

\$1.90 A DAY: WHAT DOES IT SAY?

The New International Poverty Line

THE WORLD BANK has good news. According to its latest calculations, global poverty has been falling. In 1999, an estimated 1,751,000,000 people suffered from extreme poverty. By 2011, the number had dropped to 983,000,000. World Bank President Jim Yong Kim tweeted excitedly that the rate of extreme poverty may drop below 10 per cent of the world population when the figures are in for 2015. ‘The international community can celebrate’, exclaimed a recent Bank report. Despite the global financial crisis, it had chalked up some ‘robust’ development successes. The UN’s Millennium Goals had played an important role in galvanizing efforts to reduce poverty, and that experience would help drive progress towards the Sustainable Development Goal of eradicating it altogether.¹ But what is poverty? How do you count the poor?

Since 1980 at least, the World Bank’s answer has been a concept known as the International Poverty Line.² In 1990 the Bank’s researchers, led by an LSE-trained Australian, Martin Ravallion, estimated this at \$1.02, the famous ‘dollar a day’, at purchasing power parity—that is, what the equivalent of a dollar could buy at local prices. Anyone falling below that line was counted as ‘extremely poor’. Ravallion and his colleagues derived the figure of \$1.02 as a ‘representative’ marker for absolute poverty: of the 33 poor countries they were studying, this was the approximate sum already set as a ‘national poverty line’—often by the World Bank itself—for Indonesia, the Philippines, Bangladesh, Pakistan, Kenya, Tanzania and Morocco. To be precise: this was the figure at which the Bank economists arrived after re-scaling these domestic-currency sums—denominated in Indonesian rupiahs, Filipino pesos,

Bangladeshi takas, and so forth—according to local inflation levels, to calculate their ‘equivalents’ for the year 1985; and then, for purposes of international comparison, converting those sums into a common unit of ‘purchasing power’: the 1985 PPP dollar. The purchasing-power values themselves were calculated from data gathered in its 1985 survey by the International Comparison Programme—itself now hosted by the World Bank in Washington, DC—which checks prices of goods in different countries. Each subsequent release of new global PPP data by the ICP, usually some years after the date—1985, 1993, 2005, 2011—to which the prices refer, has led to a corresponding reassessment of the International Poverty Line by the Bank’s economists.³ With the release of the 2011 PPP figures in 2014, the World Bank economists have raised the IPL to \$1.90, using methods justified at length in a new report by Bank economist Francisco Ferreira and his colleagues.⁴

The figures on global poverty extrapolated from the Bank’s poverty lines enjoy enormous international legitimacy. Canonized in the Bank’s World Development Reports, they have been used to determine priorities for resource allocation and to assess the relative success of poverty-reduction programmes. They formed the reference of the UN’s first Millennium Development Goal, to reduce poverty by half between 1990 and 2015—and were the measure of how well it had succeeded. In September 2015, world leaders at the UN’s headquarters in New York pledged to eradicate extreme poverty altogether, one of a set of Sustainable Development Goals adopted to guide their future policies—and once again based on the Bank’s assessment. Since poverty reduction is claimed to be the

¹ See Marcio Cruz, James Foster, Bryce Quillin and Philip Schellekens, ‘Ending Extreme Poverty and Sharing Prosperity: Progress and Policies’, World Bank Group Policy Research Note 15/03, October 2015, pp. 6, 2–3.

² The first attempts to project a global poverty line from national data began in the 1970s, using India’s 45th percentile of per capita income distribution as a baseline. See M. S. Ahluwalia, N. G. Carter and H. B. Chenery, ‘Growth and Poverty in Developing Countries’, *Journal of Development Economics*, vol. 6, no. 3, 1979. This analysis was the basis of the poverty analysis in the Bank’s 1980 World Development Report.

³ Thus the 1993 PPPs were used by Ravallion and Shaohua Chen to estimate an International Poverty Line of \$1.08 in 2001; similarly, the 2005 PPPs were used by Ravallion, Chen and Prem Sangraula to calculate an IPL of \$1.25 in 2009.

⁴ Francisco Ferreira, Shaohua Chen et al., ‘A Global Count of the Extreme Poor in 2012: Data Issues, Methodology and Initial Results’, World Bank Policy Research Working Paper 7432, October 2015.

Bank's central aim, the data it provides are also a key metric for assessing its own performance. Gratifyingly, the Bank has consistently been able to claim that its figures prove it is 'on the right track', even if 'much remains to be done'.

Yet the Bank's money-metric approach is open to a host of objections, even on its own terms. The criticisms that have been levelled at the Bank's methodology since the 1990s over key technical questions—PPPs, inflation measures, price variations within countries (in particular the differing costs in rural and urban areas) and the merits of income versus consumption data—apply in spades to its latest iteration, as we will show below. More broadly, it might be asked whether an approach focused only on *extreme* poverty, or 'absolute deprivation', as the Bank terms it, is potentially self-serving. Poverty lines may be set so that hundreds of millions are only just below them, with a change of only a few cents sufficient for large numbers to be 'lifted out of poverty', without any substantive change in their position. The Bank assumes *a priori* that there is no poverty in high-income countries, even though it acknowledges this may not be 'fully supported' by the facts.⁵ Indeed, as discussed below, alternative data show the assumption to be false, especially at more adequate poverty lines.

