Review of Income and Wealth
Series 37, Number 4, December 1991

QUANTIFYING ABSOLUTE POVERTY IN THE DEVELOPING WORLD

BY MARTIN RAVALLION, GAURAV DATT AND DOMINIQUE VAN DE WALLE

World Bank

We estimate that in 1985 about one in five persons in the developing world lived in poverty, judged by the standards of the poorest countries. This rises to one in three at a common, more generous, poverty line. The aggregate consumption short-fall of the poorest fifth is about one half of one percent of world consumption, while that of the poorest third is a further one percent. The shape of the distribution of consumption suggests that aggregate poverty would fall fairly rapidly if moderate growth in average consumption levels can be sustained, *and* the poor share at least proportionally in that growth. However, it would take only small adverse shifts in the world distribution of consumption to eliminate the gains to the poor from growth.

1. INTRODUCTION

In counting the poor, and measuring the severity of absolute proverty, one faces a number of difficult questions. What poverty line should be used? Should one use the same poverty line across all countries? How should one adjust for differences across countries in the purchasing power of their currencies at official exchange rates? How should one interpolate from the available grouped data on the distribution of income or consumption? How should one extrapolate to countries for which distributional data are unavailable, or are highly imperfect? And, after answering these questions: What can we really learn from the static picture of poverty about the prospects for future poverty reduction?

In this paper we propose a methodology for addressing these questions, and give aggregate results for 86 developing countries in the mid-1990s. Our aim is to make a necessarily rough but methodologically consistent assessment of the magnitude and severity of absolute poverty, based on recent available data.

In the following section we discuss possible interpretations of an "absolute poverty line" which might be considered appropriate for our purpose. This is followed in Section 3 by an empirical examination of poverty lines for a number of countries, both developing and developed. This is used to identify two poverty lines for the subsequent analysis. In Section 4 we discuss issues which arise in measuring poverty from readily available data on income and consumption distributions, and our approach to measuring poverty in countries for which such data are not available. In Section 5 we present and discuss our estimates of the prevalence and severity of absolute poverty in the developing countries in the mid-1980s. In that section we also give an alternative, largely independent, estimate based on a previous estimate of the world Lorenz curve. In Section 6

Note: This paper is based on background research done in the preparation of the World Bank's *World Development Report 1990*. Its contents should not, however, be attributed to the World Bank, or any affiliated organization. The authors are grateful to the Review's referees for their helpful comments.

we discuss some of the implications of these results, particularly their bearing on the prospects for future poverty alleviation. In the final section we offer some conclusions.

2. Approaches to Defining an "Absolute Poverty Line"

Different societies have different perceptions of what constitutes "poverty," reflecting (in part) different overall levels of living. Our aim here is only to quantify the extent of *absolute poverty* in the developing world, interpreted as the inability to attain consumption levels which would be deemed adequate in only the poorest countries. This will leave out many persons who are clearly deprived *relative* to others around them.

There are a number of possible interpretations of an "absolute poverty line" for cross-country comparisons. Two possibiltiies can be suggested:

i) One can pick the cost of a bundle of goods which is reasonably well recognized as constituting an absolute minimum by international standards. The poverty line of India has been widely used for this purpose, and that alone makes its continued use compelling. But why India's poverty line, and not that of some other country? We should at least know how sensitive poverty counts may be to that choice and how India's poverty line compares to that of other countries.

ii) In principle, one can think of the real poverty line as comprising an "absolute" component which is constant across all countries, and a "relative" component, which is specific to each country. In seeking to measure the extent of absolute poverty one might simply ask: What is the *lowest* real poverty line observed in any country? This would seem to be a good indicator of the minimum acceptable poverty line in assessing absolute poverty. However, the answer may be quite sensitive to the particular countries surveyed and the inevitable measurement errors in assessing local poverty lines, and in comparing them across countries. It will also be influenced by inter-country differences in non-income factors; a country with good public services benefiting the poor, or a relatively low-cost climate, will naturally have a lower income poverty line. In the light of these considerations, a better approach is to try to assess a "typical" poverty line amongst the poorest cou tries. To implement this approach empirically in the next section, we assume that the relative component for any country is largely determined by its mean private consumption, though we allow the possibility of other factors (such as access to public services) which may also influence the poverty line. We will then be able to estimate the poverty line to be found in the poorest country, controlling for these other factors.

3. An International Comparison of Poverty Lines

Many countries now have reasonably well established (though rarely uncontentious) local poverty lines. From a wide range of sources within and outside the World Bank, we have compiled local poverty lines for 33 countries, both developing and developed. These should not be considered as "official" poverty lines, either of the governments or the Bank. Many are the estimates of independent researchers. Nor has our survey been exhaustive; there are undoubtedly credible poverty lines we do not know about. When more than one poverty line was found (such as for urban and rural areas), the lowest was used. When poverty lines are differentiated by family size we have used the per capita value at roughly average family size. Ravallion *et al.* (1991) gives the detailed results and sources.

There are very likely to be measurement errors in our series on poverty lines, either because of the incompleteness of our survey, or errors in the primary data. However, to the best of our knowledge, those errors are random and, in particular, uncorrelated with average consumptions. Thus we do not believe that these errors will bias our econometric estimates.

In converting local poverty lines to a common currency, and indeed for all such comparisons in this study, we have used the estimates presented by Summers and Heston (1988) of the adjustments to official exchange rates needed to give purchasing power parity (PPP). Ideally one would like to construct new PPP rates for the prices most relevant to the absolute poor, in which the prices of food-staples would clearly carry a high weight.

In Figure 1, the results of our survey of poverty lines are plotted against mean private consumption per capita, and both variables are for 1985 and are measured at PPP, using local CPI's when necessary. In Figure 2 we give a "blow-up" of the part of Figure 1 for developing countries only. India's poverty line is $23.14 per person per month.

