

Toward better global poverty measures

Martin Ravallion¹

Received: 17 February 2016 / Accepted: 18 February 2016 / Published online: 6 April 2016
© Springer Science+Business Media New York 2016

Abstract While much progress has been made over the last 25 years in measuring global poverty, there are a number of challenges ahead. The paper discusses three sets of problems: (i) how to allow for social effects on welfare, recognizing the identification issues involved; (ii) the need to monitor progress in raising the consumption floor above its biological level, in addition to counting the number of people living near the floor; and (iii) addressing the longstanding concerns about prevailing approaches to making inter-country comparisons of price levels facing poor people. Some suggestions are offered for operational solutions, building on past research.

Keywords Absolute poverty · Relative poverty · Consumption floor · Purchasing power parity · International comparison program

1 Introduction

Assessments of progress against poverty at the country level, and most decisions about how best to fight poverty within countries, do not require global poverty measures. Nonetheless, such measures are important to public knowledge about the world as a whole, and they help inform the work of development agencies, including in setting targets for overall progress.

This paper discusses some current issues that are specific to *global* poverty monitoring. The paper leaves aside some important issues shared with national poverty measurement,

Electronic supplementary material The online version of this article (doi:10.1007/s10888-016-9323-9) contains supplementary material, which is available to authorized users.

✉ Martin Ravallion
mr1185@georgetown.edu

¹ Department of Economics, Georgetown University, Washington, DC, 20057, USA

including the quality of household survey data.¹ The paper focuses on three sets of issues. The first is the inadequacy of past absolute income poverty lines in reflecting social effects on welfare, notably relative deprivation and the country-specific costs of social inclusion. By this view, the standard “\$1 a day” or (\$x a day) measures of absolute poverty are unsatisfactory on their own and may well overstate progress against poverty when overall living standards rise. However, there are some serious but poorly understood challenges in how to implement defensible global relative measures.

The second issue is our ignorance about how the poorest are doing – whether they are being left behind in growing, and generally less poor, developing economies. Knowing how much progress has been made in reducing the counts of people living below one or more international poverty lines may tell us rather little about whether we have left the poorest behind. Addressing this problem calls for a credible and operational measure of the consumption floor. Here too there is a challenge given the limitations of survey data.

The third set of problems relate to cross-country comparisons of price levels for the purpose of global poverty measurement. The data for this task have improved, but there are still unresolved issues. The International Comparison Program (ICP) collects the price data and estimates the Purchasing Power Parity (PPP) exchange rates (World Bank 2014). These PPPs have been widely used in measuring global poverty.² There are puzzling changes in PPPs from one ICP round to the next and concerns about the appropriateness of the methods used. Lack of public access to the primary data has not helped.

After providing an overview of the World Bank’s approach to measuring global poverty using national poverty lines as the data, the paper discusses the three sets of problems in the above order, and suggests some possible solutions.

2 National and international poverty lines

National poverty lines have long provided the data used by the World Bank for setting global lines. Figure 1 plots the national lines for 95 countries against private consumption per capita from the national accounts. These are either official national poverty lines or (when these could not be found) they are the lines set by Eurostat or the World Bank, as part of its analytic work at country level.³ Both the poverty lines and consumption levels are converted to \$US values using the PPP exchange rates for consumption from the 2005 round of the ICP (World Bank 2008).⁴

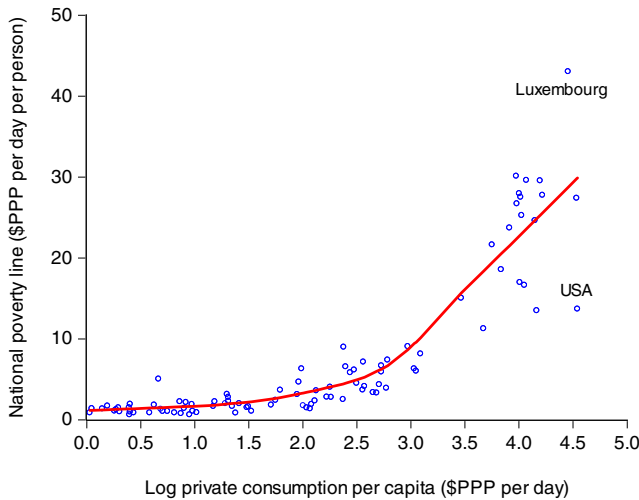
The range in Fig. 1 is large, from \$0.62 to \$43 per day. Of course there are comparability problems and measurement errors in the national lines. The range narrows to \$1.25 to \$30

¹Ravallion (2016b, Part 2) reviews these issues. Other issues not addressed here include the quality of census-based population data and the quality of national accounts data, as used in extrapolation/interpolation methods to line up surveys in time for global comparisons.

²Consumption PPPs are typically used, though sometimes with adaptations aiming to make them more appropriate to global poverty measurement (see, for example, the adaptations discussed in Chen and Ravallion 2010a).

³A similar graph for a larger set of national lines can be found in Jolliffe and Prydz (2015), who used implicit lines from national poverty measures, as well as the explicit lines used in Fig. 1. The pattern of rising lines with mean income, with an increasing slope, is still evident in this extended data set.

⁴All poverty lines are for specific years (often tied to specific survey dates) and consumption data are for that year or as close as possible; both poverty lines and consumption were then converted to 2005 prices using the country’s consumer price index, and then converted to PPP \$’s using the 2005 PPP for consumption.



Note: The figure plots national poverty lines across countries against private consumption per capita from national accounts. The lines for developing countries are found in Ravallion et al. (2009); the extra lines for high-income countries are from Ravallion (2012).

Fig. 1 Poverty lines across countries of the world

per day if one compares the empirical average line, conditional on mean consumption (as estimated by the locally smoothed scatter plot in Fig. 1). National lines tend to rise with mean consumption. The mean for the poorest 15 countries (in terms of consumption per capita) is \$1.25 while the mean for the richest 15 is 20 times higher, at \$25 a day. But even at the top of the income ladder there are marked differences; for example, the Figure points to the difference between Luxembourg's line and that in the U.S.

It is no surprise that we see higher national lines in richer countries. In identifying who is considered "poor" within its borders, a rich country tends to use a more generous allowance for basic needs. The allowances for food are almost always anchored to food-energy requirements, although these vary, with higher mean food energy requirements in places and times with better nourished populations and often with higher activity levels. (Many of us have discovered personally that there are multiple equilibria at which body weight can be maintained.) The food menus identified in practice for attaining these requirements also vary, and are typically more generous (such as with larger allowances for protein) in less poor places. However, the bulk of the gradient in Fig. 1 is due to the allowances for non-food needs, as shown by Ravallion et al. (2009). Practices differ. One common method divides the food poverty line by the typical food share (food spending divided by total spending) of people with total expenditure or income near the poverty line. Another method finds the total expenditure at which food-energy requirements are met on average. Both methods tend to impart a positive income effect on national lines (Ravallion 2012).

Even so, the relationship in Fig. 1 is quite flat at low per capita consumption levels. This is also intuitively plausible, as it is unlikely that the lines used by countries could fall to zero in the limit as mean consumption falls. The empirical elasticity of the poverty line to mean consumption rises from an average that is close to zero up to unity. We will return to this relationship in the following discussion.

As an aside, it should be noted that the poverty lines in Fig. 1 are explicit national lines, as used by the government concerned or by the World Bank in its own national poverty assessments (Ravallion et al. 2009). Jolliffe and Prydz (2015) provide a compilation of *implicit* poverty lines, which they have backed out from national poverty measures in World Bank *Poverty Assessments*, using a set of fitted distributions. (In other words, they solve their fitted distribution for the line that yields the national poverty rate, taken as given.) This provides a much larger sample of over 600 lines, with multiple observations over time for many countries. Jolliffe and Prydz confirm that the poverty line rises with mean consumption as in Fig. 1, but their data do not suggest that the relationship is flat at low mean consumption levels, as in Fig. 1. Their data suggest that the elasticity is positive, less than unity, but that the elasticity is roughly constant over the range of consumption in their data. However, as explained in the Online Appendix, when we consider the effects of measurement errors in the Jolliffe-Prydz method it is likely that their implicit poverty lines hide the flat segment – that their method overestimates the elasticity of the poverty line to the mean in poor countries. The way that their implicit poverty lines shift over time within countries is consistent with this claim (Online Appendix). Nor does this require large errors in their method; the overall fit may well be very good and yet the implicit lines turn out to be positively correlated with the mean even when the explicit lines are not.

