual effect, arising

from the fact that if one has under (over) estimated relative inequality then one will (in all likelihood) have under (over) estimated poverty. This is a serious concern if one is looking for a causal interpretation of the empirical relationship between measures of poverty and measures of inequality. However, this discussion will be concerned more with the descriptive relationship seen in the available data, which is essentially the same data that those who believe in the existence of a PIT are drawing on.

The present discussion will focus on a specific concept of absolute poverty, namely whether people in a given country live in households that consume less than an international poverty line of \$1.08 a day at 1993 Purchasing Power Parity [5].

Regarding the concept of inequality, I will consider two measures. The first is the most widely used single measure of relative inequality, the Gini index based on the same set of surveys used for measuring poverty. The measure of absolute inequality is the absolute Gini index, based on absolute differences in incomes (not normalized by the mean).

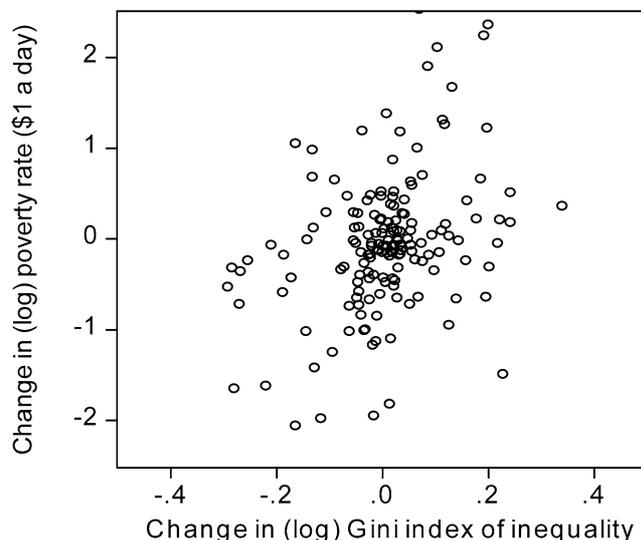
To see if the experience of developing countries is consistent with the existence of such a trade off I have assembled data on about 170 “spells” spanning two surveys for each of about 70 developing countries in the 1990s (with more than one spell for most).¹¹ For each survey one calculates a measure of poverty and a measure of inequality.

Figure 1 gives the results for relative inequality. Far from there being a negative relationship, there is a reasonably strong *positive* (unconditional) correlation. The correlation coefficient is 0.31. The regression coefficient of the proportionate rate of change (difference in logs) in poverty on the rate of change in inequality is 2.84 (with a *t*-ratio corrected for heteroscedasticity of 3.84). The intercept is not significantly different from zero ($t = -0.43$). Thus, poverty incidence does not change (on average) if inequality does not change.

If one focuses on the subset of growing economies, the relationship is even steeper. For this subsample (defined by whether there is positive growth in private consumption per capita) the regression coefficient is 4.16 (t – ratio = 3.95). For the growing economies the intercept is -0.26 and is significantly different from zero ($t = -2.92$).

It might be conjectured that the transition economies would exhibit the PIT. On repeating the above tests, one again finds a significant *positive* correlation between changes in poverty and changes in inequality; across the 50 observations (20 for the 1980s) for Eastern Europe and Central Asia (EECA) the correlation coefficient is 0.62 – even higher than for the full sample. (The regression coefficient of changes in poverty on changes in inequality is 7.72, with a *t*-ratio of 5.48.) The countries for which poverty rose the most in EECA were those for which inequality rose the most.

There is likely to be considerable noise in the short-term movements of these measures, which may make it difficult to detect longer-term trends. Suppose instead that we focus on the longest available spell for each country spanning the 1980s



Note: Each data point represents two household surveys for each country, with multiple observations over time for most countries. The calculations were done from the World Bank's *PovcalNet* data tool, which provides income and expenditure distributions from about 400 household surveys; see <http://iresearch.worldbank.org/povcalnet>. The international poverty line is \$1.08 a day at 1993 Purchasing Power Parity. For further discussion see Chen and Ravallion [5].