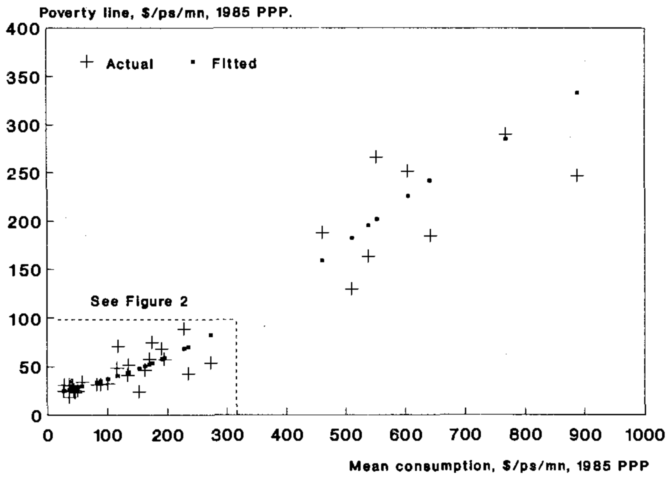

Figure 1

There is a clear tendency for the local poverty line to increase with mean consumption, though dispersion in poverty lines at most consumption levels is also evident, presumably reflecting non-income factors and/or measurement errors. The poverty line is below the mean in all cases. The figure also gives our fitted values of the poverty lines. Inspection of the figure suggests a semi-log functional form. The ordinary least squares estimate is as follows (with standard

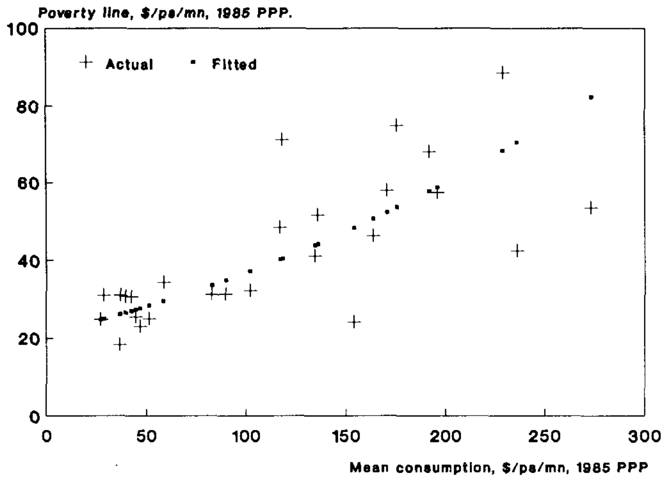

Figure 2

errors in parentheses):[1]

(1) $$\log(z_i) = 3.077 + 0.00334\mu_i - 0.0000011\mu_i^2 + \text{residual}_i$$
$$(0.105)\ (0.00056)\ \ (0.0000004)$$

$R^2 = 0.90$; SEE $= 0.292$; Mean d.v. $= 4.04$;

LM tests: NORM(2) = 1.3; HETERO(1) = 0.21; RESET(1) = 0.22.

where z_i is the poverty line in country i and μ_i is the mean private consumption per capita at constant *PPP*. The residual is assumed to pick-up other factors influencing local poverty lines, and measurement error in the latter. The actual and predicted values of z implied by the above regression are plotted in Figures 1 and 2. A test was also made for a structural break in the model between the developing and industrialized countries, but this failed to reject the null-hypothesis that the same model is valid for both.

The lowest mean consumption amongst the 86 countries studied in the *World Development Report* is Somalia at $22 per person per month in 1985 *PPP* prices. At this point, equation (1) predicts a poverty line of $23, only slightly different from that of India. Thus, India's poverty line is very close to the poverty line we would predict for the poorest country, and as such, can be considered a reasonable lower bound to the range of admissible poverty lines for the developing world.

However, it is also clear from Figure 1, that many low-income countries have more generous poverty lines. The $23 line is certainly on the low side of the range found amongst the poorest dozen or so countries in Figure 2. A more generous, and more representative, absolute poverty line for low-income countries is $31, which (to the nearest dollar) is shared by six of the countries in our sample, namely Indonesia, Bangladesh, Nepal, Kenya, Tanzania, and Morocco,

[1] Standard errors are given in parentheses below each coefficient. NORM, HETERO, and RESET are the Lagrange multiplier tests for normality, heteroscedasticity, and functional form respectively, distributed as Chi-square, with degrees of freedom in parentheses. The data are available in Ravallion *et al.* (1991).

and two other countries are close to this figure (Philippines and Pakistan). We shall use both these poverty lines, interpreting the lower line as defining "extreme absolute poverty."

A further question which the above results can illuminate is the extent to which the poverty lines of poor countries might be expected to increase in the future, with increases in average living standards. We do not have time-series data on how poverty lines have evolved over time, but the cross-sectional evidence in Figure 1 is suggestive. Here we find that, although higher mean consumption levels are associated with higher poverty lines, the poverty line tends to be less responsive to increases in the mean at low levels of consumption. At the mean over all countries of \$234 the elasticity of z with respect to μ implied by equation (1) is 0.66. However, at the mean consumption of India, the elasticity is very much lower, at 0.15. At the lowest consumption level it falls to 0.07. In short: the cross-country comparison does suggest that real poverty lines will tend to increase with growth, but they will do so very slowly for the poorest countries. Notions of "absolute poverty" appear to be relevant to low income countries, while "relative poverty" is of more relevance to high income countries.

For our purposes, it is not unreasonable to assume the same real absolute poverty line across all developing countries; two persons with the same real (*PPP* adjusted) income living in different countries will have the same measured poverty. This does not allow for differences between countries in relevant "non-income factors," such as access by the poor to public services.