In assessing poverty globally, the Bank has long argued that one should use a line with constant purchasing power, as best can be determined, and that it should be set at a level that is reasonably representative of low-income countries, making it an explicitly frugal line. Early examples used India's poverty line for global comparisons.⁵ But why India? Any one country's line can be questioned. Furthermore, the purpose is not to measure poverty in any one country (for which existing national measures should be used) but to measure *global* poverty by a common real-income standard. Some averaging is clearly called for.

From 1990, the Bank switched to an average line for a broader set of poor countries (World Bank 1990; Ravallion et al. 1991). However, it has long been recognized that this is a frugal line, sometimes called “extreme poverty.” More generous international lines can also be defended. For example, Chen and Ravallion (2010a) provide poverty measures for multiple lines from \$1.00 to \$2.50, with dominance tests over a range up to the US line of \$13 per day. But the key point is that whatever line is used, it is held constant in terms of its purchasing power. Two people with the same real consumption are treated the same way no matter where they live.

In a substantial revision to this approach, Ravallion et al. (2009) compiled a much larger sample of national lines, including the 75 observations for developing countries in Fig. 1. On this basis they set a line of \$1.25 at 2005 PPP, with a 95 % confidence interval of (\$1.05, \$1.45).⁶ This is an un-weighted mean, given that poverty lines are observed at the national level. Deaton (2010) argues instead that all national lines should have been used, weighted by the number of poor. (This led Deaton to a line of \$0.99 a day in 2005.) However, the case for such population weighting is unclear; it is questionable that an observation of the

⁵See for example (Ahluwalia et al. 1979). Ravallion (2016b, Part 1) reviews the history of thought on this topic.

⁶Ravallion et al. (2009) used Hansen's (2000) estimator for a piece-wise linear (“threshold”) model in estimating the relationship between national poverty lines and consumption per person. The standard error of the intercept is \$0.10 per day. The group of countries for which the \$1.25 line is the mean is thus endogenous. Klasen et al. (2016) derive the same line using different methods.

poverty line for a small country is any less important in this context than that for a large country.⁷

While Ravallion et al. (2009) followed the spirit of the original “\$1 a day” line of Ravallion et al. (1991) they made no other attempt to assure comparability of the \$1.25 a day line for 2005 with past lines. Given the large new sample of national lines as well as the new PPPs for 2005 a full revision to all past estimates was deemed necessary. Once the new international line had been converted back to local currency units and applied to the available national household surveys, a new series of global measures could be constructed, back to 1981. The new estimates were consistent with prior estimates indicating substantial progress in reducing the incidence of global poverty since the early 1980s Chen and Ravallion, Chen and Ravallion (2010a, 2010b). This consistency is not surprising given that the PPP conversion is only done at one date, with national price indices used in adjusting for inflation over time.

3 Taking social effects on welfare seriously

One of the main sources of dissatisfaction with poverty measures that use a constant real line is that they do not take account of the concerns people face about relative deprivation, shame and social exclusion; these can be termed *social effects on welfare*. It can be argued that such effects are weaker for poor people, for whom the externality from better off friends and neighbors may be positive (Ravallion 2008). However, it is plain today that social effects on welfare are not confined to people living in rich countries; there is evidence consistent with the existence of such effects in countries at all levels of development.⁸ This section discusses the implications for poverty measurement.

It should be noted that there can also be *biological effects* with bearing on poverty measurement, notably when better nutritional status (evident as less stunting and wasting) due to higher real income calls for higher food-energy requirements for maintaining body weight, and also greater activity levels. Measures of income poverty that are anchored to variable food-energy requirements over time and across countries (such as Moatsos 2015) are implicitly relative in the real income space, since food energy requirements natural rise with higher real incomes. Thus these measures tend to show less progress against poverty than standard absolute measures of global poverty such as discussed in Section 2. Here I focus on social effects.

In thinking about the implications of social effects on welfare for global poverty measurement, the key guiding premise here is that the international comparisons of welfare required for a global measure must be anchored to a defensible *and* common concept of individual welfare, which we can dub “welfare consistency.” When forming a global poverty

⁷Also see the thoughtful discussion in Klasen et al. (this issue).

⁸Early contributions were Duesenberry (1949) and Runciman (1966). Anthropologists describe behaviors consistent with this idea; see, for example, Geertz (1976) and Fuller (1992). Rao (2002) describes the role of celebrations to maintaining the social networks crucial to coping with poverty in rural India. Banerjee and Duflo (2007) report seemingly high expenditures on celebrations and festivals by very poor people in survey data for a number of countries. Hirsch (1977) and Frank (1985) discuss how the evaluation of certain consumption goods depends on consumption relative to others. Data on subjective welfare have suggested that relative positions matter (see, for example, Luttmer 2005) although there are concerns that the regression used may well over-estimate the income elasticity of the poverty line (Ravallion 2008).

measure, it would be morally unacceptable to be judging (possibly implicitly) that a person is better off living at the international line in a rich country than a poor one. To the extent feasible with the data available, everyone's poverty status must be judged by a consistent welfare concept. This has long been accepted by economists. For example, Sen (1983) argued that a person's capabilities should be seen as the absolute standard but that "... an absolute approach in the space of capabilities translates into a relative approach in the space of commodities" (Sen 1983, p. 168).⁹ Whether the absolute standard is an index of "utility" (as favored by most economists) or an index of "capabilities" (as in Sen) may be important for the implementation, but the first-order issue is to demand welfare consistency in some defensible sense when measuring global poverty.

That does not, however, mean that where or when one lives is irrelevant to welfare at given own consumption. Social effects imply that absolute lines in the income space do not correspond to a common level of welfare, i.e., they are not welfare-consistent. To the extent that "poverty" means a low level of welfare *and* welfare depends on relative consumption as well as own consumption, higher monetary poverty lines will be needed in richer countries to reach the same level of welfare. Similarly, poverty measures using a fixed real line will overstate progress against poverty in growing economies. While national poverty lines are rarely revised quickly – there is clearly political resistance – they have risen over time with sustained gains in overall living standards. That has happened in the rich world over the last 100 years, and it has happened in recent times in growing developing countries including China and India.¹⁰

There have been various approaches to measuring relative poverty.¹¹ The sociological literature has emphasized the idea of relative deprivation, namely that the sufficiency of a person's income must be judged relative to the society in which she lives. Townsend (1962, p.219) put it this way: "... individuals and families whose resources over time fall seriously short of the resources commanded by the average individual or family in the community in which they live... are in poverty." In operationalizing this idea, the most common approach is exemplified by the relative poverty measures popular in Western Europe and at the (non-US) OECD in which the poverty lines are set at a constant proportion (typically around half) of the current mean or median.¹² Most of the high-income countries represented in Fig. 1 use such lines. Let us call these *strongly relative measures*.

These measures have features that lead one to question their global relevance. In particular, they ignore the fact that the costs of avoiding relative deprivation and social exclusion cannot fall to zero, but must have a positive minimum. Consider the classic example of

⁹Sen was commenting in particular on the sociological approach to measuring poverty in Britain taken by Townsend (1979). There were a subsequent comment by Townsend (1985) and reply by Sen (1985).

¹⁰ China's official poverty line doubled over a period when average incomes increased by a factor of four, and India's official line has also increased in real terms (Ravallion 2012). The U.S. is one of the few rich countries for which the official poverty line has been fixed in real terms (since the 1960s), although this has been much debated in the literature; see Citro and Michael (1995), Blank (2008) and Johnson and Smeeding (2012), which also discusses the supplementary measures that have been introduced in recent years by the U.S. Census Bureau. While the U.S. did not have an official poverty line 100 years ago, the most credible estimate at the time by Hunter (1904) was only a small fraction of the current official line; indeed, the Hunter line appears to be close to the "\$1 a day" international line (Ravallion, 2016, Chapter 1).

¹¹Examples include Townsend (1979), Mack and Lansley (1985), Atkinson (1998), Eurostat (2005), OECD (2008) and Garroway and de Laiglesia (2012).