Further problems arise from inadequate data. For entire regions, including most saliently the Middle East and North Africa, the Bank presently reports no results at all, because of poor survey coverage. In a number of other cases, it appears to have based country estimates on figures deriving from other countries, to account for missing data. Judgements of this type are an unavoidable aspect of applied work in data-poor environments, but while one can sympathize with the necessity to make such calls—and indeed, to be forthright in defending them—the resulting uncertainties must be adequately recognized. There have also been serious institutional problems with the International Comparison Programme, on whose PPP estimates the Bank's entire poverty count depends. Originating in 1968 as a project of the UN Statistical Division, after its 1993 report the ICP was the subject of a highly critical UN review, which identified failures of management and resources at all levels; hence the long delay before the appearance of its 2005 figures. Although the ICP's latest price collection exercise is its most comprehensive yet,

⁵ Ferreira et al., 'A Global Count', p. 28.

there remain fundamental questions as to how to collect information and undertake comparisons, and also as to what is being measured.

At the root of most of the seemingly unrelated problems with the Bank's global poverty estimates, however, there lies a single conceptual failure: the lack of a criterion for identifying the poor that is appropriate to the task and has a consistent, substantive interpretation. It is a problem that cannot be solved within the existing approach, but requires an altogether new one. In what follows, we first examine problematic issues in the Bank's most recent poverty estimates, focusing in particular on purchasing-power parity, inflation calculations and rural–urban variations. We then demonstrate that following the Bank's own stated approach could lead to an alternate—and much higher—set of poverty estimates, which we report. However, we do not present these as a last word, but rather as a demonstration of the seriousness of the uncertainties involved—and of the need for an altogether new framework.

Updating relative to what?

The Bank's economists justify their decision to set the latest International Poverty Line at \$1.90 a day, in 2011 PPP dollars, with the argument that this will 'preserve the real purchasing power of the previous line'—that is, \$1.25 a day, in 2005 PPP dollars—'in the world's poorest countries.'⁶ But in what sense do the lines in fact correspond? The question applies equally to the Bank's previous 'updates', replacing IPLs specified in the base years of 1985 with 1993, and 1993 with 2005. In looking for an answer, we might begin by examining the Bank's own view of the matter. Its customary argument as to why the new poverty line is equal to the previous one in purchasing power is that the poverty headcount ratio is very similar in both cases. However, as with the Bank's earlier estimates, such an argument is little more than a *non sequitur*.⁷ We may think of the problem this way: suppose that an arbitrary set of new PPP dollars were chosen to translate the IPL into local currencies—or, indeed, an especially perverse set, deliberately chosen, for example, to misrepresent the real level of purchasing power in each country. By starting at a low

⁶ Ferreira et al., 'A Global Count', p. 39. A second poverty line of \$3.10 in 2011 PPPs has also been identified, as an equivalent to \$2 in 2005 PPPs, used for middle-income countries: p. 4.

⁷ See Sanjay Reddy, 'The World Bank's New Poverty Estimates: Digging Deeper into a Hole', *Challenge*, vol. 51, no. 6, 2008.

enough value and then creeping up, one could always find an IPL that would suffice to generate exactly the same headcount as did the previous PPPs. Since this argument can be used to ‘rationalize’ any set of PPPs, it cannot justify any one choice thereof. What it resoundingly does *not* show is that the new PPP dollars maintain the purchasing power of the old ones, anywhere—let alone everywhere. In any case, it is evident that this argument cannot be used to justify the original choice of the IPL.

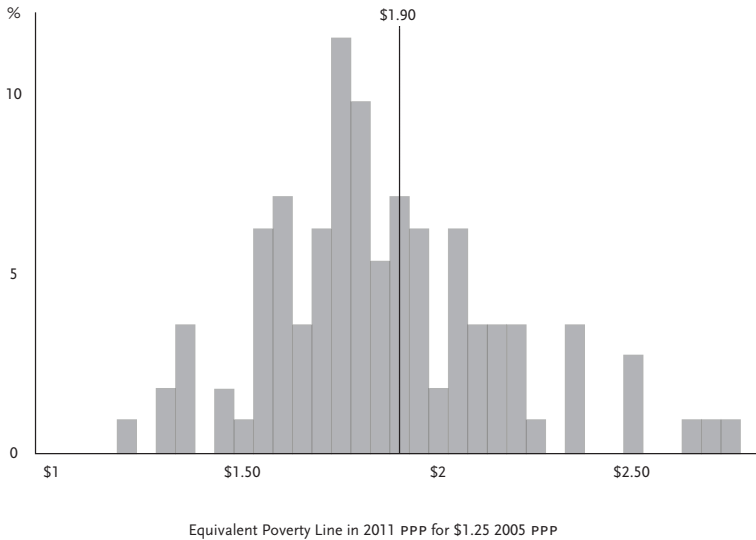
A second possibility is to ask whether the purchasing power of the 2011 IPL corresponds to that of the 2005 IPL. Unfortunately, it does not. This is because when the 2005 IPL is translated into local currencies, and then updated using the consumer price indices (CPIs) of the individual countries, this leads to amounts of local currency in 2011 values which are generally very different from those implied by converting any given IPL directly using the 2011 PPPs.⁸ This is a deep-seated problem that is intrinsic to the way in which PPPs are calculated: purchasing-power ratios are designed for spatial, not spatio-temporal, comparisons. The evidential basis of the spatial price indices used by PPPs is determined by the structure of the world economy in the year in which they are calculated; but the reference point of a temporal price index (CPI) is the pattern of consumption in a given country. The resulting diversity of reference points leads to sizable inconsistencies. The Bank’s own data shows that there is no way of choosing an IPL that will maintain its purchasing power within all countries, even remotely, using its current methods.

