WPS 1694

What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty?

Is it true that the poor have lost ground, even as average living standards have risen? No. Poor people typically share in rising average living standards.

Martin Ravallion
Shaohua Chen

FILE COPY

Summary findings

It has been claimed that in recent times the poor have lost ground, both relatively and absolutely, even as average standards of living were rising. Ravallion and Chen test that claim, using more than 100 household surveys for more than 40 countries.

Overall, there was a small decrease in poverty incidence in 1987–93, though experiences differed across regions and countries.

There was no general tendency for inequality or polarization to increase with growth. Distribution improves as often as it worsens in growing economies, and negative growth often appears to be highly detrimental to distribution.

Poor people typically do share in rising average living standards. This holds in all regions.

Produced by the Policy Research Dissemination Center

What can new survey data tell us about recent changes in distribution and poverty?

Martin Ravallion and Shaohua Chen

Acknowledgement

Is the incidence and depth of poverty rising? Does income inequality increase with rising average standards of living? Do richer societies become more "polarized"? Do the poor share in the benefits of higher average levels of living? How much do they lose from aggregate contraction? These questions are often asked. But they are hard questions to answer convincingly.

In principle, household surveys can address such questions. But coverage and quality are uneven. As a rule, the poorer a country, the less likely we know just how poor it's people are, and whether or not they are seeing any improvement over time. Other factors influence data availability and quality, such as the openness of the polity, and the size of the country (since the average cost of a representative household survey falls with the size of the population represented). Data on poor people has historically been wanting relative to most other data. For example, the 1979 World Development Report (WDR), and the WDR's for many years after, only gave distributional data from household surveys for 20 or so developing countries. Yet macroeconomic aggregates were available for almost all countries.

Estimates of distributional statistics (such as the well-known Gini index of inequality) for the 1960s and '70s have been widely used as both dependent and independent variables in cross-country regressions. Yet some of these statistics were not even based on nationally representative household surveys, but were synthetic estimates built up from other sources, including non-survey data (Fields, 1994). Even amongst the survey-based estimates, the surveys have varied greatly in terms of (for example) the measure of living standards used, with implications for summary statistics on distribution such as the Gini index.

The availability of distributional data for developing countries has improved over the last 10 years. For example, the latest WDR has distributional data for 67 low and middle income countries (World Bank, 1996a). The timeliness of data has also improved. In the 1985 WDR, the average lag was 11 years (so the average survey date was 1974!). The lag is now five years. Nationally

representative household surveys underlie all the distributional data given in recent WDRs. Efforts at improving data quality and country coverage have been made by many countries and international agencies, including the World Bank. There is a long way to go before we can even say that the all poor countries have a good quality survey for poverty monitoring, and even further before we can be confident of data comparability across countries. But there has been progress.

This paper aims to provide a "broad brush" picture of how measures of distribution and poverty have been evolving since the mid 1980s, and what correlation these changes have had with growth and contraction in average levels of living. Our approach is largely descriptive. While distributional data have improved, we remain skeptical of attempts to use these data to test seemingly sophisticated multivariate models. We will however risk drawing out some of the simple bivariate relationships, though taking some care to test their robustness to the underlying measurement problems. By care in assembling the data set, and in choosing econometric methods for estimating the relationships of interest, we hope to be able to extract the "signal" from the "noise" in these data.

The following section discusses the data and appropriate econometric methods for estimating the main relationships of interest. Section II discusses the study's results concerning how distribution has been changing over time. In section III we examine what has been happening to measures of poverty. Our conclusions can be found in section IV.

I   Data and methods

While data have improved, international comparisons of distributional statistics are still plagued with both conceptual and practical problems. We first survey some of the issues, and then discuss implications for estimating the main relationships of interest.

International comparisons of statistics on poverty and distribution

It is well recognized that official exchange rates are deceptive in making international comparisons of absolute levels of living. But the problems of making purchasing-power-parity (PPP) currency conversions should not be understated. Estimates of the PPP exchange rate have varied widely, with implications for (amongst other things) international comparisons of poverty rates.

Given that we want to include the countries of Eastern Europe in this study, absolute level comparisons of poverty across countries pose an extra problem. Applying a developing country poverty line to Eastern Europe will imply very low poverty rates in that region, while an Eastern European poverty line will give very high poverty rates in many low-income countries. Measurements at extremes of the distribution are problematic in conventional sample surveys.

A further issue is that of comparing different survey-based measures of living standards. For example, some surveys only obtain incomes and others only obtain consumptions. So one must compare an income Gini index with one for consumption. An income-based measure is bound to show higher "inequality" than one based on consumption. (At one survey date, incomes will be unusually low for some households and unusually high for others; with some opportunities for saving and/or borrowing, consumptions will be less unequal.) Also, in developing countries particularly, measurement errors are thought to be greater for incomes, which will tend to add to measured inequality. Differences between countries in measured inequality may thus reflect in part differences in the welfare indicators used.

Survey questionnaires can also differ widely in terms of (for example) the number of distinct categories of consumer goods that they identify, and the order in which questions are asked. Some income surveys still rely on questions such as "What is your income from self-employment?" which would clearly be very difficult to answer; a convincing questionnaire requires a careful and complete

3

accounting of revenues and costs in the household enterprises (recognizing that these may be tangled up with other activities). Survey quality varies, and it is also possible that even seemingly similar surveys are not comparable. This could be a serious problem for cross-country comparisons of the levels of incomes and of summary measures based on their distribution. Most of the empirical literature has compared the levels of summary measures (such as inequality measures or poverty rates) across countries; the existence of country-level fixed effects in distribution—arising from (inter-alia) survey design—can make such comparisons deceptive.

Comparisons across countries at different overall levels of development also pose a potential problem given that the relative importance of consumption of non-market goods will vary. The local market value of all consumption in kind (includes consumption from own production, which is particularly important in relatively underdeveloped rural economies) should ideally be included in the measure of total consumption expenditure; similarly, the imputed profit from production of non-market goods should be included as part of income. This is not always done. However, we think that this was a far bigger problem in the surveys prior to 1980 or so than since then. From our experience, it has become routine for survey data for developing countries to include valuations for consumption or income from own production, following guidelines of the U.N. Household Survey Capability Programme or advice from the World Bank or elsewhere. Nonetheless, the methods of valuation do vary; for example, some current surveys use the price at the nearest market, while others use the average farm-gate selling price.

On econometric methods for cross-country regressions using survey data

The data problems summarized above clearly throw doubt on simple cross-country comparisons of the measured levels inequality and poverty. However, it can still be possible to

4

detect the true relationship between (say) poverty and aggregate affluence. Indeed, as we argue in this section, there are even quite simple econometric methods which can retrieve the true relationship of interest, provided that the structure of measurement errors satisfies certain assumptions.

We want to know whether a measure of inequality or poverty responds systematically to growth in average levels of living. (For concreteness we focus on poverty in the following discussion.) However, the data are riddled with measurement errors and non-comparabilities. To some extent these will behave like country-level fixed effects, though they will also induce artificial variation over time. So there is latent heterogeneity in distribution, reflecting in part differences in the type of data. There may also be a common time trend. Combining these features, let measured poverty in country i at date t be given by:

$$\log P_{it} = \alpha_i + \beta \log \mu_{it}^* + \gamma t + \epsilon_{it} \qquad (i=1,..,N;\ t=1,..,T_i) \qquad (1)$$

where α_i is a fixed effect reflecting the time-persistent differences between countries in distribution, β is the "growth elasticity" of poverty with respect to mean consumption given by μ_{it}^*, and ϵ_{it} is a white noise error process,[1] including errors in the poverty measure.

Notice that β is not the same as the "growth elasticity" which can be derived analytically under the assumption that the Lorenz curve does not change (Kakwani, 1993). The latter elasticity must be negative, and indeed it has a unique (non-stochastic) value for any poverty measure, mean and distribution. By contrast, β is an empirical elasticity in which the Lorenz curve has shifted consistently with the data. In principle this could take any sign or magnitude, depending on how distribution changes with growth, and it has its own distribution. In estimating β our interest is whether actual growth processes typically reduce poverty.

