THE WORLD BANK ECONOMIC REVIEW, VOL. 11, NO. 2: 357-82

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

Martin Ravallion and Shaohua Chen

It has been claimed that in recent times the poor have lost ground, both relatively and absolutely, even when average levels of living have risen. This article tests that claim using household surveys for 67 developing and transitional economies over 1981–94. It finds that changes in inequality and polarization were uncorrelated with changes in average living standards. Distribution improved as often as it worsened in growing economies, and negative growth was often more detrimental to distribution than positive growth. Overall, there was a small decrease in absolute poverty, although with diverse experiences across regions and countries. Almost always, poverty fell with growth in average living standards and rose with contraction.

Are the incidence and depth of poverty rising? Does 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 falling average living standards? These questions are often asked, but they are hard to answer convincingly.

In principle, household surveys can address such questions, but the coverage and quality of surveys are uneven. As a rule, the poorer a country, the more difficult it is to know just how poor its people are and whether their living standards are improving over time. Other factors, such as the openness and size of the country, influence the availability and quality of data. For example, the average cost of a representative household survey falls with the size of the population represented. Data on poor people have historically been wanting relative to most other data. For example, the World Bank's *World Development Reports* for 1979 and for many years after only give distributional data from household surveys for 20 or so developing countries. Yet macroeconomic aggregates are available for almost all countries.

Analysts have used estimates of distributional statistics (such as the well-known Gini index of inequality) for the 1960s and 1970s as both dependent and inde-

Martin Ravallion and Shaohua Chen are with the Policy Research Department at the World Bank. This project received financial assistance from the World Bank's Poverty and Social Policy Department and the joint British-Dutch-Swedish trust fund for studying the social and environmental consequences of growth-oriented policies. For discussions on this topic and comments on the paper, the authors are grateful to Gaurav Datt, Jyotsna Jalan, Emmanuel Jimenez, Michael Lipton, Oey Meesook, Binayak Sen, Lynne Sherburne-Benz, Dominique van de Walle, Quentin Wodon, participants at various presentations, and three anonymous referees.

pendent variables in cross-country regressions. Yet some of these statistics are not based on nationally representative household surveys; rather they are synthetic estimates built up from other sources, including nonsurvey data (Fields 1994). Even among the survey-based estimates, the surveys have varied greatly in, 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 past 10 years. For example, the *World Development Report 1996* presents distributional data for 67 low- and middle-income countries (World Bank 1996b). The timeliness of data has also improved. In the *World Development Report 1985*, the average lag is 11 years, so the average survey date is 1974 (World Bank 1985). The lag is now five years. Nationally representative household surveys underlie all the distributional data given in the *World Development Reports* for recent years. Many countries and international agencies, including the World Bank, have sought to improve the quality of data and the coverage of survey data. Despite these efforts, we have a long way to go before all poor countries have a good-quality survey for monitoring poverty and even further before data can be compared with confidence across countries. But we have made progress.

This article aims to provide a broad picture of the evolution of measures of distribution and poverty since the mid-1980s and to analyze the correlation of these changes with growth and contraction in average levels of living. Our approach is largely descriptive. Although distributional data have improved, we remain skeptical of attempts to use these data to test seemingly sophisticated multivariate models. We draw out some of the simple bivariate relationships and test their robustness to the underlying measurement problems. By carefully assembling the data set and choosing appropriate econometric methods for estimating the relationships of interest, we hope to extract the signal from the noise in these data.

Section I discusses the data and econometric methods for estimating the main relationships of interest. Section II discusses the study's results concerning how distribution has changed over time. Section III examines progress in reducing poverty. Section IV presents our conclusions.

I. Data and Methods

Although data have improved, international comparisons of distributional statistics are still plagued by both conceptual and practical problems. We survey some of the issues and discuss their implications for estimating the main relationships of interest. We then describe the data set developed for this study.

International Comparisons of Statistics on Poverty and Distribution

Official exchange rates are clearly deceptive in making international comparisons of absolute levels of living. But the problems of making purchasing-

power-parity currency conversions should not be understated. Estimates of the purchasing-power-parity exchange rate have varied widely, with implications for (among 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 would imply very low poverty rates in that region, while applying an Eastern European poverty line would 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 income and others only obtain consumption. An income-based measure is bound to show higher inequality than one based on consumption. (At one survey date, income will be unusually low for some households and unusually high for others; with some opportunities for saving or borrowing, consumption will be less unequal.) Also, in developing countries particularly, measurement errors are thought to be greater for income, which tends to inflate 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, for example, the number of distinct categories of consumer goods that they identify and the order in which they ask questions. Some income surveys still rely on questions such as "What is your income from self-employment?" that are clearly very difficult to answer. A convincing questionnaire requires a careful and complete accounting of revenues and costs in the household enterprises (recognizing that these may be tangled up with other activities). Survey quality varies, and even seemingly similar surveys might not be 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 compares 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, among other things, survey design—can make such comparisons deceptive.

Comparisons across countries at different overall levels of development also pose a potential problem given variations in the relative importance of consumption of nonmarket goods. The local market value of all consumption in kind (including 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 nonmarket goods should be included as part of income. This is not always done. However, this is a far bigger problem in the surveys conducted prior to 1980 or so than in those conducted since then. 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.

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 of inequality and poverty. However, it can still be possible to detect the true relationship between (say) poverty and aggregate affluence. Indeed, some quite simple econometric methods 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 noncomparabilities. To some extent these behave like country-level fixed effects, although they 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, P, in country i at date t be given by:

$$\log P_{it} = \alpha_i + \beta \log \mu^*_{it} + \gamma t + \varepsilon_{it} \quad (i = 1,..,N;\ t = 1,..,T_i) \tag{1}$$

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}, γ is trend rate of change over time t, and ε_{it} is a white-noise-error process that includes errors in the poverty measure.[1]

Notice that β is not the same as the growth elasticity that 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 (nonstochastic) value for any poverty measure, mean, and distribution. By contrast, β is an empirical elasticity in which the Lorenz curve shifts consistently with the data. In principle it 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, not whether some hypothetical growth process does so.