As a check, we can calculate the ‘equivalent poverty line’ in 2011 local currency units for any given country by updating the 2005 IPL using the country’s own CPI. This shows that of the 117 developing countries, 70 countries would have had EPLs below \$1.90 in 2011 PPP, among them China, Afghanistan, Iran, Turkey, Ukraine, Brazil and much of Sub-Saharan Africa and Latin America.⁹ Indeed, nearly half the world’s

⁸ Ferreira and his colleagues write that ‘we updated the line so as to keep its real value constant, in terms of the purchasing power of the poorest countries. Since the real poverty line has not changed much in real terms, overall poverty levels (for a given year) don’t change much either’. See *World Bank ‘Let’s Talk Development’ blog*, 4 October 2015. However, the poverty lines have indeed shifted in every country, in real terms, according to the countries’ own consumer price indices.

⁹ Countries with EPLs over \$1.90 include India, Pakistan, Bangladesh, Indonesia and Nigeria. The highest EPL on this calculation is Yemen’s, at \$2.76; the lowest is Belarus’s at \$1.29. For all figures, see the working paper version of this article on the Global Consumption and Income Project website, www.gcip.info.

FIGURE I: *Equivalent Poverty Line in 2011 PPP for \$1.25 2005 PPP for developing countries*



Source: GCIP, WDI. Data is for 117 developing countries, excluding Sudan, Turkmenistan, El Salvador and Tajikistan, as their equivalent poverty lines are outliers (>10 or <1.2).

population (49 per cent) lives in countries with EPLs below this threshold. Figure 1 indicates the distribution of these EPLs.¹⁰ This is a sort of median, but not necessarily a happy one, insofar as the new chosen IPL is ‘wrong everywhere’, if to an extent that varies in sign and magnitude. The Bank calculates its IPL by using its preferred CPIs to update the IPLs of the fifteen countries used in its last updating exercise, themselves chosen rather arbitrarily from a longer list. But as other scholars have shown, the results depend, to a not inconsiderable degree, on the consumer price indices used for these countries.¹¹ Although, as the Bank’s

¹⁰ This graph presents in a different way the information contained in Figure 4 of Ferreira et al. ‘A Global Count’.

¹¹ Stephan Klasen et al., ‘International Income Poverty Measurement: Which Way Now?’, Courant Research Centre Discussion Papers, No. 184, Göttingen 2015, p. 15 (current draft). The authors demonstrate, for example, that Tajikistan’s national poverty line varies between \$3.18 and \$1.72, depending on the method of calculating price inflation that is used.

economists note, there is more than one way of arriving at the \$1.90 IPL, this hardly makes it unassailable. Indeed, there can be more than one way of attempting to justify any line at all.

There is a basic conceptual issue here. One can seek to maintain purchasing power, in which case one should, within the Bank's money-metric approach, fix the base line and use national CPIs for updating, as argued by Angus Deaton over a decade ago.¹² Alternatively, one can 'update' the prices used for spatial comparison. But one cannot do both. The Bank, as in the proverb, is precisely trying to have its cake and eat it.


Moreover, the proportion of the population deemed poor greatly depends on the specific choice of IPL. According to our own estimates, based on the Global Consumption and Income Project,¹³ choosing an IPL of \$2.50 in 2011 PPP would raise poverty across the world by 38 per cent, compared to choosing one of \$1.90, increasing the headcount ratio from 21 to 29 per cent. This would not just alter the global poverty level but also affect the regional composition of poverty, with South Asia contributing a significantly higher proportion of the world's poor (see Figure 2).

A third possible notion of 'equivalent' purchasing power involves the idea that successive IPLs refer to the 'same' substantive meaning in terms of basic human requirements. The Bank has made exactly this argument in relation to both its current 'update' and to previous ones, in all cases referring to a small set of poverty lines through which the IPL was constructed—the fifteen lowest, chosen from a much larger set by establishing a rather arbitrary cut-off point—held constant between the last two IPL-setting exercises, and ostensibly reflecting standards of identification of the poor in poorer countries themselves. (In fact, many of these national poverty lines were produced by World Bank consultants, giving the impression of 'Bank preconceptions in, Bank preconceptions out'.) However, the selection of poverty lines, the means used to convert them into common units, the method of identifying one poverty line by averaging or otherwise aggregating information from the rest of the set used, and even the claim that they have a meaningful

¹² Angus Deaton, 'How to Monitor Poverty for the Millennium Development Goals', *Journal of Human Development*, vol. 4, no. 3, 2003.

¹³ For further information, see the working-paper version of this article on gcip.info.

FIGURE 2: *World and regional headcount ratios for various 2011 PPP poverty lines for 2010*



Source: GCIP.

reference in terms of basic human requirements are all highly questionable. Moreover, the arbitrary and shifting methods undermine the claim of consistency over time.