Figure 1. Change in poverty plotted against changes in relative inequality.

and 1990s. Then the positive correlation between changes (difference in logs) in poverty and changes in inequality is even higher, at 0.52. The regression coefficient is 3.85 (t -ratio = 4.44).¹²

One cannot dismiss the possibility that the positive correlations found above stem in part from the aforementioned problem of correlated measurement errors in the data on poverty and inequality. There is no obvious way to deal with this concern. All we can say is that the measures we have do not offer any support for the idea of a PIT for relative inequality. The deeper causal interpretation of the relationship seen in the data remains unclear.

Figure 2 repeats the calculations underlying Figure 1 except this time the absolute Gini index is used. Now we do find evidence of a PIT. The correlation coefficient is -0.35 and the regression coefficient is -1.18 ($t = -2.87$). The intercept is not significantly different from zero. One also finds signs of a trade off for the EECA region, though it is only statistically significant at the 6% level (a regression coefficient of -1.62 with a t -ratio of -1.94). Note, however, that measurement errors may still be playing a role here, though this time there is an extra offsetting source of correlated measurement error in the measures of (absolute) poverty and (absolute) inequality, namely the fact that both depend in part on mean income or consumption.

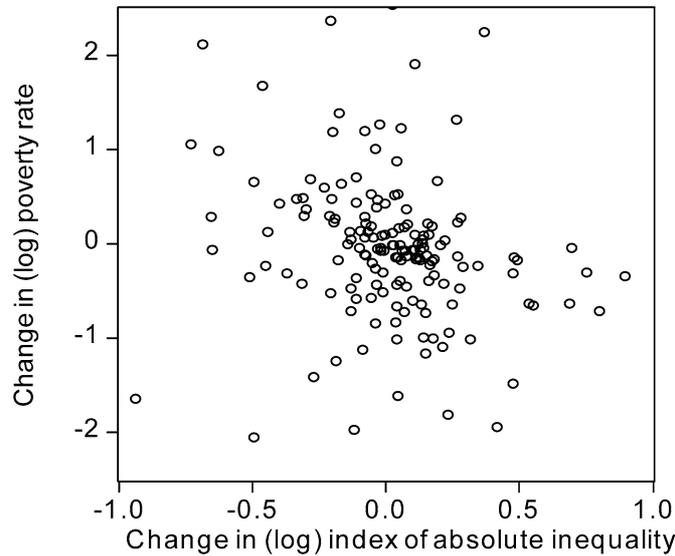


Figure 2. Change in poverty plotted against changes in absolute inequality.

4. Interpretations and implications

Where does the PIT for relative inequality break down empirically? The negative correlation between growth and poverty reduction in these data is strong, confirming past findings on earlier versions of the data set.¹³ For example, in the same data set used in Figure 1, the simple correlation coefficient between growth rates in mean household consumption or income per person as measured in the surveys and the proportionate rates of poverty reduction in the 1990s is -0.56 ;¹⁴ the correlation is weaker using national accounts consumption data as the source of growth rates (a correlation of -0.26), though this could reflect the imperfect matching of survey periods to national accounts data.

Where the trade-off breaks down is in the supposed relationship between rising relative inequality and economic growth. There is only a weak positive correlation between growth in per capita consumption and the proportionate changes in relative inequality; using growth rates in survey means the correlation is 0.16 and using growth rates from the national accounts it is 0.03. In periods of positive growth, inequality *falls* about 40% of the time. Nor is there any sign in the data of a stronger relationship between growth and rising inequality amongst the countries with initially lower income; for example, the simple correlation coefficients given above are even lower if one focuses on developing countries with below average income.¹⁵

This also makes it clear why one finds much stronger evidence of a trade off between poverty reduction and *absolute* inequality. Essentially what is driving this PIT is the correlation between growth and absolute poverty reduction. If relative inequality does not change on average with growth then absolute inequality will