¹²For example, the OECD used 50 % of median equivalized income (based on their assumed equivalence scale).

a “social inclusion” need found in Adam Smith’s (1776, Book 5, Chapter 2, Article 4) description of the role of a linen shirt in late eighteenth century Europe, whereby “a creditable day-labourer would be ashamed to appear in public without a linen shirt.”¹³ A socially acceptable linen shirt cannot cost less for the poorest person (let alone zero in the limit), so it cannot be that the relative line is a constant proportion of the mean.

The relevance of strongly relative measures for capturing what poverty means in developing countries is questionable. For example, Ravallion (2012) points out that if one uses a strongly relative line set at half the mean then its average value for the poorest 15 countries is a very low \$0.64 a day, while the value for the country with the lowest mean would be only \$0.38 per day. Similarly, the Garroway and De Laiglesia (2012) relative line, set at 50 % of the median, gives lines that are well below the poverty lines typical of even low-income countries.

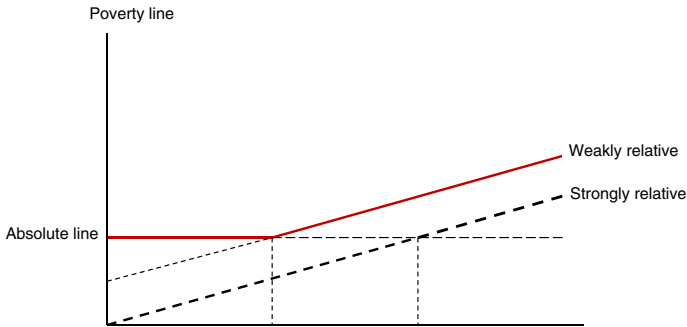
The likelihood of understating the social inclusion needs of globally poor people using strongly relative measures also comes with a seemingly perverse implication for how these measures respond to economic growth and contraction. When the poverty line is set at a constant proportion of the mean, the resulting poverty measure depends solely on the distribution of relative incomes in the population, which can be loosely interpreted as “inequality.” If all income levels grow (or contract) at the same rate then the poverty measure will remain unchanged when the poverty line is set at a constant proportion of the mean or median. Seemingly perverse poverty comparisons have been found using strongly relative measures.¹⁴

These concerns can be addressed by *weakly-relative* measures for which the poverty line rises with the overall mean but with an elasticity less than unity. An example is Foster’s (1998) “hybrid line,” given by the geometric mean of an absolute line and a strongly relative line. While this is weakly relative, it has a constant elasticity, which does not seem plausible, and is inconsistent with how national poverty lines vary across countries (Fig. 1). From this point of view, the hybrid measure proposed by Atkinson and Bourguignon (2001) is more attractive as it has an elasticity of zero at low incomes, with the elasticity then rising above some point. However, it has the undesirable feature that the relative component goes to zero at zero mean. As already noted, this almost certainly understates the costs of social inclusion in poor countries.

These problems are avoided by the weakly relative line in Fig. 2, which is the absolute line up to some critical income level, but then rises with the mean after that. Notice that the relative component of the weakly relative line does not go to zero at zero income. Thus it can allow for a positive minimum cost of social inclusion in the poorest countries. (As that positive minimum goes toward zero one gets closer to the Atkinson-Bourguignon proposal.) One is not poor if one is neither absolutely poor – as judged by the common international standard – nor relatively poor, as determined by the mean consumption of the country of residence. Again, national lines provide a basis for identification. Chen and Ravallion (2001) provided poverty measures that are relative between countries but absolute over time; the poverty line for any country was its predicted line based on its mean consumption and the

¹³Clothing often has an important social role in poor countries; see, for example, Friedman (1990). In today’s urban China or urban India (say), a similar role may well be a cell phone.

¹⁴See, for example, the UNDP (2005, Box 3) (based on Nolan et al. 2005) and Easton (2002). At some risk of understatement, the UNDP (2005, p. 334) writes: “It is clear that when economic conditions change rapidly, relative poverty measures do not always present a complete picture of the ways that economic change affects people’s lives.”



Note: The figure gives stylized examples of different types of global poverty lines. The bold line represents a set of weakly relative lines, while the dashed line represents strongly relative lines.

Fig. 2 Strongly and weakly relative poverty lines

original set of national lines used by Ravallion et al. (1991).¹⁵ Ravallion and Chen (2011) and Chen and Ravallion (2013) calibrated their weakly relative lines to the national lines for developing countries (the subset of 75 lines in the 95 lines given in Fig. 1). Chen and Ravallion find that the piece-wise linear functional form in Fig. 2 gives a very good fit to the data on national lines. Ravallion and Chen (2013b) provide global relative measures constructed on this basis. In their measures the line is both relative between countries and over time within countries.

But are such weakly relative poverty lines welfare-consistent? This depends on why we see higher lines in richer countries, as in Fig. 1. If the positive gradient in poverty lines seen in Fig. 1 only reflects social effects on welfare then a welfare-consistent global absolute poverty line must similarly rise with the mean. However, if the gradient reflects higher, socially determined, minimum levels of welfare for not being considered poor then an absolute measure in the space of welfare requires a fixed real poverty line in the income space.

We do not know which model is correct. Compare someone living on \$25 a day in a rich country (the mean for the richest 15 countries in the Figure) with someone living on \$1.25 in a poor country. A factor of 20 may seem high, but we cannot rule out the possibility that it is no more than suffices to compensate for the extra social effect on welfare of living in a rich country. Nor can we reject with any confidence the claim that richer countries may use higher reference levels of welfare in determining their poverty lines. Social standards for defining poverty may well differ between rich and poor societies, and evolve over time in growing economies.

One way of thinking about the national poverty line is that it is the lowest income in the socially optimal income distribution for a given country and date, where that optimum maximizes a weighted aggregate of all utilities constrained only by total income. However, this “first-best” optimum is not attained everywhere at all dates, given other constraints, including on public information, administrative capacities and political-economy frictions

¹⁵ This is similar to the method proposed by Reddy and Pogge (2010). However, unlike Reddy and Pogge, it is not claimed by Chen and Ravallion (2001) that these lines are absolute between countries. Also see the comments on Reddy and Pogge in Ravallion (2010).

on changing distribution. Given these constraints, the actual distribution in a country at any given time can be far more unequal than the first-best. Thus some people fall below the lower bound of the socially-optimal distribution. That lower bound (and indeed all its quantiles) will almost certainly rise with the overall mean. Thus we can expect higher poverty lines in richer countries even without social effects on welfare; rather, it is the social standard for welfare – the reference level of welfare deemed necessary to not be considered “poor” – that rises with the mean.

To give an example, consider the Official Supplementary Benefit (OSB) levels in Britain, which were once widely used for measuring poverty there. The OSB levels were selected through a political process, which led them to be revised upwards in real terms over the post-World War 2 period. We can interpret this rise in two ways. We can argue that a higher line was needed to attain the same level of welfare, reflecting relative deprivation in a growing economy. Or one can contend that with the extra public resources, the political process led to a more generous view of the socially desirable minimum level of individual welfare – just as the quality of public services generally tends to rise with sustained growth. The latter interpretation suggests that the relative lines based on the OSB are probably not welfare-consistent in that they judge poverty by a common absolute standard of welfare.

This is a serious but rarely acknowledged identification problem, which clouds the interpretation of all relative poverty lines. Social effects on welfare are no doubt at play, but so are differences in underlying welfare standards. Relative poverty lines make sense if one thinks that it is the social effects on welfare that yield higher lines in richer countries. One would be less inclined to accept such lines if the differences are seen to stem largely from different welfare standards for defining poverty. While rich countries are free to use higher reference welfare levels for defining poverty, that does not mean we should also do so in making global poverty comparisons, which should presumably apply a common *welfare* standard on ethical grounds.

Absent a credible solution to this identification problem, the only intellectually defensible approach is to use two international lines, one of which has constant purchasing power while the other is date and place specific. The constant line can be thought of as the *lower bound* to the extent of global poverty; this lower bound is relevant if one assumes that the national lines only vary according to prevailing social standards for the minimum level of welfare needed to not be considered poor. The relative lines fitted to national lines can be interpreted as providing an *upper bound*, in which the national lines are taken to reflect the costs of attaining a common level of welfare. In other words, the upper bound for a given country and date would be the predicted line at the mean consumption found at that country and date, based on the observed relationship between national lines and mean consumption.¹⁶ The true welfare-consistent measure is somewhere between the two bounds.

