



Growth, Inequality and Poverty: Looking Beyond Averages

MARTIN RAVALLION*

World Bank, Washington, DC, USA

Summary. — The available evidence suggests that the poor in developing countries typically do share in the gains from rising aggregate affluence, and in the losses from aggregate contraction. But there are large differences between countries in how much poor people share in growth, and there are diverse impacts among the poor in a given country. Crosscountry correlations are clouded in data problems, and undoubtedly hide welfare impacts; they can be deceptive for development policy. There is a need for deeper micro empirical work on growth and distributional change. Only then will we have a firm basis for identifying the specific policies and programs that are needed to complement growth-oriented policies. © 2001 Elsevier Science Ltd. All rights reserved.

Key words — economic growth, inequality, poverty

1. INTRODUCTION

The recent backlash against globalization has given new impetus to an old debate on whether the poor benefit from economic growth. The following quotes from *The Economist* represent well the two main opposing views on the matter:

Growth really does help the poor: in fact it raises their incomes by about as much as it raises the incomes of everybody else...In short, globalization raises incomes, and the poor participate fully (*The Economist*, May 27, 2000, p. 94).

There is plenty of evidence that current patterns of growth and globalization are widening income disparities and hence acting as a brake on poverty reduction (Justin Forsyth, Oxfam Policy Director, Letter to *The Economist*, June 20, 2000, p. 6).

Here we seem to have irreconcilable positions about how much the world's poorest benefit from the economic growth that is fueled by greater openness to foreign trade and investment. *The Economist's* own article is adamant that such growth is poverty reducing, drawing on a recent study by Dollar and Kraay (2000) which found that average incomes of the poorest quintile moved almost one-for-one with average incomes overall. In commenting on *The Economist's* article, Oxfam's Policy Director seems equally confident that rising inequality is choking off the potential benefits to the poor, in seeming contradiction to the

Dollar and Kraay results and earlier results in the literature pointing in the same direction.¹

As this paper will argue, however, there is some truth in both the quotes above. Indeed, it is not difficult to reconcile these two views, with important implications for development policy. In critically reviewing the arguments in this debate, I will draw heavily on evidence from a new compilation of household-level data for developing countries. The following section discusses these data. Section 3 looks at what they show about how much the poor have benefited from rising average living standards in developing countries, and how much they have lost from contractions. Section 4 looks at how distribution has been changing, to see if there is evidence to support the second quote above. The section first looks at how aggregate distribution in the developing world has been changing in the 1990s, and then it looks at what has been happening at country level. The paper then considers in more detail the ways in which

* Helpful comments on this paper were received from Nancy Birdsall, Bill Easterly, Gary Fields, Paul Isenman, Ravi Kanbur, Aart Kraay, Branko Milanovic, Giovanna Prennushi, Dominique van de Walle, Nicolas van de Walle, Michael Walton and participants at presentations at the World Bank and the World Institute of Development Economics Research. These are the views of the author, and need not reflect those of the World Bank or any affiliated organization. Final revision accepted: 15 June 2001.

distribution matters to the outcomes for the poor—both as an impediment to growth (Section 5) and as an impediment to poverty-reducing growth (Section 6). Section 7 then points to some potential pitfalls in drawing policy implications from the evidence of a weak correlation between growth and distributional changes across countries. Section 8 concludes with some observations about directions for future research.

2. NEW EVIDENCE ON AN OLD DEBATE

Data on poverty and inequality are obtained from household surveys, in which random samples of households are interviewed using a structured questionnaire. The main data I draw on here relate to “spells” defined by the periods of time spanning two successive household surveys for a given country. From the latest update of the data base on which the World Bank’s tabulations of income distribution are based (Chen & Ravallion, 2001), one can assemble two or more household surveys over time for about 50 developing countries, to create 120 such spells, mostly in the 1990s.² The estimates of poverty and inequality measures were done from the primary data (rather than using secondary sources), so that it was possible to eliminate obvious inconsistencies in existing compilations from secondary sources. Comparisons over time between any two surveys use the same indicator of economic welfare, which was either income or expenditure per person; half the time it is expenditure, which is taken to be the preferred indicator. Imputed values are included for income or consumption-in-kind from own-farm output. All measures are population weighted (taking account of household size and sample expansion factors). The underlying household surveys are nationally representative.

The data are not without problems.³ Among the concerns about the data used here, there are clearly underlying differences (between countries and over time) in the original household surveys that were the source of the data on household incomes and expenditures. There are also concerns about how best to deflate nominal values for changes in the cost-of-living; the available consumer price indices do not always reflect well the spending behavior of the poor. On top of these problems, there is likely to be underestimation of incomes and spending in household surveys, particularly (but probably

not only) by the rich, who often do not want to participate, or are hard to reach, or deliberately understate their incomes or spending. Nothing much can be done to fix these problems. One can still take partial account of the data problems by using methods of analysis that are not likely to be too sensitive to the errors in the data.

In examining the effect of growth on poverty there is also a question: “growth of what?” We want to know whether the poor are sharing in the growth in average living standards. There are, however, two quite distinct, and largely independent, sources of data on a country’s average welfare, as measured by households’ command over commodities. The level of private consumption expenditure (PCE) per capita from the national accounts (NAS) is widely used for this purpose. On the other hand, measures of average household living standards are available from the same household surveys used to measure poverty.

