

HOW NOT TO COUNT THE POOR? A REPLY TO REDDY AND POGGE

Martin Ravallion¹

A lay reader of the Reddy and Pogge chapter in this volume might be forgiven for suspecting that the World Bank's data producers and researchers are real scoundrels. Reddy and Pogge (2002, 2008; see www.socialanalysis.org) assert that there are “*deep flaws*” in the Bank's methods of measuring global poverty, and that the Bank has probably under-stated the extent of poverty in the world and over-stated global poverty reduction.² Furthermore, they contend that it is no accident that these flaws have gone unchecked for nearly 20 years; this reflects (they claim) the “*low priority that has hitherto been attached to the global problem of persistent severe poverty.*” Given that eliminating global poverty is the Bank's self-declared goal, surely only scoundrels could have imparted deep flaws into the institution's measures of progress in attaining that goal? (Yet others claim that we have systematically over-stated the extent of poverty in the world to keep ourselves employed fighting poverty.³ It seems that measuring global poverty is a treacherous business for us scoundrels!)

However, as this chapter argues, the Reddy and Pogge (RP) critique collapses under even moderate scrutiny. They do not provide anything approaching a sound basis for believing that there are “*deep flaws*” in the Bank's estimates of the extent of global poverty.

The World Bank's poverty measures

Some years ago a consensus emerged in the international development community on the idea of an international poverty line of around \$1 a day. This became the basis of the first of the U.N.'s Millennium Development Goals (MDG1), which calls for a halving of the 1990 “\$1 a

¹ The author is Director of the World Bank's research department. These are the views of the author, and need not reflect those of the World Bank or any affiliated organization. The author is grateful to Shaohua Chen for her comments on this paper.

² In this reply to RP I will refer to both their chapter in this volume and their (unpublished but widely seen) paper of 2002, under the same title, which goes into more detail on some points.

³ See Bhalla (2002). This raises different issues; see (for example) Bourguignon et al. (2008).

day” poverty rate by 2015.⁴ Reddy and Pogge (2008, p.2) declare that “\$1 a day” is a “*meaningless poverty line*,” which presumably implies that MDG1 is a meaningless goal in their view. RP call for a new approach to measuring global poverty. However, before following their advice, one should take a closer look at how exactly the “\$1 a day” line arose. This should convince anyone that the line has more meaning than RP are willing to acknowledge.

Conceptually, one can think of any poverty line as the monetary cost of a reference level of “welfare” deemed necessary to not be considered poor. As in any true cost of living index (or money-metric utility function), the reference is a matter of choice. When measuring poverty, prior information on the nutritional requirements for good health and normal activity levels is often used to guide that choice.⁵ However, it must be acknowledged that there is ample scope for different people to form different judgments on the key parameters in setting a poverty line, including the composition of the food bundle and the allowance made for non-food needs.

The (unsurprising) reality is that people at different levels of living tend to hold different views about what the reference level of economic welfare should be for defining “poverty.” The critical level of spending that a poor person would deem to be adequate in order to escape poverty is undoubtedly lower than the level that a rich person would deem adequate. (This has long been recognized in the literature on poverty measurement.)

The same point holds between countries, as well as within a given country. In a background paper produced for the 1990 *World Development Report* (WDR) (World Bank, 1990), Ravallion et al. (1991) studied how national poverty lines varied with the mean consumption of a country, when both were converted to a common currency at purchasing power parity (PPP) (meaning that the currency conversion rate is intended to assure a common purchasing power over commodities). Amongst poor countries, they have found that poverty lines tend to be low, and there is also only a modest income gradient across countries in their poverty lines—absolute consumption needs naturally dominate in a poor country. Nutritional requirements for good health and normal activity levels tend to be fairly similar between people in poor countries and rich countries. However, as living standards rise people tend to buy more expensive calories (more meat and higher quality, or more highly processed, foodgrains) and tend to have more varied diets. And prevailing notions change concerning what non-food needs

⁴ See <http://www.un.org/millenniumgoals/>.