4. Poverty Measures and Their Estimation

A wide range of existing poverty measures can be characterized by a function $P(z/\mu, L)$ which denotes the measured level of poverty in an economy in which the absolute poverty line is z, the mean "income" (or other suitable measures of living standards) is μ, and the distribution of income has the Lorenz curve L, which can be interpreted as a vector of parameters which fully characterize the structure of relative inequalities within the economy. The poverty line z may also be viewed as a function of μ (and possibly L), reflecting "relative poverty" considerations, though for the purposes of this study we will treat z as fixed, as discussed above. The function P is homogeneous of degree zero in z and μ; this is a very common property of poverty measures, and is unrestrictive.[2]

One quite straightforward approach to estimating $P(z/\mu, L)$ is to combine an estimate of the "world Lorenz curve" with our poverty lines, and an estimate of mean world income or consumption. There have been a number of estimates of the world Lorenz curve.[3] Early attempts ignored inequality within countries, but the more recent studies have used distributional data for a sub-set of countries, and assumed homogeneity within incomes groups (usually deciles) in aggregation. The most recently available world Lorenz curve is for 1986 (Berry *et al.* 1989), though this is based on 1970 Lorenz curves for individual countries (a sample

[2] Notice that the use of a poverty line which is a fixed proportion of the mean makes the poverty measure solely a function of the Lorenz curve; it might then be interpretable as a measure of inequality.

[3] See, for example, Kravis *et al.* (1988), Berry *et al.* (1983, 1989), Grosh and Nafziger (1986), Yotopoulos (1989).

of about 20 countries), up-dated solely according to country specific growth rates in national income.

Nonetheless, this may still give a satisfactory estimate, if relative inequalities *within* countries have changed little over the period 1970–86, or if the world Lorenz curve is insensitive to whatever changes have occurred. Berry *et al.* (1983, 1989) defend their estimates along these lines, and in their earlier paper they present evidence indicating that their measures of world inequality would respond little (in elasticity terms) to changes in inequality within most countries. However, their results do indicate quite high elasticities (over one for the world Theil index with respect to the national index) for some of the larger countries, including China and India. Some up-dating of at least these key distributions is clearly called for, and is feasible; for example, we now have quite recent estimates of distributions for both China and India. It may also be noted that we are interested in measuring poverty here, not inequality *per se*, and aggregate poverty measures can be quite sensitive to certain shifts in the Lorenz curve, a point which we will return to later.

In the light of these considerations, in this study we will exploit the fact that much more recent, and probably more reliable, distributional data have become available since 1970; indeed, all of the distributions we shall use are for the 1980s. We shall thus start afresh. It will be of interest to compare the results with those obtained using the Berry *et al.* world Lorenz curve for 1986.

When our data for a given country include a usable distribution of household consumption or income, we proceed by first estimating the parameters of the Lorenz curve L, and then we combine this with an estimate of z/μ to calculate P according to the explicit functional form appropriate to the specific poverty measure being considered (as discussed later).

This still leaves many countries for which adequate distributional data are unavailable. When we do not have an empirical distribution, we use an econometric extrapolation based on available correlates of poverty over the sub-set of countries with such data. This is probably a better practice than assuming that the available distributions are representative.

Given an explicit functional form for P, the contribution of changes in average incomes (for given distributions) to poverty alleviation can be estimated. A useful "benchmark" case is that of *distributionally neutral growth*, defined as an increase in μ holding L constant. However, this should only be viewed as a benchmark; there can be no presumption that growth is distributionally neutral, and from past experiences the departures from neutrality could go either way.

The specific functional form of $P(z/\mu, L)$ depends on (i) the functional form of the poverty measure and (ii) the parameterization of the Lorenz curve. In the rest of this section we discuss the assumptions we have made about functional forms.

Two specific poverty measures are considered: the *headcount index* (H) given by the proportion of the population with consumption below the poverty line, and the *income gap ratio* (G) given by the average consumption distance of the poor from the poverty line, expressed as a proportion of the poverty line. Thus, letting F denote the cumulative distribution of consumption, we have that $H = F(z)$ and $G = 1 - \mu^p/z$, where μ^p is the mean consumption of the poor.

There are advantages in following Foster *et al.* (1984) in normalizing the latter index by population size, rather than the number of poor. We then obtain the *poverty gap ratio* defined as $PG = G \cdot H$.[4] This is a more computationally convenient normalization, as it implies that the aggregate measure across any number of sub-groups is (like H) simply the population weighted mean of the sub-group values of PG.[5]

There are a number of options in specifying a parametric form for the Lorenz curve. We considered three possibilities: the original Kakwani and Podder (1973) specification, the later Kakwani specification (Kakwani, 1990), and the class of elliptical Lorenz curves (Villasenor and Arnold, 1989). All three can be readily estimated econometrically. The relative performance of these alternative specifications are discussed in Ravallion *et al.* (1991).

The calculation of the two poverty measures, H and PG, proceeds as follows. Let $L(p)$ denote the Lorenz curve, giving the estimated proportion of consumption by the poorest p% of the population. H can be obtained using the well-known fact that $x = \mu L'(p)$ is the inverse function of the distribution function $p = F(x)$, and so $L'(H) = z/\mu$. This has to be solved numerically for certain Lorenz curve specifications, though the problem is straightforward (Newton's method was used here). The poverty gap is then obtained readily using the fact that $\mu^p = \mu L(H)/H$.