We do not, however, observe the true mean μ_{it}^*, but the following estimate:

$$\log\mu_{it} = \log\mu_{it}^* + v_{it} \tag{2}$$

This contains a country-specific, time varying, error term (v_{it}) which is assumed to be white noise, as in the standard "errors-in-variables" model (EVM) (see, for example, Greene, 1991, Chapter 9). However, unlike the standard EVM, v_{it} is allowed to be contemporaneously correlated with ϵ_{it}, recognizing that both the poverty measure and mean consumption are derived from a common household survey. Using (2), equation (1) takes the form:

$$\log P_{it} = \alpha_i + \beta\log\mu_{it} + \gamma t + \epsilon_{it} - \beta v_{it} \tag{3}$$

Taking first differences,[2] we can eliminate α_i, and obtain:

$$\Delta\log P_{it} = \gamma + \beta\Delta\log\mu_{it} + \Delta\epsilon_{it} - \beta\Delta v_{it} \tag{4}$$

(where $\Delta x_{it} \equiv x_{it} - x_{it-1}$). So (roughly speaking) the rate of poverty reduction is regressed on the rate of growth in mean consumption.[3]

However, the standard Ordinary Least Squares (OLS) regression method will not in general give unbiased estimated of either β or γ even in very large samples; in other words OLS will in general be inconsistent, under the above assumptions. It can be shown that, as the number of countries (N) approaches infinity, the OLS estimate of β converges to:[4]

$$plim\,\hat{\beta} = \beta + \frac{2[Cov(\epsilon_{it},v_{it}) - \beta Var(v_{it})]}{Var(\Delta\log\mu_{it})} \tag{5}$$

The second term on the right hand side is the asymptotic bias in the OLS estimate. This is made up of the usual attenuation bias when an explanatory variable is measured with error, plus an extra

6

"common-survey bias" due to the correlated measurement errors. Surveys which have overestimated (underestimate) mean consumption will presumably tend to underestimate (overestimate) poverty measures; so it is plausible that $Cov(\epsilon_{it}, v_{it}) < 0$. Thus, as long as growth does in fact reduce poverty ($\beta < 0$), both $Cov(\epsilon_{it}, v_{it})$ and $\beta Var(v_{it})$ will be negative, and hence offsetting. Whether on balance there will be over- or under-estimation of the true value of the growth elasticity cannot be determined without imposing further structure on the measurement errors.

One way to add structure is by noting that the error term in equation (1) includes effects of measurement errors in both the mean and the Lorenz curve, for both can induce errors in measured poverty. A natural assumption to make is that overestimating the mean by (say) 10% has the same effect on measured poverty as a 10% increase in the true mean. On also allowing for other (distributional) errors in measured poverty, we can write the error term in (1) as

$$\epsilon_{it} = \beta v_{it} + \xi_{it} \tag{6}$$

where ξ_{it} is another white noise process, interpretable as the error in the poverty measure due to miss-measurement of distribution. Then the asymptotic bias in the OLS estimate simplifies to:

$$plim\,\hat{\beta} - \beta = \frac{2Cov(\xi_{it}, v_{it})}{Var(\Delta \log \mu_{it})} = 0 \tag{7}$$

as long as the "distributional error" (ξ_{it}) is uncorrelated with the "growth error" (v_{it}). Then the common-survey bias will exactly offset the attenuation bias. There is no obvious reason why the "growth" and "distributional" errors will be correlated. Overestimation of the mean might be due to overestimation of the incomes of the nonpoor in one survey (such as by over-sampling a rich area); but it does not seem plausible that this would typically be the case—sometimes the problem would be due to overestimation of the incomes of the poor.

7

So under these assumptions about the structure of measurement errors in this setting, and allowing for latent heterogeneity due to lack of strict data comparability across countries, we will obtain consistent estimates (unbiased as N approaches infinity) of the growth elasticity by simply applying OLS to equation (4). That is the approach we follow here.

But that will not give us the correct standard errors. Notice that the difference transformation used to obtain (4) has also changed the properties of the error term. In addition to eliminating the unobserved "fixed effects", the transformation has introduced a first difference in the original error term (ϵ_{it}). If the latter is white noise, then the new error process in (4) will be correlated within countries and over time, though not between countries. Successive spells for a given country are not statistically independent, since they have one survey in common. Conventional methods of calculating standard errors then have to be modified. Specifically, the variance-covariance matrix of the error process $\Delta\epsilon_{it}$ has a block diagonal structure (with a separate block for each country) in which non-zero off-diagonal elements only appear within the blocks, due to the common surveys for adjacent spells. All standard errors and t-ratios quoted in this paper have been corrected to take account of the structure of the error covariance matrix of this specification. They have also been corrected for any general type of heteroscedasticity that might be present, after first correcting for the block diagonal structure of the covariance matrix.

Would it not be better to replace μ_{it} by the private consumption component of the National Accounts? Of course, this too will be measured with error; in addition to the existing error in the National Accounts' estimate of consumption for a given year, there will be new errors in matching with the survey period used to measure poverty. Those errors would presumably be uncorrelated with the error in measured poverty. However, as we have shown above, that correlation actually works in our favor, by counter-balancing the usual attenuation bias arising from the measurement

error in the explanatory variable. Replacing the survey mean with mean consumption from National Accounts would thus create an inconsistent estimate of the growth elasticity; the attenuation bias would remain, but we could no longer rely on the off-setting common-survey bias.

The data

The data set developed for this study is a greatly expanded version of that documented in Chen, Datt and Ravallion (CDR) (1994), which was used national household surveys for 44 countries, 19 for more than one point in time. The present paper uses data for 67 countries of which 42 have at least two surveys during the period since 1980.[5] Table 1 gives the countries and dates covered by region in the new data set. We include as many surveys as available which satisfy our comparability standards (discussed below). Relative to CDR, there have been gains in coverage for all regions. Overall, 85% of the population is represented by at least one survey. The coverage varies though; the thinnest coverage is for Middle-East/North Africa (47% of the population represented) followed by Sub-Saharan Africa (66%).

All measures of household living standards are normalized by household size. The distributions are also household-size weighted. So, for example, we estimate the percentage of people living in households with consumption per person below the poverty line, not the percentage of households. Similarly the empirical Lorenz curves are household-size weighted, so they correspond to fractiles of persons not households.

In all cases we have estimated our measures from the primary data source (tabulations or household level data) rather than relying on existing estimates. The estimation from tabulations requires an interpolation method. We have mainly used parameterized Lorenz curves with flexible

functional forms; these have proved reliable in past work (see, for example, Ravallion et al., 1991). Also, we only use nationally representative surveys.

Two surveys for one country define what we term a "spell". Both measures used in a given spell are estimated the same way from the source data. In particular, we use the same living standards indicator—either expenditure or income per person—over time in constructing the spells. So we shall not compare an income measure at one date with an expenditure measure for the same country at another date. In some cases, different sub-periods use different measures for a given country; for example, surveys may switch from income to consumption. We then swap the measure at one survey date. (If this is impossible, then the spell is dropped.) When there is a choice we use consumption in preference to income.

The data set allows us to construct 64 spells for 42 countries between 1981 and 1994 (using 109 surveys). Table 1 gives the distribution of the spells across regions, and details on the specific countries and periods of each spell. The coverage deteriorates markedly for Sub-Saharan Africa (SSA) when we construct the spells; though we have 28 surveys spanning 19 countries in SSA, only seven spells were possible, for four countries. So we are less confident about results for that region.

One third of all spells are for EECA, reflecting in part the breakup of the Soviet Union. The EECA data should probably be treated differently to other regions. For one thing, the EECA countries are undergoing major structural changes which also have implications for the comparability over time and across countries of data on household living standards. For example, standard welfare measures do not allow for the rationing of consumer goods; relaxing rationing in the transition to a market economy will entail welfare gains which are not easily captured by conventional surveys. Also there are goods which were previously not market goods but have become so during the transition. The survey data and Consumer Price Index may not properly reflect this fact. And the

10

methods used for valuing consumption in kind may not have changed so as to properly reflect the changes in the economy; old "planning prices" may now bear little relationship to opportunity cost. There are also sampling biases in a number of these surveys; for example, some are likely to have undersampled (growing) "informal" segments of the economy (Atkinson and Micklewright, 1992). It is beyond our scope here to fix these problems. We will, however, take some care to note differences between the EECA data and that for other regions.

