How *Not* to Count the Poor! — A Reply to Ravallion

August 15, 2002

Sanjay G. Reddy[1] and Thomas W. Pogge[2]

We are grateful for Martin Ravallion's detailed and sincere attempt to respond to our criticisms of the World Bank's methodology for measuring global income poverty.[3] However, in this reply to Ravallion, we show that his attempt to defend the World Bank's methodology is unconvincing. It seriously misstates a number of our criticisms, responds inadequately to others, and fails altogether to mention still more. Although we would like to be able to accept the global estimates of severe income poverty the Bank has produced, we continue to find that they do not stand up to serious scrutiny, and regretfully hold to our recommendation that they cease to be used.

We respond to Ravallion's comments paragraph by paragraph below, referring to them according to their order of appearance, and italicizing quotations from Ravallion for ease of recognition.

Paragraph 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 is an inadequate description of our view. First, we believe that there is evidence to suggest that the Bank's estimates have misrepresented the extent of poverty in the world in *all* years. Second, there is one unambiguous sense in which evidence suggests that the Bank has underestimated the extent of world income poverty. This is that the use of more realistic purchasing power parity conversion factors, more closely related to the real costs of living of the poor, would have led to higher national poverty lines in most countries, even if the Bank's '$1/day' international poverty lines were retained. We argue however, that the '$1/day' concept should be abandoned because of the deep and irremediable problems attached to it. The estimates that would result if a more *meaningful* concept of poverty were used are entirely *unknown*, and we have never suggested otherwise. Third, we argue that the picture presented by the Bank of the *change* in world poverty over time is not *necessarily* incorrect, but is certainly unjustified. The Bank's methods are so flawed that it is impossible

[1] sr793@columbia.edu

[2] tp6@columbia.edu

[3] We are responding to the version of Ravallion's comments sent to us on May 10, 2002 and posted on www.socialanalysis.org . We wrote to Ravallion on June 27, 2002 to confirm that this was the final version of his comments, but we received no response.

to judge from them what the true trend has been. There are two distinct problems. First, because the Bank's estimates of the extent of global income poverty at *any* point in time are not credible, its methodology offers no evidential basis for judgements regarding the trend of poverty. Second, the Bank's method for "updating" purchasing power parity conversion factors (used to convert a constant international poverty line into national currencies) over time generates tendencies for poverty lines to fall in an unjustified manner, falsely suggesting a decrease in the extent of poverty in instances in which there may have been no actual decrease. These tendencies result from the nature of the methods used to calculate PPPs for general consumption, which are extensively discussed in our paper. This is why we say that the Bank's methodology has given rise to an *incorrect inference* that global income poverty has declined, without taking a view on whether the conclusion inferred is *actually* true or false.

Paragraph 2: "*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*". Ravallion advances this claim in order to suggest that there is a high degree of consensus among poor countries as to the level of the poverty line, so that the choice of the international poverty line in the vicinity of this common level is unproblematic. This is not true, however, and begs the central questions raised in our paper. Whether the domestic poverty lines of poor countries are clustered together or not depends on the type of PPP conversion factor used to translate them in to a common measure (US$ PPP). We have shown that when PPPs based on consumption items more closely related to the needs of the poor are used, this apparent clustering breaks down. We found that a 1 percent increase in consumption per capita is associated with a 0.64 percent increase in the official domestic poverty line, when the cost of bread and cereals is used to determine PPPs (or a 0.43 percent increase when the cost of food in general is used). Although this elasticity is less than that which exists in rich countries, it is hardly negligible. Scrutiny of the list of official poverty lines used by Ravallion and Chen reveals a wide divergence between the lowest and the highest poverty lines, even when PPPs based on general consumption are used.

Paragraph 3: This paragraph is rather irrelevant to the argument of our paper. It concerns the role of 'relative' poverty lines meant to capture the idea that more resources are required for an individual to avoid poverty as others in the society become richer. This is an important and interesting issue, but not one which we address or with which we are concerned. Following the Bank's own reports, we treat the $1/day measure as referring to 'absolute' poverty, which is defined in terms of a set of core deprivations (such as the presence or absence of hunger) that are independent of such relative considerations.