During earlier IPL setting exercises, the Bank had also drawn frequently on the auxiliary argument that the poverty lines supposedly set by the poorest countries themselves were very similar, in contrast to those of less poor countries, which were deemed to increase with income.¹⁴ This claim, depending in part on a visual trick, becomes even more strained when subsequent PPP base years are used, which is perhaps why it has not been relied upon when attempting to justify the latest revision.¹⁵ In any case, it would be difficult to maintain that the poverty lines in question have a common substantive meaning, in light of the demonstrably different standards and methods of construction used and the resulting variability.

Translation into local currencies

The lack of consistent and substantive meaning in the Bank's poverty criteria not only undermines attempts to 'update' the IPL, but also infects the translation of the poverty line into local currency units. There is, it bears repeating, no such thing as purchasing power in the abstract; rather, it must be defined in relation to a specific purpose, which in turn can be translated into an account of the specific commodities required to achieve that purpose. For example, if we consider purchasing power over tradable necessities, such as food, rather than over all goods and services, then the local currency equivalent of a given US dollar amount is found by the ICP to be around one-third higher.¹⁶ Similarly, PPPs are broad aggregates of price levels over goods and services that reflect, in practice, the influence of the overall pattern of consumption in the world in a given year. This leads to 'irrelevant commodities' and 'irrelevant countries' affecting the PPP of a given country in irrational ways. For

¹⁴ For one instance, see Martin Ravallion, Shaohua Chen and Prem Sangraula, 'Dollar a Day Revisited', *World Bank Economic Review*, 2009.

¹⁵ It depended on, among other things, using a log-scale for the visual appearance of a 'flat portion' of the relevant curve.

¹⁶ See Sanjay Reddy and Thomas Pogge, 'How Not to Count the Poor', in Sudhir Anand, Paul Segal and Joseph Stiglitz, eds, *Debates on the Measurement of Global Poverty*, Oxford 2010, p. 52. A first draft of this paper was made available in 2002.

example, under the current method, data on Japanese real-estate prices may impinge on whether a household in India is deemed to be living in extreme poverty or not.

This is the static analogue of the dynamic problem created by changes in base year, which raise or lower a country's PPP relative to its CPI change, to an extent that is different and hard to predict or interpret, from country to country. In a departure from its earlier practice, this central deficiency is acknowledged in the Bank's justification of the new IPL—the authors indeed admitting that the results of the updates 'can be jarring and cast into doubt the reliability of the global poverty estimates'—but it is not addressed. The Bank's new procedure of continuing to use the old 2005 IPL and PPPs for countries with very high or low discrepancies ('delta') in this regard is merely an *ad hoc* attempt to mitigate an intrinsic consequence of its own method.¹⁷

From this point of view, the notion that the latest set of PPPs is always the best one to use, as the Bank's public statements have suggested, deserves closer scrutiny.¹⁸ On the one hand, the most recent set of PPPs aims to reflect the latest global consumption patterns. On the other hand, and for the very same reason, it miscasts the patterns that prevailed in previous years. In a world economy subject to rapid and large-scale shifts in composition, this creates serious difficulties, especially when dealing with trends over long periods of time. Even within the conceptual framework of money-metric poverty assessment, it is far from obvious why 2011 offers a better base year for examining trends between 1980 and 2015, for instance, than does 1990. (The same arguments, of course, extend well beyond poverty assessment to a range of other economic analyses.) It is a different and additional matter that the coverage and quality of price surveys have arguably improved, although

¹⁷ Ferreira et al., 'A Global Count', pp. 3, 21.

¹⁸ See, for example, the *Financial Times* interview with Jim Yong Kim, who explains that the decision to adjust the poverty line was 'a necessary update due to new data on purchasing power': 'We don't think we moved the goalposts. We think we simply updated the goalposts to 2015.' Shawn Donnan, 'Earth's poor set to swell as World Bank moves poverty line', *Financial Times*, 23 September 2015. The misleading headline was based on a proposed IPL of \$1.92, which produced a 148,000,000 surge in the number of the extreme poor. As we have seen, the results were far more 'satisfactory' when the new IPL was set at \$1.90.

the methodological changes introduced with the 2011 ICP survey have attracted a degree of controversy.

From village to slum

A further problem in specifying purchasing power arises when prices differ sharply within a country, most typically between the cities and the countryside. The International Comparison Programme defines PPP conversion factors only at national level. For three particularly large and heterogeneous countries, China, India and Indonesia, the Bank has made an attempt to define sector-specific PPPs for rural and urban areas. Again, however, the approach is strikingly inconsistent. In China, which did not participate in earlier ICP surveys, the 2005 ICP data-collection exercise was limited to eleven cities and hence was considered to provide an urban PPP. For India, although ICP data was collected in both rural and urban areas, Bank economists suggested that the 2005 survey was more representative of city prices.¹⁹ In Indonesia, although Bank economists judged that there was no survey bias, they nevertheless made an adjustment to account for differences in price levels between rural and urban areas. Similar adjustments to purchasing power estimates for these three countries have been incorporated in the latest update based on the ICP's 2011 PPPs.

