How *Not* to Count the Poor? A Reply to Reddy and Pogge

Martin Ravallion[1]

Reddy and Pogge (2002) claim that the World Bank has systematically underestimated the extent of poverty in the world in the 1990s, and overestimated its rate of decline. This note argues that their claims do not stand up to close inspection and that there are serious concerns about their proposed alternative method of counting the world's poor.

The World Bank's poverty measures

Different people naturally have different ideas of what "poverty" means. This is true between countries as well as within a given country. The 1990 *World Development Report* (WDR; World Bank, 1990) explicitly recognized that richer countries tend to have higher poverty lines when converted to a common currency at exchange rates that attempt to assure purchasing power parity (PPP). Amongst poor countries, there is very little income gradient across countries in their poverty lines — absolute consumption needs dominate in a poor country. But the gradient rises as incomes rise, with poverty lines being roughly proportional to mean consumption for rich countries (World Bank, 1990; Ravallion et al., 1991).

Recognizing this feature of how poverty lines vary, how should we measure poverty in the world as a whole? One might use the poverty lines that prevail in each country. But then one would not be treating people with the same level of real consumption the same way, and so the resulting measures would lose meaning as measures of absolute consumption poverty. Relative poverty lines can still be defended if one believes that relative deprivation matters to a person's welfare. For comparison purposes, the Bank has also been producing poverty measures that take this approach. (Chen and Ravallion, 2001, provide such relative poverty measures for the developing world and by region.)

In the bulk of its efforts at global poverty monitoring, the Bank has taken the position that to measure absolute consumption poverty on a consistent basis across countries one must use a common poverty line. But if one wants a single number for the poverty count in the world, whose poverty line should it be? Since the WDR 1990, the Bank chose to measure global poverty by the standards of what poverty means in poor countries, which gave the "$1/day" line. It is fully acknowledged that this is a conservative definition; one could hardly argue that the people in the world who are poor by the standards typical of the poorest countries are not in fact poor.

This poverty line is then converted to local currency using the latest available PPP exchange rates for consumption. The best available consumer price indices are then used to convert the international poverty line in local currency to prices prevailing at about the time

---

[1] The author is Research Manager for Poverty in the World Bank's Development Research Group. These are the views of the author, and need not reflect those of the World Bank.

of the surveys. These poverty lines are applied to distributions of consumption per person (or income if consumption is not available) constructed from nationally representative household surveys. Adjustments to the data are often required for consistency, such as assuring that population weights are used to obtain an unbiased estimate of the individual distribution of household consumption per person. Calculations are done from the primary data (either micro data or appropriate tabulations).

However, it should be noted that in the vast bulk of the Bank's analytic and operational work on poverty, the "$/day" line is ignored, and with good reason. When one works on poverty in a given country, or region, one naturally tries to use a definition of poverty appropriate to that setting. Most of the time, the Bank's poverty analysts don't need to know what the local poverty line is worth in international currency at purchasing power parity. The main annual tabulation of the "$/day" poverty numbers, the Bank's *World Development Indicators* gives estimates based on national poverty lines side by side with the international lines, and has done so since these data were first published. Behind every one of these country numbers is a body of work as part of the Bank's Country *Poverty Assessments* and (more recently for low-income countries) the country's own *Poverty Reduction Strategy Paper*.

Nor does the Bank claim that the popular $1/day line is the only line that can be used for international comparisons. Indeed, the Bank has regularly produced estimates for a line set at twice this value — giving a poverty line more typical of middle-income countries. Past estimates have covered a range from $21 to $60 per month at 1985 PPP (see, for example, Chen et al., 1994).

Purchasing power parity and poverty measurement

Since the 1990 WDR, the Penn World Tables (PWT; see, for example, Summers and Heston, 1991) have been the main source of the PPP rates for consumption, though the latest estimates used the 1993 PPPs estimated by the Bank's Development Data Group.