4 Monitoring progress in raising the floor

The traditional approach to poverty measurement sets a line, such as \$1.25 a day in 2005 prices, and then counts how many people live in households with consumption or income

¹⁶ Note that one has no choice but to use the predicted lines since we have incomplete data on the actual lines across all countries and dates. Indeed, the 95 data points in Fig. 1 represent only 8 % of the relative lines that are needed to estimate global poverty for 125 countries in each of 10 years, as in Chen and Ravallion (2013).

per person less than this line.¹⁷ But that is not how everyone thinks about progress against poverty. Some focus instead on how the poorest are faring. For example, at the launch of the United Nations' (2011) *Millennium Goals Report*, the U.N.'s Secretary-General Ban Ki-moon said that: "The poorest of the world are being left behind. We need to reach out and lift them into our lifeboat." Similarly, the President of the World Bank, Jim Yong Kim, talks about "eliminating" global poverty.

These popular views have deep roots. The idea that justice is only served when every individual is covered by its precepts – none are left behind – is implicit in the objectives of social policies that strive for a minimum level of living in society. For example, the two largest antipoverty programs in the world (in terms of population coverage) aim to do just that, namely China's *Dibao* program and India's *National Employment Guarantee Scheme*.¹⁸ Similarly, in prominent discussions of distributive justice, a society's progress is judged in part by its ability to enhance the economic welfare of the least advantaged group ("maximin"). Famously, this is one of the principles of justice proposed by Rawls (1971). In social choice theory, a lexicographic version of maximin (often called "leximin") has been derived as the rule for choice under certain assumptions (Hammond 1976; Fleurbaey and Maniquet 2011, Chapter 3).¹⁹

The traditional counting approach does not explicitly address these concerns about whether the poorest are left behind. Logically, for the poorest to not be left behind there must be an increase in the lower bound to the distribution of levels of living. The lower bound can be called the *consumption floor*, which we can think of as the typical level of living of the poorest stratum. An appealing concept of "level of living" is permanent consumption (Friedman 1957).²⁰ If the poorest see a gain in permanent consumption then we can say that (by definition) the consumption floor has risen. Human physiology makes a positive floor plausible, given the nutritional requirements for basal metabolism and normal activity levels. This can be called the "the biological floor." Given economic growth and (private and public) redistribution, the consumption floor may well be above the biological floor. But it is almost certainly lower than prevailing poverty lines (national or international). To my knowledge, the only estimate to date of the biological floor indicates that it is about \$0.67 per day at 2005 PPP (Lindgren 2015).

The count of those living below \$1.25 a day (say) can fall without a higher floor. Indeed, none of the prevailing measures of economic progress are likely to put much weight on progress in raising the floor. The growth rate in the overall mean income will have a low (possibly very low) implicit weight on the growth rate in the floor, given that the share of total income going to the poorest is likely to be very low. The same is true of standard poverty measures, as illustrated in Fig. 3, each panel of which gives two cumulative distribution functions (CDFs). In each case, the upper CDF is the initial one and the lower CDF is for a later date. The drop in the incidence of poverty is similar in panels (a) and (b). In (a), the counting approach can claim that many of the poorest have been reached even though

¹⁷ The counts may or not be equally weighted; more sophisticated "distribution-sensitive" poverty measures give higher weight to gains to the poorest.

¹⁸ How well they do this in practice is another matter; on these schemes see Ravallion (2014a) and Dutta et al. (2014).

¹⁹ By the lexicographic version, if the worst off are equally well off in two states then one looks to the next worst off, and so on.

²⁰ Past approaches to measuring poverty have often called for inter-temporal averaging; recall, for example, Townsend's (1962, p.219) reference to "resources over time." For further discussion see Ravallion (2016b, Ch. 3).

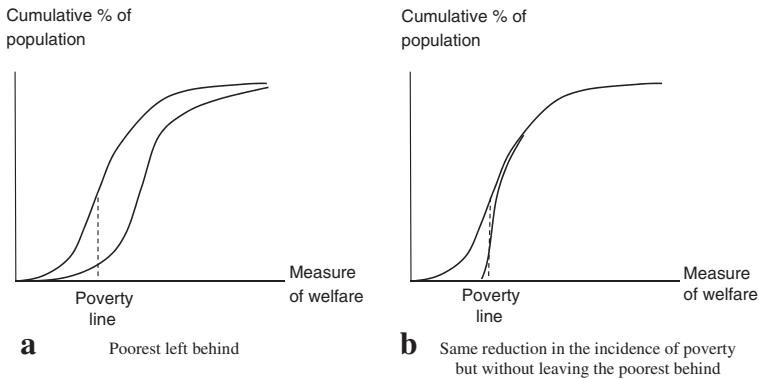


Fig. 3 Same reduction in the poverty count but different implications for the poorest

the floor has not risen, so some people still remain living at the same very low level. In panel (b), the same reduction in the poverty rate has come with a rising floor – implying that the poorest are not left behind.

This neglect of explicit attention to the poorest is due, at least in part, to the difficulties in identifying the floor. The lowest observed level of living in a survey may differ greatly from that household’s normal level of living. For example, food consumption in a survey typically relates to the last week or two, and may be subject to large transient effects. Given the current interest in assuring that no person is left behind, this is a gap in the “dashboard” of development indicators.

Elsewhere I have argued that there is a relatively easy fix for this inadequacy of our existing measures (Ravallion 2016a). Success in assuring that no person is left behind can be monitored from existing data sources under certain assumptions. The approach recognizes that there are both measurement errors and transient consumption effects in the observed survey data. However, the data are assumed to be reliable enough to assure that the person with the lowest observed consumption is more likely to be living at the floor than anyone else.

To make this approach operational with readily available data, Ravallion (2016a) assumes (as a first approximation) that the probability of any observed consumption being the floor falls linearly up to an upper bound.²¹ Then the ratio of the squared poverty gap to the poverty gap relative to that bound – two readily-available poverty measures – emerges as the key (inverse) indicator for assessing progress in raising the floor.²² For various parameter values, the evidence suggests that the developing world has so far had very little success at raising the floor above the biological level, despite the progress in reducing the number of people living near the floor (Ravallion 2016a). The world’s poorest have not been (literally) “left behind,” but they have gained disappointingly little over the last 30 years.

²¹This assumption can be relaxed to allow for nonlinearity in the form of a quadratic probability function, although then an approximation is required to allow estimation from the available data (Ravallion 2016a). The results for this nonlinear function are found to be quite similar to those for the linear probability model reported here.

²² This is proved in Ravallion (2016a). The squared poverty gap was proposed by Foster, Greer and Thorbecke (1984) and is available in the *PovcalNet* output along with the poverty gap index.

It is hard to defend the view that we should look solely at the level of the floor in assessing global poverty. The counting approach remains relevant. However, the level of the floor in society does appear to have normative relevance independently of success in reducing the numbers of people living near that floor. While progress against poverty should not be judged solely by the level of the floor, the latter should not be ignored.

5 Price level comparisons across countries

If strongly relative poverty measures are used then we do not, of course, need to make international comparisons of price levels. The paper has argued, however, that strongly relative measures are hard to defend for the purpose of global poverty measurement. As long as such measures are rejected for that purpose (although they may still be deemed relevant to high-income countries), international price-level comparisons are called for.

The inadequacy of using market (or official) exchange rates for this purpose is well understood. There is no obvious economic mechanism to assure price parity across borders for those commodities that are not internationally traded. Poorer countries tend to have lower wage rates and (hence) lower prices of these non-traded goods relative to traded ones. Thus the PPP exchange rate differs systematically from the market exchange rate (MER). This is the prediction of the classic Balassa-Samuelson model of a competitive market economy with mobile factors of production between the sectors producing traded and non-traded-goods.²³

Motivated by this argument, the ICP collects prices in each country. Field workers report the prices of selected commodities to the designated country authority (often the government's statistics office), under the guidance of the ICP's relevant regional authority.²⁴ Overall guidelines are provided by the global office (housed in the World Bank), although regional and country units have a degree of independence in implementation, such as in setting commodity lists for pricing. The PPPs for each country are then estimated by the regional office, with the global office doing the final linkage across regions. In this final step, relative PPPs across countries are kept fixed within each of the ICP's regions.²⁵ The ICP does not provide sub-national PPPs. Nor does it provide open access to the raw price data.