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

$$\log \mu_{it} = \log \mu^*_{it}, + \nu_{it}. \tag{2}$$

Equation 2 contains a country-specific, time-varying error term (ν_{it}) that is assumed to be white noise, as in the standard errors-in-variables model (see, for example, Greene 1991, chap. 9). However, unlike the standard errors-in-

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

variables model, ν_{it} is allowed to be contemporaneously correlated with ε_{it} in equation 1, recognizing that both the poverty measure and mean consumption are derived from a common household survey. Using equation 2, equation 1 takes the form:

$$\log P_{it} = \alpha_i + \beta \log \mu_{it} + \gamma t + \varepsilon_{it} - \beta\nu_{it}. \tag{3}$$

Taking first differences, we can eliminate α_i and obtain:

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

(where $\Delta x_{it} \equiv x_{it} - x_{it-1}$).[2] 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 does not in general give unbiased estimates of either β or γ even in very large samples; in other words OLS is 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]

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

The second term on the right-hand side of equation 5 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 common-survey bias caused by the correlated measurement errors. Surveys that overestimate (underestimate) mean consumption presumably tend to underestimate (overestimate) poverty measures; so it is plausible that $Cov(\varepsilon_{it}, \nu_{it}) < 0$. Thus, as long as growth does in fact reduce poverty ($\beta < 0$), both $Cov(\varepsilon_{it}, \nu_{it})$ and $\beta Var(\nu_{it})$ are negative and hence offsetting. Whether on balance there is over- or underestimation 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 percent has the same effect on measured poverty as a 10 percent increase in the true mean. Also allowing for other

2. Alternatively, we could 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, we 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 with 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 gives biased estimates of equation 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.

(distributional) errors in measured poverty, we can postulate the following decomposition of the error term in equation 1:

$$\varepsilon_{it} = \beta\nu_{it} + \xi_{it} \tag{6}$$

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

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

as long as the distributional error (ξ_{it}) is uncorrelated with the growth error (ν_{it}). Then the common-survey bias exactly offsets the attenuation bias. There is no obvious reason why the growth and distributional errors are correlated. Overestimation of the mean might be due to overestimation of the incomes of the nonpoor in a survey (such as by oversampling a rich area), but it does not seem plausible that this is typically the case. Sometimes the problem is caused by overestimation of the incomes of the poor.

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

But that does not give us the correct standard errors. Notice that the difference transformation used to obtain equation 4 also changes the properties of the error term. In addition to eliminating the unobserved fixed effects, the transformation introduces a first difference in the original error term (ε_{it}). If the latter is white noise, then the new error process in equation 4 is correlated within countries and over time, although not between countries. Successive spells for a given country are not statistically independent, because 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 $\Delta\varepsilon_{it}$ has a block diagonal structure (with a separate block for each country) in which nonzero off-diagonal elements only appear within the blocks, because of the common surveys for adjacent spells. In this article we correct all standard errors and t-ratios to take account of the structure of the error covariance matrix of this specification. We also correct them for any general type of heteroscedasticity that might be present, after first correcting for the block diagonal structure of the covariance matrix.

Would it be better to replace μ_{it} by the private consumption component of the national accounts? This component, too, is measured with error; in addition to the existing error in the national accounts' estimate of consumption for a given year, there are new errors in matching the survey period used to measure poverty. Those errors are presumably uncorrelated with the error in measured pov-

erty. However, as we have shown above, that correlation actually works in our favor, by counterbalancing the usual attenuation bias arising from the measurement error in the explanatory variable. Replacing the survey mean with mean consumption from national accounts data thus creates an inconsistent estimate of the growth elasticity; the attenuation bias remains, but we can no longer rely on the offsetting common-survey bias.

The Data

We developed a data set for this study that greatly expands the data set documented in Chen, Datt, and Ravallion (1994), which uses national household surveys for 44 countries, 19 for more than one point in time. The present article 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 that satisfy our comparability standards (discussed below). Relative to Chen, Datt, and Ravallion, the new data set gains in coverage for all regions. Overall, 85 percent of the population (in the countries included) is represented by at least one survey. The coverage varies, though; the thinnest coverage is for the Middle East and North Africa (47 percent of the population represented), followed by Sub-Saharan Africa (66 percent).

All measures of household living standards are normalized by household size. The distributions are also weighted by household size. 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 weighted by household size, so they correspond to fractiles of persons, not households.

In all cases we estimate our measures of living standards 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 mainly use parameterized Lorenz curves with flexible functional forms, which have proved reliable in past work (Ravallion, Datt, and van de Walle 1991 and Datt and Ravallion 1992). Also, we only use nationally representative surveys.

Two surveys for one country define what we term a "spell." Both measures of living standards used in a given spell are estimated the same way from the source data. In particular, in constructing the spells we use the same living standards indicator—either expenditure or income per person—over time. So we do not compare an income measure at one date with an expenditure measure for the same country at another date. In some cases, different subperiods use different measures for a given country; for example, surveys may switch from income to