These two measures do not agree in general, either in the levels or in their growth rates. This is not surprising, given the differences in coverage, definitions and methods. There are the aforementioned problems in survey data. But national accounts have their own data problems. For example, PCE is typically determined residually in the NAS, after accounting for other uses of domestic output and imports at the commodity level. In developing countries, there are concerns about how well both output and consumption by unincorporated (“informal sector”) businesses is measured, though it is not clear how this would affect NAS consumption. A further problem is that it is not generally possible to separate the spending by nonprofit institutions (such as nongovernmental organizations, religious groups, and political parties) from that of households. In many developing countries, the nonhousehold sectors that are implicitly lumped together with households appear to be sizable and possibly growing, so PCE may well overstate the growth rate in household welfare. There are also consistency problems between the two sources, such as arising from imperfect matching between survey dates (which also vary between types of commodities, according to assumed recall periods) and the accounting periods used in the NAS.

There are differences in the extent of these data concerns both between regions and between types of surveys. India stands out as an unusual case in the 1990s. The growth rates in

consumption that we have seen in the national accounts for India in the 1990s have not been reflected in the main national household survey of expenditures on consumption (the National Sample Survey). This divergence is naturally putting a brake on how much poverty reduction we are seeing in the survey data during this period of economic growth (Datt, 1999). At the same time, there are signs that measured inequality is increasing, which is also slowing the rate of poverty reduction given the rate of growth (Ravallion, 2000).

How one interprets the data for India depends critically on why we are seeing this rising divergence between the two data sources on consumption. One interpretation assumes that all consumptions are being underestimated by the surveys, and so concludes that poverty is falling faster than the survey data suggest (Bhalla, 2000). While agreeing that the surveys are probably missing a share of the aggregate consumption gains, an alternative interpretation is that the problem is more likely to be due to underestimation of consumption by the nonpoor. The latter interpretation would appear to accord better with our limited knowledge of the problems of underreporting and noncompliance in consumption and income surveys (see, for example, Groves & Couper, 1998). The fact that the divergence is correlated with growth (over time and across states) of India is also consistent with an income effect on survey underestimation, which one expects to hold also between households (Ravallion, 2000). If the problem is entirely due to underreporting of consumption by the nonpoor, who are nonetheless correctly weighted in the survey design, then one will still get the poverty measures right. But, there could well be problems of sample weighting and underestimation of consumption by the poor, leading to an underestimation of the rate of poverty reduction.

If one is willing to discount income (rather than expenditure) surveys for measuring average levels of economic welfare, and if one puts aside the (highly problematic) data from the transition economies of Central and Eastern Europe for growth rates, then the tests for bias reported in Ravallion (2001) do not point to a systematic overall discrepancy between national accounts and survey-based estimates of aggregate consumption. (This holds in the aggregate across countries; large discrepancies can still be found for specific countries, in both directions.) Nonetheless, it is notable that in the aggregate, and for most regions, the elasticity

of the survey mean to NAS consumption growth is less than one (even though the difference is often not statistically significant). This could well be an attenuation bias due to measurement error. By implication, elasticities of measured poverty to NAS growth will be less than those implied by the measured elasticities of poverty to growth in the survey mean.

The fact that the mean from the surveys is consistent with the data used to calculate poverty measures makes it an appealing candidate for measuring the growth rate. This creates a further problem, however, namely that survey measurement errors can create a spuriously high correlation between poverty measures and the means of the distributions on which those measures are based. The fact that there is measurement error in the surveys (probably creating a spurious negative correlation between measured poverty and the measured mean) speaks to the use of econometric methods that are robust to this type of problem. Examples will be given later.

3. POVERTY REDUCTION AND GROWTH IN AVERAGE LIVING STANDARDS

There is little or no correlation in these data between growth in average household income per person and the change in measured inequality. The correlation coefficient between the annualized change in the log of the Gini index and the annualized change in the log of the survey mean is -0.09 ($n = 115$). The correlation is even lower if one uses growth rates in consumption from the national accounts (a correlation coefficient of 0.01). This finding is consistent with previous research. Earlier versions of this data set also indicated that growth in average household income per person and the change in measured inequality are virtually orthogonal (Ravallion, 1995; Ravallion & Chen, 1997). Similarly, Dollar and Kraay (2000) find that, across countries, log mean income of the poorest quintile (inferred from distributional shares and GDP per capita) changes one-to-one with the overall log GDP per capita. This is equivalent to saying that the share of the poorest quintile is uncorrelated with log GDP per capita.

It does not follow, however, that growth raises incomes of the poor "...by about as much as it raises the incomes of everybody else" (in the quote from *The Economist* at the

beginning of this paper). Finding that the share of income going to the poor does not change on average with growth does not mean that growth raises the incomes of the poor as much as for the rich. Given existing inequality, the income gains to the rich from distribution-neutral growth will of course be greater than the gains to the poor. For example, the income gain to the richest decile in India will be about four times higher than the gain to the poorest quintile; it will be 19 times higher in Brazil.⁴ The fact that, on average, the rich will tend to capture a much larger share of the increment to national income from growth than the poor is directly implied by the empirical results in the literature, including Dollar and Kraay (2000).

Of course, if distributional shares do not change on average then the poor will gain in absolute terms: growth is poverty reducing, and contraction is poverty increasing. Figure 1 plots the proportionate changes in the poverty rate against the growth rate in average income. The poverty measure is the proportion of people living below \$1/day (using 1993 Purchasing Power Parity exchange rates), though other poverty lines show a similar pattern.⁵

The figure also gives the regression line that fits the data best. The line virtually passes through the origin, implying that the average rate of poverty reduction at zero growth is

zero—consistent with the pattern of zero change in inequality on average. The line has a slope of -2.50 with a (heteroskedasticity corrected) standard error of 0.30 ($R^2 = 0.44$). This can be thought of as an overall “growth elasticity” of poverty, since the two variables are proportionate changes. Thus for every 1% increase in the mean, the proportion of the population living below \$1/day (at 1993 Purchasing Power Parity) falls by an average of 2.5%. For example, in a large enough sample of countries for which exactly half of the population lives below \$1/day, a 3% increase in the mean will bring that proportion down to about 0.46. A 3% drop in mean income will push the poverty rate up to about 0.54 on average.