II   Changes in distribution

We now use these data to address the set of questions posed at the beginning of this paper. But we must first be more precise about the distributional measures we will be using.

What do we mean by "distribution" and how should it be measured?

Conventional measures of inequality satisfy the "transfer principle" whereby inequality is said to have fallen if the new distribution can be obtained from the old one by a set of transfers in which the gainers are poorer than the losers. There are numerous measures which satisfy the transfer principle (for a survey of standard measures and their properties see Sen, 1973). Here we shall use the most common measure found in practice, namely the Gini index.

However, there are aspects of distribution of interest which are not captured well by conventional inequality measures. Popular perceptions of whether "distribution" has improved or not may well be in discord with the usual axioms of social welfare comparisons used to justify specific inequality measures (Amiel and Cowell, 1992).

In democracies, impacts on the middle strata can be important to the political feasibility of policy reform. However, conventional inequality measures may not capture well the gains and losses

to the middle stratum. This calls for a measure of <u>polarization</u> i.e., the extent to which the society is divided into the "haves" and "have-nots". Roughly speaking, distribution A is said to be more polarized than B if the incomes in A tend to be more bimodal, in that there are more "poor" and "rich", but fewer people in the middle. For example, if one makes transfers amongst the poorest half such that the gainers are initially poorer than the losers, and does the same amongst the richest half, then polarization will have increased; yet inequality will have decreased.

To illustrate the ways in which inequality and polarization can diverge in a developing country context, consider the effects of a shift in the domestic terms of trade in favor of the rural sector. Suppose (to simplify the exposition) that there are four income groups: when ranked from lowest to highest income they are the rural poor, the urban poor, the rural rich and the urban rich. The rural poor and the rural rich gain from the shift in the terms of trade (at least in the long run), while both the urban groups lose. To simplify the exposition, let us assume that the gain to the rural poor is roughly equal to the loss to the urban poor; similarly the gain to the rural rich is about equal to the loss to the urban rich. Also assume that the rankings of the four groups are preserved. The pro-rural shift in the terms of trade will <u>reduce</u> inequality by any measure satisfying the transfer principle—the new distribution can be obtained from the old one by a set of transfers in which the recipient is poorer than the donor. But the change will <u>increase</u> polarization, by the above definition; the overall distribution will become more bimodal, due to the lower inequality both amongst the poor (due to the convergence in incomes between the rural and urban poor) and amongst the rich (with "rural rich" gaining relative to the better off "urban rich").

Thus an analysis which is concerned solely with inequality as conventionally defined may miss relevant aspects of how distribution has changed. It can be argued that some of the claims that have been made about how inequality may impinge on a growth process are really more to do with

12

polarization. It is possible, for example, that in our attempts to understand the political-economy of distributional impacts of policy reform we have been looking at the wrong measures; inequality may well fall with reform—and the change would be judged a social welfare improvement by conventional ethical criteria used in economics—and yet the society has become more polarized, with heightened social tensions arising from the polarizing effects of diverse impacts among middle income groups, whereby some become poorer while others prosper.

To measure polarization we use the index proposed by Wolfson (1994). Like the Gini index, this is between zero (no polarization) and one (complete polarization). When there is complete equality there is also zero polarization. However, while maximum inequality entails that the richest person has all of the income, maximum polarization occurs when half the population has zero income and the other half has twice the mean. The Wolfson polarization index (W) can be written as:

$$W = 2(\mu^* - \mu^L)/m \tag{8}$$

where μ^* is the distribution-corrected mean (given by the actual mean times one minus the Gini index), μ^L is the mean of the poorest half of the population, and m is the median.

Changes in inequality and polarization

Table 2 gives a regional summary of the changes in distribution. Inequality rose in 37 of the 64 spells, while polarization rose in 40. Both measures indicate a worsening in 6 out of 9 East Asian spells, and 18 out of 21 spells for EECA. In Latin America and the Caribbean (LAC) and SSA distribution improved more often than it worsened; inequality fell in 10 out of 14 spells for LAC and polarization fell in 8 out of 14 spells, while for SSA the Gini fell 4 times out of 7, and polarization fell for 5 spells. In South Asia (SA) inequality fell in 6 out of 10 spells, while polarization fell in

4 spells. Of the three spells for Middle East and North Africa (MENA), distribution worsened in two.

Combining all the spells, the average rate of increase in both the Gini index and the polarization index was significantly positive; for the Gini index, the mean rate of increase was 1.6% per year with a standard deviation of 0.48%; for the polarization index the mean rate of increase was 1.4% with a standard deviation of 0.52%. However, this worsening of distribution on average is largely due to the experience of EECA. If we exclude that region from the calculations, neither the mean rate of change in the Gini index nor the polarization index is significantly different from zero.

While there is a clear conceptual distinction between "inequality" and "polarization", there is a surprisingly close correspondence between them for these data. The relationship is quite strong and significant (the overall correlation coefficient is 0.83). In all but 7 of the 64 spells the two measures of distribution moved in the same direction. In 32 cases both inequality and polarization increased, while both fell in 23 cases. In the largest deviation from the least squares regression line (estimated on the full sample of spells), the Gini index fell 2.6% while the polarization index rose by 8.3%. The bulk of the EECA points are in the region of both increasing polarization and increasing inequality. And the measured rates of increase in inequality and polarization in EECA are high by any standards.

Growth and distributional change

Recognizing that we are concerned about how the benefits of growth in aggregate incomes are distributed, the question arises as to whether there is any systematic tendency for distribution to change in the process of rising average household income. This is a long standing issue in development economics (for a recent overview of the arguments see Bruno et al., 1996). In recent

14

times there has been much concern about the distributional implications of the types of growth processes we are seeing in poor countries. It is difficult to predict the effect of economic growth on distribution on a priori grounds. We turn instead to empirical evidence.

Figure 1 plots the changes in (log) Gini index against the changes in (log) real household consumption (or income) per person. (The picture looks very similar for the polarization index.) Over the 64 spells, the correlation is negative. On regressing the change in log Gini index on that in mean consumption and allowing a trend (by adding the number of years between surveys as an additional explanatory variable) we get the results reported in Table 3. There is a significant negative effect of higher mean consumption on inequality. There is also a significant underlying trend increase in inequality. However, when we separate out the EECA spells, both effects vanish from the remaining non-EECA (Figure 1). When we tried adding a complete set of regional dummy variables (both slope and intercept), we found no other significant regional differences.

So there is no sign in these data that higher average consumption tends to be associated with higher inequality, and there is no sign that inequality tends to increase independently of growth. For EECA, there is still a negative effect of growth on inequality, but it is not significant. There is a trend increase in inequality in the EECA countries. The same conclusions hold for polarization (Table 3). There is no evidence here that some middle income households have become worse off during spells of growth, while others have gained.

III   Progress in reducing poverty

Assessing and comparing progress in reducing poverty

All our poverty comparisons over time use poverty lines which have constant real value, according to country-specific consumer price indices. When we want to also make comparisons of

15

the level of poverty between countries we shall use Purchasing Power Parity exchange rates. These are not, however, available for a number of the countries in our data set (particularly, but not only, in the Former Soviet Union). We can thus expand the number of data points considerably by using poverty lines which are relative across countries, but absolute over time; since we are only comparing rates of change, the lack of absolute comparability of the levels is not too worrying. However, we do test robustness to this practice, by also comparing rates of change in level-comparable poverty measures and mean consumptions.

We first examine our results using poverty lines which are absolute over time, but relative between countries. To begin, the initial value (at the beginning of the first spell) is set at a common proportion of the initial survey mean in each country (i.e., the mean from the first survey). The poverty line is then up-dated over time using the local Consumer Price Index. We present summary results in Table 4 for three such poverty lines, set at 50%, 75% and 100% of the mean for the first survey date in each country. Poverty lines for European countries are typically around 50% of the mean, and this is also a common figure in middle income developing countries, while a figure closer to 75-100% of the mean is more common in low-income countries (Ravallion, Datt and van de Walle, 1991). The range 50-100% would appear to embrace the range of poverty lines found in practice.