Paragraph 4: "*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*". We are in complete agreement with this evidently sensible view. "*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*". We do not accept this description of the "$1/day" line, for two distinct reasons. The first is that, as described in our reply to paragraph 2 above, the international poverty line does not correspond consistently or closely to the official domestic poverty lines of poor countries. Whether a cluster of official domestic poverty lines around a common level appears depends on the year in which the comparison is made and on the PPP conversion factors used. The use of PPP conversion factors more relevant to poverty assessment seems to increase the divergence. Moreover, some poor countries' official domestic poverty lines are far below the level chosen and others are far above. Second, the international poverty line determined 'inductively' in the way Ravallion describes does not correspond to any single underlying conception of poverty, defined in terms of human needs or elementary capabilities. It is clear that it cannot correspond to "what poverty means" since it means nothing in particular. There is no evidence that the domestic poverty lines of poor countries are in fact closely grouped together in the sense of representing the same real purchasing power over goods and services.

Paragraph 5: "*This poverty line is then converted to local currency using the latest available PPP exchange rates for consumption*". Ravallion refers here (as he has elsewhere) to his determination to convert the international poverty line into national currencies using the purchasing power parity conversion factors calculated from the most recent data available, in order to attempt to maintain the salience of the international poverty line in terms of national currency units. Although this proposal periodically to "update" the purchasing power parity factors used is well-intentioned, it leads to a host of inadvertent and deep difficulties. As we have discussed in our paper, such 'updating' will result in national poverty lines in poor countries falling, even though nothing may have changed in those countries, *simply* because of changes in the structure of the world economy. People will appear to be less poor even though nothing in their situation has changed. If the Bank is indeed committed to such 'updating' as a feature of its methodology, then our reason for concern about the possibility of mistaken inferences regarding the trend in global income poverty is deepened rather than lessened.

Paragraph 6: We are in full agreement with the sentiment of this paragraph, which is that most work concerning poverty reduction at the national level need not make any reference to an international poverty line. However, the comparative assessment of national poverty levels and poverty reduction strategies does require a common standard of poverty. Moreover, the Bank does produce figures concerning the global incidence, trend, and geographical distribution of severe income poverty. We assume that the Bank produces these figures because it is important to do so. This is also why we have voiced our concerns about their quality.

Paragraph 7: "*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*". *All* of the criticisms we make of the '$1/day' poverty line apply equally to the '$2/day' poverty line and any other arbitrarily defined international poverty line converted via general-consumption PPPs. We do not emphasize the problems with the $2/day poverty line because our primary concern is with assessments of the level of *absolute* poverty (a concept that the Bank has linked to the $1/day poverty line in the past), although we could easily have done so.

Paragraph 8: No need for comment.

Paragraph 9: "*Reddy and Pogge … appear to be proposing to price a single bundle of goods in each country relative to a reference country*". This is a misunderstanding of our view, as we outline further in our response to paragraph 10.

Paragraph 10: " *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/day concept"*. As we made very clear in our paper, our suggestion is to specify the 'fixed' bundle of commodities at an adequately abstract level, so as to accommodate relevant differences in tastes, requirements, and substitution opportunities, while maintaining the uniformity of the underlying poverty concept. In this sense, the 'fixed' bundle of commodities would not be *wholly* fixed. For example, the food component of a global standard for poverty avoidance could be defined in terms of the calorie and micro-nutrient content of foods, without specifying these foods exactly. This would permit wheat to be consumed in wheat eating regions, and rice in rice eating regions, so long as the calorie standard was met. This approach, isolating the *characteristics* of commodities rather than commodities themselves, has been widely used in poverty assessments at the national level, especially in large federal countries in which such variations are significant.[4] The specification of the characteristics of these commodities should in turn be motivated by a fundamental concern for the particular elementary capabilities that people must possess if they are to count as non-poor. The global standard, once specified, should be given specific regional and local content in a transparent, participatory and rigorous way. We see no reason why this approach should be more controversial than the decidedly controversial $1/day approach. As we have argued, the seeming uniformity of this approach masks hidden variations in the meaning of the $1/day "standard". The fact that people undertake substitution of goods to make adjustments for relative prices is well known. Whether and how to take account of such variations is the subject of controversy even in rich countries [see e.g. Jorgenson and Slesnick (1999)] which otherwise pursue an approach of the kind we advocate. The difficulty of taking account of such substitution is not therefore particularly embarassing to the method we suggest. In any case, by specifying the characteristics of commodities rather than commodities themselves, we fully allow in principle for such substitution to be taken into account.