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

Table 1. *Coverage of the Data Set*

Region	*Percentage of 1993 population represented*	*Country*	*Survey dates*	*Welfare indicator*
East Asia	88.0	China	1985, 1990, 1992, 1993	Income
		Indonesia	1984, 1987, 1990, 1993	Expenditure
		Malaysia	1984, 1989	Income
		Philippines	1985, 1988	Expenditure
		Thailand	1981, 1988	Income
			1988, 1992	Expenditure
Eastern Europe and Central Asia	85.9	Belarus	1988, 1993	Income
		Bulgaria	1988, 1992	Income
		Czech Republic	1988, 1993	Income
		Estonia	1988, 1993	Income
		Hungary	1989, 1993	Income
		Kazakstan	1988, 1993	Income
		Kyrgyz Republic	1988, 1993	Income
		Latvia	1988, 1993	Income
		Lithuania	1988, 1993	Income
		Moldova	1988, 1992	Income
		Poland	1985, 1987, 1989, 1993	Income
			1990, 1992	Expenditure
		Romania	1989, 1992	Income
		Russia	1988, 1993	Income
		Slovak Republic	1988, 1992	Income
		Slovenia	1987, 1993	Income
		Turkmenistan	1988, 1993	Income
		Ukraine	1988, 1992	Income
		Yugoslavia	1985, 1989	Income
Latin America and the Caribbean	83.9	Bolivia	1990	Income
		Brazil	1985, 1989	Income
		Chile	1990, 1992	Income
		Colombia	1988, 1991	Income
		Costa Rica	1981, 1989	Income
		Dominican Republic	1989	Income
		Ecuador	1994	Expenditure
		Guatemala	1986/87, 1989	Income
		Honduras	1989, 1992	Income
		Jamaica	1988, 1989, 1990, 1991, 1992, 1993	Expenditure
		Mexico	1984, 1992	Expenditure
		Nicaragua	1993	Expenditure
		Panama	1989	Income
		Peru	1985/86, 1994	Expenditure
		Venezuela	1981, 1987, 1989, 1991	Income
Middle East and North Africa	46.7	Algeria	1988	Expenditure
		Egypt	1991	Expenditure

Table 1. *(continued)*

Region	*Percentage of 1993 population represented*	*Country*	*Survey dates*	*Welfare indicator*
		Jordan	1986/87, 1992	Expenditure
		Morocco	1984/85, 1990	Expenditure
		Tunisia	1985, 1990	Expenditure
South Asia	98.4	Bangladesh	1983/84, 1985/86, 1988/89, 1991/92	Expenditure
		India	1983, 1986/87, 1987/88,1988/89, 1989/90, 1990/91, 1992	Expenditure
		Nepal	1984/85	Income
		Pakistan	1991	Expenditure
		Sri Lanka	1985, 1990	Expenditure
Sub-Saharan Africa	65.9	Botswana	1985/86	Expenditure
		Côte d'Ivoire	1985, 1986, 1987, 1988	Expenditure
		Ethiopia	1981/82	Expenditure
		Ghana	1987, 1988, 1992	Expenditure
		Guinea	1991	Expenditure
		Guinea-Bissau	1991	Expenditure
		Kenya	1992	Expenditure
		Lesotho	1986/87	Expenditure
		Madagascar	1993	Expenditure
		Mauritania	1988	Expenditure
		Niger	1992	Expenditure
		Nigeria	1985, 1992	Expenditure
		Rwanda	1983/85	Expenditure
		Senegal	1991/92	Expenditure
		South Africa	1993	Expenditure
		Tanzania	1991, 1993	Expenditure
		Uganda	1989/90, 1992	Expenditure
		Zambia	1991, 1993	Expenditure
		Zimbabwe	1990	Expenditure
Total	85.0			

Note: Income denotes household income per person, and expenditure denotes household consumption expenditure per person. The 1991–92 and 1991–93 spells for Jamaica and Tanzania, respectively, were not used because of serious comparability problems. The 1993 China survey, the 1992 Honduras survey, and the 1992 Uganda survey arrived too late to be used in constructing the spells but were used for other calculations (tables 2, 4, and 5).

Source: Household surveys done for individual countries, mostly by government statistical agencies.

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 67 countries between 1981 and 1994 (using 109 surveys). Table 1 gives the distribution of the spells across regions and presents details on the specific countries and periods for each spell. The coverage deteriorates markedly for Sub-Saharan Africa when we construct

the spells; although we have 28 surveys spanning 19 countries in Sub-Saharan Africa, only 7 spells are possible for 4 countries. So we are less confident about results for that region.

One-third of all spells are for Eastern Europe and Central Asia, reflecting in part the breakup of the Soviet Union. The data for Eastern Europe and Central Asia should probably be treated differently than the data for the other regions. For one thing, the countries in Eastern Europe and Central Asia are undergoing major structural changes that also have implications for the comparability of data on household living standards over time and across countries. For example, standard welfare measures do not allow for the rationing of consumer goods; relaxing rationing in the transition to a market economy has entailed welfare gains that are not easily captured by conventional surveys. Similarly, subsidies on publicly provided goods are often ignored and may have changed during the transition. Some nonmarket goods have become market goods during the transition. And the methods used for valuing consumption in kind may not have changed so as to reflect properly the changes in the economy; old planning prices may now bear little relationship to opportunity cost. The survey data and the consumer price index may not properly reflect these facts. 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 do, however, take some care to note differences between the data for Eastern Europe and Central Asia and those for other regions.

II. Changes in Distribution

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

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. Several measures satisfy the transfer principle (for a survey of standard measures of inequality and their properties see Sen 1973). Here we use the most common measure of inequality found in practice, namely the Gini index.

However, a conventional inequality measure may not pick up distributional changes of concern to policymakers. Impacts on the middle strata can be important to the political feasibility of policy reform, yet an inequality measure such as the Gini index may not capture changes in the share of income held by the middle stratum. This calls for a measure of *polarization*, that is, 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 (Wolfson 1994). For example, if we transfer income within the poorest half such that the gainers are poorer than the losers, and we do the same within the richest half, then polarization will have increased—there will be fewer people at middle incomes—yet inequality will have decreased. (Suppose that there are four people with incomes \$1, \$2, \$3, and \$4. We take \$0.50 from the person with \$2 and give it to the person with \$1, and we take \$0.50 from the person with \$4 and give it to the one with \$3. The new distribution is \$1.50, \$1.50, \$3.50, and \$3.50. Inequality has fallen, because the gainers were poorer than the losers, but polarization has risen, because the new distribution is more sharply divided into "rich" and "poor.")