Reddy and Pogge are concerned that existing PPPs (from either PWT or the Bank) do not correspond to the cost of a *"well-defined basket of commodities"* which leads them to claim that *"...existing PPPs are generally inappropriate for identifying the real incomes of poor households and hence the incidence of absolute poverty"* (p.10). They go on to argue that *"...the only way to avoid this problem is to <u>start</u> from a particular reference basket of commodities and to construct PPPs that accurately reflect the relative costs of purchasing <u>this</u> basket in different countries."* So they appear to be proposing to price a single bundle of goods in each country relative to a reference country.[2]

The idea of basing PPPs on a fixed bundle of goods is problematic for well-known reasons. People consume very different things in different countries, reflecting in part the differences they face in relative prices. I would be surprised if any kind of consensus could

---

[2]   I say "appear" here because Reddy and Pogge seem to back away from this position somewhat in the paper's conclusions, possibly in recognition of the likely problems. However, it is difficult to understand what exactly they are proposing from their conclusions, so I will confine my comments to the parts of the paper that seem reasonably clear to me.

be reached on what should be included in the single global bundle of goods, comparable to the consensus that has been established around the $1/day concept.

In fact, the deficiencies of the idea of using a single bundle of goods led to the types of price indices currently in use for constructing PPPs. (And not just for the Bank's global poverty monitoring task, which is just one of many applications for existing PPP exchange rates.) Ideally the underlying price index would only reflect differences in the cost of a reference level of welfare, fixed across all countries. This means that the reference bundle of goods cannot be the same across countries, given that relative prices vary (particularly, but not only, for non-traded goods) and hence that consumers can substitute amongst goods to achieve the same level of welfare — moving along their demand curves. The PPPs underlying the Bank's global poverty measures are based on the Fisher index, which gives a true cost of living index (reflecting differences in relative prices consistently with consumer preferences) if the utility function representing consumer preferences is quadratic. (This has been known since the 1920s.)

That is not to say that existing PPPs are ideal. Indeed, as is also well recognized in the literature, no index found in practice can claim to be perfect. However, the specific criticisms of the Fisher index made by Reddy and Pogge are based on a questionable analysis of the properties of that index. They claim that *"simple differentiation"* reveals a number of dubious properties of the index. To show this, they take the derivative of the Fisher index with respect to quantity holding prices constant. However, the quantities are functions of prices, as given by the underlying demand functions. (Equivalently, one can write the prices as functions of quantities, as given by the inverse demand functions.) So their differentiation entirely misses the demand effects that the index captures, to assure that it gives valid cost-of-living comparisons (at least for quadratic utility functions). The subsequent criticisms of the Fisher index based on this differentiation are thus ill-founded.

While it would not be progress to follow Reddy and Pogge's approach, it would be better to have PPPs designed for poverty measurement, weighted to say the consumption bundle of the poor, which one might set at say the 25[th] percentile from the bottom or the median. Such "PPPs for the poor" have been proposed often (including by the people doing the poverty measures that Reddy and Pogge criticize). But convincing implementation is not going to be easy. Unlike the mean consumption bundle (derived from national accounts), we do not have data for many countries on the consumption bundles at each income level, and there are comparability problems amongst the countries for which we do have these data. An alternative approach is proposed by Deaton (2001), whereby one would fix seemingly "plausible" PPPs at a base date, and make no further revisions beyond deflation for price changes over time within each country.

Biases?

The fact that we judge the extent of consumption poverty in the world by the standards typical of low-income countries clearly does not mean that we are underestimating the extent of the world poverty problem. Obviously if you use a higher standard you will get a higher poverty count. For example, the "$2/day" poverty count is published side-by-side with the "$1/day" count in the Bank's *World Development Indicators* for all years in which the numbers have been published. The "$1/day" line does not claim to be anything else than a poverty line typical of poor countries. To say that we are underestimating poverty by this

method is like saying that one underestimates length using a ruler calibrated in inches rather than centimeters. Surely if one knows how the ruler is calibrated there can be no confusion?

While the aims of the exercise have never changed, there have been improvements in data, and some corresponding changes in methods. There has been a huge increase in the country coverage of the Bank's global poverty aggregates; from one national survey for each of 22 countries in the WDR 1990 to 300 surveys for 90 countries now. A substantial effort by the Bank has gone into this expansion in the data base on household living standards.

The PPPs have changed too with new and better price data. The 1993 PPPs are a great improvement in terms of country coverage, though there are continuing (and possibly increasing) concerns about data quality. However, the two sets of PPPs are not comparable. So there is no straightforward way to convert the old $1/day line at 1985 PPP to a new line with base 1993. Instead, the only defensible approach was to go back to the original poverty line for the WDR 1990, and recalculate them with the new set of PPPs, and re-estimate the relationship between national poverty lines and mean consumption which led to the original $0.75 and $1/day lines used in the WDR 1990. The 1990 WDR $1/day line had been picked by eye-balling the scatter of points in the relationship between national poverty lines and national mean consumption. For the revision we used instead the median of the lowest 10 poverty lines amongst the 33 countries, which gave the figure of $1.08 at 1993 PPP.