⁵ For an overview of the main approaches to setting poverty lines see Ravallion (2008a).

should be met if one is to not be deemed “poor.” Poverty is a socially-specific concept. Thus, above a critical level of mean consumption, the national poverty line tends to rise sharply with mean consumption (Ravallion et al., 1991).

This issue has recently been revisited by Ravallion et al. (2008), who have compiled an entirely new data base of national poverty lines across 75 developing countries. (In another chapter of this volume, Chen and Ravallion, 2008, summarize the results from this new compilation of national poverty lines.) Their results indicate that the pattern found in the original Ravallion et al. (1991) is quite robust: national poverty line rise with mean consumption, though with a low elasticity at low consumption. Both the food and non-food components of the national poverty lines rise with mean consumption. As one would expect, the “income” elasticity tends to be higher for the non-food component of the national poverty line, there is still a significantly positive elasticity for the food component (Ravallion et al., 2008).

In this light, the key question is: *By whose definition of what “poverty” means should we judge its extent 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 welfare, as measured by real consumption, the same way. By treating absolutely poor people similarly to relatively poor people one would risk diverting the focus away from what is surely the highest priority: to raise the living standards of the poorest in the world. The resulting measures would lose meaning as measures of absolute poverty. Relative poverty lines can still be defended if one believes that relative deprivation matters to a person’s welfare.⁶ For comparison purposes, the Bank has also produced poverty measures that take this approach (Chen and Ravallion, 2001, 2004). However, in the bulk of its efforts at global poverty monitoring, the Bank has taken the position that to measure absolute poverty on a consistent basis across countries one should use a poverty line with the same real value.

The “\$1 a day” line aims to judge poverty in the world as a whole by the standards of what poverty means in poor countries. The latest available estimates indicate that about one billion people live below this line, representing about one fifth of the population of the developing world (Chen and Ravallion, 2007). This is an explicitly conservative definition; one could hardly argue that the people in the world who are poor by the standards typical of poor

⁶ For further discussion of the theory and evidence on this point see Ravallion (2008b).

countries are not in fact poor. The point is that one cannot reasonably argue that there is less poverty in the world as a whole than is indicated by this calculation. Chen and Ravallion (2001) also argue that the \$1 a day line is a defensible lower bound to relative poverty lines.

We do not claim that the “\$1 a day” line is the only line that one would want to use for international comparisons. Indeed, in Chen et al (1994) we provided estimates spanning a fairly wide range. We regularly publish estimates for a line set at twice the level found in the poorest countries. The “\$2/day” poverty count is published alongside the “\$1/day” count in the Bank’s *World Development Indicators* for all years in which the numbers have been published. (Pritchett, 2006, has proposed an “upper bound” poverty line of around \$10 a day; about 95% of the developing world’s live below this line.)

As one would expect, there are measurement errors and idiosyncratic differences between countries in how poverty lines are constructed, which can be interpreted as noise in the mapping from the underlying welfare space into the income space. So it would not make sense to pick the lowest poverty line in the world; in fact that is well below \$1 a day. Some averaging is called for, as is normally the case in economic measurement. 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. Since then we have taken an average of the lines for the poorest countries, and provided tests of sensitivity to alternative methods of forming that average.

Having set an international poverty line, we then convert it back to local currency using the same PPPs. The best available consumer price indices (CPIs) are then used to convert the international line in local currency to prices prevailing at the time of the surveys. Next, these poverty lines are applied to distributions of consumption per person (or income if consumption is not available) constructed from nationally representative household surveys—we currently use well over 500 surveys for over 100 countries. 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). The latest estimates are found in Chen and Ravallion (2007).⁷

⁷ We have also created a web site, *PovcalNet* that allows users to replicate our calculations and try alternative assumptions; see <http://econ.worldbank.org/povcalnet>.

As an aside, the vast bulk of the Bank's analytic and operational work on poverty does not use the "\$/day" line, and with good reason. When one works on poverty in a given country, or region, one naturally tries to use a definition of poverty considered 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. In its annual tabulation of the "\$/day" poverty numbers, the Bank's *World Development Indicators* (for example, see World Bank, 2007) 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*.