To illustrate how 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: 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, we 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. We also assume that the rankings of the four groups are preserved. The prorural shift in the terms of trade *reduces* 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 *increases* polarization, by the above definition; the overall distribution becomes more bimodal, due to the lower inequality both among the poor (due to the convergence in incomes between the rural and urban poor) and among the rich (with the rural rich gaining relative to the better-off urban rich).

Thus an analysis that is concerned solely with inequality as conventionally defined may miss relevant aspects of how distribution has changed. Claims about how inequality changes during a growth process could well have more to do with 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 decrease with reform—and the change would be judged a social welfare improvement by conventional ethical criteria used in economics—and yet the society may 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, it is between 0 (no polarization) and 1 (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, maxi-

Table 2. *Regional Summary of Changes in the Distribution of Income or Consumption*

Region	Number of spells	Real survey mean per capita household income or consumption: Number of spells for which it Fell	Rose	Mean rate of change (percent per year)	Inequality[a]: Number of spells for which it Fell	Rose	Mean rate of change (percent per year)	Polarization[b]: Number of spells for which it Fell	Rose	Mean rate of change (percent per 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 and the Caribbean	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
Total excluding Eastern Europe and Central Asia	43	17	26	0.4	24	19	−0.1	21	22	−0.2

Note: See table 1 for countries and survey dates.
a. Measured by the Gini index.
b. Measured by the Wolfson (1994) polarization index.
Source: Authors' calculations.

mum 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 income (given by the actual mean times 1 minus the Gini index), μ^L is the mean income of the poorest half of the population, and m is the median income. Like inequality, this is not the only available measure of polarization, but it appears adequate for the present purpose.

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 spells for East Asia and in 18 out of 21 spells for Eastern Europe and Central Asia. In Latin America and the Caribbean and in Sub-Saharan Africa distribution improved more often than it worsened; inequality fell in 10 of 14 spells for Latin America and the Caribbean, and polarization fell in 8 of 14 spells, while for Sub-Saharan Africa the Gini index fell in 4 of 7 spells, and polarization fell in 5 spells. In South Asia inequality fell in 6 of 10 spells, while polarization fell in 4 spells. Of the 3 spells for the Middle East and North Africa, distribution worsened in 2.

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 percent a year with a standard deviation of 0.48 percent; for the polarization index, the mean rate of increase was 1.4 percent with a standard deviation of 0.52 percent. However, this worsening of distribution on average is largely due to the experience of Eastern Europe and Central Asia. If we exclude that region from the calculations, neither in the Gini index nor in the polarization index was the mean rate of change significantly different from 0.

Although there is a clear conceptual distinction between our measures of 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 between the rates of change 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 percent, while the polarization index rose 8.3 percent. The bulk of the points for Eastern Europe and Central Asia were in the region of both increasing polarization and increasing inequality. And the measured rates of increase in inequality and polarization in Eastern Europe and Central Asia were high by any standards.

Growth and Distributional Change

Is there any systematic tendency for distribution to change in the process of rising average household living standards? The distribution of the benefits of

economic growth is a long-standing issue in development economics (for a recent overview of the arguments see Bruno, Ravallion, and Squire 1996). In recent times economists have been much concerned about the distributional implications of the types of growth processes in poor countries. It is difficult to predict the effect of growth on distribution on a priori grounds. We turn instead to empirical evidence.

Figure 1 plots the changes in the (log) Gini index against the changes in (log) real household consumption (or income) per person. (The picture looks very similar for the polarization index, which is to be expected given their high correlation.) Over the 64 spells, the correlation is negative. On regressing the change in the 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 obtain the results reported in table 3. Higher mean consumption has a significant negative effect on inequality. We also find a significant underlying trend increase in inequality. However, when we remove the spells for Eastern Europe and Central Asia, both effects vanish (figure 1). When we try adding a complete set of regional dummy variables (both slope and intercept), we find no other significant regional differences.

Figure 1. *Inequality and Growth*

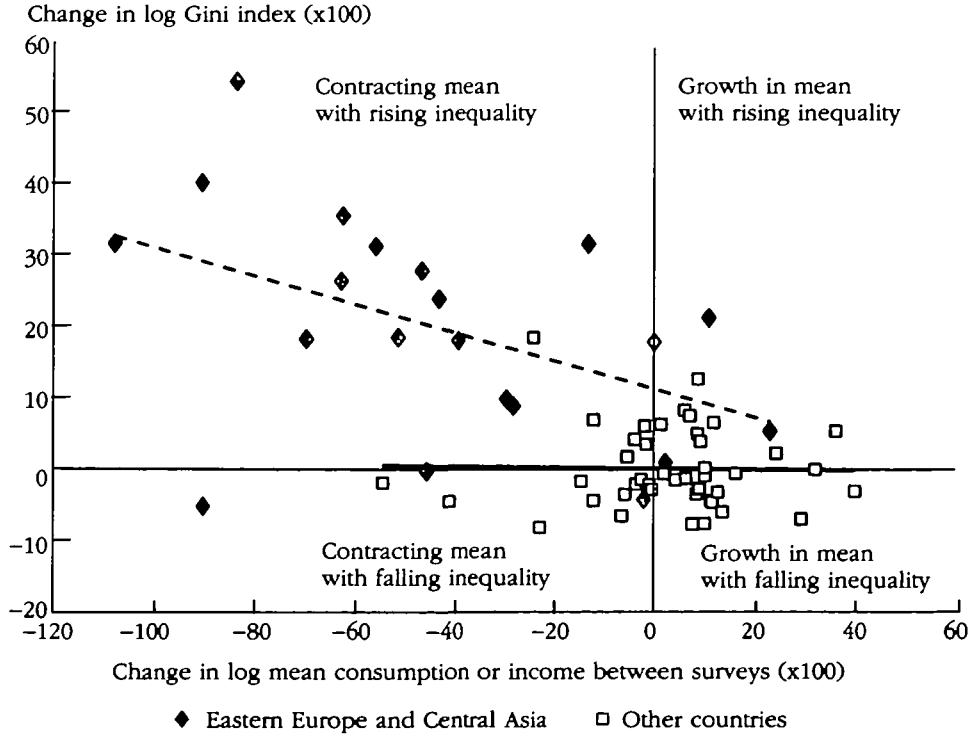

Note: See table 1 for countries and survey dates.
Source: Authors' calculations.