Poverty fell in 24 of the 64 spells for all three poverty lines, while it increased for 34 spells for all three lines; in only 6 spells was the trend ambiguous (poverty increasing for some poverty lines and decreasing for others).

Table 4 also gives the results by region. The regions in which poverty fell unambiguously for half or more of the spells were East Asia (7 of the 9 spells) and Latin America (7 of the 14 spells). The regions in which poverty rose for half or more of the spells were EECA (17 of 21

16

spells showing an unambiguous increase) and Sub-Saharan Africa (5 out of 7). In South Asia an unambiguous increase in poverty was as common as a decrease (4 of the 10 spells in each case, with two ambiguous spells). While there seem to be some regional patterns, it is notable how much variation there is within regions; indeed, in no case do all spells for a region indicate the same direction of change.

The sharp increase in the poverty measures for most of Eastern Europe and Central Asia is striking. It is known that this has been happening (Milanovic, 1995). We find that the impact is particularly pronounced at the lower end of the distribution, which we comment further on below. However there was one glaring outlier in the EECA spells; poverty measures for Poland fell sharply in 1987-89; indeed, this is the spell with the largest drop in poverty amongst all 64 spells. The Poland spells are, however, erratic; for example, the (income-based) spell 1989-93 shows a sharp increase in poverty, while the (expenditure-based) spell 1990-92 shows little change. There may be comparability problems here.

Next we consider the results in which we attempt to fix the absolute value of the poverty line across countries. Table 5 gives our estimates of the percentage of the population living below $1/day at 1985 international prices. This is a typical poverty line amongst low income countries (World Bank, 1990; Ravallion, Datt and van de Walle, 1991). We also give the poverty gap index, given by the mean shortfall below the poverty line (counting the non-poor as having zero shortfall) expressed as a percentage of the poverty line. The table up-dates past estimates available for 1990 including in CDR.[6] Half of the 122 surveys used are household consumption surveys, and consumption expenditure (including the imputed value of consumption in kind) per person is used as the indicator of household welfare. When only an income survey is available, mean income per person has been re-scaled according to the estimated consumption share from National Accounts.

As in CDR, adjustments have been made to line up the surveys in time. Of the 67 countries represented, 22 had only one survey in the period 1984-94; two surveys were available for 36 countries, and three or more surveys were available for nine. If there is a survey within one year of the "target" date then that survey is used. If there is not, then the closest survey is used, adjusting the survey mean according to the rate of growth in real private consumption per person from the National Accounts. When the target date is between two surveys this is done for both and a time-weighted average is taken. This could not be done for EECA in 1990 due to the substantial amount of missing data in the Bank's data base.

Like past estimates, a dollar is not converted into local currencies at official exchange rates, but rather at rates which attempt to assure purchasing power parity—so that $1 is worth roughly the same in different countries. For currency conversions, the Purchasing Power Parity rate for consumption in 1985 in Penn World Tables 5.6 (PWT5.6) has been used. This is the latest available comprehensive set of consumption PPP rates, and are widely considered to be the most reliable source for consumption PPPs. However, PWT5.6 has entailed some important revisions to past PPPs. The main change is a substantial increase in the estimated proportion living below $1/day in East Asia, mainly arising from an upward revision in the number for China. This is due entirely to the revision in the PPP rate for China. If one uses instead Penn World Table 5.0, the East Asia percentages fall to 14.0 (1987), 14.0 (1990) and 11.6 (1993). Other changes due to the revised PPP rates include a lower estimate for India, bringing down the South Asia aggregate, and lower rates for MNA. Holding the survey data set constant, the numbers for Latin America and SSA are affected very little by the revisions to PPPs.

The results indicate a small drop in poverty between 1987 and 1993, though a small increase is indicated for the poverty gap index between 1990 and 1993 (in the aggregate index, excluding

EECA). A fall in poverty is indicated for East Asia, MNA, South Asia (though signs of a slight reversal from 1990 to 1993), while increases in poverty are indicated for EECA, Latin America and SSA. In 1993, the regional ranking from highest to lowest percentages living below $1/day is South Asia, SSA, East Asia, Latin America, MNA and EECA; for the poverty gap index the ordering is SSA, South Asia, Latin America, East Asia, EECA, and MNA. So, for example, while South Asia has the highest overall poverty incidence, SSA has the highest depth of poverty (so that at some lower poverty line, poverty incidence is highest in SSA). The share in the total number of people living below $1/day is falling in East Asia, staying roughly constant in South Asia, but rising elsewhere (Table 5).

Table 5 also gives the mean poverty gap of the poor as a percentage of the poverty line (which is simply the poverty gap index divided by the proportion of the population who is poor). While the proportion living under $1/day is falling the aggregate, it can be seen that the average distance below $1/day amongst the poor has remained close to $0.31 over the period.

Poverty and growth

The extent to which poor people share in a rising average standard of living has been much debated. A still common view is that the poor are generally left behind, though this has been challenged by a number of recent studies suggesting that a rising (falling) overall mean is typically associated with falling (rising) absolute poverty (Fields, 1989; World Bank, 1990, 1995; Squire, 1993; Ravallion, 1995). Here we see what further light our up-dated and expanded data set can throw on the issue.

Figure 2 gives a scatter plot of the change in log poverty rate between surveys (vertical axis) against that in average consumption. We have set the poverty line at 75% of the initial mean; the

pattern is similar for other poverty lines. Higher rates of growth in average living standards are associated with higher rates of poverty reduction. Unlike the distributional measures, the slope is similar between EECA and the rest.

To estimate the overall growth elasticities and distributional trends for various poverty measures, we use the spells data to estimate equation (4). On allowing for the uneven spacing of the surveys, the constant term in (4) is replaced by the lapsed time in years between surveys (and the usual constant term is suppressed). OLS gives consistent estimates under our assumptions about the structure of measurement errors, though the standard errors have to be corrected (section I). The results are given in Table 6.

Regressing the first difference of the log of the proportion of the population living on less than 50% of the initial mean against the difference in the log of the real value of the mean for the 64 spells we obtained a growth elasticity -2.6; thus a 10% increase in the mean can be expected to result in a 26% drop in the proportion of people living on less than half the initial mean. Moving to higher poverty lines, the growth elasticity falls (in absolute value). Regressing the rates of change in the proportion of the population living on less than 75% of the initial mean against the percentage change in the real value of the survey mean the regression coefficient is -1.3. At 100% of the initial mean the elasticity falls to -0.7 (Table 6).

If we use instead the international "$1/day" poverty line then we find a larger variance across countries in both the levels and rates of poverty reduction. The estimate growth elasticity of the proportion of the population living under $1/day is -3.1 (Table 6). We obtain a slightly higher elasticity for the poverty gap index.[7]

Thus the relationship between rates of poverty reduction and rates of growth in average consumption becomes flatter and more precisely estimated as one moves to higher poverty lines. It

is not the case that the incidence of "extreme" poverty tends to be less responsive to growth in average living standards than does the incidence of only "moderate" poverty. If anything these data point to the opposite conclusion. Similarly the "depth" of poverty, as reflected in the poverty gap index, is more responsive to growth than the "incidence".

There is no sign of a significant distributional trend overall, except for the poverty line set at 50% of the initial mean (Table 6). That is due to EECA, where there is a strong trend increase in poverty, independently of growth as seen in Section II; distribution is clearly worsening in these transitional economies. For the developing countries there is no sign of a trend independently of growth; zero is our best estimate of the rate of change in poverty at zero growth.

Are there other significant regional differences in the impact of a given rate of growth on poverty? We added a set of intercept dummy variables for the regions. (We also tried an intercept dummy variable for whether the survey data for a given spell was for incomes or expenditures, but this was insignificant.) At a given rate of growth in the mean, the only region which had a significantly different rate of poverty reduction to East Asia (taken as the arbitrary reference) was EECA, where the rate of increase in poverty was significantly higher than one would have expected given the rate of change in average living standards.