The naïve approach of simply adjusting the old line upwards for inflation in the US would ignore the fact that there has been (in effect) a PPP devaluation of poor countries relative to the US over the period. For example, China's and Indonesia's poverty lines at 1985 PPP are almost identical to their poverty line at 1993 PPP; India's poverty line at 1993 PPP is only 17% higher than its poverty line at 1985 PPP. Yet adjusting the 1985 $1/day line for US inflation would entail an upward increase of roughly 50%. In other words, if we had simply adjusted the $1/day line for inflation in the US between 1985 and 1993 we would have obtained a poverty line which is well above the median of the 10 lowest poverty lines at 1993 PPP, and so could no longer claim to be the poverty line that is typical of poor countries. That would certainly entail a re-calibration of the ruler.

Coincidentally, the new $1.08 line with the new PPPs gave a very similar global poverty count to the old $1/day line and old PPPs for the common reference year, 1993 — both give a poverty count of around 1,300 million people for that year. This was not done by design, but it was judged to be a desirable feature. I agree with Reddy and Pogge that there is no reason why the new method should give the same poverty count for the same year as the old one. But it did. I don't think any neutral observer would go so far as to say that it was a *"serious conceptual error"* on our part to check if we got the same count for the reference year; it is surely a natural thing to look at the resulting estimate of ostensibly the same thing when you change your data and methods, and to draw comfort from their similarity. But we did not choose the method to make the aggregate poverty count similar for the reference year.

"Methodological revisions" and "erroneous estimates"?

Reddy and Pogge criticize the Bank for *"methodological poverty revisions"*. They give a number of examples (Tables 2 and 3), drawing on the Bank's published estimates at different dates.

It can hardly be surprising that the numbers will change as a result of new data, even for the same country and year. This can arise from changes in the underlying estimate of the PPP exchange rate, revisions to the Consumer Price Indices at country level and changes in the processing of the underlying survey data (a more consistent consumption or income aggregate may have been formed, for example). For example, quite a few of the *"pure methodological revisions"* they cite in their Table 2 between the poverty counts using 1985 PPP and 1993 PPP are for the Former Soviet Union (FSU). For 1985 there was only one PPP rate for the FSU, while with the new 1993 price data from the International Comparisons Project it was possible to estimate separate PPP rates for all countries within the FSU. So naturally we revised the estimates for all countries within the FSU. Reddy and Pogge chastise us for making such changes. A knowledgeable external consumer of these numbers would surely be far more inclined to criticize us if we had <u>not</u> made these revisions. The fact that Reddy and Pogge can see all these changes speaks for itself about our openness in making the necessary revisions in the light of new data.

Reddy and Pogge also confuse *"methodological revisions"* with real effects when they also compare our estimates for the same country at different dates (Table 3). They acknowledge the possibility that these changes are real, but assert that this *"seems unlikely"* (p.7) though they give us no basis for their judgment. Against their interpretation, the substantial increase in the measured poverty rate in Indonesia (for example) between 1996 and 1999, which they identify as a likely methodological revision, is more plausibly attributable to the severe macroeconomc crisis Indonesia faced in 1998, compounded by a poor agricultural year. There are other examples of the same confusion of real effects with revisions in the light of better data.

Reddy and Pogge assert that our methods systematically overstate the rate of poverty reduction for another reason, namely the method we use to line surveys up in time. Naturally, different countries do their surveys at different times. As I have noted, we currently have an average of about three surveys per country, but these are at different dates, and for about 20 countries out of 90 there is only one usable survey. So there has to be some interpolation method to get an estimate for 1998, say. Again, our methods are well documented. If there is only one survey for a country, then we estimate measures for each reference year by applying the growth rate in real private consumption per person from the national accounts to the survey mean — assuming that the distribution does not change. When the reference date is between two surveys, we interpolate from each survey to the reference date and take a weighted mean (Chen and Ravallion, 2001).