Table 3. *Trends and Growth Elasticities of Inequality and Polarization*

Measure of distribution and sample	*Trend (γ) (x100)*	*Growth elasticity (β)*	R^2
Gini index of inequality			
Full sample	1.10	−0.24	0.54
	(3.21)	(6.07)	
Excluding Eastern Europe and Central Asia	0.13	−0.01	0.01
	(0.58)	(0.23)	
Eastern Europe and Central Asia	3.71	−0.11	0.75
	(3.18)	(1.21)	
Wolfson polarization index			
Full sample	1.00	−0.21	0.40
	(2.55)	(4.51)	
Excluding Eastern Europe and Central Asia	0.00	−0.01	0.00
	(0.22)	(0.12)	
Eastern Europe and Central Asia	3.82	−0.05	0.68
	(3.08)	(0.56)	

Note: Estimates were obtained using OLS, 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 are in parentheses, based on robust standard errors corrected for heteroscedasticity and serial correlation due to common surveys across sequential spells. Sample sizes are 64 spells for the full sample, 43 spells for the sample excluding Eastern Europe and Central Asia, and 21 spells for Eastern Europe and Central Asia. See table 1 for countries and survey dates.

Source: Authors' calculations.

So these data do not indicate that higher average consumption tends to be associated with higher inequality or that inequality tends to increase independently of growth. For Eastern Europe and Central Asia, growth still negatively affects inequality, but the effect is not significant. There is a trend increase in inequality in the Eastern Europe and Central Asia countries. The same conclusions hold for polarization (table 3). We find 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

This section looks at how poverty measures have been changing and what relationship those changes have had with changes in average living standards.

Assessing and Comparing Progress in Reducing Poverty

All our poverty comparisons over time use poverty lines that have constant real value, according to country-specific consumer price indexes. When we also want to compare the level of poverty between countries we use purchasing-power-parity exchange rates. However, these are not available for several countries in our data set (particularly, but not only, in the Former Soviet Union). Therefore, we expand the number of data points considerably by using poverty lines that are relative across countries, but absolute over time; because we only compare

Table 4. *Regional Summary of Changes in Poverty*

Region	Number of spells: Total	Poverty fell for all three poverty lines	Trend is ambiguous	Poverty rose for all three poverty lines	Mean rate of change in poverty[a] (percent per year): 50 percent	75 percent	100 percent
East Asia	9	7	1	1	–6.1	–4.6	–2.7
Eastern Europe and Central Asia	21	2	2	17	109.2	25.4	9.4
Latin America and the Caribbean	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
Total excluding Eastern Europe and Central Asia							
Before 1990	30	16	4	10	–0.6	–0.7	–0.4
After 1990	13	6	0	7	1.7	1.2	0.9

Note: See table 1 for countries and survey dates.

a. The three poverty lines are set at 50, 75, and 100 percent of the mean household income or consumption expenditure per person for the first survey date in each country.

Source: Authors' calculations.

rates of change, the lack of absolute comparability of the levels is not too worrying. However, we do test the robustness of this practice by also comparing rates of change in level-comparable poverty measures and mean consumptions.

We first examine poverty lines that are absolute over time, but relative between countries. The initial value for the poverty line (at the beginning of the first spell) is set at a common proportion of the mean living standards indicator from the first survey. The poverty line is then updated 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 percent of the initial survey mean in each country. Poverty lines for European countries are typically around 50 percent of the mean, and this is also a common figure in middle-income developing countries, while a figure closer to 75–100 percent of the mean is more common in low-income countries (Ravallion, Datt, and van de Walle 1991). The range of 50–100 percent appears to embrace the range of poverty lines found in practice.

As shown in table 4, poverty fell in 24 of the 64 spells for all three poverty lines, while it increased in 34 spells for all three lines; in only 6 spells was the trend ambiguous (poverty increased for some poverty lines and decreased for others). The table also gives the results by region. The regions in which poverty fell unambiguously in 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 in half or more of the spells were Eastern Europe and Central Asia (17 of 21 spells showed 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). Although there seem to be some regional patterns, the variation within regions is notable; indeed, in no case did 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. (See Milanovic 1995 for further discussion.) We find that the impact was particularly pronounced at the lower end of the distribution. However, there is one glaring outlier for Eastern Europe and Central Asia. Poverty measures for Poland fell sharply in 1987–89; indeed, this is the spell with the largest drop in poverty among all 64 spells. However, the Poland spells were erratic; for example, the (income-based) 1989–93 spell showed a sharp increase in poverty, while the (expenditure-based) 1990–92 spell showed little change. There may be comparability problems here.