Purchasing power parity and poverty measurement

PPPs are derived from the country-level price surveys that have been done since 1968 by the International Comparison Program (ICP). Estimating global poverty is only one of many applications of these PPPs in economic research. (For example, they have also been widely used in the vast literature on measuring and explaining differences in real incomes across countries.)

Prior to 2000, the Penn World Tables (PWT) were our main source of the PPP rates derived from the ICP; see, for example, Summers and Heston (1991). In 2000 we switched to the 1993 PPPs estimated by the Bank's Development Data Group. New PPPs have recently become available for 2005, again based on ICP price surveys (World Bank, 2008). There were numerous data improvements over time and various methodological differences between these sets of PPPs.⁸

RP clearly do not like any of these PPPs. Their main concern seems to be that 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*" (Reddy and Pogge, 2002, p.10). They go on to argue in their 2002 paper that "*...the only way to avoid this problem is to start from a particular reference basket of commodities and to construct*

⁸ The PWT used the Geary-Khamis (GK) method, while the Bank used the EKS method, which is the multilateral extension of the bilateral Fisher index. On the differences between the GK and EKS methods and implications for global poverty measures see Ackland et al. (2006).

PPPs that accurately reflect the relative costs of purchasing this 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.⁹ 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 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 a 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. 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 and hence that consumers can substitute among goods to achieve the same level of welfare—moving along their indifference 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) under certain assumptions.¹⁰

While it would certainly not be progress to follow RP’s recommendations, it can be agreed that it would be better to have PPPs designed for poverty measurement, weighted to the consumption bundle of people near the poverty line, using an appropriate iterative estimation method.¹¹ This was argued by Ravallion et al. (1991) in the first paper estimating the “\$1 a day” global poverty measures, although there was little or no progress toward that goal until the latest round of the ICP. An effort is underway at the Bank to estimate “PPP for the poor,” by re-weighting the 2005 ICP prices to accord more closely with consumption patterns of poor people. Preliminary results reported in Deaton and Dupriez (2007) do not suggest that the re-weighting needed to derive a PPP for the poor will make an appreciable change to the aggregate

⁹ I say “appear” here because they seem to back away from this position in their 2008 paper, possibly in recognition of the problems readers of their 2002 paper pointed out to them.

¹⁰ Notably that the utility function is quadratic. For a recent discussion see Ackland et al. (2006).

¹¹ Given that one needs to set a poverty line to determine the relevant consumption bundle, but only then can one determine the poverty line; an iterative method for this problem is proposed by Ravallion (1998), in the context of setting national poverty lines.

consumption PPP. However, further work will be needed before we can be confident about the implications for global poverty measurement.

Does the Bank underestimate the extent of global poverty?

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 world poverty. Obviously if you use a higher standard you will get a higher poverty count. The “\$1/day” line does not claim to be anything other 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. If one knows how the ruler is calibrated there should be no confusion.

RP question whether the national poverty lines are reliable as a basis for setting an international poverty line. There appear to be two concerns. First, they claim that the national poverty line may be “*influenced by political and other considerations*” (RP, 2008, p.15). They appear to be implying that this would lead us to underestimate the extent of global poverty, although this is evidently little more than a casual conjecture on their part, and they give no reason to expect a bias. Our new data set on 75 national poverty lines was formed from the World Bank’s *Poverty Assessments* and the governments’ own *Poverty Reduction Strategy Papers*. In every case, these poverty lines are constructed by professional teams, often comprising both staff of the governmental statistics office and economists and/or statisticians working for the Bank. There are (as already noted) idiosyncratic differences in how poverty lines are constructed, which is why we take an average of the poverty lines found amongst low-income countries. Errors certainly cannot be ruled out, but there is no obvious reason to expect systematic bias one way or the other.