We also tested whether or not the impact of growth was any different amongst regions, by adding to our regressions the interaction effects between the rate of change in the mean and the regional dummy variables. None of these dummy variables were significant. Thus, for the set of countries in our data set, we could find no significant differences between regions in how responsive the poverty measures are to growth.

In summary, we find strong evidence that higher rates of growth in average living standards are associated with higher rates of poverty reduction. The adverse distributional effect of recent

growth in a number of the developing countries has not been strong enough to change the conclusion that growth has benefited the poor. For the developing countries as a whole, there is no significant trend distributional effect for or against the poor. So at zero growth, the expected rate of poverty reduction is also zero. For EECA there is an adverse distributional effect independently of the rate of growth in average levels of living.

IV   Conclusions

The main body of our analysis used distributional data from 109 household surveys done since 1980 to construct 64 spells of distributional change for 42 developing and transitional economies. The two surveys we have used to construct each spell must satisfy minimal criteria for comparability, including being nationally representative, and using ostensibly the same indicator of welfare. We have then estimated various summary statistics on how distribution and poverty have changed. We have mainly looked at rates of change. However, we have also offered an overall assessment of the absolute levels of poverty (at constant international prices) and how this has changed over the period 1987-93. That assessment used 122 surveys (including countries with only one survey) and extrapolated over time when necessary.

There are numerous sources of measurement error and comparability problems in these data, even after the quality controls we have applied. This is particularly worrying for the comparisons of absolute levels of poverty. While only comparing changes will avoid some of the difficulties of making level comparisons, the measures of change over time will undoubtedly include noise due to errors or inconsistencies of measurement. We have argued, however, that the main sources of bias in our estimation methods are likely to be offsetting, and (under certain assumptions about the

structure of measurement errors) will cancel each other out, leaving an unbiased estimate of the relationship of interest. So we can reasonably hope to have extracted the signal from the noise.

Our results suggest that both inequality and polarization increased more often than they decreased amongst the 64 spells. However, the experience of Eastern Europe and Central Asia is not typical; if we exclude this region from the analysis then both inequality and polarization fell more often than they rose. Distribution deteriorated more often than not in East Asia and it improved more often than not in Africa and Latin America.

For the sample as a whole, we find no support for the view that higher growth rates in average living standards tend to come with worsening distribution. Indeed, over the whole sample, rising average consumption was associated with lower inequality and polarization. However, this conclusion is not robust to excluding the countries of Eastern Europe and Central Asia, where there has been a tendency for both inequality and polarization to increase during a time of overall economic contraction. Excluding this set of countries from the analysis, we find that neither inequality nor polarization are correlated with growth in average consumption; nor do either have an underlying trend, in either direction.

Turning to performance at reducing absolute poverty, we have calculated rates of change in the proportions of the population living on less than 50%, 75% and 100% of the initial survey mean for each country. For all three of these cut-off points, poverty fell in 24 of the 64 spells, and it rose for all three cut-offs in 34 spells (the remaining six being ambiguous according to which cut-off is used). In East Asia, poverty fell in all except one spell, while it rose in almost all cases in Eastern Europe and Central Asia. Poverty rose during five of the seven African spells. In South Asia and Latin America poverty rose about as often as it fell.

23

When we force level comparability, we find that the overall percentage living below one dollar per day (at 1985 international prices) has fallen between 1987 and 1993, from 31% to 29%. The depth of poverty, as measured by average distance below the poverty line, has remained static in the aggregate over this period. Progress has been uneven across regions, with falling poverty incidence in East and South Asia, in Middle-East and North Africa, but rising poverty incidence in Eastern and Central Europe, Latin America and Sub-Saharan Africa.

There is a strong association between the rate of growth in average living standards and the rate at which absolute poverty falls. In terms of elasticities, the response of the poverty measures to changes in average consumption is even stronger for lower poverty lines. The benefits of higher total consumption appear to be spread quite widely, on average. Structural changes going on in the transitional economies entail rising poverty even at zero growth. But for the developing economies as a whole, stagnation in average living standards entails stagnation for the poor too. We could find no significant regional differences in how responsive the poverty measures are to growth.

Notes

1.   A "white-noise" error is one which has zero mean, is independent over time and between countries, and has constant variance.

2.   Alternatively, one can take deviations from the means over time (giving the "within" or "fixed effects" estimator).   However, this requires stronger assumptions for consistency under the present structure of measurement errors.   Under certain conditions, one can assure consistency by combining the estimates obtained from the two methods of transforming the data (Hsiao, 1986).   However, those conditions include that the time-varying measurement error in the right hand side variable is uncorrelated to that in the left-hand-side variable, which is implausible in this setting.

3.   Note however that using growth rates rather than changes in logs will give biased estimates of (4) for all except small changes.

4.   This is proved by taking the probability limit ("plim") of the formula for the OLS regression coefficient as N approaches infinity.

5.   The data set has been used for various recent compilations of regional and country-level distributional and poverty data, including World Bank (1996a,b, 1997).   The data set overlaps that used by Deininger and Squire (1996).   The latter data set focuses solely on inequality, but goes back further in time.

6.   There are a number of differences between these numbers and previous estimates published in World Bank (1990, 1992, 1993).   Aside from new data, the main difference is that, unlike past estimates, no model-based extrapolations have been used for countries without survey data.   The numbers used here are only based on household surveys

7.   The poverty gap index is the average distance in cents below $1/day (averaged over the whole population, with zero for the non-poor).

References

Alesina, Alberto and Dani Rodrik. 1994. 'Distributive Politics and Economic Growth', *Quarterly Journal of Economics* 108: 465-90.

Amiel, Yoram and Frank Cowell. 1992. 'Measurement of Income Inequality: Experimental Evidence by Questionnaire', *Journal of Public Economics* 47: 3-26.

Atkinson, Anthony B., and John Micklewright. 1992. *Economic Transformation in Eastern Europe and the Distribution of Income*, Cambridge: Cambridge University Press.

Bruno, M., M. Ravallion and L. Squire. 1995. 'Equity and Growth in Developing Countries: Old and New Perspectives on the Policy Issues', Policy Research Working Paper 1563, World Bank, Washington DC.

Chen, Shaohua, Gaurav Datt, and Martin Ravallion. 1994. 'Is Poverty Increasing or Decreasing in the Developing World?', *Review of Income and Wealth*, 40: 359-376.

Clarke, George R.G. 1995. 'More Evidence on Income Distribution and Growth', *Journal of Development Economics* 47: 403-428.

Datt, Gaurav and Martin Ravallion. 1992. 'Growth and Redistribution Components of Changes in Poverty Measures: A Decomposition with Applications to Brazil and India in the 1980s', *Journal of Development Economics* 38: 275-295.

Deininger, Klaus and Lyn Squire. 1996. 'A New Data Set Measuring Income Inequality', *World Bank Economic Review* 10: 565-592.

Fields, Gary. 1989. 'Changes in Poverty and Inequality in Developing Countries', *World Bank Research Observer*, 4:167-186.

_____. 1994. 'Data for Measuring Poverty and Inequality Changes in the Developing Countries', *Journal of Development Economics*, 44:87-102.

Greene, William H. 1991. *Econometric Methods*, New York: Macmillan Publishing Company.

Hsiao, Cheng. 1986. *Analysis of Panel Data*, Cambridge: Cambridge University Press.

Kakwani, Nanak. 1993. 'Poverty and Economic Growth with Application to Côte D'Ivoire'. *Review of Income and Wealth* 39, 121-139.

Milanovic, Branko. 1995. 'Poverty, Inequality and Social Policy in Transition Economies', Research Paper No.9, Transition Economics Division, Policy Research Department, World Bank.

Persson, Torsten and Guido Tabellini. 1994. 'Is Inequality Harmful for Growth?' *American Economic Review* 84: 600-621.

Ravallion, Martin. 1995. 'Growth and Poverty: Evidence for the Developing World', *Economics Letters* 48: 411-417.

Ravallion, Martin, Gaurav Datt, and Dominique van de Walle. 1991. 'Quantifying Absolute Poverty in the Developing World', *Review of Income and Wealth* 37: 345-361.

Sen, Amartya. 1973. *On Economic Inequality*, Oxford: Oxford University Press.