Next 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 on less than $1 a day at 1985 international prices. This is a typical poverty line among low-income countries (World Bank 1990 and Ravallion, Datt, and van de Walle 1991). The table also gives the poverty gap index, the mean shortfall below the poverty line (counting the nonpoor as having zero shortfall) expressed as a percentage of the poverty line. The table updates past estimates available for 1990, including those in Chen, Datt, and Ravallion (1994). There are a number of differences between these numbers and previous estimates published in World Bank (1990, 1992,

Table 5. *Poverty Measures Using an International Poverty Line of $1 a Day per Person at 1985 Purchasing Power Parity, 1987–93*

Region	*Percentage of population consuming less than $1 a day*			*Poverty gap index (percent)*			*Mean poverty gap of the poor (cents)*		
	1987	*1990*	*1993*	*1987*	*1990*	*1993*	*1987*	*1990*	*1993*
East Asia	29.7	28.5	26.0	8.3	8.0	7.8	27.9	28.1	29.9
Eastern Europe and Central Asia	0.6	—	3.5	0.2	—	1.1	27.1	—	30.8
Latin America and the Caribbean	22.0	23.0	23.5	8.2	9.0	9.1	37.2	39.3	38.8
Middle East and North Africa	4.7	4.3	4.1	0.9	0.7	0.6	18.3	15.9	15.7
South Asia	45.4	43.0	43.1	14.1	12.3	12.6	31.1	28.6	29.1
Sub-Saharan Africa	38.5	39.3	39.1	14.4	14.5	15.3	37.3	37.0	39.1
Total	30.7	—	29.4	9.5	—	9.2	30.9	—	31.3
Total excluding Eastern Europe and Central Asia	33.9	32.9	31.9	10.8	10.3	10.5	31.7	31.2	32.8

— Not available.

Note: The poverty measures are population-weighted means over all countries in the data set within each region. See table 1 for countries and survey dates.

Source: Authors' calculations.

1993). Aside from new data, the main difference is that, unlike past estimates, no model-based extrapolations are used for countries without survey data. The numbers used here are only for countries with appropriate household surveys.

Half of the 122 surveys used are household consumption surveys. We use consumption expenditure (including the imputed value of consumption in kind) per person as the indicator of household welfare. When only an income survey is available, we rescale mean income per person according to the estimated consumption share from the national accounts. As in Chen, Datt, and Ravallion (1994), we make adjustments to line up the surveys in time. Of the 67 countries represented for 1981–94, 22 have only one survey; 35 have two surveys, and 10 have three or more surveys. If there is a survey within one year of the target date, then we use that survey. If there is not, then we use the closest survey, adjusting the survey mean consumption or income according to the rate of growth in real private consumption per person from the national accounts. When the target date is between two surveys, we make the adjustment for both and use a time-weighted average. We cannot make the adjustment for Eastern Europe and Central Asia in 1990 because the World Bank's data base is missing a substantial amount of data.

As for past estimates, we do not convert dollars into local currencies at official exchange rates, but rather at rates that attempt to assure purchasing power parity (PPP)—so that $1 is worth roughly the same in different countries. For currency conversions, we use the PPP rate for consumption in 1985 in the Penn World Tables 5.6. (Summers and Heston 1991 describe the Penn World Tables, Mark 5.) This is the latest available comprehensive set of consumption PPP rates and is widely considered to be the most reliable source for consumption PPPs. However, Penn World Tables 5.6 entails some important revisions to past PPPs. The main change is a substantial increase in the estimated proportion of people living on less than $1 a day in East Asia, mainly arising from an upward revision in the number for China. This increase is due entirely to the revision in the PPP rate for China. If we use instead Penn World Tables 5.0, the East Asia percentages fall to 14.0 (1987), 14.0 (1990), and 11.6 (1993). Other changes caused by the revised PPP rates include a lower estimate for India, bringing down the South Asia aggregate, and lower rates for the Middle East and North Africa. Holding the survey data set constant, the numbers for Latin America and Sub-Saharan Africa are affected very little by the revisions to the PPP values.

From table 5, the results indicate a small drop in aggregate poverty between 1987 and 1993. This holds for both the headcount index (percentage of the population consuming less than $1 a day) and the poverty gap index (average distance in cents below $1 a day, when averaged over the whole population, with 0 for the nonpoor). The regional breakdown indicates a fall in poverty for East Asia, the Middle East and North Africa, and South Asia (with signs of a slight reversal from 1990 to 1993) and increases in poverty for Eastern Europe and Central Asia, Latin America, and Sub-Saharan Africa. For 1993 the regional ranking from highest to lowest percentages of the population living on less than

$1 a day is South Asia, Sub-Saharan Africa, East Asia, Latin America, the Middle East and North Africa, and Eastern Europe and Central Asia; for the poverty gap index the ordering is Sub-Saharan Africa, South Asia, Latin America, East Asia, Eastern Europe and Central Asia, and the Middle East and North Africa. So, for example, while South Asia has the highest overall poverty incidence, Sub-Saharan Africa has the highest depth of poverty (so that at some lower poverty line, the incidence of poverty is highest in Sub-Saharan Africa).

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 headcount index). Although the aggregate proportion living on less than $1 a day is falling, the average distance below $1 a day among 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. Several recent studies have challenged this view, suggesting that a rising (falling) overall mean standard of living is typically associated with falling (rising) absolute poverty (Fields 1989; World Bank 1990, 1995; Squire 1993; Ravallion 1995). Here we apply our updated and expanded data set to this issue.

Figure 2 plots the change in the log poverty rate between surveys against that in average consumption. We set the poverty line at 75 percent of the initial mean standard of living; 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 for Eastern Europe and Central Asia and the other regions.

To estimate the overall growth elasticities and distributional trends for various poverty measures, we use the data on spells to estimate equation 4. To allow for the uneven spacing of the surveys, the constant term in equation 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, although the standard errors have to be corrected (see 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 percent of the initial mean standard of living against the difference in the log of the real value of the mean for the 64 spells, we obtain a growth elasticity of −2.6. Thus, a 10 percent increase in the mean standard of living can be expected to result in a 26 percent drop in the proportion of people living on less than half the initial mean. For 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 percent of the initial mean standard of living against the percentage change in the real value of the survey mean, the regression coefficient is −1.3. At 100 percent of the initial mean, the elasticity falls to −0.7 (table 6).