In principle, if such an approach is appropriate for China or India, it is not clear why it should not be used for other large countries, such as Nigeria, or for smaller but still cumulatively populous ones, such as Bangladesh or Pakistan. The more serious issue, however, is that these sector-specific PPPs have been constructed by the Bank using highly questionable, back-of-an-envelope assumptions. The idea of having a single price level for the whole of rural India is only a little less absurd than that of having a single price level for all India, or all China, or any other regionally variegated country. Moreover, the specific methods used to construct the rural price level were poorly justified and open to serious criticisms. In its rough-and-ready approach, the Bank assumed for India (and followed a parallel procedure for other countries), firstly, that the ratio of rural to urban prices, and thus of sectoral PPPs, could be derived from the ratio of previously defined rural and urban poverty lines; and

¹⁹ Martin Ravallion, 'A Global Perspective on Poverty in India', *Economic & Political Weekly*, 25 October 2008. This view had not been put forward by the Bank in regard to earlier rounds.

secondly, that the national price level was a weighted average of the (unknown) rural and urban price levels, with the number of price points sampled by the ICP in its national PPP-determination exercise defining the weights.²⁰ The first equation was used to specify the ratio of the sectoral PPPs, the second their absolute level. The resulting two-equation system was used to generate the Bank's estimates.

If the goal was to construct a robust, rural-versus-urban price index, a better approach might have been built by referring directly to prices paid by households as inferred from household surveys, combined with a hypothetical basket of goods and services. Alternatively, the Bank might have applied the price data used to generate official domestic price indices for different categories of workers, which are available in India for both agricultural and industrial workers. The official Indian rural and urban poverty lines used in the Bank's first equation have also come in for severe criticism.²¹ Insofar as they may be relied upon, they represent the presumed costs of obtaining basic human requirements, in the village or the city. The variations between them thus reflect differences *both* in prices and in assumptions about the commodities that must be purchased in the two sectors to meet human requirements. The weights used in determining the relationship between national and sectoral price levels are only for the specific categories of goods—food, clothing, footwear—for which price information is collected in rural areas by the ICP, whereas the national PPP reflects prices of *all* goods. It is therefore far from obvious that the exercise leads to a correct estimate of the price levels in each sector.

The Bank employed questionable methods in determining Indian, Chinese and Indonesian rural and urban poverty lines in its latest IPL construction exercise. For India, it chose to adopt the highly controversial Tendulkar Committee poverty lines, which were not taken up officially. Indeed, the Government of India felt obliged to appoint the Rangarajan Committee to re-examine the problem, due to the perception that it had been inadequately addressed.²² This apparently arbitrary choice would seem to have required some justification. If we take the ratio of the

²⁰ Ravallion, 'A Global Perspective on Poverty in India', p. 35 and footnote 9.

²¹ S. Subramanian, 'The Poverty Line: Getting It Wrong Again', *Economic & Political Weekly*, 26 November 2011.

²² See for example Yogima Seth Sharma, 'India has 100 million more poor: C. Rangarajan Committee', *Economic Times*, 7 July 2014.

urban–rural poverty lines as a measure of price variations alone, and not of differences in requirements as implicitly supposed by the Bank, then rural prices must be judged to have risen by almost 30 per cent more than urban prices over a two-year period, as the ratio of urban to rural poverty lines declined from 1.51 to 1.22. In fact, according to sectoral price indices reported for India by the Bank’s own Povcalnet website, the two sets of prices rose at a comparable rate: 76 per cent in rural areas as against 70 per cent in towns. The conclusion to be drawn is that the ratio of rural to urban poverty lines is unlikely to be a very good guide to relative price levels. In the recent literature, methods proposed for estimating rural and urban price levels produce widely varying differentials, which in turn give rise to widely varying poverty estimates.²³ For India, the Bank estimates a 51 per cent difference between rural and urban price levels, with the rural poverty headcount accordingly low, whereas the ICP puts the difference at only 3 per cent; other sources report estimates in between. Here it will suffice to say that it is not only paying attention to sectoral variation that matters, but also which method of inter-sectoral adjustment is used. The Bank’s chosen approach leads to the most optimistic portrayal of rural purchasing power, and thus of the rural headcount. A measure of the impact of using sectorally adjusted PPPs for these three important countries, China, India and Indonesia, is given below (Tables 1 and 2).

Estimates of global poverty levels are thus enormously affected by this single, very questionable choice. Using ICP national PPPs for China, India and Indonesia would substantially increase their poverty rates, raising the estimated number of poor people in the world on 2011 figures by an additional 290,000,000. The trend-rate of global poverty reduction from 1990 to 2011 also appears more favourable using the sectoral PPPs the Bank’s economists have chosen in the recent period. The Bank has offered no sensitivity analysis nor discussed the impact of this choice, leaving open the question of why it made the particular decisions that it did. Although taking note of intra-national specificities, including rural–urban differences, is in principle desirable, it is vital that this should be done in a manner that is both justifiable and consistent across countries. Tellingly, two Bank researchers had argued

²³ See for example Angus Deaton and Olivier Dupriez, ‘Spatial Price Differences Within Large Countries’, manuscript, Princeton University, 2011; and Yuri Dikhanov, ‘Income Effect and Urban–Rural Price Differentials from the Household Survey Perspective’, ICP Global Office, 2010.