There is no indication in the data that the elasticity is any different when the mean is increasing versus decreasing; one cannot reject the null hypothesis that the elasticity is the same in both directions (the t -statistic is 0.11). So there is no sign that distributional changes help protect the poor during contractions in average living standards.

The relationship looks similar if one uses PCE per capita from the national accounts instead of the mean from the survey, although then the correlation is not as strong, and the elasticity is -1.96 with a considerably higher standard error of 0.89 (though still statistically




Figure 1. Poverty tends to fall with growth in mean household income or expenditure. (Note: Based on data for 47 developing countries in the 1980s and 1990s (multiple spells for most countries). The horizontal axis is the annualized change in the log of the real value of the survey mean; the vertical axis is the annualized change in the log of the percentage of the population living below \$1/day at 1993 Purchasing Power Parity. The figure has been trimmed of extreme values, but this does not alter the line of best fit indicated.)

significant at the 3% level). This is partly because of measurement problems, such as the fact that survey periods do not match exactly the periods used in National Accounts. It is also partly because changes in PCE can arise solely from the nonhousehold sector of the economy (notably spending by nonprofit organizations). Nor is there any sign that the elasticity to growth from the national accounts is any different in expansions versus contractions.

A possible concern about this estimate of the average growth elasticity of poverty reduction is that there may be negatively correlated measurement errors in the rate of poverty reduction and the rate of growth in the survey mean. If the second survey overestimated the mean for some reason (relative to the first survey) it will probably overestimate the rate at which poverty is falling. To check for a bias due to this problem, I used the growth rate in private consumption per capita from the national accounts as the instrumental variable for estimating the regression line in Figure 1, i.e., as the predictor of the growth rate in the survey mean. This assumes that the errors in the national accounts growth rate are uncorrelated with the errors in the growth rates based on the survey means. This instrument is not valid, however, for the countries of Eastern Europe and Central Asia for which there is no correlation between the growth rates from the surveys and those from the national accounts (Ravallion, 2001). So I dropped the data for that region. This estimation method gave a growth elasticity of -2.07 with a standard error of 0.72 (significant at the 1% level).

Therein lies the truth in the first quote at the start of this paper. The incidence of absolute poverty in developing countries tends to fall with growth. This is not a new point; indeed, the empirical relationship has been well known for some time (Bruno *et al.*, 1998, provide a survey). But it is worth emphasizing in the context of the recent debate.

Looking behind the averages, however, the experience is diverse. Even ignoring extreme values, the 95% confidence interval of the last estimate above of the growth elasticity implies that a 1% rate of growth in average household income or consumption will bring anything from a modest drop in the poverty rate of 0.6% to a more dramatic 3.5% annual decline. We will now try to better understand this variance in growth elasticities of poverty.

4. IS RISING INEQUALITY IMPEDING POVERTY REDUCTION?

Let us look first at distribution in the developing world as a whole in the 1990s. In the same way that rising inequality in one country can clearly put a brake on prospects for poverty-reducing growth, rising inequality in the developing world as a whole can inhibit overall poverty reduction. Has that been happening?

The proportion of the population of the developing world living in households with consumption per capita less than about \$1/day in 1998 (at 1993 Purchasing Power Parity) is estimated to be 23% which was only five percentage points lower than in 1987 (Chen & Ravallion, 2001). The total number of poor by this standard was about the same in 1998 as in 1987, with roughly 1.2 billion people living below \$1/day. Chen and Ravallion (2001) try to assess what role worsening distribution played in aggregate poverty reduction during the 1990s. They simulate what would have happened if there had been no change in the overall interpersonal distribution for developing and transitional countries during 1987–98. In other words, all household consumptions and incomes grow at the same rate, given by the growth rate in the (population-weighted) survey mean over their entire data set. The 1987 Lorenz curve of interpersonal consumption for the developing world as a whole would thus remain fixed over the period. If it were true that distribution is worsening over time in the developing world as a whole then this distribution-neutral simulation would give lower poverty in 1998 than actually observed. Chen and Ravallion find, however, that the poverty rate in 1998 would have been 24.4% in the distribution neutral case, instead of 23.4% as calculated from the data.

It can be inferred from this that there was no worsening in the overall interpersonal distribution from the point of view of the poor. Indeed, the actual distributional changes were slightly pro-poor, since the measured poverty rate in 1998 is slightly lower than the simulated rate without any change in distribution. On investigating this finding more closely, one finds that the difference is almost entirely attributable to growth in China. If one takes China out of the calculation then the actual poverty rate in 1998 of 25.2% is almost exactly the same as the simulated rate (25.9%). So income distribution has not been deteriorating overall in the 1990s, from the point of view of the poor.⁶

This aggregate picture hides more than it reveals. The previous section pointed to the heterogeneity in the gains to the poor from a given rate of growth. Underlying this heterogeneity lies the fact that during spells of growth or contraction one sees changes in inequality over time within most developing economies—changes in both directions.

Table 1 divides the data points of Figure 1 (each spell representing two surveys for a given country) into four groups, according to whether the mean is increasing or not, interacted with whether inequality is increasing or not. Even in the countries in which inequality is rising with growth in average living standards, poverty is falling on average. But it typically falls at a much slower rate than in countries experiencing more equitable growth. The median rate of decline in the proportion of the population living below \$1/day among countries with both rising average income and rising inequality was 1.3% per year (Table 1). By contrast, the median rate of poverty reduction was seven times higher, at about 10% per year, among the countries that combined growth in average living standards and falling inequality. Among contracting economies it also mattered greatly what was happening to inequality; when inequality was rising while average living standards fell, the poverty rate was rising by a dramatic 14% per year on average, while with falling inequality the poverty rate rose by less than 2%.