Figure 2. *Poverty and Growth*

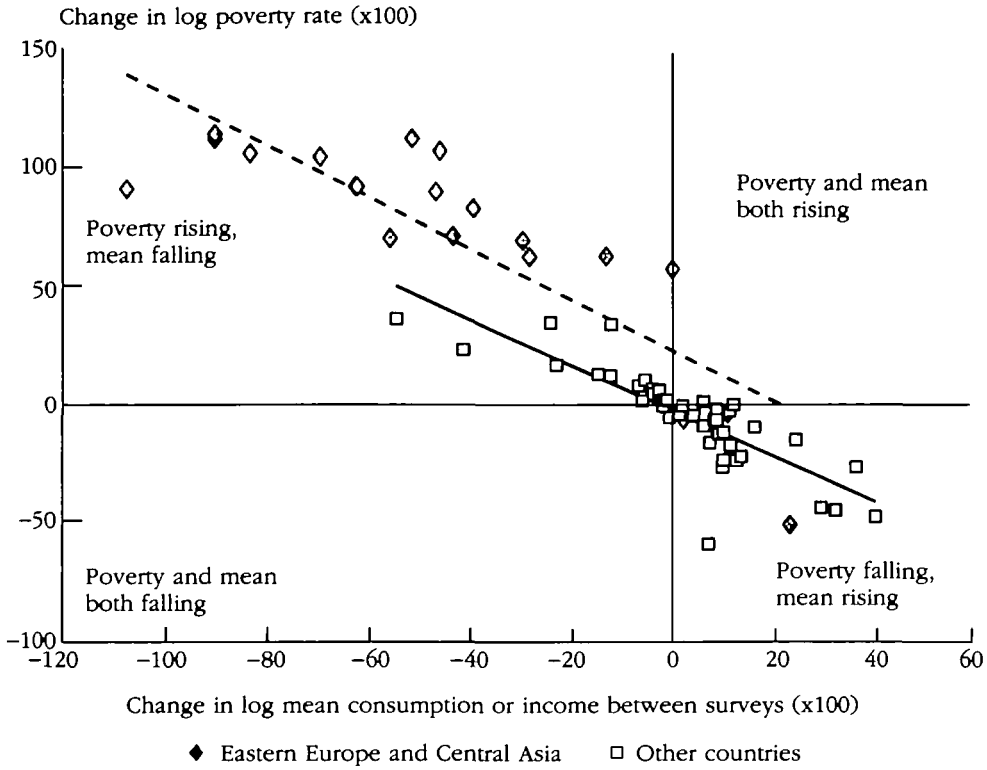

Note: See table 1 for countries and survey dates.
Source. Authors' calculations.

If we use instead the international $1 a day poverty line, then we find a larger variance across countries in both the levels and rates of poverty reduction. The estimated growth in the elasticity of the proportion of the population living on less than $1 a day is −3.1 (table 6). We obtain a slightly higher elasticity for the poverty gap index.

Thus the relationship between rates of poverty reduction and rates of growth in average consumption becomes flatter and more precisely estimated for higher poverty lines. The incidence of extreme poverty does not tend 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 is the incidence of poverty.

There is no sign of a significant distributional trend overall, except for the poverty line set at 50 percent of the initial mean standard of living (table 6). The trend for the 50 percent poverty line is due to Eastern Europe and Central Asia, where there is a strong trend increase in poverty independent of growth, as seen in section II; distribution is clearly worsening in these transi-

Table 6. *Distributional Trends and Growth Elasticities of Various Poverty Measures*

Poverty measure[a]	*Distributional trend* (γ) *(x100)*	*Growth elasticity* (β)	R^2
Poverty line at 50 percent			
Full sample	3.52	−2.59	0.84
	(2.37)	(15.01)	
Excluding Eastern Europe and Central Asia	−0.95	−1.57	0.58
	(0.87)	(6.37)	
Eastern Europe and Central Asia	16.66	−1.91	0.93
	(2.88)	(4.43)	
Poverty line at 75 percent			
Full sample	0.87	−1.29	0.83
	(1.40)	(13.24)	
Excluding Eastern Europe and Central Asia	−0.87	−0.95	0.72
	(1.54)	(10.23)	
Eastern Europe and Central Asia	6.75	−0.97	0.92
	(2.46)	(4.05)	
Poverty line at 100 percent			
Full sample	0.15	−0.69	0.84
	(0.51)	(11.81)	
Excluding Eastern Europe and Central Asia	−0.38	−0.64	0.85
	(1.38)	(10.50)	
Eastern Europe and Central Asia	2.68	−0.53	0.88
	(1.64)	(3.59)	
Proportion consuming less than $1 a day, 1985 PPP	−3.86	−3.12	0.37
	(1.40)	(2.62)	
Poverty gap index in cents per day	−6.04	−3.69	0.36
	(1.63)	(2.61)	

Note: Estimates were obtained using OLS, 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 are in parentheses, based on robust standard errors corrected for heteroscedasticity and serial correlation due to common surveys across sequential spells. Sample sizes for the poverty lines are 64 spells for the full sample, 43 spells for the sample excluding Eastern Europe and Central Asia, and 21 spells for Eastern Europe and Central Asia. Sample sizes are both 42 for the proportion living below $1 a day and for the poverty gap index. See table 1 for countries and survey dates.

a. The proportion of the population with income or consumption below 50, 75, or 100 percent of the mean household income or consumption expenditure per person for the first survey date in each country.

Source: Authors' calculations.

tional economies. For the developing countries there is no sign of a trend independent 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 on poverty of a given rate of growth in average living standards? We add 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 were for incomes or expenditures, but this was insignificant.) At a given rate of growth, the only region that has a rate of poverty reduction significantly different from that of East Asia (taken as the arbi-

trary reference) is Eastern Europe and Central Asia, where the rate of increase in poverty is significantly higher than we would expect given the rate of change in average living standards.