Let us now take a closer look at why Reddy and Pogge think we have overestimated the rate of poverty reduction. Though it is not entirely clear from their paper, one reason is that they appear to think that inequality is increasing within countries, thus leading us to overestimate the rate of poverty reduction by the above method. Yet, as we have established in other work and re-established in the latest up-date to our global poverty numbers, inequality within developing countries is falling about as often as it is increasing (Ravallion and Chen, 1997; Ravallion, 2001). And this is true during spells of growth too; indeed, the sample data for growing economies are almost exactly split between inequality-increasing cases and inequality-decreasing cases (Ravallion, 2001). Furthermore, even if Reddy and Pogge were right that inequality tends to increase as poor countries grow, note that for all except the 20 countries with only one survey, they would again be wrong since we are doing

the interpolation in both directions. This much could have been readily verified from a careful reading of the technical documentation they cite (notably Chen and Ravallion, 2001).

However, it is important to note that in our published regional and global aggregates we have re-calculated all numbers back in time in the light of improved survey data, revised price indices and new PPPs. Reddy and Pogge ignore the fact that in all updates of the Bank's global and regional aggregates all the numbers have been revised back in time on a consistent basis. So at whatever line one chooses, "$1" or $2" per day, the aggregate comparisons are consistent over time.

Another argument they make is that the PPP for food is *"..a more appropriate PPP concept"* for poverty measurement, and that this gives a higher poverty count. The food PPP is only available for about one third of the countries in the global poverty monitoring data set, so the scope for comparison is rather limited. Moreover, Reddy and Pogge provide no argument, and it is far from obvious, that putting zero weight on nonfood goods would give you a better PPP than that based on all consumption, even recognizing that the latter PPP is anchored to the mean consumption bundles. I am not surprised that using a PPP that ignores about half of consumption (probably about 40% for the poor on average) gives different poverty counts (though they do not present any evidence to suggest that it would give different trends). But this is hardly a convincing basis saying that the estimates based on consumption PPPs are *"erroneous"* as they claim so forthrightly.

There are some points where I agree with Reddy and Pogge, though they understate the difficulty of doing a better job. One example is the need to have better data on the prices paid by the poor; as Reddy and Pogge note, there is some evidence that the poor pay higher prices (though it does not follow that we underestimate global poverty since there are also indications that the poor pay more in rich countries too; the PPPs may not then change much). Another example is the calculation of believable standard errors for these estimates is a much more difficult problem than Reddy and Pogge appear to realize. We do have an internal proposal for doing this, which we will be attempting to implement in due course.

In conclusion

Reddy and Pogge begin their paper as follows: *"How many poor people are there in the world? This simple question is surprisingly difficult to answer."* I would argue instead that there is nothing simple about the question, and nothing surprising about how difficult it is to answer it. Reddy and Pogge have oversimplified the problem of measuring poverty in the world, and exaggerated the supposed faults in the Bank's methods. That is not to deny that there are problems galore in the data and methods of measuring poverty — problems that the Bank has taken the lead in exposing and addressing. There is still much more to do. While I have serious doubts about their proposed alternative method, the attention that Reddy and Pogge have given this issue is welcome.

References

Chen, Shaohua, Gaurav Datt and Martin Ravallion, 1994, "Is Poverty Increasing in the Developing World?" *Review of Income and Wealth*, 40(4), 359-376.

Chen, Shaohua and Martin Ravallion, 2001, "How Did the World's Poor fare in the 1990s?" (with Shaohua Chen), *Review of Income and Wealth*, 47(3), 283-300.

Deaton, Angus, 2001, "Counting the World's Poor: Problems and Possible Solution," *World Bank Research Observer* 16, 125-148.

Ravallion, Martin, 2001, "Growth, Inequality and Poverty: Looking Beyond Averages," *World Development*, 29(11), 1803-1815.

Ravallion, Martin and Shaohua Chen, 1997, "What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty?," *World Bank Economic Review*, 11(2), 357-82.

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

Reddy, Sanjay G., and Thomas W. Pogge, 2002, "How *Not* to Count the Poor," (Version 3.0), mimeo, Barnard College, New York.

Summers, Robert and Alan Heston, 1991, "The Penn World Tables (Mark 5): An Extended Set of International Comparisons, 1950-1988," *Quarterly Journal of Economics,* 106, 327-368.

World Bank. 1990, *World Development Report: Poverty*, Oxford University Press for the World Bank.