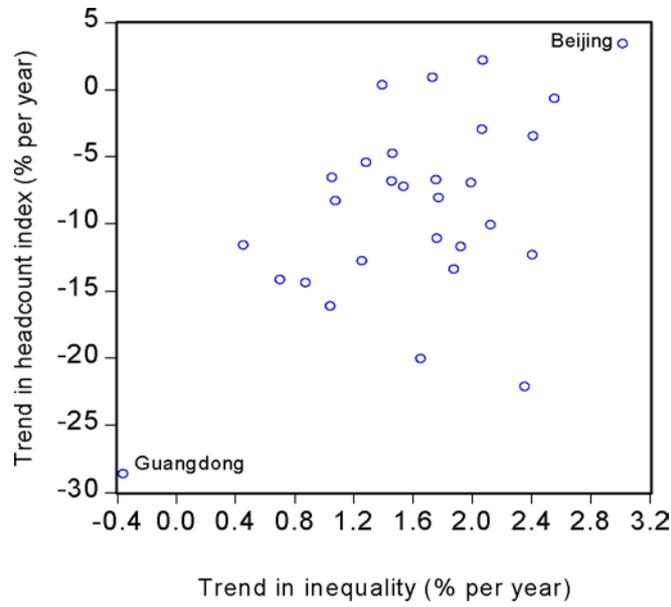
tend to rise with growth. Thus we see a negative relationship between changes in poverty and changes in absolute inequality.

There is clearly a lot more to the changes in poverty seen in these data than can be explained by changes in inequality. There are also changes in mean incomes that are not accountable to changes in inequality. And there are interaction effects between changes in inequality and initial conditions. A potentially important example in the present context is that higher inequality tends to have more impact on poverty when the incidence of poverty is lower. However, even if one confines the analysis to countries with above average initial poverty rates (greater than 15%) there remains a significant positive correlation between rising relative inequality and rising poverty; the regression coefficient of the proportionate rate of change in the poverty measure on the proportionate change in the inequality measure drops to 0.88 but is still significant (a t -ratio of 2.55 with $n = 65$). That remains true if one also confines attention to growing countries with above average initial poverty rates.

The cross-country correlations reported in the previous section do not rule out the possibility that rising inequality is (more or less) unavoidable in some specific growing economies, such as when the growth comes from freeing up prior controls on economic activity that had been used to keep inequality low. There is likely to be heterogeneity in various dimensions (including policies) that can yield exceptions to the generalizations suggested by these aggregate comparisons.

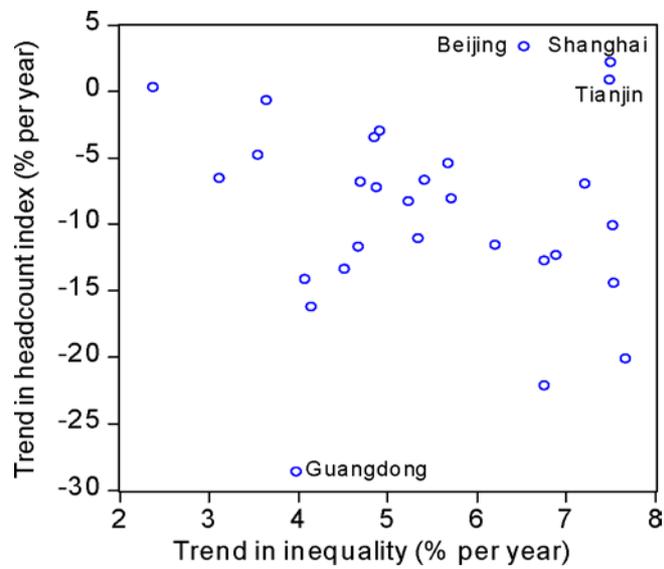
Here it is of interest to also point to some recent evidence for China, drawing on Ravallion and Chen [20]. China is often cited as an example of a country that achieved rapid poverty reduction with rising inequality. Over the last two decades of the twentieth century, the proportion of the population of China living under \$1 a day fell by about two percentage points per year on average [5]. As is well known, China has also enjoyed a period of sustained (and high) economic growth. And the available evidence suggests that inequality has been rising since the mid-1980s though not continuously [20].