These methods will not always work well. For example, if there are very few class intervals in the available data for the poor, then any estimate of poverty will be subject to error. Also, certain groupings of the data can yield distorted estimates of the Lorenz curve; for example, econometric estimates of the parameters may not satisfy the theoretical conditions required of a Lorenz curve.[6] We have checked those conditions and used quadratic or linear interpolation as options.[7]

To help assess the accuracy of these methods, in Table 1 we present an experiment we conducted in which we estimated poverty in Indonesia for 1984 using $z = \$31$ and alternatively calculated using:

(i) the primary data tapes of Indonesia's socio-economic survey (SUSENAS) of consumption by each of 50,000 sampled households,

[4] In general, we can write the FGT class of poverty measures for any distribution of income or consumption $\{y_i\}$, $i = 1, \ldots, n$, as follows:

$$P_\alpha = \sum_{y_i < z} [(z - y_i)/z]^\alpha / n$$

where α is a non-negative parameter. H is then obtained when $\alpha = 0$, while PG is obtained when $\alpha = 1$.

[5] That is not the only advantage. The additivity property also guarantees that measured poverty is *sub-group monotonic*, meaning that if poverty increases (decreases) within a sub-group then (*ceteris paribus*) aggregate poverty will also increase (decrease) (Foster *et al.*, 1984).

[6] A valid Lorenz curve $L(p)$ must be monotonic increasing and strictly convex for p in the $(0, 1)$ interval. When not implied for all such p, this was tested in the neighborhood of the estimated headcount index. A valid Lorenz curve should also have the limiting properties that $L(0) = 0$ and $L(1) = 1$. However, we have not imposed or tested these conditions (and they do not hold for the Kakwani model), on the grounds that we are mainly interested in obtaining a good interior fit. Villasenor and Arnold (1989) compare various Lorenz specifications, including the Kakwani-Podder model, on Australian data and find the elliptical model to be preferable. However, their comparison does not include the Kakwani (1990) specification used here.

[7] Though the Kakwani specification can go astray at very low values of p due to the limiting properties of its slope. This does not appear to be a practical problem here.

TABLE 1

ALTERNATIVE ESTIMATES OF POVERTY IN INDONESIA, 1984

Poverty Measure	Calculated Directly from Unit Record Data (sample size = 50,000)	Number of Intervals Used in Estimation from Grouped Data:			
		50 (18 "poor")	15 (8 "poor")	10 (4 "poor")	5 (2 "poor")
H: Headcount index (%)	33.02	33.74	33.64	33.76 33.88*	33.50 33.63*
PG: Poverty gap index (%)	8.52	9.10	9.04	9.09 9.17*	9.07 9.10*

* Poverty line in the middle of the class interval.

(ii) parameterized Lorenz curves (using the Kakwani specification) calibrated to a detailed description of the frequency distribution based on 50 class intervals formed from the unit record data, of which 18 intervals are below the poverty line,

(iii) Lorenz curves calibrated to far more "coarse" frequency distributions using 15, 10, and 5 class intervals, of which 8, 4, and 2 respectively are below the poverty line. These are typical of the detail usually available from published sources. With the more aggregated data it is also very common for the poverty line to fall well within a class interval. We also present results for this case, where we have constructed grouped distributions in which the poverty line is in the middle of a class interval.

The results confirm a loss of accuracy in using grouped data, though it is not large; the headcount index, for example, is estimated within three-quarters of a percentage point using the grouped data. Furthermore, the loss of accuracy is remarkably unaffected by contraction in the number of class intervals in the grouped data. We do not know how specific this result is to these data (as we rarely have access to the unit record data). However, it does suggest that there need not be much loss of accuracy in using even highly aggregated data for measuring poverty.[8]

Problems of comparison across countries are another source of error in the aggregate estimates. Available distributions generally pertain to either household income or consumption at one point in time, and are either household aggregates or are normalized by the number of persons in the household. Our first preference was to use a distribution of household consumption per person. This was possible for 12 of the countries. Income distributions were used for the others, adjusted pro rata according to an average propensity to consume estimated from national accounts.[9] If the distribution was in household form, we adjusted it according

[8] In Ravallion *et al.* (1991), we further discuss these methods and the various experiments we conducted to assess their accuracy.

[9] Income distributions will generally exhibit higher inequality, though this need not mean that the poverty measures are biased upward relative to a consumption based measure. What we need to know is whether the differences in the extent of inequality associated with whether one is using incomes or expenditures yield different poverty measures at a given mean (since we have controlled for the latter, at least according to differences in average propensity to consume). To see if this is

to either income/consumption specific household sizes (when available) or average household sizes, for which purpose one has no choice but to assume that the ranking of households is the same.

All currency conversions use constant purchasing power exchange rates, based on Summers and Heston (1988). For two countries, Burma (Myanmar) and China, *PPP* rates are not given in this source. We estimated these from an extrapolation model of the *PPP* rates across the other 84 countries. The variables used for extrapolation were the official exchange rate, the mean private consumption at official exchange rates, and various social indicators (Ravallion, *et al.* 1991). For both countries, the constant *PPP* exchange rate is estimated to be about one-quarter of the official rate. We estimate the *PPP* for China as a proportion of the official exchange rate to be 23 percent, which is very close to the estimate for 1975 reported by Kravis (1981) of 24 percent (using the geometric mean of *PPP* estimates based on U.S. and China expenditure weights).

However, imprecision in the estimated *PPP* for China is likely to be an important source of imprecision in the final poverty estimates, given the country's population share and (as we will see later) the high elasticity of the country's distribution function in the region of the poverty lines. We shall return to this point.

Further information on the data and the estimation methods used for the 22 countries for which adequate distributional data were available is given in Ravallion *et al.* (1991). Data are known to exist for other countries, but were either not available, or deemed inadequate for our purposes. The average population size of this set of countries is much higher than the average for the 86 countries, so that these 22 countries represent 76 percent of the total population covered.