Second, RP (2008) claim another fault in our method, stemming from the fact that different low-income countries have different (national) poverty lines. They present results for a truncated sample of the domestic poverty lines for “the poorest 14 countries,” which appears to be the 14 countries with lowest mean consumption at PPP in the original sample of 33 countries used by Ravallion et al. (1991). RP show that there is a variance in poverty lines found amongst

these countries. However, this has never been at issue.¹² Again, some averaging is called for, as we have always argued. RP do not present anything that would lead one to question that the mean poverty line (conditional on consumption) rises with consumption. Note that if RP had not chosen to truncate the Ravallion et al. (1991) sample—and no reason is given for this odd truncation—their readers would no doubt have seen the same strong positive relationship between the poverty line and mean consumption reported in Ravallion et al. (2001), and in Ravallion et al. (2008), for the new sample of 75 national poverty lines.

While the aims of the Bank's global poverty measurement effort have not changed over time, there have been many improvements in the underlying data. 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 World Bank (1990) and Ravallion et al. (1991) to well over 500 surveys for 100 countries now. Indeed, the Bank has put substantial effort into expanding the data base on household living standards.

The PPPs have changed too, with new and better price data. The 1993 PPPs were an improvement over those for 1985 in terms of country coverage, although there were 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 is to go back to the original poverty lines for the WDR 1990, and recalculate those lines with the new set of PPPs, and re-estimate the relationship between national poverty lines and mean consumption which led to the original \$1/day lines used in the WDR 1990. The same basic approach has been proposed by Ravallion et al. (2008), based on the updated data set of national poverty lines referred to above, using the PPPs derived from the 2005 ICP price surveys. (The Chen-Ravallion chapter in this volume provides a summary.)

Why not just update the international poverty line for inflation in the US? This would be a valid method if the purchasing power parity principle held (whereby the PPP for a given country evolves over time according to differences in that country's rate of inflation and that for the US) and one deemed 1993 PPPs to be beyond question; indeed, under these conditions one would not have needed to do the 2005 ICP. However, neither condition holds. Ravallion et al.

¹² They also show that the variance is higher, and the relationship with mean consumption is steeper, if one uses food PPPs. Given that the underlying national poverty lines were based on both food and non-food needs, it would seem more appropriate to use the full consumption PPP.

(2008) show that the joint implications of the purchasing power parity principle and comparability of the 1993 and 2005 ICP data can be convincingly rejected statistically.

Thus the naïve approach of simply adjusting the old line upwards for inflation in the US ignores some key features of how PPPs have evolved over time, including problems of data comparability over time. 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 that is well above the median of the 10 lowest poverty lines at 1993 PPP, and so we could no longer claim to be using a poverty line that is typical of poor countries. That would certainly entail a re-calibration of the ruler. The same point applies to the switch from the 1993 to 2005 PPPs, as shown in Ravallion et al. (2008).

In the light of these observations, one should not accept the claim by RP, echoed by Wade (2004), that we have devalued the poverty line over time, and hence over-estimated the extent of the world's progress against poverty.

We have long recognized these problems in switching PPP base years and data sources. The latest version of RP's critique is more cognizant of the problems than their earlier paper. However, they still do not properly acknowledge that our practice has always been to revise all of our estimates back in time (currently back to 1980) when new PPPs become available. The PPP currency conversion is only done at the base data; then the comparisons over time for a given country depend on the best available CPI for that country. The country-level CPIs are not always ideal, but they are the best data we have for making such comparisons over time.

The key point is that, in assessing progress against poverty in the world, we do not need to make comparisons across different (non-comparable) sets of PPPs. So the entire discussion of this issue in RP (2002, 2008) is irrelevant.

As an aside, it was noted in Chen and Ravallion (2001) that the \$1.08 line using the 1993 PPP gave a very similar global poverty count to the old \$1/day line at 1985 PPP for the common reference year in the series, namely 1993. RP assert that it was a "*serious error of reasoning*" (RP, 2008, p.3) on our part to have made this check for whether the poverty counts matched for the same year. This is surely over-stated. It is natural to look at how new data and methods

affect one's final estimate of (ostensibly) the same thing, and to draw some comfort from their similarity. RP are right that they could have come out at very different poverty counts; that is obvious enough since (as explained above), we did not choose the new poverty line to make the aggregate poverty count similar for any year.