There have been plenty of cases of rising inequality during spells of growth. Indeed, inequality increases about half the time (Table 1; also see Ravallion & Chen, 1997). Therein lies the truth in the second quote at the beginning of this paper. The first quote implicitly averages over this diversity; the second looks not at the averages, but the cases in which the poor share little in the gains from growth.

The fact that we are seeing plenty of cases of rising inequality during spells of growth does not, however, imply that the rising inequality is putting a brake on the rate of poverty reduction. It cannot be concluded from the information in Table 1 that the growing economies with rising inequality could have achieved something like a 9.6% rate of poverty reduction—instead of 1.3% on average—if only inequality had been falling. For that to hold one requires the assumption that the growth rate would have been no lower with falling inequality. Possibly there is an aggregate tradeoff between growth and inequality reduction. That depends critically on exactly how the reduction in inequality is achieved. The next section considers this point further.

5. INEQUALITY AS AN IMPEDIMENT TO GROWTH

One way inequality can matter to the rate of poverty reduction is through the rate of growth in average income. There are a number of arguments that have been made as to why greater equality can actually be good for growth, belying the presumption of an aggregate tradeoff. A seemingly plausible argument points to the existence of credit market failures such that people are unable to exploit growth-promoting opportunities for investment in (physical and human) capital (see, for example, Aghion, Caroli, & Garcia-Penalosa, 1999; Benabou, 1996). 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.⁷

Crosscountry comparisons of growth rates provide some support for the claim that countries with higher initial inequality in incomes

Table 1. *Diverse impacts on poverty underlie the fact that changes in inequality are uncorrelated with economic growth*^a

		What is happening to average household income between the surveys?	
		Falling	Rising
What is happening to inequality between the surveys?	Rising	16% of spells Poverty is rising at a median rate of 14.3% per year	30% of spells Poverty is falling at a median rate of 1.3% per year
	Falling	26% of spells Poverty is rising at a median rate of 1.7% per year	27% of spells Poverty is falling at a median rate of 9.6% per year

^aBased on 117 spells between two household surveys covering 47 developing countries in the 1980s and 1990s. Poverty is measured by the % of the population living below \$1/day at 1993 Purchasing Power Parity. Inequality is measured by the Gini index.

experienced lower rates of growth controlling for other factors such as initial average income, openness to trade and the rate of inflation.⁸ The robustness of this finding has been called into question in some studies. There are difficult problems in identifying this relationship empirically, and the results in the literature have not been robust to alternative specifications, such as allowing for country fixed effects (Barro, 2000; Forbes, 1997; Li & Zou, 1998).

Again, there are a number of concerns about the data and methods used. There are measurement errors in both the levels and changes in measured income inequality, including comparability problems between countries and over time arising from survey error (sampling and non-sampling) and heterogeneity in survey design and processing (see, for example, Atkinson & Brandolini, 1999). One expects that this will matter more to tests which allow for country fixed effects than to standard growth regressions, since the signal-to-noise ratio could well be quite low for changes in measured inequality in existing data sets. Greater attenuation bias should be expected in the fixed-effects regressions of growth on inequality. Using a pooled regression of growth on inequality, Knowles (2001) finds that trimming the data set to reduce the comparability problems changes the results obtained in important ways. Knowles finds, however, that using more recent and more comparable measures of inequality in consumption expenditures does indicate significant negative effects of inequality on growth.

Another concern is that spurious inequality effects in an aggregate growth regression can arise from the assumptions made in aggregating across micro relationships, given credit market failures (Ravallion, 1998). The direction of this bias could go either way, though results for China indicate that regional aggregation hides the adverse effect of inequality on growth. The validity of the common assumption that initial inequality has a linear effect on aggregate growth is also questionable; Banerjee and Duflo (1999) find evidence that changes in income inequality are bad for growth, which ever way the changes go. The choice of control variables in identifying the relationship is also open to question; for example, past tests of the effect of inequality on growth have controlled for the human capital stock, yet reducing investment in human capital is presumably one of the ways that inequality matters to growth.

On balance, the existing evidence using crosscountry growth regressions appears to

offer more support for the view that inequality is harmful to growth than the opposite view, which was the prevailing view in development economics for decades. That does not imply, however, that any reduction in inequality will enhance growth; indeed, it can have the opposite effect if it comes at the expense of other factors that are also known to matter to growth. Reducing inequality by adding further distortions to external trade or domestic economy will have ambiguous effects on growth and poverty reduction.

Given the concerns about past tests based on crosscountry aggregates, it is of interest to ask if there might be some other way of testing for an effect of initial distribution on growth. Returning to the various theories about why initial distribution might matter, one finds that many of the proposed models share some strong and testable implications for micro data. An example is the common feature of a number of the theoretical models based on credit-market failures that individual income or wealth at one date is an increasing concave function of its own past value. This implication of the class of models of distribution-dependent growth based on credit market failures is testable on micro panel data; Lokshin and Ravallion (2001) provide supportive evidence in panel data for Hungary and Russia.⁹

As with macro tests of whether inequality is bad for growth, finding the appropriate non-linearity in household-level income dynamics would not constitute a case for public redistribution as a means of stimulating aggregate growth. But, with the right data, dynamic micro models of income or consumption can be augmented to allow for (possibly endogenously placed) public programs.¹⁰ Micro structural modeling of growth in the presence of specific redistributive interventions may offer hope of a deeper understanding of the policy implications.