Each new round of the ICP has brought its surprises. Many developing countries saw substantial changes to their real incomes relative to expectations using the ICP's prior round for 2005 (World Bank 2008). The ICP and its supporters typically tell users that the unexpected changes are due to faults in the prior round. (Based on past experience, we can expect the ICP to be telling us some time in the future about the defects of their 2011 round.) Users are asked to fully revise their economic map of the world with each new round. Some users, including the author, have felt a bit like they are on an ICP roller-coaster ride, as the relative sizes of economies rise and fall unexpectedly. We might like to get off this roller coaster, but that is not easy given that the ICP remains the best source of data on price levels across

²³See Balassa (1964) and Samuelson (1964). Some concerns about the relevance of this model to developing countries are noted in Ravallion (2013a). An alternative model was proposed by Bhagwati (1984) based on factor endowments, leading (labor-intensive) services to be cheaper in poor countries.

²⁴For example, the Asian Development Bank in Manila runs the ICP for all of Asia, including South and West Asia.

²⁵Thus the original PPPs within a given region are scaled up or down by the same proportion when switching to the global PPPs.

countries. Rather than switch fully to a new round, ignoring the history, it would make more sense to use a moving average of PPPs across ICP rounds, making the measures less vulnerable to errors in the ICP data and methods.

The PPPs have entered the calculation of the World Bank's global poverty measures in two ways.²⁶ First, they are used to convert national poverty lines into a common currency. Second, the international line derived from the national lines is converted to local currencies at PPP. It is then converted to the prices prevailing at the time of the relevant household survey using the best available price index for that country.²⁷ Then the poverty rate is calculated from that survey using the micro data or specially commissioned tabulations. As we will see, the fact that the PPPs also matter in determining the (endogenous) international poverty line has not been properly appreciated in some of the literature.

Assessments of the impact of the PPP revisions have sometimes confounded changes due to the PPPs with other changes in methodology, notably in the set of national lines used to determine the international line. For example, it has been claimed that the prior set of PPP revisions in 2005 substantially increased the poverty count. But in fact there was a small downward revision once one isolates the effect of PPP revisions from changes in the sample of national poverty lines. Chen and Ravallion (2010a) show that the effect of the 2005 ICP revisions on their own was to bring down the poverty count; they show that the fact that the poverty count rose was due instead to their new and better sample of national poverty lines (as reported in Ravallion et al. 1991).

5.1 PPP revisions and global poverty measures

The ICP's estimates from the 2011 round were released in World Bank (2014). Since then we have seen various attempts to update the \$1.25 a day line in the light of the price changes from 2005 to 2011, and there has been much debate. It may be worth laying out the differences as I see them.

The first updating attempt was by Dykstra et al. (2014), in a blog post entitled "Global Absolute Poverty fell by Almost Half on Tuesday," which came out just a few days after the release of the 2011 ICP. They simply adjusted the \$1.25 a day line in 2005 prices for US inflation, which gave them a line of \$1.44 for 2011. Using this line, Dykstra et al. argued that the new PPPs imply almost half the global poverty rate for 2011 as the old PPPs. A debate ensued about the new ICP and its implications for the global economic landscape. The calculations of the impact of the new PPPs were found to be quite sensitive to the level of the poverty line, as shown by Edward and Sumner (2015).

It was realized by many observers that fixing the US purchasing power of the international line is very hard to defend given the generally higher inflation rates in developing countries.²⁸ Essentially two approaches emerged in the subsequent literature. In the first,

²⁶This follows from the approach of basing international lines on national lines proposed by Ravallion et al. (1991).

²⁷PPP are rarely used to make intertemporal comparisons for a given country. The tradition in applied work has been to use national price indices for such comparisons, while PPP conversions are only done in the ICP benchmark year. This has been the approach in the literature on growth empirics as well as in global poverty monitoring.

²⁸Thus, while \$1.44 a day in 2011 has the same purchasing power in the US as \$1.25 in 2005, when \$1.44 is expressed in local currencies of developing countries using the 2011 PPPs it has lower purchasing power in many of those countries than when the prior \$1.25 line in local currency is adjusted for inflation in those countries.

Klasen et al. (this issue) calculated international lines for 2011 that broadly follow the methods of Ravallion et al. (2009) in setting the \$1.25 line, although Klasen et al. argued for some methodological changes, which I will not go into for brevity. Their preferred estimates indicate an international line of around \$1.70 a day in 2011, which gives a lower overall poverty count for 2011 than the \$1.25 line, but not (of course) as low as Dykstra et al. (2014). Second, Ferreira et al. (2016) propose a line of \$1.90 a day for 2011, which has been adopted by the World Bank. This line was obtained by updating the national poverty lines underlying the \$1.25 a day line proposed by Ravallion et al. (2009) for 2005, where the updating adjusted for inflation in the countries concerned.²⁹ Kakwani and Son (2016) use the same method but do so for all developing countries instead of just the 15 poorest from the Ravallion et al. (2009) study. On taking a population-weighted mean of these 2011 lines at PPP Kakwani and Son get an international line of \$1.93 a day – very similar to Ferreira et al. (this issue), although (as noted) the case for population weighting in this context is unclear.

Using a line around \$1.90 a day for 2011 one finds much less change in the aggregate poverty count when compared to the old numbers using the 2005 ICP, although there are some large differences for individual countries. It is clear that the initial assessments suggesting a halving of the poverty count using the 2011 ICP were deceptive. This illustrates the importance of the aforementioned point that the PPPs matter to setting both the international line and the global poverty counts for a given line.

The changes in PPPs underlying these comparisons between poverty measures anchored to different ICP benchmarks have been the subject of critical concern in the development community, and not just among those measuring global poverty. The real exchange rate revisions implied by the 2011 ICP have been puzzling to some observers (the author included). The real exchange rate (as often measured) is the inverse of what is more often called in the PPP literature the Price Level Index (PLI) – the ratio of the PPP to the MER. (One can interpret the PLI as a measure of how cheaply one can live in a country with the \$US.) Here I will focus on the PLI rather than the real exchange rate.

The main concern has been why we have not seen more sign of rising PLIs in the growing developing countries. Recall why PPPs are collected in the first place, namely the fact of lower wage rates in poorer countries and (hence) cheaper non-traded goods. This must presumably also happen over time, so that one expects the PLI to rise in developing economies experiencing sustained growth and structural change. This can be called the Dynamic Penn Effect (DPE) (Ravallion 2013a). The DPE is a strong and stable feature of the changes in PLIs between the ICP rounds (Ravallion 2013a, b, 2015; Majumder et al. 2015).³⁰ Thus we would expect a rise in the PLI in a country such as India over 2005–11. Instead, the 2011 PPP implies a *fall* in India's PLI despite the country's recent high growth rate and rising real wages (Ravallion 2014b). Possibly the 2005 PPP was overestimated, or the 2011 PPP has been underestimated. The puzzles are not just for India. There has been a general downward

²⁹Specifically, they update the local currency equivalent of \$1.25 at 2005 PPP for inflation in that country and then convert to \$US in 2011 using the new 2011 PPP.

³⁰The term "Penn effect" stems from the Penn World Tables (Summers and Heston 1991), which provided the data that were used to establish this effect empirically, indicating a higher PLI in countries with higher GDP per capita. The DPE simply says that the same relationship should also hold over time in growing economies. Inklaar (2013) questions whether the DPE is present; Ravallion (2013b) shows that the DPE is present in a model encompassing (Ravallion's 2013a) DPE test and Inklaar's alternative. Majumder et al. (2015) confirm the presence of the DPE in Ravallion's augmented regression, but also find that inequality within countries matters, consistently with their assumption that the poor have a higher propensity to consume nontraded goods than do the rich.

drift in the PLIs in the 2011 ICP, undoing the ICP's prior trend in the opposite direction (Ravallion 2015). Controlling for the GDP growth rate, one finds a downward drift of about 1 % per annum in the PLIs over 2005-11, reversing the upward trend in the 2005 ICP relative to 1993. The downward drift is concentrated in the ICP's Asia regions. We can call this the Asia drift in the 2011 ICP.