“Methodological revisions,” “erroneous estimates” and “false precision”?

As in virtually all aspects of socio-economic data, there is still scope for improving the data underlying the global poverty measures, namely the survey-based distributional data and the price data (both CPIs and PPPs). However, data are improving in the developing world, thanks to the efforts of international agencies such as the Bank, as well as the governmental statistical offices in developing countries. There is no doubt that both data sources have improved enormously in terms of coverage and quality in the time since our estimates of those poverty measures began, around 1990.

As data improve, it is not too surprising that our knowledge gets revised as well. In the light of better data, we have always revised our global poverty estimates accordingly, including back in time (as have other data sources, such as the national accounts).

These revisions to our past estimates provide fuel for many of RP's criticisms, though I expect most people would agree that to not revise knowledge in the light of new data would be far worse. In particular, RP criticize the Bank for “*methodological poverty revisions*” (RP, 2002, p.7). They give a number of examples (Tables 2 and 3 of RP, 2002), drawing on the Bank's published estimates at different dates.

It can hardly be surprising that the numbers 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 CPIs 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 RP, 2002, 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. RP 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 not made these revisions. The fact that RP can see all these changes speaks for itself about our openness in making the necessary revisions in the light of new data.

RP also confuse “*methodological revisions*” with real effects when they also compare our estimates for the same country at different dates (see Table 3 of RP, 2002). They acknowledge the possibility that these changes are real, but assert that this “*seems unlikely*” (p.7) though they give us no justification for this judgment. Against their interpretation, the substantial increase in the measured poverty rate in Indonesia (for example) between 1996 and 1999, which RP identify as a “methodological revision,” is more plausibly attributable to the severe macroeconomic crisis Indonesia faced in 1998, compounded by a poor agricultural year (Ravallion and Lokshin, 2007). While it may “seem unlikely” to RP that such a crisis could have substantially increased poverty, it is very clear from the evidence that it did do so. There are other examples of the same confusion of real effects with revisions in the light of better data.

RP (2002) assert that our methods systematically overstate the rate of poverty reduction for yet another reason, namely the method we use to line surveys up in time. In the latest available estimates at the time of writing, Chen and Ravallion (2007) used over 500 surveys spanning 100 countries. But these surveys do not (of course) line up neatly in time across different countries, so an interpolation method is needed to obtain an aggregate estimate for any given reference year in the aggregate time series of regional or global estimates. Again, our methods are well documented. The reference years chosen lie comfortably within the range of the data. 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. However, for 94 of the 100 countries we had two or more surveys. 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, 2004).

Let us now take a closer look at why RP think we have overestimated the rate of poverty reduction. Though it is not entirely clear from their papers, 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, 2007). 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. Furthermore, even if RP were right that inequality tends to increase as poor countries grow, note that for all except the six countries with only one survey, they would again be wrong since we interpolate in both directions and then take the average. This much could have been readily verified from the documentation they cite (notably Chen and Ravallion, 2001, 2004).

However, it is again 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. RP 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; they also assert that this gives a higher poverty count. However, RP 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 gives different poverty counts for a fixed poverty line (though they do not present any evidence to suggest that it would give different trends). But this is hardly a convincing basis for saying that the estimates based on consumption PPPs are “*erroneous*” as they claim. RP's calculations are also deceptive given that they ignore the fact that switching to a food PPP would also change the poverty line; Ravallion et al. (2008) show that using the food PPPs implies an appreciably lower poverty line.

RP also accuse us of what they call “*false precision*” (RP, 2008, p.9) in the poverty estimates reported in the various technical papers by Chen and Ravallion, documenting their methods for the Bank's global poverty estimates. “False precision” refers to the fact that the estimates of the global poverty measures (in millions of people and percentages of the population) are given to two decimal places in Chen and Ravallion (2001, 2007) (and one decimal place in Chen and Ravallion, 2004). RP believe that greater rounding off of the estimates would have better revealed their true precision. We choose to give as much accuracy as

we could reasonably fit into our published tables, leaving it to the data users to do further rounding off. Maybe RP would be happier to round off the poverty counts to the nearest billion people before publishing them; then the count would have been unchanged (at one billion) between 1981 and 2004, even though the number of poor fell by 500.80 million (Chen and Ravallion, 2007). Readers can judge for themselves the merit of RP's claims of "false precision."