Paragraph 11: "*In fact, the deficiencies of the idea of using a single basket of goods has 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.)"* This statement is both untrue and irrelevant. The classical concept of PPP, first associated with Gustav Cassel, exactly concerns the costs in different countries of purchasing a single basket of goods. Modern PPP calculation methods take a more agnostic approach concerning the basket to be priced, and therefore compute a kind of average of relative prices for different kinds of goods. This agnosticism may be justified for tasks such as determining the relative size of different countries' GDP, and possibly even for the comparative assessment of living standards in general, but it *cannot* be appropriate for the measurement of absolute poverty. It *is* possible to give specific meaning to a conception of *absolute* poverty, although there will necessarily be some controversy over the exact content of this conception. It is a great irony that for the *one* task for which the specification of a particular basket (allowing for a sensible degree of abstraction in this specification) is quite appropriate, Ravallion insists that it is not appropriate. There is certainly room for controversy about whether some particular commodity should be included among the requirements for avoiding poverty and about what relative weights the included commodities

[4] India, for instance, is one of many countries in which the poverty line has been defined in relation to a calorific standard.

should be assigned. However, there can be little doubt about the most significant needs of the poor (especially for food, which typically accounts for 70 percent of poor households' expenditures) and no doubt that the vast majority of commodity expenditures by the non-poor are **ir**relevant to the poverty estimation exercise. The controversies that do exist should be acknowledged and responded to, both by producing a *band* of estimates of global poverty corresponding to distinct assumptions (rather than point estimates corresponding to a single assumption) and by identifying the core conception of poverty that is ultimately used in a transparent and globally consultative way.

"*Ideally, the underlying price index would only reflect differences in the cost of a reference level of welfare, fixed across all countries.*" We do not share this particular ideal. As discussed by Sen (1992) and others, poverty (and in particular absolute poverty) is best thought of in terms of deprivation of specific elementary capabilities (or even basic needs), rather than in utilitarian 'welfarist' terms. This is the conception that we too adopt. There are many reasons that subjective 'welfare' is the inappropriate space in which to identify the poor. An obvious reason is that many unhappy people are not unhappy because they are poor, and many poor people are not necessarily unhappy.

"*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)*". The PPPs used by the Bank are in fact derived using the EKS method, which is quite different from the Fisher index, although it incorporates it into its formula. The EKS method is a response to a more complex problem than that set by Fisher, and the PPPs which result from it cannot be interpreted in the same way. In any event, there is no particular reason to think that 'consumer preferences' are quadratic. We suspect that this too has been known since the 1920s. Finally, all of the arguments made are rather beside the point, since absolute income poverty assessment should properly proceed from a non-welfarist normative criterion (concerning, for instance, capabilities or basic needs), whereas all of the statements made concern the ability of particular measures to capture a constant level of *utility* (under specific assumptions).

Paragraph 12: "*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…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)*". There are two misunderstandings here. First, as noted above, our criticisms are of the use of general consumption PPPs derived from the EKS approach in absolute poverty assessment, and not of the use of the EKS approach for other purposes, nor of the Fisher index (which is merely a component of the EKS approach) as such. Second, Ravallion has seriously misunderstood the purpose of our formal analysis. Our intent was to examine the nature of the EKS index by examining various *counterfactuals*. Specifically, we wished to examine the impact of a change in the pattern of demand in rich countries on poverty lines in poor countries, *holding all other factors constant*. We showed that a changing pattern of demand in rich countries could lower poverty lines in poor countries, though nothing had changed in the latter. This entry of irrelevant alternatives into the construction of purchasing power parities used for poverty assessment is highly inappropriate. For the purpose of demonstrating this fact, our differentiation exercise was entirely appropriate, and concern with equilibrium demand effects (based on a hypothetical relation between prices and quantities demanded) would have been quite inappropriate.

Paragraph 13: *"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 the 25th percentile from the bottom of the median….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."* Defining the consumption bundle of the poor as that which is consumed by the 25th percentile from the bottom of the median of the income distribution is arbitrary, and would not permit of any consistent interpretation. Moreover, individuals at the 25th percentile of a country's income distribution might be extremely poor in absolute terms in some countries and rather well off in others, undermining the uniformity of assessment that must underpin a meaningful exercise of attempting to estimate absolute poverty in the world. Finally, the informational requirements (collection of data on the consumption data of this particular group, and of the prices of the goods they consume) are more intensive than for the alternative we propose. Prof. Deaton's proposal has the virtue of providing a consistent basis for assessment, but suffers from its lack of connection to a clear and ongoing underlying conception of what poverty means, defined in terms of basic needs or capabilities.