It is important to understand these PPP changes if one is to make comparisons between global poverty measures using different benchmark rounds of the ICP. The changes in PPPs reflect in part the changes in domestic prices implied by national Consumer Price Indices (CPIs). There are also economic changes that can influence the evolution of PPPs relative to CPIs. The differences in the weights and, in all likelihood, the prices used, can readily entail that the economic changes have differing effects on PPPs versus CPIs. Thus the DPE can play a role in explaining PPP changes controlling for the domestic inflation rates implied by the CPIs.

The literature has identified a number of other factors, both substantive and methodological. One difference is in how the PPPs were linked across regions. In the 2005 round the linkage used a set of 18 "ring countries" spanning the ICP regions. These countries did their own pricing exercise for a common list of goods. In 2011, the ring method was replaced by a common global core list (GCL), priced in all countries. Deaton and Aten (2014) and Inklaar and Rao (2014) argue that this was an improvement, and that the deficiencies of the prior ring method account for much of the change in PPPs in 2011. Elsewhere I have argued that neither economic effects nor the observed methodological changes fully account for the unexpected changes in PPPs (Ravallion 2015). While the literature offers some clues, the source of the Asia drift in the PLIs for 2011 remains unclear. The Asian Development Bank (2014) indicates that they did a better job in 2011 of including rural areas in their ICP price surveys for Asia, although this is less evident in the available meta data for other regions. A differential improvement in sampling may help explain the otherwise puzzling Asia drift in the 2011 ICP relative to macroeconomics covariates (Ravallion 2015). Further research is needed to properly test this hypothesis, which raises concerns about the inter-regional comparability of the 2011 ICP.

The methods used to collect the primary prices have received less attention than other aspects. ICP users do not always appreciate how important those methods are to the PPPs. In poor countries, many of the goods described in the global or regional lists will be missing in the sampled markets. These should be treated as missing values. To some extent an effort is probably made to fill in the missing values by finding substitutes, which (in poor areas) will almost certainly be of lower quality. (Incentive payments for more complete price reports encourage this effort.) The risk of underestimating price levels in poorer places is plain. The 2005 ICP introduced more detailed product descriptions (World Bank 2008). It is not clear how much that effort was sustained in 2011; more information from the ICP's field reports would be welcome. Past ICP methods have used simple averages of the price quotations within each of the (quite aggregated) "basic headings" of the national accounts and do not appear to have made any adjustments for the likely quality bias when replacing missing prices in the field.³¹

In the efforts of the user community to understand the PPP revisions it has not helped that the raw price data have not been publicly available, to allow researchers to do their own

³¹The country-product dummy (CPD) method has been widely used, and appears to have been the main method used for filling in missing values in the 2005 and 2011 ICP rounds. On the CPD method see Silver (2009).

calculations – both to replicate and to assess the implications of alternative assumptions.³² While it can be credibly claimed that confidentiality limits access for some types of data, this can hardly be so for consumer prices, which should be public knowledge if markets work well. The micro price data from the 2011 ICP, and future ICP rounds, should be made public.

5.2 Going forward

How should the PPPs be estimated? It must first be acknowledged that the existing PPPs from Penn World Tables and the World Bank have never been ideal for global poverty measurement. They are designed for comparing national accounts not for measuring poverty, as was pointed out in the first paper on the “\$1 a day” poverty measures (Ravallion et al. 1991). Past practice in using the existing PPPs for global poverty measurement has entailed some ex post adjustments to enhance their relevance to global poverty measurement. For many countries, the ICP’s price collection has favored urban areas as it is less costly to obtain the price reports.³³ Yet prices differ, and often substantially, between urban and rural areas. The metadata for the 2011 ICP suggest that substantial urban bias is still present, with rural prices not being collected for about half of the developing countries.³⁴ Systematic ex post adjustments for urban bias are called for in using the PPPs for global poverty measurement.

There has also been a long-standing concern that the ICP under-weights food. To assess this, Fig. 4 compares the food shares implied by national household survey data with the shares in the 2011 PPPs from the ICP. For as many countries as possible, the Figure gives the food Engel curve, in the form of the regression of the food share on log total consumption per person (though this is not constrained to be linear). This is constructed in two ways. One is the Engel curve implied by the expenditure share on food and non-alcoholic beverages in the household consumption PPPs from the 2011 ICP.³⁵ The other is the average food share in the closest available household survey, with no survey older than 2006. The median year of the surveys is 2011.³⁶

We see that there is a sizeable gap between the two Engel curves in Fig. 4. For example, the ICP food share for India is 30 %, as compared to 52 % and 44 % for rural and urban areas respectively (based on India’s National Sample Survey). As expected, the gap falls as mean consumption rises. For the poorest quarter of the sample – up to log consumption per capita of 5 – the gap is an average of 11.3 percentage points (with a standard error of 1.7 %). By contrast the gap for the other three quarters of countries is 5.6 percentage points (standard error of 0.9 %).

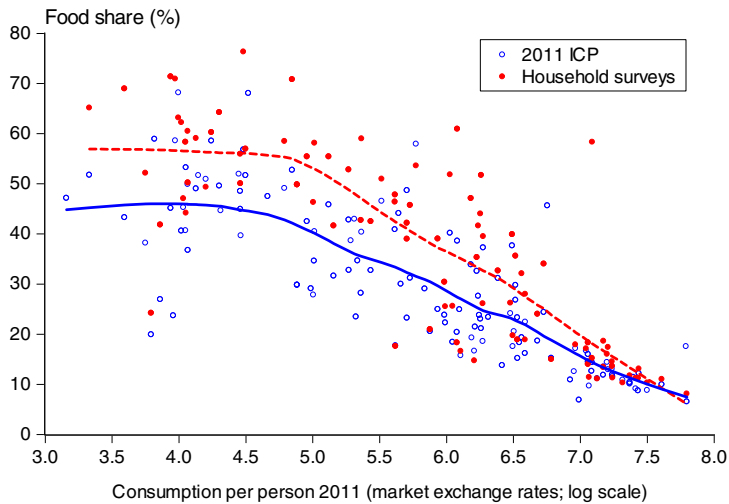
³² The prices have only been available at the aggregate basic heading level of the national accounts. ICP staff appear to take a simple average of the raw prices under each heading.

³³ For example, there was an urban bias in the sampling for the 2005 PPP for China, which required a correction using supplementary data on the cost-of-living differentials between urban and rural areas, as described in Chen and Ravallion (2010b). Failing to correct for this bias entailed a large over-estimation of the global poverty count.

³⁴ For the 156 countries for which the metadata are available at the time of writing, my count indicates that the price surveys only included rural areas (beyond the capital city) for 55, about one third. Excluding the OECD and CIS, the proportion rises to about half. (Details are available from the author.)

³⁵ The household food share in the ICP can be derived from the public tables released by the ICP on the World Bank’s ICP web site. The ICP food share in household consumption is obtained by dividing the GDP share attributed to household expenditures on food and non-alcoholic beverages by the GDP share accounted to individual household consumption, both of which are found in Table D2 in the [Online Supplementary Material](#) on the Bank’s ICP web site.

³⁶ The calculations were done by World Bank staff mainly using the micro data; see Ravallion (2015).



Note: The figure plots the food share (share of household expenditure devoted to food) against consumption per person for food shares implied by household surveys and those implied by the 2011 ICP. The bold line is the estimated (nonparametric) regression function for the 2011 ICP data while the dashed line is for the national household surveys. The source is Ravallion and Chen (2015).

Fig. 4 Two Engel curves

In 2008 the ICP's global office (based in the World Bank) recognizing that existing PPPs do not have appropriate weights for measuring poverty. The global office initiated work on the 2005 ICP to create alternative PPPs re-weighted consistently with household surveys. This work is documented in Deaton and Dupriez (2011), which constructed a set of "PPP for the poor" that accord better with the consumption patterns of people living near the international poverty line, based on household surveys. As it turned out, there was not much difference between the PPPs for the poor and the standard PPPs based on national accounts for 2005; Chen and Ravallion (2010a) give their global poverty estimates using the Deaton-Dupriez PPPs, and show that they are very similar to the estimates obtained using the regular PPPs. However, while Deaton and Dupriez (2011) provide a good starting point, their exercise was limited by the fact that they did not have access to the micro price data (only the aggregated data at the basic-heading level of the national accounts). We could do better with those data.