How would Reddy and Pogge measure global poverty?

In their conclusions, Reddy and Pogge claim that there is a better way of counting the world's poor. They are rather vague about what this would entail, but refer to a paper by Reddy et al. (2006). This paper measures poverty in three countries (Nicaragua, Tanzania and Vietnam) using a method that will be recognized by specialists on poverty measurement as a version of what is termed the "cost of basic needs" (CBN) method (Ravallion, 1994, 2008a).¹³ By this method one calculates the cost of a bundle of goods deemed to be nutritionally adequate and conforming to local tastes—to give the food-poverty line—and one adds to this an allowance for non-food spending consistent with the spending patterns of those near the food-poverty line. The Reddy et al. (2006) version of this method is that used by the government of Vietnam (following recommendations from the World Bank). They then repeat the method for the other two countries, using the same caloric cut-off point for all three countries (2100 calories per person per day) but different (country-specific) food bundles and different allowances for non-food spending (anchored to the spending behavior in each country for the quintile at which 2100 calories is reached on average).

What is being proposed here is essentially the method most developing countries use to set their own, national, poverty lines; indeed, virtually all of the 75 countries in our new compilation of national poverty lines have used some version of the same CBN method (Ravallion et al., 2008). There are of course differences; caloric cut-off points vary somewhat, as do valuation methods and the allowances made for non-food needs. However, to a first order

¹³ Reddy et al. (2006) call their method a "capability-based approach." However, the relationship to Sen's (1985) "capability approach" is unclear. The fact of being able to afford a diet that yields (say) 2100 calories per day does not assure that the functionings that come with being adequately nourished will actually be met, even on average, let alone for each individual. A true "capability approach" would presumably look rather different to what Reddy et al. (2006) propose. For further discussion see Ravallion (2008a).

approximation, one expects that the poverty lines generated by Reddy et al. (2006) are more like national poverty lines. In that light, the fact that (as Reddy et al. show) the resulting poverty measures differ from those obtained using an international (\$1 or \$2 a day) line is hardly surprising. As noted already, the purchasing power over commodities of the national poverty lines is demonstrably not constant, as best that can be measured. So two people with the same absolute standard of living in terms of their command over commodities will be treated differently, depending on where they live. Typically the person living in the poorer country will be less likely to be deemed poor.

RP may object that the Reddy et al (2006) measures can still be considered “absolute” because they have used the same caloric cut-off point for all three countries. However, this response would ignore an important lesson from the literature on nutrition and poverty (and from common sense), namely that a given food energy intake can be attained in multiple ways, requiring very different levels of income. As already noted, there is a strong income effect on both the food and non-food components of national poverty lines (though stronger for non-food component); for example, one obvious reason why poorer countries tend to have lower poverty lines is that they consume cheaper calories.

So RP have not solved the problem of setting an international poverty line with constant purchasing power over commodities, but rather they have sidestepped that problem. Arguably, RP have not taken the discussion of how best to set an international poverty any further than its starting point in the 1990 WDR (World Bank, 1990).

Conclusion

Reddy and Pogge begin their chapter in this volume as follows: “*How many poor people are there in the world? This simple question is surprisingly difficult to answer at present*” (RP, 2008, p.1). 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, have greatly exaggerated the supposed faults in the Bank’s methods, and their proposed alternative method does not take us very far in the goal of setting an international poverty line.

Postscript in response to Thomas Pogge's rejoinder

I thought I was just a scoundrel in their eyes, but it seems I am “Nixonesque” to boot. (Although given that Pogge is “not questioning the integrity of the Bank’s researchers” I am left wondering in what sense he sees me as Nixonesque.) Much of what I say above can be repeated in response to Pogge’s rejoinder.¹⁴ I will avoid doing so. But there are a couple of points that do require clarification, so readers are not misled by Pogge’s rejoinder.