However, the evidence for China does not offer much support for the view that rising relative inequality has facilitated the country's rapid poverty reduction. Two empirical facts lead one to question the view that China has faced a serious trade-off between poverty and (relative) inequality. Firstly, 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. Secondly, the provinces that saw a more rapid rise in rural inequality saw *less* progress against poverty, not more. This is illustrated by Figure 3, which plots the provincial trend rates of change in the headcount index of poverty for rural areas against the trends in the Gini index using time series data for each province spanning the period from the mid 1980s to 2001. (The trends are calculated by regressing the log of each measure on time.) One finds that the provinces with lower increases in inequality had higher rates of poverty reduction; the correlation coefficient is 0.52 ($t = 3.14$; $n = 29$).



Source: Ravallion and Chen [20].

Figure 3. Trend in rural headcount index against trend in relative Gini index for provinces of China, mid-1980s to 2001.



Note: Author's calculations based on Ravallion and Chen [20].

Figure 4. Trend in rural headcount index against trend in absolute Gini index for provinces of China, mid-1980s to 2001.

What about absolute inequality? Figure 4 gives the corresponding scatter plot using instead the trends in the absolute Gini index. There is no correlation; the correlation coefficient is a mere -0.07 . However, there are four provinces that stand out as untypical. The correlation rises to -0.20 if one drops Guangdong, though this is still not statistically significant. There are also three provinces where rural poverty rose (Beijing, Shanghai and Tianjin) that are unusual for being the rural hinterlands of large and thriving urban areas. If one drops these four provinces then evidence of a significant PIT does emerge. The correlation coefficient is -0.56 ($n = 25$) and the regression coefficient is -2.11 with a t -ratio of 3.20. So the bulk of the sub-national data for China does offer support for the existence of a trade off between poverty and absolute inequality.

5. Conclusions

Looking at the experience of 70 developing and transition economies in the 1990s, this paper finds no sign of a systematic trade off between absolute poverty incidence and relative inequality. Indeed, lower (higher) poverty tends to come hand in hand with lower (higher) relative inequality. The main reason why the trade off is not found in these data is that economic growth shows little correlation with changes in relative inequality. There is clearly a lot more to the changes we see in measures of inequality and poverty than is captured by the simple idea of a poverty-inequality trade off.

The story changes dramatically if one switches to the concept of absolute inequality. Then the evidence suggests that rising (falling) inequality is associated with falling (rising) poverty. This highlights the crucial importance to development policy discussions on this topic of the concept of inequality one uses. While those who think about inequality in relative terms will reject the idea of a poverty-inequality trade off as a generalization of experience across countries, those who are more inclined to think about inequality in absolute terms will tend to see such a trade off.

Notes

¹ See for example the cover story of *The Economist* magazine, March 13–19, 2004 (p. 13, pp. 69–72), entitled “A question of justice?”

² For a survey of alternative theoretical approaches to the measurement of inequality see Cowell [6].

³ For further discussion of the implications of absolute inequality for development policy debates on growth and inequality see Ravallion [18].

⁴ See again the aforementioned cover story of *The Economist* magazine, March 13–19, 2004; it is unclear whether the characterization of “inequality” is absolute or relative. Also see the article in the May 27, 2000, issue of the same magazine (p. 94).

⁵ For a more precise formulation of the KH, and necessary and sufficient conditions for the inverted U for various inequality measures see Anand and Kanbur [2].

⁶ On India see Datt and Ravallion [8], on China see Ravallion and Chen [20].

⁷ Ravallion [16], Ravallion and Chen [19] and Bruno et al. [4]. Also see the discussion in Fields [9].

⁸ This follows from the well-known fact that H is the point at which the slope of the Lorenz curve equals the ratio of the poverty line to the mean. Depending on precisely how the Lorenz curve shift, the headcount index could fall or rise.

⁹ Examples include Ravallion and Huppi [21], Datt and Ravallion [7], Kakwani [11], Besley and Burgess [3], Kraay [12]; for an overview see Lipton and Ravallion [15].

¹⁰ This is an unnecessary assumption. Instead one can use the empirical distribution.