Paragraph 14: *"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?".* The $1/day line is not what Ravallion claims it to be, namely typical of poor countries, as we have discussed in our response to paragraph 2 above. To be meaningful, a concept of poverty must be associated with a clear standard that can be interpreted in terms of human needs, elementary capabilities, or related notions. The methodology employed by the Bank is *not* meaningful in this sense. Moreover, as we have argued, it also fails to retain whatever meaning it might have in terms of dollars when translated in to other currencies. In this sense, the comparison between the Bank's method and the alternative is *not* like that between inches and centimetres, since the Bank's method, properly speaking, measures nothing at all. Finally, as we have expressed above, it is unclear ex ante what estimates an alternative global income poverty estimation exercise would generate, and we have not taken a view in this regard.

As we (unlike the Bank) have not so far had the resources with which to carry out a full-fledged alternative global income poverty estimate, we conducted two other analyses. In the first we asked what would be the impact of taking the existing Bank poverty line of $1/day (PPP), and translating it into national currencies by using food-based PPPs rather than general consumption PPPs. We found that national poverty lines in poor countries would invariably be higher than suggested by the Bank. In the second exercise we 'endogenized' the international poverty line by using food-based PPPs and Ravallion's own procedure for calculating the international poverty line to create, from data about the official domestic poverty lines used by poor countries, a new international poverty line and subsequently to translate that international poverty line in to national currencies. Here too we found that national poverty lines were higher than suggested by the Bank. It is only in the sense of these two *specific* exercises — consistent with the letter and the spirit of Ravallion's own favored approach — that we say that poverty has been underestimated.

Paragraph 15: This paragraph simply notes the improvements in the Bank's household survey database over the 1990s. We are, of course, pleased about this development.

Paragraph 16: *"However, the two sets of PPPs [1985 and 1993] 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[s] 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 lack of comparability of PPPs from different years is precisely the reason that there is no clear sense in which the $1993 poverty line contains equivalent purchasing power to the $1985 poverty line, as claimed by Ravallion elsewhere. Moreover, there is another defensible alternative to the procedure adopted by the Bank, as we have outlined at length in this response and in our original paper. This is to specify a poverty line in terms of the requirements of fulfilling basic capabilities and to hold to this criterion consistently over time.

"The 1990 WDR $1/day 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." As we have discussed in our paper, the cluster of poverty lines that appears when 1985 PPP conversion factors are used to put official poverty lines on a comparable basis breaks down when 1993 PPP conversion factors are used. Both of the methods used to calculate the international poverty line are arbitrary. Neither has been adequately motivated, and they are mutually inconsistent.

Paragraph 17: *"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 … This was not done by design, but was judged a desirable feature."* As we discuss in our paper, there is no reason at all that this should be thought a desirable feature of a methodology for poverty assessment, as it is achievable by any methodology whatsoever! The idea that this similarity is a 'desirable' feature continues to reflect a serious conceptual error.

Paragraph 18: No need for comment.

Paragraph 19: This paragraph notes that many of the changes in poverty estimates that we attribute in the paper to methodological revisions may in fact be due to the availability of improved data. We accept that this is the case. However, we continue to be struck by the magnitude of the changes (see Table 2 of our paper). It seems reasonable to conclude that the shift in the Bank's poverty measurement methodology and improvements to data were jointly responsible for these large changes. The data remain interesting and indicative for this purpose, though certainly not conclusive.

Paragraph 20: This paragraph notes the possibility that sharp changes in national circumstances in a short period of time may have been responsible for some of the fluctuations in recorded poverty estimates. We agree that this is wholly possible, but hold to the view that such drastic changes do not explain all of the cases in which substantial changes were reported. For many of these, the methodological shift is likely to be the most important factor. Only a detailed country by country analysis can reveal the respective roles of these factors.

Paragraph 21: This paragraph describes the method used by the Bank to 'line surveys up in time'. It is purely factual, and corresponds to Ravallion's documentation of this method in other publications.