One such issue is Pogge’s assertion that the national poverty lines for developing countries (that we use in identifying an international poverty line) are unreliable in that they do not reflect “a level of income or consumption sufficient to meet basic human needs.” The national poverty lines we use are all founded on reasonably well-specified “basic needs,” and in this respect they are no more “questionable” than (say) the official poverty line of the US. Indeed, they are typically set essentially the same way that Reddy and Pogge advocate that an international line should be set!

Of course, as I pointed out in my reply to Reddy and Pogge, what constitutes a “basic human need” is socially specific—it depends on the standards of living in a specific setting. And national poverty lines in poor countries are clearly not “survival points” below which nobody can live (for if they were there would not be so many people living below them!). Rather they reflect the prevailing notions of what “poverty” means in those societies. I have no objection when Pogge pleads for a higher line, reflecting what he thinks constitutes “basic human needs.” The “\$1 a day” line has not claimed to be anything more than a defensible lower bound, by which poverty in the world as a whole is judged by what are seen to be basic human needs in the poorest countries in the world.

Pogge goes on to repeat some textbookish points about the properties of PPPs. The original “\$1 a day” paper in 1990 fully acknowledged that standard PPPs (designed for

¹⁴ In a footnote, Pogge tells us that Sanjay Reddy chose not to co-author the rejoinder because two conditions were (apparently) made by the editors: that I had the last word on the exchange, and that Reddy and Pogge did not have access to another chapter I wrote (with Shaohua Chen) in this volume, which Pogge claims I “repeatedly cite” above. I doubt if this postscript will be the last word, but at least both sides had two turns in addressing the issues in debate, which seems fair. I don’t speak for the Editors, but I would have given Reddy or Pogge my other chapter in this volume if they had asked. Indeed, the working paper version of that chapter has been downloadable on the Bank’s external research web site since May 19, 2008—almost three months before I received Pogge’s rejoinder. So too has the main unpublished paper referred to above—which is not the one in this volume (cited twice) but another paper, “Dollar a Day Revisited” (cited seven times).

comparing national accounts across countries) were not strictly appropriate for measuring poverty. It took 12 years to get to a new round of the ICP that could do a better job in this respect. The analysis is still underway, but the preliminary results suggest that the “PPP’s for the poor” are quite similar to the standard PPPs, and that the global poverty estimates are reasonably robust to this change. I have no disagreement in principle with Pogge on the importance of weighting PPPs and CPIs appropriately to the task of measuring poverty.

However, near the end of his rejoinder, Pogge comes to what he sees as “the most compelling evidence one can have that the method (used by the Bank) is no good,” namely that one gets different poverty counts if one changes the base year for the PPPs. Yes, new rounds of the ICP have generated new PPPs that imply changes to our poverty estimates. New ICP rounds bring new and better data on the cost-of-living in developing countries. Consider the latest ICP round for 2005—almost certainly the biggest global statistical exercise in history, involving numerous international agencies and the government statistics offices of almost 150 countries. This is clearly far superior to past rounds in terms of country participation, survey data collection and processing methods and the quality and comparability of the price data collected. Nobody who bothers to look into the history of the ICP—from 1970 (crude price surveys for 10 countries) to 2005 (state of the art price surveys for 150 countries)—could contend otherwise. It cannot be too surprising that the 2005 ICP has changed our estimates of real volumes for all international economic comparisons, including poverty measures. (My chapter with Shaohua Chen in this volume shows the impacts for China.) There has also been considerable progress in improving the coverage and quality of the household survey data required for measuring poverty. This too has entailed some significant changes to our poverty measures.

The key point here is that the data have improved and that this has led to new and better estimates. That hardly means that “the method is no good.”