Poverty counts for the remaining countries have been based on extrapolations. The extrapolations used a set of variables which are available in a complete series across the 86 countries. These variables include mean consumptions for 1985 as estimated from national accounts and converted to constant $U.S. using the 1985 *PPP* deflators, and a broad set of social and economic indicators, including life expectancy, infant and child mortality, primary and secondary school enrollment rates, proportion of the population in the labor force, female labor force participation, and the share of the population living in urban areas. Logit transformations of both poverty measures were used to guarantee that the predicted values are within the theoretical bounds (0, 1) and that the error term is theoretically unbounded. Initial runs of the model were based on a wide set of social indicators and their multiplicative interactions, and only variables with standard errors greater than the estimated coefficients were deleted, thus minimizing the regression's error variance.[10]

---

the case, we regressed each poverty measure (in logit transform) against mean consumption per capita in constant *PPP* dollars (in log form) and a dummy variable for whether or not the distribution is one of income or expenditure (1 for income, 0 for expenditure). The dummy variable had a positive sign for both H and PG, and for both poverty lines, but it was not significantly different from zero in any of the four cases (highest t-ratio was 0.96). We conclude that, for these data, the use of the income based surveys is unlikely to lead to a significant over-estimation of consumption poverty.

[10] The extrapolation models used are discussed further in Ravallion *et al.* (1991).

It should also be noted that there are many potentially important country specific effects which will not be captured by the extrapolation model. This is of no concern for the sample of 22 countries (as we shall, of course, use the actual values for those countries). However, it is worrying for the other countries.

On the whole, within the sample the predictive performance of the extrapolation models seems adequate for the purpose of estimating *aggregate* poverty in the countries for which distributional data are unavailable. However, the individual country extrapolations should be treated as very approximate. We shall calculate standard errors for the aggregate measures.

5. The Estimates of Aggregate Poverty

Our estimates of both poverty measures for both poverty lines for the 86 developing countries are given in Table 2. About one in three persons in the developing countries fails to attain our upper poverty line, while one in five does not reach our lower line, defining extreme poverty. The combined population

TABLE 2
Estimates of Poverty Measures for the Developing Countries, 1985

Region	Poverty Line ($/ps/mn)	Headcount Index		Poverty Gap Index	
		Point estimate	95% confidence interval*	Point estimate	95% confidence interval*
Aggregate	31	33.0	(27.9, 39.2)	10.2	(8.4, 14.3)
estimates	23	18.8	(14.9, 25.5)	5.1	(4.0, 8.9)
By region:					
South Asia	31	50.9	(49.8, 52.6)	15.9	(15.5, 16.9)
	23	29.4	(28.7, 31.1)	7.2	(7.0, 8.1)
of which India:	31	55.0		17.6	
	23	32.7		8.2	
East Asia	31	21.2	(21.1, 21.5)	6.0	(5.9, 6.0)
	23	9.7	(9.6, 9.9)	2.8	(2.8, 2.9)
of which China:	31	21.1		5.9	
	23	9.2		2.9	
Sub-Saharan	31	46.9	(18.6, 75.7)	15.0	(5.4, 36.5)
Africa	23	30.5	(8.9, 65.0)	7.7	(2.1, 27.5)
Middle-East and	31	31.0	(13.3, 50.9)	9.6	(3.3, 24.0)
North Africa	23	20.6	(6.4, 42.4)	5.5	(1.5, 19.3)
Eastern Europe	31	7.8	(7.3, 9.7)	2.6	(2.4, 3.3)
	23	4.3	(4.0, 5.7)	1.4	(1.3, 2.0)
Latin America	31	19.1	(14.0, 28.9)	6.9	(5.2, 11.5)
	23	12.5	(9.0, 21.0)	5.1	(4.0, 9.6)

* Allowing solely for imprecision due to the need to predict the poverty measures for those countries for which suitable distributional data are unavailable; see text.

size of the countries covered is 3,442 million, so the estimated total number of poor is 1,137 million, of which 645 million are deemed to live in extreme poverty. We shall consider some implications of these results in the following section.

It should go without saying that these are rough estimates only. There are a number of sources of imprecision in the estimation methods. Most are familiar, such as measurement errors in the underlying distributions, or errors in the estimated *PPP* rates. One source of error, on which we can readily comment further, is our use of extrapolations for those countries for which adequate distributional data were unavailable. We have calculated the 95 percent confidence interval around the point estimate for each of the countries for which extrapolations have been used. The poverty measures for the 22 countries for which we have used distributional data are assumed to be measured without such error. The implied confidence intervals around the aggregate point estimates are given in parentheses in Table 2.

The aggregate headcount index is found in the range 28–39 percent with 95 percent confidence using the upper poverty line and 15–25 percent using the lower line. The corresponding confidence intervals for our estimates of the numbers of poor are 960–1,348 millions and 513–876 millions.

Another source of error is in the estimation of a *PPP* rate for China's currency. This is worrying for two reasons. Firstly, although poverty is below average in China, given the country's size it still accounts for a large share of aggregate poverty. Secondly, the consumption distribution function for China is unusually steep in the region of the poverty lines; our estimates of the headcount index at the upper and lower poverty lines imply an elasticity of 3.7 (as compared to an elasticity of about two for the developing countries as a whole). This undoubtedly reflects distributional policy in China. By implication, aggregate povery measures will be quite senstive to errors in measuring the *PPP* rate of exchange for China; the elasticity of the aggregate headcount index at the lower poverty line with respect to China's *PPP* rate will be about 0.54. Suppose, for example, that our *PPP* rate has been over-estimated by 10 percent. this would imply about a 5 percent overestimation in the aggregate headcount index, which would fall by one percentage point, or about 35 million people.