6. INEQUALITY AS AN IMPEDIMENT TO PRO-POOR GROWTH

Even when inequality is not rising, a high initial level of inequality can stifle prospects for pro-poor growth. In an economy where inequality is persistently low, one can expect that the poor will tend to obtain a higher share of the gains from growth than in an economy in which inequality is high. To put this another way, an important determinant of the rate of

poverty reduction is the *distribution-corrected rate of growth* in average income, given by a measure of initial equality (100 minus the measure of inequality) times the rate of growth. Indeed, the distribution-corrected growth rate knocks out the ordinary growth rate when both are used in a regression for the rate of poverty reduction (Ravallion, 1997). It is not the rate of growth that matters, but the distribution-corrected rate of growth.

One can represent this in the form of a very simple model in which the proportionate rate of change in the incidence of poverty (P) between surveys is directly proportional to the distribution-corrected rate of growth; on adding an error term, this can be written as:

$$\Delta \ln P_{it} = \gamma(1 - G_{i,t-\tau})\Delta \ln Y_{it} + \varepsilon_{it}, \quad (1)$$

where the difference is taken between surveys that are τ years apart (which varies between countries and over time), $G_{i,t-\tau}$ is the Gini index (between zero and one) for country i at the beginning of the spell, Y_{it} is real value of the survey mean at date t , and γ is a parameter to be estimated. Using the same data as in Figure 1, I obtained an estimate of -3.74 for γ , with standard error of 0.68 (this is very close to the estimate in Ravallion, 1997, on an earlier and smaller data set). A possible concern about this estimate is that there may be (negatively) correlated measurement errors in the changes in P and Y . Using the growth rate in PCE per capita from the national accounts as the instrumental variable for the growth rate in the survey means and dropping the observations for Eastern Europe and Central Asia (where the instrument fails) I found a lower estimate of γ , namely -2.94 , with standard error of 1.18.

The elasticity of poverty to growth declines appreciably as the extent of initial inequality rises. Consider a per capita growth rate of (say) 2% per annum (roughly the mean for low-income countries in the 1990s). With $\gamma = -3$ a country with high inequality (a Gini index of 60% say) can expect to see a rate of poverty reduction of 2.4% per year. By contrast, a relatively low-inequality country, with a Gini of 30%, can expect a rate of poverty reduction of 4.2% per year.

The above results are unrevealing about what specific aspects of inequality matter. The theoretical arguments based on credit-market failures point to the importance of asset inequality, not income inequality *per se*. There is evidence of adverse effects of asset inequality in growth (Birdsall & Londono, 1997; Deininger & Ol-

into, 2000, both using crosscountry data; Ravallion, 1998, using regional data for China).

Some clues have been found by comparing rates of poverty reduction across states of India, for which we can compile a long series of reasonably comparable survey data back to about 1960. The analyses of these data confirm that economic growth has tended to reduce poverty in India. Higher average farm yields, higher public spending on development, higher (urban and rural) nonfarm output and lower inflation were all poverty reducing (Ravallion & Datt, 2001). But, the response of poverty to nonfarm output growth in India has varied significantly between states. The differences reflected 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 nonfarm sector. Rural and human resource development appear to be strongly synergistic with poverty reduction though an expanding nonfarm economy.

7. "NO CORRELATION" DOES NOT MEAN "NO IMPACT"

We have seen that the data suggest little or no correlation between growth and changes in inequality across countries. The same holds for indicators of growth promoting policies for which significant correlations have rarely been found, one way or the other. This is confirmed by Dollar and Kraay (2000), who find negligible correlation between changes in inequality and indicators of policy reform, including greater openness.¹¹ If there is no effect on inequality then the outcomes for the poor depend solely on the growth effects.

There are three main reasons to be cautious in drawing implications for policy from this lack of correlation between growth and changes in inequality. First, this apparent distribution-neutrality of reform (on average) could simply reflect the fact that changes in inequality are not well measured as a rule. For example, it should be emphasized again that although the main data set used above has been constructed to try to eliminate as many of the problems as possible, there are still changes in survey design that can add considerable noise to the measured changes in inequality.

Second, the data relate to averages within countries. Aggregate inequality or poverty may

change relatively little over time, and yet there are both gainers and losers at all levels of living. Indeed, in cases in which the survey data have tracked the same families over time (“panel data”), it is quite common to find considerable churning under the surface; Baulch and Hoddinott (2000) provide evidence of this for a number of countries. Some of this reflects measurement error, but probably not all, since the changes seen in the data are partially explicable in terms of observable characteristics and measurable shocks (see, for example, Jalan & Ravallion, 2000, using data for rural China).

One can find that many people have escaped poverty while others have fallen into poverty, even though the overall poverty rate may move rather little. For example, comparing household incomes immediately after the 1998 financial crisis in Russia with incomes of the same households two years earlier, one finds a seemingly small 2% point increase in the poverty rate. But, this was associated with a large proportion of the population (18%) falling into poverty, while a slightly smaller proportion (16%) escaped poverty over the same period (Lokshin & Ravallion, 2000).

People are often hurting behind the averages. Panel data and observations from the ground can reveal this, but the aggregate statistics cannot. It is important to know the aggregate balance of gains and losses, but it will be of little consolation to those suffering to be told that poverty is falling on average.

A third reason why the low correlations found between policy reform and changes in overall inequality can be deceptive is that starting conditions vary a lot between reforming countries. Averaging across this diversity in initial conditions can readily hide systematic effects. This argument warrants further elaboration since it holds lessons for policy.