¹¹ The spells were formed from the World Bank's *PovcalNet* data tool, which provides income and expenditure distributions from about 400 household surveys; <http://iresearch.worldbank.org/povcalnet>.

¹² I deleted two extreme outliers that are almost certainly large measurement errors. However, a highly significant positive correlation remains without trimming the data.

¹³ See, for example, Ravallion [16], Ravallion and Chen [19] and Ravallion [17].

¹⁴ The regression coefficient of the difference in log poverty rate on the difference in log mean is -2.13 with a standard error of 0.42 .

¹⁵ Focusing on countries with below median consumption per capita at the start data, the correlation coefficient is 0.01 using growth rates in survey means and -0.05 using growth rates from the national accounts.

References

1. Amiel, Y. and Cowell, F.: *Thinking about Inequality: Personal Judgment and Income Distributions*, Cambridge University Press, Cambridge, MA, 1999.
2. Anand, S. and Kanbur, S.M.R.: The Kuznets process and the inequality-development relationship, *J. Development Economics* **40** (1993), 25–52.
3. Besley, T. and Burgess, R.: Halving global poverty, *J. Economic Perspectives* **17**(3) (2003), 3–22.
4. Bruno, M., Ravallion, M. and Squire, L.: Equity and growth in developing countries: Old and new perspectives on the policy issues, in V. Tanzi and Ke-young Chu (eds), *Income Distribution and High-Quality Growth*, MIT Press, Cambridge, MA, 1998.
5. Chen, S. and Ravallion, M.: How have the world's poorest fared since the early 1980s?, *World Bank Research Observer* **19**(2) (2004), 141–170.
6. Cowell, F.: Measurement of inequality, in A.B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*, North-Holland, Amsterdam, 2000.
7. Datt, G. and Ravallion, M.: Growth and redistribution components of changes in poverty measures: A decomposition with applications to Brazil and India in the 1980s, *J. Development Economics* **38** (1992), 275–295.
8. Datt, G. and Ravallion, M.: India's checkered history in the fight against poverty: Are there lessons for the future?, *Economic and Political Weekly* **31** (1996), 2479–2486.
9. Fields, G.S.: *Distribution and Development*, Russell Sage, New York, 2001.
10. Honohan, P.: Inequality and poverty, *J. Economic Perspectives* **18**(2) (2004), 271–272.
11. Kakwani, N.: Poverty and economic growth with application to Côte D'Ivoire, *Review of Income and Wealth* **39** (1993), 121–139.
12. Kraay, A.: When is growth pro-poor? Evidence from a panel of countries, *J. Development Economics*, forthcoming.
13. Kolm, S.: Unequal inequalities I, *J. Economic Theory* **12** (1976), 416–442.
14. Kuznets, S.: Economic growth and income inequality, *American Economic Review* **45** (1955), 1–28.
15. Lipton, M. and Ravallion, M.: Poverty and policy, in J. Behrman and T.N. Srinivasan (eds), *Handbook of Development Economics*, Vol. 3, North-Holland, Amsterdam, 1995.

16. Ravallion, M.: Growth and poverty: Evidence for developing countries in the 1980s, *Economics Letters* **48** (1995), 411–417.
17. Ravallion, M.: Growth, inequality and poverty: Looking beyond averages, *World Development* **29**(11) (2001), 1803–1815.
18. Ravallion, M.: Competing concepts of inequality in the globalization debate, in S. Collins and C. Graham (eds), *Brookings Trade Forum 2004: Globalization Poverty and Inequality*, Brookings Institution, Washington, DC, 2004.
19. Ravallion, M. and Chen, S.: What can new survey data tell us about recent changes in poverty and distribution?, *World Bank Economic Review* **11**(2) (1997), 357–382.
20. Ravallion, M. and Chen, S.: China's (uneven) progress against poverty, Working Paper 3408, World Bank, 2004 (<http://econ.worldbank.org/resource.php?type=5>).
21. Ravallion, M. and Huppi, M.: Measuring changes in poverty: A methodological case study of Indonesia during an adjustment period, *World Bank Economic Review* **5** (1991), 57–82.