Some observers have asked whether it would not be better to give up on the ICP. One can simply stay with the old PPPs, continuing to update for inflation within each country, consistently with what is done between ICP rounds; this is one of the proposals in Klasen et al. (this issue). One can also estimate PPPs without ICP prices, or indeed any data on prices. Under certain conditions, one can identify the PPPs from the estimable parameters of empirical Engel curves rather than using data on prices. Majumder, Ray and Santra (MRS) (2015) estimate PPPs using this method.³⁷ Essentially their PPPs become functions of country-specific parameters of the Engel curves relating budget shares to incomes. Against

³⁷ Antecedents using empirical Engel curves to estimate price indices include (Hamilton 2001) and Almås (2012).

the advantage of not requiring data on prices, there is the concern that there may well be other reasons why Engel curves shift geographically besides price differences³⁸. The key assumption of a geographically stable Engel curve can be questioned, although it is somewhat less problematic when one confines attention to developing countries as in MRS. And, of course, all methods of estimating PPPs make assumptions that can be questioned. The value of the MRS calculations is in providing estimates under a different set of assumptions. And it is clear that the PPPs are not robust to this difference. While the MRS estimates using the Engel curve method are highly correlated with those from the ICP – a correlation coefficient over all countries of 0.91 – there are some marked differences in the levels when compared to the PPPs from the 2011 ICP, with the MRS PPPs being substantially higher for many developing countries including much of the Asia (East, South and West).³⁹ The Engel curve method can also deliver PPPs for the poor (or any level of income) without data on prices, as demonstrated by MRS. And that method suggests that the PPPs vary more across income groups than found by Deaton and Dupriez.

We do not yet know how stable these alternative PPPs are over time. Similarly to the concern about the geographic stability of the Engel curve, its parameters may well shift over time in ways that have little to do with prices. Similarly to past practice using ICP PPPs it is probably preferable to only update over time using domestic deflators after anchoring the country comparisons to the chosen benchmark year.

6 Conclusions

The paper has discussed three deficiencies in prevailing approaches to monitoring global poverty. The first relates to a deep, but poorly recognized, identification problem in long-standing efforts to set international poverty lines based on national lines while striving for global welfare consistency – that one judges poverty by a globally consistent welfare standard no matter where someone lives. The identification problem stems from the fact that we do not know to what extent the higher lines found in richer countries reflect social effects on welfare or more generous welfare standards for defining poverty. Thus the paper has argued that two global poverty lines are needed – a familiar lower line with fixed purchasing power across countries and a new upper line given by the poverty line that one would expect given the country's level of average income, based on how national poverty lines vary across countries. The true welfare-consistent absolute line lies somewhere between the two bounds. By this approach, to be judged “not poor” one needs to be neither absolutely poor (independently of where and when one lives) nor relatively poor (depending on where and when one lives).

The second problem is an evident disconnect between how poverty is measured in practice and the emphasis given in social policy and moral philosophy to leaving none behind. To assess whether the poorest are being left behind one needs a measure of the consumption floor. Here there is a severe data constraint, namely that a low observed consumption or income in a survey could be purely transient, and so unrepresentative of permanent consumption. However, a more reliable estimate of the consumption floor can be derived from

³⁸In what appears to be the only test of the Engel curve method to date Gibson et al. (2016) found that it performed poorly in Vietnam when compared to reliably-observed geographic price relatives. Further tests are needed.

³⁹ Both PPPs are set to unity for India in the MRS calculations.

existing measures of poverty under certain assumptions. This can be readily implemented from existing poverty data, and it provides a rather different vantage point on progress against poverty.

Finally, the paper has reviewed concerns about the current PPP exchange rates from the International Comparison Program. The days are (thankfully) gone when the community of users simply accepts without question the aggregate statistics produced by publicly-funded statistical organizations like the ICP. Recurrent debates about the ICP's results have been fueled in part by poorly-understood methodological changes and in part by the ICP's long-standing lack of openness, notably in access to primary data. Calculating PPPs for global poverty measurement using ICP price data is not exactly easy, but nor is it the hardest task imaginable as long as researchers have access to the data. There are also options to using ICP prices, although further testing is needed on their performance. Even staying with the ICP, a moving average of PPPs across ICP rounds would probably make the measures less sensitive to the likely errors in this data source – smoothing the “ICP roller coaster.” Adjustments of the PPPs will also be called for, such as to deal with urban bias in the price surveys. Going forward, better price-level comparisons for the purpose of measuring poverty, including sub-national analysis, require re-estimating the PPPs from the primary data. If the ICP is to continue to be a valuable resource, it needs to make public the primary data to facilitate such calculations. Greater central control over the data collection methods used by the regional ICP offices would also be desirable.

Each of the paper's proposals for addressing these problems could undoubtedly be improved upon and refined if there is enough agreement that effort is needed to develop better global poverty measures along these lines. That effort is justified if our global measures are to continue to have relevance in global public knowledge, and to international policy making and poverty monitoring.

Acknowledgments The author is grateful to Nora Lustig and Jacques Silber for encouraging him to write this paper and for their comments. For helpful comments the author is also grateful to Caitlin Brown, Shaohua Chen, Yuri Dikhanov, John Gibson, Stephan Klasen, Ranjan Ray and Dominique van de Walle.

References

- Ahluwalia, M.S., Carter, N.G., Chenery, H.B.: Growth and poverty in developing countries. *J. Dev. Econ.* **6**, 299–341 (1979)
- Almås, I.: International income inequality: measuring PPP bias by estimating Engel. *American Econ. Rev.* **102**(2), 1093–1117 (2012)
- Asian Development Bank: Purchasing Power Parities and Real Expenditures: A Summary Report. Manila: Asian Development Bank (2014)
- Atkinson, A.B.: Poverty in Europe. Blackwell Press, Oxford (1998)
- Atkinson, A.B., Bourguignon, F.: Poverty and inclusion from a world perspective. In: Stiglitz, J.E., Muet, P.-A. (eds.) *Governance, Equity and Global Markets*. Oxford University Press, Oxford (2001)
- Balassa, B.: The purchasing power parity doctrine: a reappraisal. *J. Politic. Econ.* **72**, 584–596 (1964)
- Banerjee, A., Duflo, E.: The economic lives of the poor. *J. Econ. Perspect.* **21**(2), 141–167 (2007)
- Bhagwati, J.: Why are services cheaper in poor countries? *Econ. J.* **94**, 279–286 (1984)
- Blank, R.M.: How to improve poverty measurement in the United States. *J. Policy Anal. Manag.* **27**(2), 233–254 (2008)
- Chen, S., Ravallion, M.: How did the world's poor fare in the 1990s? *Rev. Income Wealth* **47**(3), 283–300 (2001)
- Chen, S., Ravallion, M.: The developing world is poorer than we thought, but no less successful in the fight against poverty. *Q. J. Econ* **125**(4), 1577–1625 (2010a)