One obvious way that countries differ is in their initial level of economic development. It has been argued that greater openness to external trade will have very different effects on inequality depending on the level of economic development—increasing inequality in rich countries and decreasing it in poor ones (Wood, 1994, makes a qualified argument along these lines). The opposite outcome is possible, however, when economic reforms, including greater openness to external trade, increase demand for relatively skilled labor, which tends to be more inequitably distributed in poor countries than rich ones.¹² Geographic disparities in access to infrastructure

also impede prospects for participating in the growth generated by reform, and these disparities tend to be correlated with incomes.

A simple test for the effect of openness on inequality is suggestive. Using the same data set as Li, Squire, and Zou (1998), I found no significant effect of exports as a share of GDP on the Gini index across 50 countries (100 observations). The regression included controls for schooling, financial sector development, urbanization and the black-market premium (the same explanatory variables used by Li *et al.*, 1998). But, I found a strong *negative* interaction effect with initial GDP per capita (with openness entering positively on its own in the same regression). This suggests that openness is associated with *higher* inequality in poor countries. Barro (2000) also reports a significant negative interaction effect between GDP per capita and openness in a regression for inequality, using different controls.

Heterogeneity might also be expected at given levels of economic development. Suppose that reforming developing countries fall into two categories: those in which pre-reform controls on the economy were used to benefit the rich, keeping inequality artificially high, and those in which the controls had the opposite effect, keeping inequality low. The reforms may well entail sizable redistribution between the poor and the rich, but in opposite directions in the two groups of countries. Then one should not be surprised to find that there is zero correlation between growth and changes in inequality or that the average impact of policy reform on inequality is not significantly different from zero. Yet there could well be non-random distributional change going on under the surface of this average impact calculation. This can arise when policy reforms shift the distribution of income in different directions in different countries. Moreover, it is not implausible that they would do so, given the diversity in initial conditions across developing countries at the time reforms begin.

There is evidence to support this interpretation. As noted already, using the same data set underlying Figure 1, one finds virtually zero correlation between changes in the Gini index of inequality and growth in mean income or consumption. But, suppose that the true relationship is one in which initial inequality interacts with growth, such that the growth attenuates high inequality, but it increases inequality when it is initially low. Using the same set of developing countries as used for Figure 1,

one finds evidence for such an interaction effect by regressing the change in the log of the Gini index on the growth rate in PCE and the product of that growth rate with initial inequality. More precisely, the test regression takes the form:

$$\Delta \ln G_{it}/\tau = (\beta_0 + \beta_1 \ln G_{it-\tau})\Delta \ln Y_{it}/\tau + \mu_{it}, \quad (2)$$

where G_{it} is the Gini index in country i at date t , and Y_{it} is the private consumption per capita of country i at date t and τ is the time between surveys. The estimate of β_0 is significantly positive (6.03, with a standard error of 2.14), while β_1 is significantly negative (a regression coefficient of -1.60 with a standard error of 0.57). (Again, the standard errors are heteroskedasticity-consistent.) Thus one finds a significant *negative* interaction effect between growth and initial inequality.¹³ At a value of the log Gini index of $-\beta_0/\beta_1$, growth has no effect; this occurs at a Gini index of 0.433, which is very close to the median Gini index in the sample of 0.425.

A possible concern about this test is that the interaction effect with initial inequality might be due to measurement error in the latter variable. If the Gini index is over- (under-) estimated this year, then the growth rate in the Gini index will tend to be under- (over-) estimated, which will be reflected in a negative interaction effect with growth in the above test (a version of a problem known as “Galton’s Fallacy”). But, the negative interaction effect remained significant (at the 2% level) when I used a higher lag of the inequality measure as the instrument for “initial” inequality.¹⁴ This will eliminate any bias due to measurement error as long as the errors are serially independent. The turning point was almost the same (a Gini index of 0.432). So for roughly the lower half of countries in terms of inequality, growth tends to come with higher inequality, while for the upper half, growth tends to attenuate inequality.

None of this denies that growth-oriented reforms have an important role in fighting poverty or that policies can intervene to alter the distributional outcomes. But these observations do point to the need for a deeper understanding of the evident diversity in the impacts on poverty, and what role other policies have played. This requires further research on the role of initial conditions (including distribution) and how they interact with policy change.

Economic reforms in developing countries can create opportunities for poor people. But

only if the conditions are in place for them to take advantage of those opportunities will absolute poverty fall rapidly. Given initial inequalities in income and nonincome dimensions of welfare, economic reforms can readily bypass the poor. The conditions for pro-poor growth are this closely tied to reducing the disparities in access to human and physical capital, and sometimes also to differences in returns to assets, that create income inequality and probably also inhibit overall growth prospects.

Policy discussions have often emphasized the need to combine policies conducive to growth with investments in the human and physical assets of poor people.¹⁵ But, many questions remain unanswered. What specific interventions should have priority in specific circumstances? Should reform be redesigned or delayed when initial conditions are not favorable, and take time to change?

8. CONCLUSIONS

The seemingly opposing positions taken in this ongoing debate are not as hard to reconcile as it might seem at first sight. The poor typically do share in the benefits of rising aggregate affluence, and they typically do suffer from economic contraction. But, there is a sizable variance around the “typical” outcomes for the poor. One source of variance is that “economic growth,” as measured in the national accounts, is not always reflected in average household living standards as measured in surveys, at least in the short run.

The sources of the heterogeneity in outcomes for the poor go deeper than that, however. Finding zero average impact on inequality of growth-oriented policy reforms does not mean that reforms are generally distribution neutral. An average is just that, and it is deceptive when one averages over large differences across countries in their starting points. There are important differences in initial inequalities, with implications for how much the poor share in aggregate growth and contraction. The churning that is found under the surface of the aggregate outcomes also means that there are often losers during spells of growth, even when poverty falls on average. While various papers in the literature have found that growth-promoting policies have little or no average impact on inequality, that finding is perfectly consistent with sizable distributional impacts in spe-

cific countries, albeit in different directions. Average neutrality is consistent with strong distributional effects at the country level. There is truth in both the quotes at the beginning of this paper, though each is potentially deceptive on its own.