TABLE 1: *Poverty headcount and ratio as defined by \$1.25 (2005 PPP) poverty line using sectoral poverty lines for 2011*

	1990		2011	
	% poor	Millions of poor	% poor	Millions of poor
<i>China</i>	60.4	686.0	6.4	86.0
<i>India</i>	49.3	454.0	24.1	294.5
<i>Indonesia</i>	54.3	97.0	16.2	39.5
<i>Total</i>		1,237.0		420.0

Source: World Bank.

TABLE 2: *Poverty headcount and ratio as defined by \$1.25 (2005 PPP) poverty line using national poverty lines for 2011*

	1990		2011	
	% poor	Millions of poor	% poor	Millions of poor
<i>China</i>	71.6	813.5	13.2	177.5
<i>India</i>	65.7	604.7	39.9	487.7
<i>Indonesia</i>	63.4	113.3	19.3	47.1
<i>Total</i>		1,531.5		712.3

Source: GCIP.

that it could be appropriate to drop the use of sectoral adjustments for China, India and Indonesia when using the 2011 PPPs, on the grounds that there was an urban bias in the 2005 ICP data-collection round that does not exist in the 2011 ICP data.²⁴ Again, this viewpoint pays no attention to the reality of systematic price differentials within large countries, but focuses merely on producing the ‘right’—by what criteria?—national PPP.

Income or consumption?

Further methodological problems arise from the fact that surveys in many countries, notably in Latin America, collect data on income, while the Bank’s poverty lines are defined in terms of consumption. In

²⁴ Dean Jolliffe and Espen Prydz, ‘Global Poverty Goals and Prices: How Purchasing Power Parity Matters’, World Bank Policy Research Working Paper 7256, 2015.

earlier rounds the Bank made an effort to convert income figures into consumption estimates, multiplying them by the ratio of consumption-to-income in the national accounts. In recent years, the Bank has switched to pooling both kinds of data and using them without further adjustment. It claims that the choice of method makes little difference, although it has admitted that the new procedures lower headcounts. Its argument is that two sets of distortions—those caused by the share of bottom quintiles being lower in income than in consumption surveys, and those caused by the means being lower in consumption than in income surveys—‘cancel out’. But the claim that the use of the pooled, unadjusted data produces a closer estimate of the underlying unknown values is implausible. The new method is inappropriate in two distinct ways, which cannot be assumed to negate each other. In the comparisons discussed below we find widely varying results, depending on whether estimated income surveys, estimated consumption surveys or the pooled approach is used.²⁵

To illustrate the impact of these methods on poverty headcount ratios, we calculate the latter for selected countries for which both types of data, income and consumption, are available from the same survey: Uganda, Angola, Nepal, Uzbekistan and Bolivia. Using a single poverty line of \$2.50 in 2005 PPP of consumption, we show the resulting percentage of the poor varying significantly between the current Bank method (income surveys included with no adjustment to distributions or to means), the old Bank method (income surveys included with adjustment to means but not to distributions) and the count derived from the actual consumption surveys (Table 3). The difference in poverty estimates across the three methods varies by country, as the magnitude of the dependence on the method used relies on various factors (in particular the income and consumption survey distributions and means). In some cases the difference is small and in others more sizable. In these cases the new Bank method leads to lower headcounts than the old one, but whether either method leads to lower or higher headcounts than

²⁵ In order to gauge the underlying and unknown true values, when we estimate consumption poverty on the basis of income surveys, we adjust both survey distribution and means to enhance comparability. This differs both from what the Bank does now and its previous practice, which was to adjust means alone. For details see Rahul Lahoti, Arjun Jayadev and Sanjay Reddy, ‘The Global Consumption and Income Project (GCIP): An Introduction and Preliminary Findings’, working paper, 2015.

TABLE 3: *Percentage of poor (defined by \$2.50 2005 PPP consumption poverty line) in select countries*

	Angola	Bolivia	Nepal	Uganda	Uzbekistan
Survey year	2009	2000	2010	2000	2001
Income survey, no adjustment	71.5	46.3	66.7	81.3	81.3
Income survey, means adjusted	76.9	48.9	71.6	83.5	89.5
Consumption survey	80.2	32.9	70.7	80.3	96.2

Source: World Bank, GCIP.

do consumption surveys—which ought to be used wherever available—depends on the country.

Finally, as noted above, for a number of countries consumer price indices either do not exist or are rejected by the Bank on the grounds that they are implausible. In these cases, the Bank's economists undertake *ad hoc* measures. Again, this choice is potentially consequential, as it includes countries with large numbers of poor people, such as Bangladesh, and may account for some of the discrepancy between the estimates of the Bank's economists and those of other scholars using its chosen IPL of \$1.90, in 2011 PPP. As demonstrated above, the choice of CPI matters for 'updating' the IPL. The risk is that this could create an impression of selective choice. Since the Bank has an entirely abstract conception of purchasing power, there is no guidepost as to what an appropriate consumer price index might be, beyond the fact that 'it looks right to us'. In sum, both the sectoral (rural–urban) adjustments and the choice of CPIs may have been of consequence in enabling the Bank to achieve, quite remarkably, a similar regional distribution using the 2011 PPPs to what it attained using the 2005 PPPs. This is not something the Bank was able to do with previous base-year changes—in particular the shifts from 1985 to 1993 PPPs, and from 1993 to 2005 PPPs, which led to sizeable changes in the regional composition of poverty.