In Table 2 we also give the regional breakdown of the poverty measures, both point estimates and 95 percent confidence intervals. The estimates are clearly far more accurate for some regions than others. In particular, while we have used distributional data for three-quarters of the aggregate population, such data were only available for 11 percent of the population of the Middle East and North Africa region and a mere 6 percent of the population of Sub-Saharan Africa. At the other extreme, the Asia coverage is excellent, being 95 percent of the population of South Asia, and 97 percent of the population of East Asia. (The Latin America coverage is 55 percent, while for Eastern Europe it is 85 percent). These regional variations in survey coverage are clearly reflected in the 95 percent confidence intervals around the point estimates given in Table 2. While we can be reasonably confident about our estimates for South Asia, for example, very wide margins for error must be allowed around those for Sub-Saharan Africa and the Middle-East/North-Africa.

Subject to this (important) caveat, we find that for the upper poverty lines the highest headcount index is for South Asia, closely followed by Sub-Saharan Africa. However, this order reverses for the lower poverty line. The same observation holds for the poverty gap index. Loosely speaking, the incidence and severity of "moderate" absolute poverty can be said to be highest in South Asia, while that of "extreme" poverty is highest in Sub-Saharan Africa.

It should be recalled that these poperty assessments will not necessarily accord with *local* perceptions of poverty, since we have used fixed absolute poverty lines across all countries. For example, our counts for Latin America will seem low to observers familiar with the results obtained using the (typically) higher poverty lines of that region. Conversely, our counts for some regions will seem high. For example, a commonly used poverty line for China translates into about $25 per person per month at our estimated *PPP* rate. This is well below our upper poverty line, and (as we have already discussed) China's distribution is very steep in this region, so our figure for the upper poverty line is much higher than local assessments. Our lower poverty line is more consistent with current perceptions of poverty in China.

As an aside, it is of interest to compare our results with an alternative method of estimating aggregate poverty using the Berry *et al.* (1989) estimate of the world Lorenz curve of consumption for 1986 (based on 1970 distributions, and national growth rates over the intervening period). From the results of Summers and Heston (1988), we estimate mean consumption per capita of the non-socialist countries to be $230 per person per month in 1985 at *PPP*.[11] In Table 3 we give our estimates of both poverty measures for both poverty lines.[12] We also give an estimate which includes our figure for China. The total population covered is then 4,278 million, and the main difference in coverage over the earlier results is that this now includes the industrialized countries. It is reasonable to assume that the number of persons in the industrialized countries who do not attain these very low poverty lines (by the standards of those countries) would be negligible. Thus the estimated numbers of poor should be comparable with those we obtained earlier.

TABLE 3

ALTERNATIVE ESTIMATES BASED ON A WORLD LORENZ CURVE

Poverty Measure	Poverty ($/ps/mn)	Non-socialist Countries (Elliptical model of Berry *et al.* 1989 (s = 0.76)	(s = 0.24)	Aggregate Poverty (Non-socialist + China)
H: Headcount	31	25.1	21.1	24.1
index (%)	23	18.4	9.2	16.2
PG: Poverty	31	10.8	5.9	9.6
gap index (%)	23	6.9	2.9	5.9

Note: s = population share, total population = 4,278 million.

We find that there are 1,031 and 691 million for the upper and lower poverty lines respectively, by this alternative estimation method, as compared to 1,137 and 645 million by our earlier method. The two methods are in reasonably close accord.

[11] This is the 1985 population weighted mean of the consumption per capita estimates at constant *PPP* implied by Summers and Heston (1988, Table 3), after correcting the typing errors for India and El Salvador.

[12] This is based on the elliptical model. The Kakwani specification violated the theoretical condtions for a Lorenz curve.

6. Some Implications

There is only so much one can learn from a single "snapshot" of poverty. However, a number of observations of possible interest can be made about our main results in Table 2.

We are clearly looking at a fairly *steep* segment of the consumption distribution function for the developing countries. In going from the lower to the higher poverty line in Table 2 the headcount index increases from 19 percent to 33 percent, representing an elasticity of 2.2.[13] This implies that, in the region of these poverty lines, the headcount index of poverty will decrease with an elasticity of 2.2 under distributionally neutral growth (recalling that the poverty measure is homogeneous of degree zero in z and μ).

For example, at a growth rate in real mean consumption of 1 percent per year, and without any change in relative inequalities, the headcount index will decline at about 2 percent per year. If the rate of population growth does not exceed this amount, then the number of poor will also decline. In fact, the average rate of population growth in the developing countries is about 2 percent per year, so the total number of poor will be in decline as long as the developing countries can attain a distributionally neutral growth rate of at least 1 percent per year.

The aggregate poverty gap is 10 percent and 5 percent of the poverty line for the upper and lower poverty lines respectively. This can also be expressed as a proportion of mean consumption. (The aggregate poverty gap as a proportion of aggregate consumption is given by *PG* times z/μ.) The mean consumption of the sample of 86 countries is $102 per person per month in 1985, at *PPP*. Thus the aggregate poverty gap represents 3.1 percent and 1.1 percentage of aggregate consumption for the upper and lower poverty lines respectively. Also recall that this calculation uses the aggregate consumption of the developing countries. The aggregate poverty gap of the developing countries represents 1.5 and 0.5 percent of the aggregate consumption of the non-socialist countries for the upper and lower poverty lines respectively.

Provided that transfers could be perfectly targeted to the poor, and without reducing mean consumption, these calculations suggest that only modest gains to the poor would be needed to eliminate poverty. For example, with less than one-half of one percent of world consumption perfectly targeted without distortion, one could guarantee that every one in the world could reach at least India's poverty line.