These observations point to the importance of more micro, country-specific, research on the factors determining why some poor people are able to take up the opportunities afforded by an expanding economy—and so add to its expansion—while others are not. Individual endowments of physical and human capital have rightly been emphasized in past work, and

suggest important links to policy. Other factors that may well be equally important have received less attention, such as location, social exclusion and exposure to uninsured risk.

While good policy-making for fighting poverty must obviously be concerned with the aggregate impacts on the poor, it cannot ignore the diversity of impacts underlying the averages, and it is here that good micro empirical work can help. That diversity also holds potentially important clues as to what else needs to be done by governments to promote poverty reduction, on top of promoting economic growth.

NOTES

1. Earlier contributions include Fields (1989), World Bank (1990, Chapter 3), Squire (1993), Ravallion (1995), Ravallion and Chen (1997) and Bruno, Ravallion, and Squire (1998).
2. The latest version of the data set can be found at <http://www.worldbank.org/research/povmonitor/>. This web site is updated regularly; the results in this paper are based on the data set at mid-2000, as used in Chen and Ravallion (2001).
3. For a critical review of the methods underlying the data set used here, see Deaton (2001).
4. These are based on income shares for Brazil in 1996 and consumption shares for India in 1997; in both cases the ranking variable is per capita (World Bank, 2000a).
5. One possible concern is that the poverty line is fixed (in PPP terms) across countries. While this is valid for making comparisons of absolute poverty in terms of command over commodities, actual poverty lines tend to rise with mean consumption or income in a country (Chen & Ravallion, 2001). But, comparisons of national poverty lines with mean consumption (both at PPP) indicate that the “income” elasticity of the poverty line is very low among poor countries, only rising at middle to upper income levels (Chen & Ravallion, 2001). Thus the growth elasticity of poverty in low- and middle-income countries is unlikely to be affected much by using relative poverty lines consistent with the way national poverty lines vary.
6. Here I am only referring to distribution in the developing world as it affects poverty. There is evidence of rising overall interpersonal inequality in the world (developing and developed); see Milanovic (2001) for 1987–93.
7. Banerjee and Duflo (1999) sketch a simple but elegant model of the intergenerational accumulation of wealth in which individuals start with an endowment from the previous generation but face a borrowing constraint. In this model, individual wealth at one date is a concave function of the individual’s endowment, given declining marginal products of capital. Thus mean wealth in the economy at one date is a quasi-concave function of the vector of endowments left over from the previous period. It follows from well-known properties of concave functions that higher initial inequality will entail lower future mean wealth for any given initial mean wealth.
8. Examples include Persson and Tabellini (1994), Alesina and Rodrik (1994), Birdsall, Ross, and Sabot (1995), Clarke (1995), Perotti (1996), Deininger and Squire (1998) and Deininger and Olinto (2000).
9. Distribution-dependent growth is possible without nonlinear income or wealth dynamics at the micro level. Such models that have been driven instead by their assumptions about political economy, notably the way initial distribution influences the balance of power over public spending (Alesina & Rodrik, 1994; Persson & Tabellini, 1994).
10. For example, research on government anti-poverty programs in China suggests that there have been dynamic consumption gains from the program at farm-household level (Jalan & Ravallion, 1998).
11. Dollar and Kraay (2000) do find that stabilizing against inflation is associated with lower inequality. This is consistent with other evidence that inflation hurts the poor (including Easterly & Fischer, 2001, using cross-country data, and Datt & Ravallion, 1998, using data for India).

12. Disparities in returns to schooling need not favor the poor either. For evidence on this point see, van de Walle's (2000) results for Vietnam.
13. This test has used the data for developing market economies. The negative interaction effect vanishes if one includes the transition economies of Eastern Europe and Central Asia. This is not too surprising given that these economies are going through a transition process to a market economy in which the economies are tending to contract, and inequality is rising, which may well be confounding the true relationship.
14. Naturally the number of observations drops considerably, to $n = 69$.
15. Arguments along these lines can be found in World Bank (1990, 2000b), Bruno *et al.* (1998) and Kanbur and Squire (1999) among others.