References

- Ackland, Robert, Steve Dowrick and Benoit Freyens, (2006). "Measuring Global Poverty: Why PPP Methods Matter," mimeo, Research School of Social Sciences, Australian National University.
- Bhalla, Surjit, (2002), *Imagine There's No Country: Poverty, Inequality and Growth in the Era of Globalization*, Institute for International Economics, Washington DC.
- Bourguignon, François, Francisco Ferreira, Branko Milanovic and Martin Ravallion, (2008), "Global Inequality," in Kenneth Reinert and Ramkishan Rajan (eds) *Princeton Encyclopedia of the World Economy*, Princeton University Press, forthcoming.
- 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.
- _____ and _____, (2004), "How Have the World's Poorest Fared Since the Early 1980s?" *World Bank Research Observer*, 19/2: 141-70.
- _____ and _____, (2007), "Absolute Poverty Measures for the Developing World, 1981-2004." *Proceedings of the National Academy of Sciences of the United States of America*, 104/43: 16757-62.
- _____ and _____, (2008), "China is Poorer than we Thought, but no Less Successful in the Fight Against Poverty," in Sudhir Anand, Paul Segal, and Joseph Stiglitz (ed), *Debates on the Measurement of Poverty*, Oxford University Press (forthcoming).
- Deaton, Angus and Olivier Dupriez, (2007), "Poverty PPPs for Latin America and Asia," mimeo, Development Data Group, World Bank.
- Pritchett, Lant, (2006), "Who is *Not* Poor? Dreaming of a World Truly Free of Poverty," *World Bank Research Observer* 21(1):1-23.
- Ravallion, Martin, (1994), *Poverty Comparisons*, Harwood Academic Press, Chur, Switzerland.
- _____, (1998), *Poverty Lines in Theory and Practice*. Living Standards Measurement Study Paper 133. Washington DC: World Bank.

- _____, (2007), “Inequality is Bad for the Poor,” in J. Micklewright and S. Jenkins (eds.), *Inequality and Poverty Re-Examined*. Oxford: Oxford University Press.
- _____, (2008)a, “Poverty Lines,” in *The New Palgrave Dictionary of Economics*, 2nd edition, Larry Blume and Steven Durlauf (eds) London: Palgrave Macmillan.
- _____, (2008)b, “On the Welfarist Rationale for Relative Poverty Lines,” in Kaushik Basu and Ravi Kanbur (eds) *Arguments for a Better World: Essays in Honor of Amartya Sen. Volume II: Society, Institutions and Development*, Oxford: Oxford University Press.
- 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, Shaohua Chen and Prem Sangraula, (2008), “Dollar a Day Revisited,” Policy Research Working Paper 4620, World Bank (<http://econ.worldbank.org/docsearch>).
- 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.
- Ravallion, Martin and Michael Lokshin, (2006), “On the Consistency of Poverty Lines,” in *Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke*, edited by Alain de Janvry and Ravi Kanbur, Springer.
- _____ and _____, (2007), “Lasting Impacts of Indonesia’s Financial Crisis,” *Economic Development and Cultural Change*, 56(1): 27-56.
- Reddy, Sanjay G., and Thomas W. Pogge, (2002), “How *Not* to Count the Poor,” (Version 3.0), mimeo, Barnard College, New York.
- _____ and _____, (2008), “How *Not* to Count the Poor,” in Sudhir Anand, Paul Segal, and Joseph Stiglitz (ed), *Debates on the Measurement of Poverty*, Oxford University Press (forthcoming).
- Reddy, Sanjay G., Sujata Visaria and Muhammad Asali, (2006), “Inter-Country Comparisons of Income Poverty Based on a Capability Approach,” Department of Economics, Barnard College. (Available at <http://ssrn.com/abstract=915406>.)

- Sen, Amartya, 1985. *Commodities and Capabilities*. Amsterdam: North-Holland.
- 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.
- Wade, Robert, (2004), "Is Globalization Reducing Poverty and Inequality?" *World Development* 32(4): 567-589.
- World Bank. (1990), *World Development Report: Poverty*, Oxford University Press for the World Bank.
- _____, (2000), *World Development Report: Attacking Poverty*, Oxford University Press for the World Bank.
- _____, 2(007), *World Development Indicators*, Washington DC: World Bank.
- _____, (2008), *2005 International Comparison Program: Tables of Final Results*, Washington DC: World Bank.