However, neither of these provisos should be taken lightly. For example, the information needed for perfect non-distortionary targeting to the poor is not readily available, and may, indeed, be very costly to obtain. The potential impact on poverty of *informationally feasible* targeting may be far more modest (see, for example, the simulations for India reported in Datt and Ravallion, 1990). To illustrate how much informational constraints could bite into the impact on poverty suppose instead that absolutely no targeting was possible i.e. each person

[13] Note that this is an "arc elasticity", which will generally differ from the point elasticities at both poverty lines. The point elasticity for the headcount index requires estimation of the density function of incomes at the poverty line. That is only possible for the sub-set of countries for which we have distributional data.

receives the same gain, irrespective of their initial position. Assuming that the lowest consumption is zero, the transfer needed to eliminate poverty is given by z times the population of the developing world. Clearly this is very much higher than the amount needed to eliminate poverty with perfect targeting; the necessary proportion of the aggregate consumption of the non-socialist countries would be 14 percent and 11 percent for the upper and lower poverty lines respectively. Thus, depending on the information available for targeting, between 0.5 and 11 percent of aggregate consumption would be needed to guarantee that nobody in the developing world falls below India's poverty line. Of course, in practice we will at least know countries of residence and this will undoubtedly be a useful indicator for targeting. The 11 percent estimate could be well above the actual cost of eliminating poverty with readily available information.

Our results can also provide information on the prospects for alleviating poverty through economic growth. The point elasticity of the poverty gap index with respect to distributionally neutral growth is given by $1-H/PG$ (also equal to $-\mu^p/(z-\mu^p)$), and can be readily calculated from the data in Table 2. The implied absolute elasticities are 2.2 and 2.7 for the upper and lower poverty lines respectively. At a growth rate in the mean of 1 percent per year, and without any change in relative inequalities, the mean poverty gap will decrease at a rate of 2–3 percent per year, over the range of poverty lines considered.

However, deviations from distributionally neutral growth can matter to the future prospects for reducing absolute poverty. This will depend on the way in which Lorenz curves shift. Following Kakwani (1990), let us assume that the shift in the Lorenz curve $L(p)$ at each value of p is directly proportional to $p-L(p)$ i.e. the Lorenz curve shifts by a constant proportion of the difference between each income group's actual share of total income, and the share it would have if there were perfect equality. We cannot know if this is plausible, though it has been found to give an excellent approximation to recent shifts in Indonesia's Lorenz curve (Ravallion and Huppi, 1989). It also yields quite a close approximation to the shifts in world Lorenz curves from 1960 to 1986, and from 1970 to 1986, reported by Berry *et al.* (1989).[14]

Under this assumption, the elasticity of the poverty gap index with respect to the Gini index is given by $1+(H/PG)[(\mu/z)-1]$ (Kakwani, 1990). Notice that this takes a value of one if the poverty line equals the mean. Thus, for example, the elasticity is found to be around one for Bangladesh (Ravallion, 1990). However, for the developing countries as a whole (and, indeed, the world distribution), the poverty line is low relative to the mean. When combined with the "skew" in the aggregate Lorenz curve associated with the disparities in means across countries, this has the effect that the elasticity of poverty with respect to the Gini index can be very much higher than one.

This is exactly what we find from our results in Table 2. The elasticity of the poverty gap is 8.4 for the higher poverty line, and 13.6 for the lower one. Thus even small deviations from distributionally neutral growth could have a significant impact on progress in reducing world poverty. Consider again the effect of a

[14] The correlation coefficients between predicted and actual Lorenz curves for 1986 over the 11 available points on the interior of the (0, 1) interval are 0.9998 based on the 1960 Lorenz curve and 0.9999 based on 1970. The standard deviations of the errors in prediction are 0.0061 and 0.0078.

1 percent growth rate in mean consumption, representing a 16 percent increase in the mean over the 15 years from 1985 till the end of this century. The entire effect of this on the poverty gap for our higher poverty line would be wiped out if it was associated with a 4.3 percent increase in the Gini index over this 15 year period. For the lower poverty line it would take a 3.1 percent increase in the Gini index.

A 3–4 percent increase in the Gini index between 1985 and 2000 is not inconceivable. For example, Berry *et al.* (1989) report that the world Gini index increased by 5 percent between 1960 and 1986, and by 3 percent between 1970 and 1986. Taking their estimate of the world Gini index for consumption in 1986 of 0.67, a 4 percent increase would bring this up to 0.70 by 2000. This would be equivalent to a lump-sum transfer of 6 percent of the world mean, from each person below the median to each person above it, over 15 years (or about 0.4 percent of the mean per year).[15]

7. Conclusions

A strong case can be made for treating a consumption level of $23 per person per month in 1985 at $U.S. purchasing power parity as a reasonable lower bound for the poverty line. This has been widely used in poverty assessments for India, and we have argued that it is also the poverty line that one would expect in the poorest country in our sample. However, for comparative purposes it is also worth considering a rather less meager criterion. We find that a consumption level of $31 per month is actually a far more common poverty line amongst the dozen or so low income countries for which poverty lines have been calculated. The range $23–31 embraces quite well the poverty lines used by low income countries.

In measuring the magnitude and severity of absolute poverty in the developing countries, we have used recently available distributions of consumption or (occasionally) income for a sub-set of countries, accounting for about three-quarters of the total population covered by the estimates. We have then used regression models calibrated to data on average private consumption and various social indicators for that subset to extrapolate to the remaining quarter of the population. Allowing for imprecision in the extrapolations, our method yields acceptably low variances in estimates of aggregate poverty. We have also considered an alternative methodology, based on recent estimates of a world Lorenz curve. Allowing for the differences in methodology and coverage between these two approaches, it is encouraging that they yield quite similar results on the magnitude of absolute poverty.

However, it should go without saying that these estimates are no better than the quality of the underlying data allows. We have made a number of tests of the reliability of our estimates. Two points stand out:

(i) China's high population share, and the fact that China's consumption distribution is particularly steep at the lower end, make our estimates of aggregate poverty particularly sensitive to errors in estimating the *PPP* rate for China.