REFERENCES

- Aghion, P., Caroli, E., & Garcia-Penalosa, C. (1999). Inequality and economic growth: the perspectives of the new growth theories. *Journal of Economic Literature*, 37(4), 1615–1660.
- Alesina, A., & Rodrik, D. (1994). Distributive politics and economic growth. *Quarterly Journal of Economics*, 108, 465–490.
- Atkinson, A. B., & Brandolini, A. (1999). *Promise and pitfalls in the use of secondary data sets: income inequality in OECD countries*. Mimeo, Nuffield College Oxford.
- Banerjee, A., & Duflo, E. (1999). *Inequality and growth: what can the data say?* Mimeo, Department of Economics, MIT.
- Barro, R. (2000). Inequality and growth in a panel of countries. *Journal of Economic Growth*, 5, 5–32.
- Baulch, B., & Hoddinott, J. (2000). Economic mobility and poverty dynamics in developing countries. *Journal of Development Studies*, 36(6), 1–24.
- Benabou, R. (1996). Inequality and growth. In B. Bernanke, & J. Rotemberg (Eds.), *National bureau of economic research macroeconomics annual* (pp. 11–74). Cambridge, MA: MIT Press.
- Bhalla, S. (2000). Growth and poverty in India—myth and reality, available from <http://www.oxus-research.com/economic.asp>.
- Birdsall, N., & Londono, J. L. (1997). Asset inequality matters: an assessment of the World Bank's approach to poverty reduction. *American Economic Review Papers and Proceedings*, 87(2), 32–37.
- Birdsall, N., Ross, D., & Sabot, R. (1995). Inequality and growth reconsidered: lessons from East Asia. *World Bank Economic Review*, 9(3), 477–508.
- Bruno, M., Ravallion, M., & Squire, L. (1998). Equity and growth in developing countries: old and new perspectives on the policy issues. In V. Tanzi, & K. Chu (Eds.), *Income distribution and high-quality growth*. Cambridge, MA: MIT Press.
- Chen, S., & Ravallion, M. (2001). How did the world's poorest fare in the 1990s? *Review of Income and Wealth* (in press).
- Clarke, G. R. G. (1995). More evidence on income distribution and growth. *Journal of Development Economics*, 47, 403–428.
- Datt, G. (1999). Has poverty in India declined since the economic reforms. *Economic and Political Weekly*, 34(December), 11–17.
- Datt, G., & Ravallion, M. (1998). Farm productivity and rural poverty in India. *Journal of Development Studies*, 34, 62–85.
- Deaton, A. (2001). Counting the world's poor: problems and possible solutions. *World Bank Research Observer*.
- Deininger, K., & Olinto, P. (2000). *Asset distribution inequality and growth*. Policy Research Working Paper 2375 Washington, DC: World Bank.
- Deininger, K., & Squire, L. (1998). New ways of looking at old issues: inequality and growth. *Journal of Development Economics*, 57(2), 259–287.
- Dollar, D., & Kraay, A. (2000). *Growth is good for the poor*. Washington, DC: World Bank.
- Easterly, W., & Fischer, S. (2001). Inflation and the poor. *Journal of Money Credit and Banking*, 33(2), 160–179.
- Fields, G. (1989). Changes in poverty and inequality in developing countries. *World Bank Research Observer*, 4, 167–186.
- Forbes, K. J. (1997). *A reassessment of the relationship between inequality and growth*. Mimeo. Department of Economics, MIT.
- Groves, R. M., & Couper, M. P. (1998). *Nonresponse in household surveys*. New York: Wiley.
- Jalan, J., & Ravallion, M. (1998). Are there dynamic gains from a poor-area development program? *Journal of Public Economics*, 67(1), 65–86.
- Jalan, J., & Ravallion, M. (2000). Is transient poverty different? Evidence for rural China. *Journal of Development Studies*, 36, 82–99.
- Kanbur, R., & Squire, L. (1999). *The evolution of thinking about poverty: exploring the interactions*. Department of Agricultural Resource and Managerial Economics Paper, 99-24. Cornell University.
- Knowles, S., (2001). *Inequality and economic growth: the empirical relationship reconsidered in the light of comparable data*. Paper prepared for the WIDER Conference on Growth and Poverty. WIDER, Helsinki.
- Li, H., & Zou, H. (1998). Income inequality is not harmful to growth: theory and evidence. *Review of Development Economics*, 2(3), 318–334.
- Li, H., Squire, L., & Zou, H. (1998). Explaining international and intertemporal variations in income inequality. *Economic Journal*, 108, 26–43.

- Lokshin, M., & Ravallion, M. (2000). Welfare impacts of Russia's 1998 financial crisis. *Economics of Transition*, 8(2), 269–295.
- Lokshin, M., & Ravallion, M. (2001). *Nonlinear household income dynamics in two transition economies*. Mimeo, Development Research Group, World Bank.
- Milanovic, B. (2001). True world income distribution: 1988 and 1993: First calculations based on household surveys alone. *Economic Journal*.
- Perotti, R. (1996). Growth income distribution and democracy: what the data say. *Journal of Economic Growth*, 1(2), 149–187.
- Persson, T., & Tabellini, G. (1994). Is inequality harmful for growth? *American Economic Review*, 84, 600–621.
- Ravallion, M. (1995). Growth and poverty: evidence for developing countries in the 1980s. *Economics Letters*, 48, 411–417.
- Ravallion, M. (1997). Can high inequality developing countries escape absolute poverty? *Economics Letters*, 56, 51–57.
- Ravallion, M. (1998). Does aggregation hide the harmful effects of inequality on growth? *Economics Letters*, 61(1), 73–77.
- Ravallion, M. (2000). Should poverty measures be anchored to the national accounts? *Economic and Political Weekly*. September.
- Ravallion, M. (2001). Measuring aggregate welfare in developing countries: How well do national accounts and surveys agree? Policy Research Working Paper, World Bank.
- Ravallion, M., & Chen, S. (1997). What can new survey data tell us about recent changes in distribution and poverty? *World Bank Economic Review*, 11(2), 357–382.
- Ravallion, M., & Datt, G. (2001). Why has economic growth been more pro-poor in some states of India than in others? *Journal of Development Economics* (in press).
- Squire, L. (1993). Fighting poverty. *American Economic Review Papers and Proceedings*, 83(2), 377–382.
- van de Walle, D. (2000). *Are returns to investment lower for the poor? Human and physical capital interactions in rural Vietnam*. Policy Research Working Paper, WPS 2425. Washington, DC: World Bank.
- Wood, A. (1994). *North–South trade employment and inequality. Changing fortunes in a skill driven world*. Oxford: Clarendon Press.
- World Bank. (1990). *World development report: Poverty*. New York: Oxford University Press.
- World Bank. (2000a). *World development indicators*. Washington, DC: The World Bank.
- World Bank. (2000b). *World development report: Attacking poverty*. New York: Oxford University Press.