Taking the Bank at its word?

In its latest update, then, the Bank has once again adopted an approach that locks in previous mistakes, making minor modifications while failing to address deeper criticisms. The upshot is a set of results which,

despite their seeming technical authority, should enjoy little credibility. As we outline below, our constructive proposal is not merely to modify the Bank's money-metric approach to global poverty estimation, but to reject it altogether. Credible assessments of poverty's level, trend and regional composition require a comprehensive new approach. For purposes of illustration, however, we will first contrast the Bank's new poverty estimates with the figures that would result from taking the Bank at its word, as far as the concepts it uses are concerned. Our claim is not that these are correct estimates, but that conceptually they are more warranted by the Bank's own method than the estimates it reports.


The framework we use relies on the following surmise: if the International Poverty Line is to reflect a reasonable level for what it means to be poor, it must correspond to some conception of adequacy for basic human requirements. Even if the IPL is meant to reflect poverty lines defined by (or for) the poorest countries, as the Bank claims, it must still be deemed appropriate according to an ordinary-language conception of what poverty is and why we care about it.²⁶ Further, the supposed interpretation of PPP conversion factors is that they preserve purchasing power across countries. In that case, the IPL chosen must suffice for purchasing the most elementary requirements in the base country for which the price indices are defined—that is, the US—especially if those requirements are conceived of in absolute terms, i.e., without deferring particularly to the contextual specificities of that country. This seems an unavoidable consequence of claiming to preserve purchasing power when one uses PPPs. Referring to differences in standards, or in the purchasing power of currencies, cannot avoid this logical implication.

A measure of what might just suffice in this respect is available. The Thrifty Food Plan produced by the US Department of Agriculture has established, with great care, the minimum cost of achieving 'Recommended Dietary Allowances' in the United States. (The Plan was initially used to set food-stamp allocations and is now used for SNAP, the Supplementary Nutrition Assistance Programme.) Taking a model family of specified size and composition, it collects 'scanner' price data

²⁶ See the recent biting expositions by S. Subramanian: 'Identifying the Poor', 'Controversy over Poverty Line' and "'Price-corrected' Poverty Lines', *The Tribune*, 24 July, 7 August, 21 August 2015.

from markets around the US and calculates the mathematical least cost of achieving the allowances at these prices, using linear-programming techniques; the result is modestly adjusted to make allowance for prevailing tastes. It then verifies that the amount suffices for cooking model recipes in a test kitchen. The Thrifty Food Plan allowance is based on the supposition of home cooking and makes no reference to the costs of the kitchen or the utensils. By definition, the allowance does not suffice for any non-food requirements—for example, shelter, clothing, transportation—and should therefore be taken as a lower bound on real American requirements. In 2011, the Thrifty Food Plan allowance was \$5.04 per person per day, based on a family of four with two children of intermediate age.²⁷ However, to take note of the possible criticism that the Thrifty Food Plan is overly generous, we consider expenditure levels corresponding both to this allowance and to *half* its value, or \$2.52. These can be thought of as food poverty lines, to which non-food requirements should be added. Further, we apply both general consumption PPPs (as does the Bank) and food PPPs, more appropriate to food requirements in particular. Combining these possibilities, and drawing on the data of the Global Consumption and Income Project (GCIP), leads to four alternative poverty lines, with resulting poverty estimates and trends (Figure 3).

FIGURE 3: *Alternate 2011 PPP poverty lines headcount estimates for the world*



Source: World Bank, GCIP, Thrifty Food Plan.

²⁷ See 'Official USDA Food Plans: Cost of Food at Home at Four Levels, US Average, June 2011', available on the US Department of Agriculture's website.

Selecting the Thrifty Food Plan's allowance as an IPL leads to a substantial increase in poverty headcount ratios, both globally and across all regions. If general consumption PPPs are used, more than 80 per cent of people in South Asia and Sub-Saharan Africa are found to live below the line of \$5.04 per capita per day. Even if only half the US level is used, the poverty headcount ratio nearly doubles in East Asia and South Asia. Using the more appropriate food PPPs increases the poverty rate across all regions; 90 per cent of South Asians consume below the Thrifty Food Plan level (Table 4). In contrast to the Bank's findings, the poverty headcount, or the absolute number of poor, increases after 1980 on these measures, peaking in 1990 for the lower lines and in 2000 for the higher ones (Figure 4).²⁸ But for the higher lines, the pace of decline is slower and poverty is still above levels seen in 1990.

This is a purely illustrative approach to generating alternate poverty estimates, which does not explore the implications of further steps, such as alternate choices of inter-sectoral price adjustments for large countries, which again might vary widely from the Bank's. However, the exercise should suffice to prove that the Bank's methods do not generate credible estimates even within its own conceptual framework.

A better approach

There is a practicable alternative for the assessment of income poverty. It would involve anchoring poverty assessment in clear identification criteria with consistent meanings and appropriate substantive interpretations. Specifically, we would advocate a conception of poverty that is absolute in the space of basic human requirements, and relative in the space of commodities needed for those requirements.²⁹ Such a requirements-based approach leaves ample room for the assessment of non-income aspects of poverty as well; the two are complements, not substitutes. At the core of such an approach is the idea of fixing one or more sets of reference requirements that a person must be able to

²⁸ For an incisive analysis of the reasons to be concerned with absolute numbers as well as relative proportions of the poor, see Nicole Hassoun and S. Subramanian, 'Variable Population Poverty Comparisons', *Journal of Development Economics*, vol. 98, no. 2, 2012.