[15] This uses a result due to Blackburn (1989).

(ii) The very limited coverage of the available distributional data for Africa has necessitated the use of extrapolations from other data for most of that region, giving greater imprecision than for other regions. This should be of concern, particularly given the current high incidence of poverty in the region, and the low expected growth rates.

In assessing the overall magnitude of poverty, we estimate that about 1,137 million people in the developing world did not attain a consumption level of $31 per month in 1985. Of these, 645 million did not attain our lower, and extremely frugal, poverty line. The *severity* of poverty amongst the poor can be gauged by their average poverty gap. We find that the average consumption of the poor in the developing world was about 27-31 percent below the poverty line (for lower and higher poverty lines respectively). This may be a very significant gap for a poor person. But, despite the large numbers of poor, the aggregate gap turns out to be a very small proportion of world consumption; for example, the aggregate poverty gap of the developing countries at our slightly more generous poverty line is about 1.5 percent of the aggregate consumption of the non-socialist countries, falling to a mere 0.5 percent at our lower poverty line. This suggests that, with perfect targeting, it would take only a modest redistribution of world income to bring everyone up to these minimal consumption levels. However, the information available to policy-makers is generally a lot less than one would need for perfect targeting. So this calculation may substantially understate the magnitude of the transfers needed to eliminate absolute poverty in practice.

Our results imply that aggregate poverty in the developing world will respond fairly elastically to distributionally neutral growth; for example, a 1 percent annual growth rate at all consumption levels will reduce the proportion of the population who are poor by about 2-3 percent per year. If annual population growth rates stay at about 2 percent or lower, then the total number of poor will decline.

However, our results also suggest that even a seemingly modest worsening in distribution could upset this progress in poverty alleviation. For example, if the same 1 percent growth rate in the mean was associated with only a 0.25 percent annual increase in the world Gini index of inequality, then the reduction in the poverty gap attainable through growth would be virtually eliminated (assuming that the Lorenz curve shifts out roughly "proportionately" at all points). Such a rate of increase in the world Gini index has been observed over recent decades, associated with the relatively low growth rates of a number of the poorest countries. In this case, the numbers of persons who do not attain even the most meager consumption levels would almost certainly increase.

On the other hand, a pattern of growth more favorable to the poor could rapidly accelerate global poverty alleviation. Consider, instead, a rate of *decrease* in the world Gini index of 0.25 percent per year, roughly equivalent to a transfer of one-third of one percent of the world's mean income from the better off half to the poorer half of the world's population. This would roughly double the rate of decrease in the aggregate poverty gap (measured against our more generous poverty line) associated with a 1 percent annual growth rate in mean consumption of the developing countries; instead of the decrease of 2.2 percent per year we could expect with distributionally neutral growth, we would see the poverty gap fall at an impressive annual rate of 4.5 percent.

REFERENCES

Berry, A., Bourguignon, F., and Morrisson, C., The Level of World Inequality: How Much Can We Say?, *The Review of Income and Wealth*, 29, 217–242, 1983.

———, Changes in the World Distribution of Income between 1950 and 1977, *The Economic Journal*, 93, 331–350, 1983.

———, The World Distribution of Income: Evolution over the Recent Period and Effects of Population Growth. Paper presented at a conference on "Consequences of Rapid Population Growth," United Nations, New York, 1989.

Blackburn, M. L., Interpreting the Magnitude of Changes in Measures of Income Inequality, *Journal of Econometrics*, 42, 21–25, 1989.

Datt, G. and Ravallion, M., Regional Disparities, Targeting, and Poverty in India, in M. Lipton and J. van der Gaag (eds.) *Including the Poor*, World Bank, Washington, D.C., 1991.

Foster, J., Greer, J., and Thorbecke, E., A Class of Decomposable Poverty Measures, *Econometrica*, 52, 761–765, 1984.

Grosh, M. E. and Nafziger, E. W., The Computation of World Income Distribution, *Economic Development and Cultural Change*, 34, 347–360, 1986.

Kakwani, N., Poverty and Economic Growth with Applications to Cote D'Ivoire, Living Standards Measurement Study Working Paper 63, World Bank, 1990.

Kakwani, N. and Podder, N., On the Estimation of Lorenz Curves from Grouped Observations, *International Economic Review*, 14, 137–148, 1973.

Kravis, I., An Approximation of the Relative Real Per Capita GDP of the People's Republic of China, *Journal of Comparative Economics*, 5, 60–78, 1981.

Kravis, I., Heston, A., and Summers, R., Real GDP Per Capita for More Than 100 Countries, *Economic Journal*, 88, 215–242, 1978.

Ravallion, M., Datt, G., van de Walle, D., and Chan, E., Quantifying the Magnitude and Severity of Absolute Poverty in the Developing World, PRE Working Paper 587, World Bank, 1991.

Ravallion, M. and Huppi, M., Poverty and Undernutrition in Indonesia During the 1980s, Policy, Planning and Research Working Paper 286, World Bank, 1989.

Ravallion, M. and Huppi, M., Poverty and Undernutrition in Indonesia During the 1980s, Policy, Planning and Research Working Paper 286, World Bank.

Summers, R. and Heston, A., A New Set of International Comparisons of Real Product and Price Levels Estimates for 130 countries, 1950–1985, *Review of Income and Wealth*, 34, 1–26, 1988.

Villasenor, J. and Arnold, B. C., Elliptical Lorenz Curves, *Journal of Econometrics*, 40, 327–338, 1989.

Yotopoulos, P. A., Distributions of Real Income: Within Countries and by World Income Classes, *Review of Income and Wealth*, 35, 357–376, 1989.