Paragraph 22: This paragraph notes that inequality does not always have a tendency to increase (the circumstance under which we argued that the Bank's method would automatically lead to the appearance of falling poverty). We are in full agreement with the empirical point that is being made here, but believe that our criticism has been misunderstood. We argued as follows. If global poverty is being estimated for a particular year (say 2002), the Bank's method involves assuming that the distribution of consumption in a country continues to be as last observed in a household survey (which may have been in, say, 1996). It is assumed, in other words, that in all social strata the change in consumption expenditure between 1996 and 2002 has exactly mirrored the change in mean consumption expenditure reported in the country's national income accounts. We pointed out that this assumption will necessarily lead to the appearance of falling poverty in countries where mean consumption expenditure has risen — regardless of whether consumption expenditure of the poorest strata has increased or declined. The point we were making pertains to the implication of this method for estimated poverty in all countries for which the last survey was before the year for which the estimate of global income poverty was produced. We were especially concerned with the latest estimates of poverty as they specially affect apparent trends and we therefore neglected to discuss cases in which the estimate of global income poverty is for years between survey years. We were fully aware, however, of the interpolation methodology used by the Bank in these cases. Although we agree that the global poverty decrease could potentially be underestimated rather than overestimated (if income distribution has improved in many countries over the period from the last survey date) this cannot be known without further data. The main point that we would like to draw out here is that there are very large errors potentially associated with the Bank's method, which must be duly acknowledged as such, and which until now have not been.

Paragraph 23: This paragraph claims that whenever the Bank has adopted a new estimation methodology, it has produced estimates running backward and forward in time, and that these estimates are therefore consistent over time. This statement is certainly true in a limited sense. In particular, the figures most recently reported by the Bank for global poverty between 1987 and the present, which consistently use the $1.08 PPP 1993 poverty line, adopt a consistent (if arbitrary and meaningless) standard. However, these estimates are not comparable with those produced using the $1 PPP 1985 poverty line for years prior to 1987. We do not disagree with the point being made here, but consider it to be of limited relevance in view of the fundamental infirmities of the Bank's methodology.

Paragraph 24: *"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…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"*. Our view here has been severely misrepresented. We have not advocated the use of food based PPPs as an alternative to general consumption PPPs. Rather, we have simply pointed out that the resulting poverty numbers change substantially when food-based PPPs rather than general consumption PPPs are used as an indicator or proxy for the costs faced by the poor in relation to their basic needs. We used this evidence to argue for the need for a fuller study, and for poverty estimates based on a comprehensive basket of commodities that are relevant for furthering elementary capabilities. We have, recently, extended our research to non-food PPPs closely related to basic capabilities (clothing, shelter, and medical care) derived from the benchmark surveys of the International Comparison Programme. We have found that for the commodities most closely related to the needs of the poor, just as with food, prices are much higher in poor countries than is suggested by general consumption PPPs. We would be pleased to supply our analysis on request. Moreover, the data on which it is based is freely available in the public domain.

Paragraph 25: *"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…Another example is the calculation of believable standard errors for these estimates"*. As we have discussed in our paper, the collection of price data relevant to poverty assessment could easily be undertaken under the existing auspices of the International Comparison Programme (ICP). This has not been done so far primarily because it is not a priority. The ICP already collects data on many other kinds of prices. Moreover, private consulting firms collect price data on luxury consumption around the world in order to determine cost of living adjustments for private sector executives. There is no reason why something like this cannot be done for the consumption of the poor, especially since the World Bank is now the organization that hosts the ICP. Further, we have urged the Bank to be more forthright about the massive errors potentially associated with its poverty assessments. We have never suggested that the Bank should '*calculate* believable standard errors' (italics added). The unmerited pseudo-scientific precision suggested by this phrase is exactly what we seek to avoid.

Paragraph 26: *"Reddy and Pogge have oversimplified the problem of measuring poverty in the world, and exaggerated the supposed faults in the Bank's methods"*. We leave it for others to judge the merit of this claim. Our concern is first and foremost to further progress toward the development of credible estimates of world poverty. It is of enormous importance that the world's institutions devote greater resources and attention to this cause than has so far been the case.

REFERENCES

Jorgenson, D. and D.Slesnick (1999). "Indexing Government Programs for Changes in the Cost of Living." *Journal of Business and Economic Statistics*, Vol. 17 No. 2.

Reddy, S. and T. Pogge (2002), "How Not to Count the Poor", unpublished manuscript, Columbia University. Available on www.socialanalysis.org .

Sen, A. (1992). *Inequality Reexamined*. Oxford: Oxford University Press.