We also test whether the impact of growth is any different among regions, by adding to our regressions the interaction effects between the rate of change in the mean standard of living and the regional dummy variables. None of these dummy variables is significant. Thus, for the set of countries in our data set, we can find no significant differences between regions in the responsiveness of the poverty measures 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 Eastern Europe and Central Asia there is an adverse distributional effect.

IV. Conclusions

The main body of our analysis used distributional data from 119 household surveys since 1980. We constructed spells of distributional change for 42 developing and transitional economies using two surveys for each spell that satisfy minimal criteria for comparability, including being nationally representative and using ostensibly the same indicator of welfare. We estimated various summary statistics on how distribution and poverty have changed. We mainly looked at rates of change. However, we also offered an overall assessment of the absolute levels of poverty (at constant international prices) and how this changed over 1987–93. For that assessment, we used 122 surveys (including countries with only one survey) and extrapolated over time when necessary.

There are numerous sources of measurement errors and comparability problems in these data, even after the quality controls were applied. This is particularly worrying for the comparisons of absolute levels of poverty. Although comparing only changes avoids some of the difficulties of making level comparisons, the measures of change over time undoubtedly include noise caused by errors or inconsistencies of measurement. We argue, however, that the main sources of bias in our estimation methods for testing the effect of growth on distribution and poverty are likely to be offsetting and (under certain assumptions about the structure of measurement errors) to 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 in these data.

Our results suggest that both inequality and polarization increased more often than they decreased among the 64 spells. However, the experience of Eastern Europe and Central Asia is not typical; if we excluded 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 Sub-Saharan Africa and Latin America.

For the sample as a whole, we found no support for the view that higher growth rates in average living standards tended to accompany 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 found that neither inequality nor polarization was correlated with growth in average consumption; nor did either have an underlying trend, in either direction.

Turning to performance at reducing absolute poverty, we calculated rates of change in the proportions of the population living on less than 50, 75, and 100 percent of the initial survey mean for each country. For all three of these cutoff points, poverty fell in 24 of the 64 spells, and it rose for all three cutoff points in 34 spells (the remaining six being ambiguous according to which cutoff 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 spells in Sub-Saharan Africa. In South Asia and Latin America, poverty rose about as often as it fell.

When we forced level comparability, we found that the overall percentage of people living on less than $1 a day (at 1985 international prices) fell between 1987 and 1993, from 31 to 29 percent. The depth of poverty, as measured by average distance below the poverty line, remained static in the aggregate over this period. Progress was uneven across regions, with falling poverty incidence in East Asia, South Asia, and the Middle East and North Africa, but with 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 fell. 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 found no significant regional differences in the responsiveness of the poverty measures to growth.

References

The word "processed" describes informally reproduced works that may not be commonly available through library systems.

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

Bruno, Michael, Martin Ravallion, and Lyn Squire. 1996. "Equity and Growth in Developing Countries: Old and New Perspectives on the Policy Issues." Policy Research Working Paper 1563. Policy Research Department, World Bank, Washington, D.C. Processed.

Chen, Shaohua, Gaurav Datt, and Martin Ravallion. 1994. "Is Poverty Increasing or Decreasing in the Developing World?" *Review of Income and Wealth* 40(December):359–76.

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(April):275–95.

Deininger, Klaus, and Lyn Squire. 1996. "A New Data Set Measuring Income Inequality." *The World Bank Economic Review* 10(September):565–92.

Fields, Gary. 1989. "Changes in Poverty and Inequality in Developing Countries." *The World Bank Research Observer* 4(July):167–86.

———. 1994. "Data for Measuring Poverty and Inequality Changes in the Developing Countries." *Journal of Development Economics* 44(June):87–102.

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

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

Kakwani, Nanak. 1993. "Poverty and Economic Growth with Application to Côte d'Ivoire." *Review of Income and Wealth* 39(June):121–39.

Milanovic, Branko. 1995. "Poverty, Inequality, and Social Policy in Transition Economies." Research Paper 9. Policy Research Department, World Bank, Washington, D.C. Processed.

Ravallion, Martin. 1995. "Growth and Poverty: Evidence for the Developing World." *Economics Letters* 48(July):411–17.

Ravallion, Martin, Gaurav Datt, and Dominique van de Walle. 1991. "Quantifying Absolute Poverty in the Developing World." *Review of Income and Wealth* 37(December):345–61.

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

Squire, Lyn. 1993. "Fighting Poverty." *American Economic Review* 83(May):377–82.

Summers, Robert, and Alan Heston. 1991. "The Penn World Table (Mark 5): An Extended Set of International Comparisons, 1950–1988." *Quarterly Journal of Economics* 106(May):327–68.

Wolfson, Michael. 1994. "Diverging Inequalities." *American Economic Review* 84(2, May):353–58.

World Bank. 1979. *World Development Report 1979.* New York: Oxford University Press.

———. 1985. *World Development Report 1985: International Capital and Foreign Trade.* New York: Oxford University Press.

———. 1990. *World Development Report 1990: Poverty.* New York: Oxford University Press.

———. 1992. *World Development Report 1992: Development and the Environment.* New York: Oxford University Press.

———. 1993. *Implementing the World Bank's Strategy to Reduce Poverty: Progress and Challenges.* Washington, D.C.

———. 1995. *World Development Report 1995: Workers in an Integrating World.* New York: Oxford University Press.

———. 1996a. *Poverty Reduction and the World Bank: Progress and Challenges in the Early 1990s.* Poverty and Social Policy Department, Washington, D.C.

———. 1996b. *World Development Report 1996: From Plan to Market.* New York: Oxford University Press.

———. 1997. *World Development Indicators 1997.* Washington, D.C.