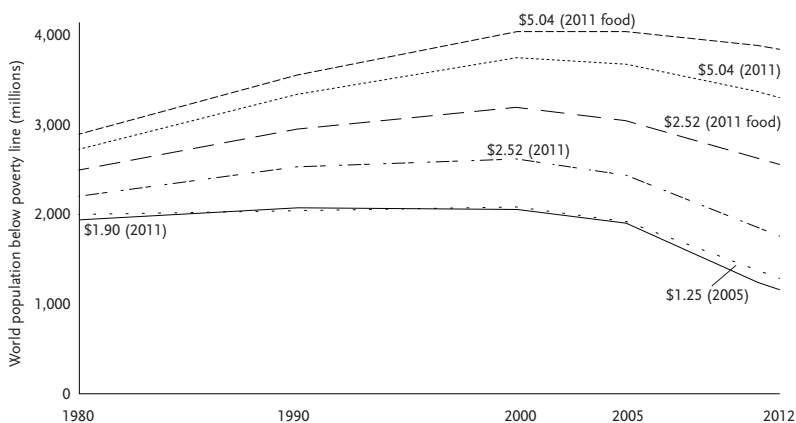
²⁹ The concept of basic human requirements is compatible with Amartya Sen's 'capabilities' approach.

TABLE 4: GCIP headcount ratio estimates for alternate 2011 PPP poverty lines (general consumption and food PPPs)³⁰, 2012

	\$1.90	\$5.04	\$2.52	\$5.04 (food)	\$2.52 (food)
<i>East Asia & Pacific</i>	12.1	42.1	19.3	54.9	30.1
<i>Europe & Central Asia</i>	1.5	7.7	2.4	9.5	2.1
<i>Latin America & Caribbean</i>	2.9	26.2	6.1	28.8	8.9
<i>Middle East & North Africa</i>	3.7	36.2	8.5	59.6	27.1
<i>North America</i>	0.0	0.0	0.0	0.0	0.0
<i>South Asia</i>	23.3	81.8	41.7	90.9	66.2
<i>Sub-Saharan Africa</i>	45.1	81.8	58.2	83.9	72.3
<i>World</i>	15.9	48.5	24.9	56.6	37.1

Source: GCIP.

FIGURE 4: GCIP estimates of the number of poor in the world for alternate poverty lines



Source: GCIP.

³⁰ User discretion is advised, keeping in mind the motivation we have presented, which is internal to the Bank's own method. We do not include results for countries which lack a food PPP when reporting regional results.

achieve in order to be deemed non-poor. The basic human requirements that are typically income-dependent, such as adequate nourishment, are of special interest in relation to income-poverty assessment, although the extent to which requirements are income-dependent would vary across contexts—determined in part, for example, by the extent to which a market economy prevails.

These reference requirements would themselves be fixed across contexts, perhaps through a coordination exercise of the kind previously undertaken by the United Nations in relation to establishing the System of National Accounts. Once the basic requirements are fixed, it will be possible to investigate which specific combinations of commodities possess the characteristics sufficient to meet them—which combinations of foods can generate the specified nutritional requirements, for example. A reference set of commodity characteristics might also be specified across contexts. Finally, the sets of commodities that meet the reference requirements can be identified and priced explicitly. (This is roughly the approach of the Thrifty Food Plan, as of initiatives to determine the cost of a decent standard of living undertaken by the Rowntree Foundation in the UK, or by the Market Basket Measure project in Canada.) In such an approach, the poverty line corresponds not to a cash sum, ‘a dollar a day’, but to ‘food in the belly’. The reference commodities, once identified, can be periodically priced or adjusted. Although there is an important role for expert judgement—for example, nutritionists’ advice—in mapping, from requirements to characteristics to commodities, this process also necessarily involves a democratic component, both in identifying basic human requirements and in validating commodities.³¹

How is this approach different from determining a sound poverty line for any given country? It isn’t, but it adds something extra, which is the element of coordination across countries. This provides a common reference for the human requirements that can be met by a given level of resources. In effect, the development of poverty lines that have a common substantive interpretation generates bottom-up comparability that does away altogether with the need for an IPL, PPPs or other

³¹ On the relation between poverty assessment and democracy, see Anthony Atkinson, ‘Promise and Performance: Why We Need an Official Poverty Report’, in Paul Barker, ed., *Living as Equals*, Oxford 1996.

artefacts of the money-metric approach.³² We believe an alternative approach along these lines is both feasible and desirable, and that it could help to catalyse debate as well as to facilitate poverty monitoring. Even those who do not share our optimism in this regard can agree that the uncertainties associated with the Bank's approach to global poverty require challenging. The problem is not beyond public understanding, and it is rather too important to be left to a small group of technicians, pretending to precision.

³² In effect, this is what the UN's International Civil Service Commission or global human resources consultancy firms implicitly do when they develop cost-of-living indices tied to specific, often explicitly identified, understandings of what constitutes an adequate level of life (albeit far above that needed to avoid poverty).