Squire, Lyn. 1993. 'Fighting Poverty', *American Economic Review* 83(2): 377-382.

Summers, R. and A. Heston. 1991. 'The Penn World Table (Mark 5): An Extended Set of International Comparisons, 1950-1988', *Quarterly Journal of Economics* 106:327-368.

Wolfson, Michael. 1994. 'Diverging Inequalities', *American Economic Review* 84 (2, May).

World Bank. 1979, 1985, 1990, 1992, 1995, 1996a. *World Development Report*. New York: Oxford University Press.

_____. 1993. *Implementing the World Bank's Strategy to Reduce Poverty. Progress and Challenges*. Washington DC: World Bank.

_____. 1996b. *Poverty Reduction and the World Bank: Progress and Challenges in the Early 1990s.* Poverty and Social Policy Department, World Bank.

_____. 1997. *World Development Indicators.* Washington DC: World Bank.

Table 1: Coverage of the data set

	% of 1993 population represented	No. countries	No. of spells (defined by two surveys)	Countries, survey dates and welfare indicator
East Asia (EA)	88.0	5	9	China 1985, 90, 92, 93 (I); Indonesia 84, 87, 90, 93 (E); Malaysia 84, 89(I); Philippines 85, 88(E); Thailand 81, 88(I), 88, 92(E)
East Europe and Central Asia (EECA)	85.9	18	21	Belarus 1988, 93(I); Bulgaria 88, 92(I); Czech 1988, 93(I); Estonia 1988, 93(I); Hungary 1989, 93(I); Kazakhstan 88, 93(I); Kyrgyz Republic 1988, 93(I); Latvia 88, 93(I); Lithuania 88,93(I); Moldova 88, 92(I); Poland 85, 87, 89(I), 90, 92(E), 89, 93(I); Romania 1989, 92(I); Russia 88, 93(I); Slovak 88, 92; Slovenia 87, 93(I); Turkmenistan 88, 93(I); Ukraine 88, 92(I); Yugoslavia 85, 89(I)
Latin America and Caribbean (LAC)	83.9	15	14	Bolivia 1990(I); Brazil 85, 89(I); Chile 90, 92(I); Colombia 88, 91(I); Costa Rica 81, 89(I); Dominican Republic 89(I); Ecuador 94(E); Guatemala 86/87, 89(I); Honduras 89, 92(I); Jamaica 88, 89, 90, 91, 92, 93(E); Mexico 84, 92(E); Nicaragua 93(E); Panama 89(I); Peru 85/86, 94(E); Venezuela 81, 87, 89, 91(I)
Middle East and Nth Africa (MENA)	46.7	5	3	Algeria 1988(E); Egypt 91(E); Jordan 86/87,92(E); Morocco 84/85,90(E); Tunisia 85, 90(E)
South Asia (SA)	98.4	5	10	Bangladesh 1983/84, 85/86, 88/89, 91/92(E); India 83, 86/87, 87/88, 88/89, 89/90, 90/91, 92(E); Nepal 84/85(I); Pakistan 91(E); Sri Lanka 85, 90(E)
Sub-Saharan Africa (SSA)	65.9	19	7*	Botswana 1985/86(E); Cote d'Ivoire 85, 86, 87, 88(E); Ethiopia 81/82(E); Ghana 87, 88, 92(E); Guinea 91(E); Guinea-Bissau 91(E); Kenya 92(E); Lesotho 86/87(E), Madagascar 93 (E); Mauritania 88(E); Niger 92(E); Nigeria 85, 92(E); Rwanda 83/85(E); Senegal 91/92(E); South Africa 93(E); Tanzania 91, 93(E); Uganda 89/90, 92(E); Zambia 91, 93(E); Zimbabwe 90(E)
Total	85.0	67	64	

Note: The dates are those of each of the surveys used. "I" denotes that household income per person is the welfare indicator, while "E" indicates that it is household consumption expenditure per person. * The 1991-93 spell for Tanzania was not used because of serious comparability problems between the surveys.

Table 2: Regional summary of changes in distribution

	Number of spells	Real survey mean per capita			Inequality			Polarization		
		fell	rose	mean rate of change (%/year)	fell	rose	mean rate of change (%/year)	fell	rose	mean rate of change (%/year)
East Asia	9	0	9	3.6	3	6	1.1	3	6	1.5
Eastern Europe and Central Asia	21	18	3	-6.9	3	18	5.0	3	18	4.6
Latin America	14	5	9	1.5	10	4	-0.3	8	6	-0.5
Middle East and North Africa	3	1	2	1.3	1	2	0.7	1	2	1.3
South Asia	10	6	4	0.2	6	4	0.0	4	6	-0.2
Sub-Saharan Africa	7	5	2	-6.0	4	3	-1.5	5	2	-2.1
Total	64	35	29	-2.0	27	37	1.6	24	40	1.4
Non-EECA	43	17	26	0.4	24	19	-0.1	21	22	-0.2

Table 3: Trends and growth elasticities of inequality and polarization

Measure of distribution		Trend (γ) (x100)	Growth elasticity (β)	R^2
Gini index of inequality	Full sample (n=64)	1.10 (3.21)	-0.24 (6.07)	0.54
	Excluding EECA (n=43)	0.13 (0.58)	-0.01 (0.23)	0.01
	EECA (n=21)	3.71 (3.18)	-0.11 (1.21)	0.75
Wolfson polarization index	Full sample (n=64)	1.00 (2.55)	-0.21 (4.51)	0.40
	Excluding EECA (n=43)	0.00 (0.22)	-0.01 (0.12)	0.00
	EECA (n=21)	3.82 (3.08)	-0.05 (0.56)	0.68

Note: OLS estimates obtained by regressing the difference between household surveys in the log of the measure of distribution on the time elapsed between the surveys and the difference in the log of the real value of the survey mean. Absolute t-ratios in parentheses, based on robust standard errors corrected for heteroscedasticity and serial correlation due to common surveys across sequential spells.

Table 4: Regional summary of changes in poverty

	Number of spells	Poverty fell for all three poverty lines	Trend is ambiguous	Poverty rose for all three poverty lines	Mean rate of change (%/year) Poverty line as % of initial mean:		
					50%	75%	100%
East Asia	9	7	1	1	-6.1	-4.6	-2.7
Eastern Europe and Central Asia (EECA)	21	2	2	17	109.2	25.4	9.4
Latin America	14	7	1	6	-1.2	-0.8	-0.4
Middle East and North Africa	3	2	0	1	1.3	-0.5	-0.9
South Asia	10	4	2	4	2.6	0.7	0.2
Sub-Saharan Africa	7	2	0	5	6.8	6.0	4.4
Total	64	24	6	34	35.9	8.3	3.1
Non-EECA countries before 1990	30	16	4	10	-0.6	-0.7	-0.4
Non-EECA countries after 1990	13	6	0	7	1.7	1.2	0.9

Table 5: Poverty measures using an international poverty line of $1/day/person at 1985 purchasing power parity

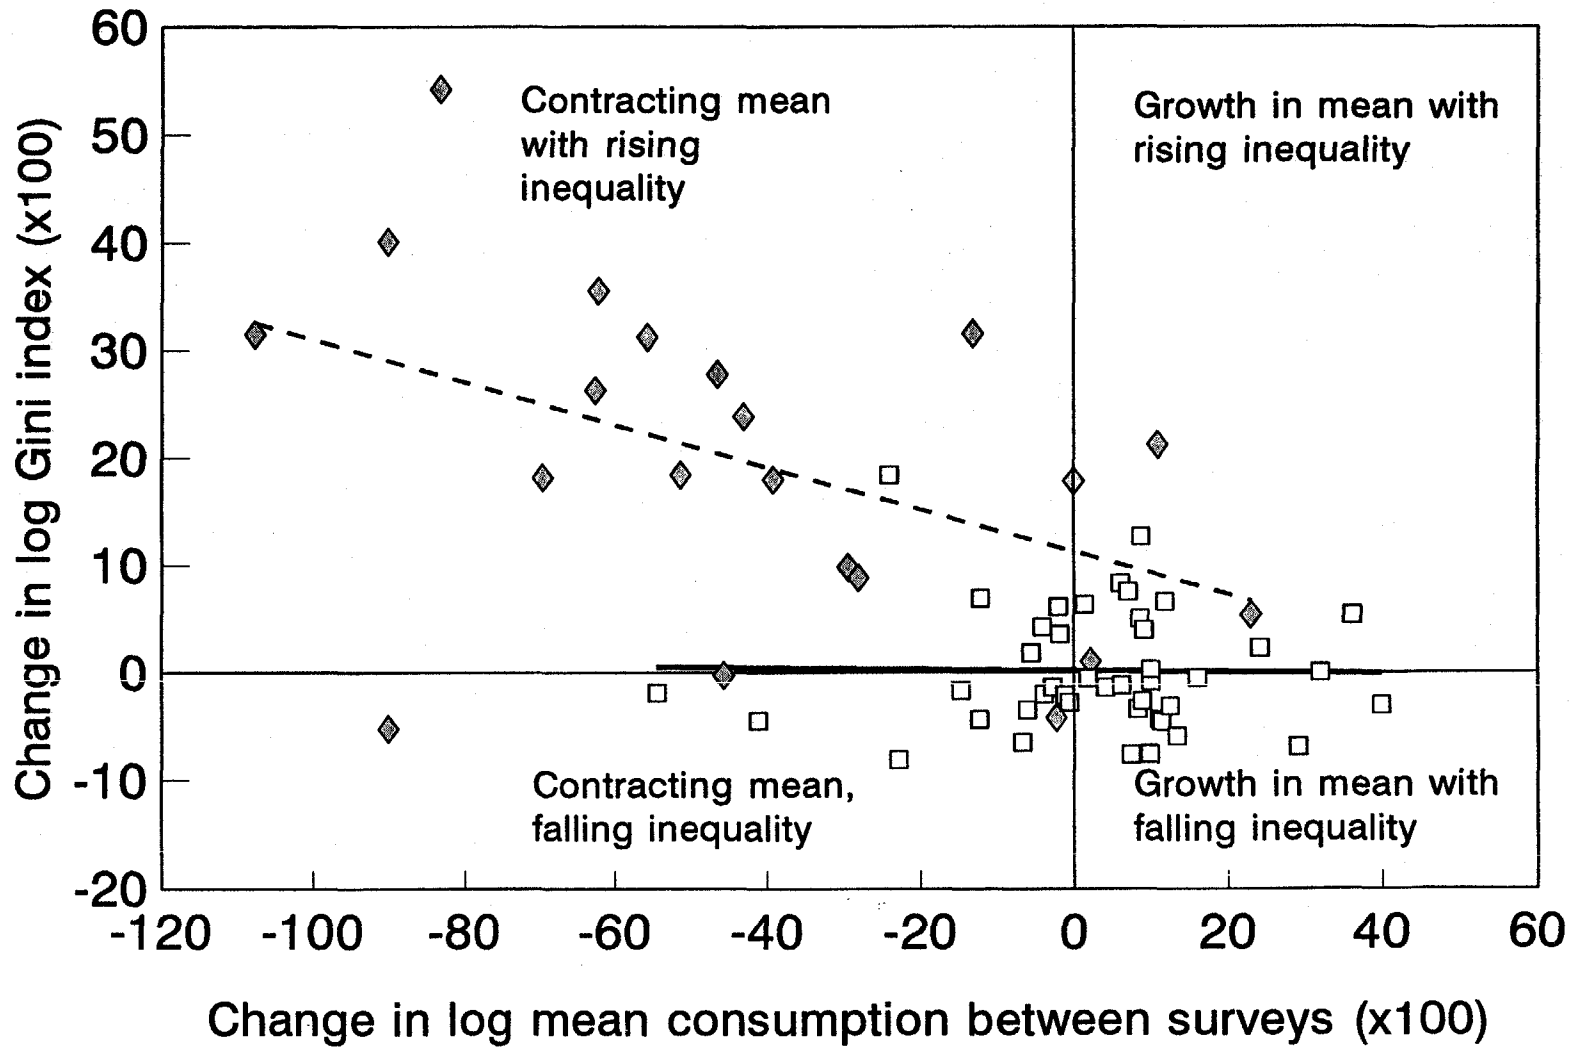| Region	Percentage of population consuming less than $1/day			Number of people consuming less than $1/day (millions; % of total in parentheses)		Poverty gap index (%) (Mean poverty gap in cents given in parentheses)		
	1987	1990	1993	1987	1993	1987	1990	1993
East Asia	29.7	28.5	26.0	464.0 (37.8)	445.8 (33.9)	8.3 (27.9)	8.0 (28.1)	7.8 (29.9)
Eastern Europe and Central Asia	0.6	n.a.	3.5	2.2 (0.2)	14.5 (1.1)	0.2 (27.1)	n.a.	1.1 (30.8)
Latin America	22.0	23.0	23.5	91.2 (7.4)	109.6 (8.3)	8.2 (37.2)	9.0 (39.3)	9.1 (38.8)
Middle East and North Africa	4.7	4.3	4.1	10.3 (0.8)	10.7 (0.8)	0.9 (18.3)	0.7 (15.9)	0.6 (15.7)
South Asia	45.4	43.0	43.1	479.9 (39.1)	514.7 (39.2)	14.1 (31.1)	12.3 (28.6)	12.6 (29.1)
Sub-Saharan Africa	38.5	39.3	39.1	179.6 (14.6)	218.6 (16.6)	14.4 (37.3)	14.5 (37.0)	15.3 (39.1)
Total	30.7	n.a.	29.4	1227.1 (100.0)	1313.9 (100.0)	9.5 (30.9)	n.a.	9.2 (31.3)
Total (excluding EECA)	33.9	32.9	31.9			10.8 (31.7)	10.3 (31.2)	10.5 (32.8)

Note: The figures given for the number of people consuming less than $1/day assume that the aggregate proportion of people who live below $1/day in the countries for which we have survey data is also representative of the countries for which we don't have such data.

33

Table 6: Distributional trends and growth elasticities of various poverty measures
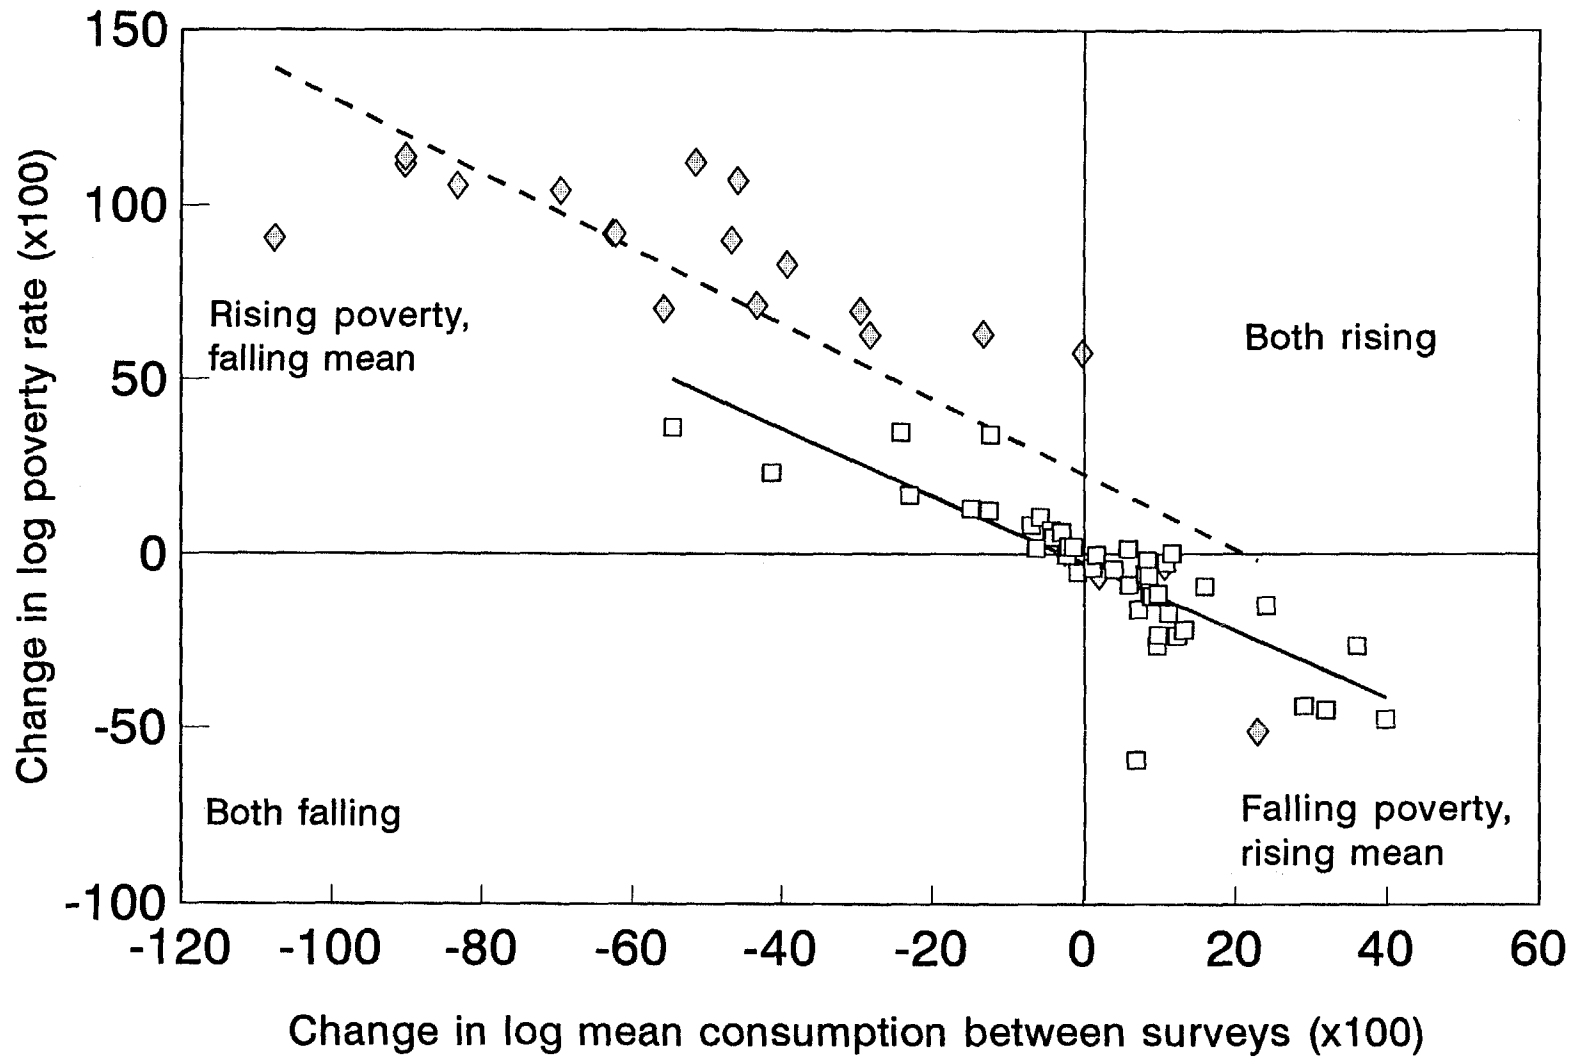
Poverty measure		Distributional trend (γ) (x100)	Growth elasticity (β)	R^2
Proportion below a poverty line set at 50% of initial mean	Full sample (n=64)	3.52 (2.37)	-2.59 (15.01)	0.84
	Excluding EECA (n=43)	-0.95 (0.87)	-1.57 (6.37)	0.58
	EECA (n=21)	16.66 (2.88)	-1.91 (4.43)	0.93
Proportion below a poverty line set at 75% of initial mean	Full sample (n=64)	0.87 (1.40)	-1.29 (13.24)	0.83
	Excluding EECA (n=43)	-0.87 (1.54)	-0.95 (10.23)	0.72
	EECA (n=21)	6.75 (2.46)	-0.97 (4.05)	0.92
Proportion below a poverty line set at initial mean	Full sample (n=64)	0.15 (0.51)	-0.69 (11.81)	0.84
	Excluding EECA (n=43)	-0.38 (1.38)	-0.64 (10.50)	0.85
	EECA (n=21)	2.68 (1.64)	-0.53 (3.59)	0.88
Proportion below $1/day, 1985 PPP	(n=42)	-3.86 (1.40)	-3.12 (2.62)	0.37
Poverty gap index in cents per day	(n=42)	-6.04 (1.63)	-3.69 (2.61)	0.36

Note: OLS estimates obtained by regressing the difference in the log of the poverty measure between household surveys on the time elapsed between the surveys and the difference in the log of the real value of the survey mean. Absolute t-ratios in parentheses, based on robust standard errors corrected for heteroscedasticity and serial correlation due to common surveys across sequential spells.

Figure 1: Inequality and Growth

The figure shows a scatter plot with:
- X-axis: "Change in log mean consumption between surveys (x100)" ranging from -120 to 60
- Y-axis: "Change in log Gini index (x100)" ranging from -20 to 60

Quadrant labels:
- "Contracting mean with rising inequality" (upper left)
- "Growth in mean with rising inequality" (upper right)
- "Contracting mean, falling inequality" (lower left)
- "Growth in mean with falling inequality" (lower right)

Legend: □ non-EECA   ◆ EECA

Side margin text: 35

Figure 2: Poverty and Growth

Policy Research Working Paper Series

	Title	Author	Date	Contact for paper
WPS1678	Financial Development and Economic Growth: Views and Agenda	Ross Levine	October 1996	P. Sintim-Aboagye 38526
WPS1679	Trade and the Accumulation and Diffusion of Knowledge	Pier Carlo Padoan	November 1996	M. Pateña 39515
WPS1680	Brazil's Efficient Payment System: A Legacy of High Inflation	Robert Listfield Fernando Montes-Negret	November 1996	T. Ishibe 38968
WPS1681	India in the Global Economy	Milan Brahmbhatt T. G. Srinivasan Kim Murrell	November 1996	S. Crow 30763
WPS1682	Is the "Japan Problem" Real? How Problems in Japan's Financial Sector Could Affect Developing Regions	Shigeru Otsubo Masahiko Tsutsumi	November 1996	J. Queen 33740
WPS1683	High Real Interest Rates, Guarantor Risk, and Bank Recapitalizations	Philip L. Brock	November 1996	N. Castillo 33490
WPS1684	The Whys and Why Nots of Export Taxation	Shantayanan Devarajan Delfin Go Maurice Schiff Sethaput Suthiwart-Narueput	November 1996	C. Bernardo 37699
WPS1685	Macroeconomic Crises and Poverty Monitoring: A Case Study for India	Gaurav Datt Martin Ravallion	November 1996	P. Sader 33902
WPS1686	Institutions, Financial Markets, and Firms' Choice of Debt Maturity	Asli Demirgüç-Kunt Vojislav Maksimovic	November 1996	P. Sintim-Aboagye 37644
WPS1687	Regionalism versus Multilateralism	L. Alan Winters	November 1996	A. Kitson-Walters 33712
WPS1688	Risk, Taxpayers, and the Role of Government in Project Finance	Michael Klein	December 1996	S. Vivas 82809
WPS1689	Is Economic Analysis of Projects Still Useful?	Pedro Belli	December 1996	K. Schrader 82736
WPS1690	Stock Markets, Banks, and Economic Growth	Ross Levine Sara Zervos	December 1996	P. Sintim-Aboagye 38526
WPS1691	Integrating the Unofficial Economy into the Dynamics of Post-Socialist Economies: A Framework of Analysis and Evidence	Daniel Kaufmann Aleksander Kaliberda	December 1996	S. Torres 39012

	Title	Author	Date	Contact for paper
WPS1692	Regulating Market Risk in Banks: A Comparison of Alternative Regulatory Regimes	Constantinos Stephanou	December 1996	P. Infante 37642
WPS1693	Famines and Economics	Martin Ravallion	December 1996	A. Ramirez 85734
WPS1694	What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty?	Martin Ravallion Shaohua Chen	December 1996	A. Ramirez 85734