- Chen, S., Ravallion, M.: China is poorer than we thought, but no less successful in the fight against poverty. In: Anand, S., Segal, P., Stiglitz, J. (eds.) *Debates on the Measurement of Poverty*. Oxford University Press, Oxford (2010b)
- Chen, S., Ravallion, M.: More relatively poor people in a less absolutely poor world. *Rev. Income Wealth* **59**(1), 1–28 (2013)
- Citro, C., Michael, R.: *Measuring Poverty: A New Approach*. National Academy Press, Washington, DC (1995)
- Deaton, A.: Inequality and the measurement of world poverty. *Am. Econ. Rev.* **100**(1), 5–34 (2010)
- Deaton, A., Aten, B.: Trying to understand the PPPs in ICP 2011: why are the results so different? NBER Working Paper 20244. National Bureau of Economic Research, Cambridge, Mass (2014)
- Deaton, A., Dupriez, O.: Purchasing power parity exchange rates for the global poor. *Am. Econ. J. Appl. Econ.* **3**(2), 137–166 (2011)
- Duesenberry, J.S.: *Income, Saving and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, Mass (1949)
- Dutta, P., Murgai, R., Ravallion, M., van de Walle, D.: *Right-to-Work? Assessing India's Employment Guarantee Scheme in Bihar*. World Bank, Washington DC (2014)
- Dykstra, S., Kenny, C., Sandefur, J.: *Global absolute poverty fell by almost half on Tuesday*. Center for Global Development (2014)
- Easton, B.: Beware the median. *Soc. Policy Res. Cent. Newslett.* **82**, 6–7 (2002)
- Edward, P., Sumner, A.: New estimates of global poverty and inequality: how much difference do price data really make? pp. 2015–365. ECINEQ Working Paper (2015)
- Eurostat: *Income Poverty and Social Exclusion in the EU25*. Statistics in Focus 03 2005, Office of Official Publications of the European Communities, Luxembourg (2005)
- Ferreira, F., Chen, S., Dabalén, A., Dikhanov, Y., Hamadeh, N., Jolliffe, D., Narayan, A., Prydz, E., Revenga, A., Sangraula, P., Serajuddin, U., Yoshida, N.: A global count of the extreme poor in 2012: data issues, methodology and initial results. *J. Econ. Inequal.* this issue (2016)
- Fleurbaey, M., Maniquet, F.: *A Theory of Fairness and Social Welfare*. Cambridge University Press, Cambridge (2011)
- Foster, J.: Absolute versus relative poverty. *Am. Econ. Rev. Paper. Proc.* **88**(2), 335–341 (1998)
- Foster, J., Greer, J., Thorbecke, E.: A class of decomposable poverty measures. *Econometrica* **52**, 761–765 (1984)
- Frank, R.H.: *Choosing the right pond: human behavior and the quest for status*. Oxford University Press, New York (1985)
- Friedman, J.: Being in the world: globalization and localization. *Theory, Cult. Soc.* **7**(2), 311–328 (1990)
- Friedman, M.: *A Theory of the Consumption Function*. Princeton: Princeton University Press (1957)
- Fuller, C.J.: *The Camphor Flame: Popular Hinduism and Society in India*. Princeton: Princeton University Press (1992)
- Garroway, C., De Laiglesia, J.R.: On the relevance of relative poverty for developing countries. OECD Development Centre Working Paper, 314, OECS, Paris (2012)
- Geertz, C.: *The Religion of Java*. University of Chicago Press, Chicago (1976)
- Gibson, J., Le, T., Kim, B.: Prices, Engel curves and time-space deflation: impacts on poverty and inequality in Vietnam. University of Waikato. World Bank Econ. Rev. forthcoming (2016)
- Hamilton, B.: Using Engel's Law to estimate CPI bias. *Am. Econ. Rev.* **91**(3), 619–630 (2001)
- Hammond, P.: Equity, Arrow's conditions and Rawls' difference principle. *Econometrica* **44**, 793–804 (1976)
- Hansen, B.E.: Sample splitting and threshold estimation. *Econometrica* **68**, 575–603 (2000)
- Hirsch, F.: *The Social Limits to Growth*. Routledge & Kegan Paul, London (1977)
- Hunter, R.: *Poverty*. MacMillan Company, London (1904)
- Inklaar, R.: Price levels and economic growth: making sense of revisions to data on real incomes—A comment. *Rev. Income Wealth* **59**(4), 614–22 (2013)
- Inklaar, R., Rao, P.: Cross-country income levels over time: did the developing world suddenly become much richer? Working Paper 151, Growth and Development Center. University of Groninger, Germany (2014)
- Johnson, D., Smeeding, T.: A consumer's guide to interpreting various us poverty measures. Fast Focus 14, Institute for Research on Poverty University of Wisconsin (2012)
- Jolliffe, D., Prydz, E.: Global poverty goals and prices. How purchasing power parity matters. Policy Research Working Paper 7256, World Bank (2015)
- Kakwani, N., Son, H.H.: Global poverty estimates based on 2011 purchasing power parity: where should the new poverty line be drawn? *J. Econ. Inequal.* this issue (2016)
- Klasen, S., Krivobokova, T., Greb, F., Lahoti, R., Pasaribu, S., Wiesenfarth, M.: International poverty measurement: Which way now? *J. Econ. Inequal.* this issue (2016)

- Lindgren, M.: The elusive quest for the subsistence line. How much does the cost of survival vary between populations? Comparative Institutional Analysis Working Paper 2015:1. Lund University, Sweden (2015)
- Luttmer, E.: Neighbors as negatives: relative earnings and well-being. *Q. J. Econ.* **120**, 963–1002 (2005)
- Mack, J., Lansley, S.: *Poor Britain*. Allen and Unwin, London (1985)
- Majumder, A., Ray, R., Santra, S.: Preferences, purchasing power parity, and inequality analytical framework, propositions, and empirical evidence. Policy Research Working Paper 7395 World Bank (2015)
- Moatsos, M.: Global absolute poverty: behind the veil of dollars. Working Paper 77, Center for Global Economic History University of Utrecht University (2015)
- Nolan, B., Munzi, T., Smeeding, T.: Two Views of Irish Poverty Trends Background Paper to Human Development Report, vol. 2005. UNDP, New York (2005)
- Organization for Economic Co-Operation and Development (OECD): *Growing Unequal? Income Distribution and Poverty in OECD Countries*. Paris, OECD (2008)
- Rao, V.: Poverty and public celebrations in rural India. *Ann. Am. Acad. Polit. Soc. Sci.* **573**(1), 85–104 (2002)
- Ravallion, M.: On the welfarist rationale for relative poverty lines. In: Basu, K., Kanbur, R. (eds.) *Social Welfare, Moral Philosophy and Development: Essays in Honour of Amartya Sen's Seventy Fifth Birthday*. Oxford University Press, Oxford (2008)
- Ravallion, M.: A reply to Reddy and Pogge. In: Anand, S., Segal, P., Stiglitz, J. (eds.) *Debates on the Measurement of Poverty*. Oxford University Press, Oxford (2010)
- Ravallion, M.: Poverty lines across the world. In: Jefferson, P.N. (ed.) *Oxford Handbook of the Economics of Poverty*. Oxford University Press, Oxford (2012)
- Ravallion, M.: Price levels and economic growth. *Rev. Income Wealth* **59**(4), 593–613 (2013a)
- Ravallion, M.: Reply to Inklaar. *Rev. Income Wealth* **59**(4), 623–628 (2013b)
- Ravallion, M.: An emerging new form of social protection in 21st century China. In: *Oxford Companion to the Economics of China*, edited by S. Fan, R. Kanbur, S. Wei, X. Zhang. Oxford University Press, Oxford (2014a)
- Ravallion, M.: India's puzzling new PPP, Center for Global Development Blog Post (2014b)
- Ravallion, M.: An exploration of the international comparison program's new global economic landscape. NBER Working Paper, 20338 (2015)
- Ravallion, M.: Are the world's poorest being left behind? *Journal of Economic Growth*, forthcoming (2016a)
- Ravallion, M.: *The Economics of Poverty: History, Measurement and Policy*. Oxford University Press, New York and Oxford (2016b)
- Ravallion, M., Chen, S.: Weakly relative poverty. *Rev. Econ. Stat.* **93**(4), 1251–1261 (2011)
- Ravallion, M., Chen, S.: More relatively poor people in a less absolutely poor world. *Rev. Income Wealth* **59**(1), 1–28 (2013a)
- Ravallion, M., Chen, S.: A proposal for truly global poverty measures. *Glob. Policy* **4**(3), 258–265 (2013b)
- Ravallion, M., Chen, S.: Rising food prices in poor countries: a new clue to those puzzling PPP revisions. Center for Global Development Blog Post (2015)
- Ravallion, M., Datt, G., Van de Walle, D.: Quantifying absolute poverty in the developing world. *Rev. Income Wealth* **37**, 345–361 (1991)
- Ravallion, M., Chen, S., Sangraula, P.: Dollar a day revisited. *World Bank Econ. Rev.* **23**(2), 163–184 (2009)
- Rawls, J.: *A Theory of Justice*. Harvard University Press, Cambridge MA (1971)
- Reddy, S., Pogge, T.W.: Debates on the Measurement of Poverty. In: Anand, S., Segal, P., Stiglitz, J. (eds.) *Oxford University Press* (2010)
- Runciman, W.G.: *Relative Deprivation and Social Justice*. Routledge and Kegan Paul, London (1966)
- Samuelson, P.: Theoretical notes on trade problems. *Rev. Econ. Stat.* **46**, 145–154 (1964)
- Sen, A.K.: Poor, relatively speaking. *Oxf. Econ. Pap.* **35**(2), 153–69 (1983)
- Sen, A.K.: A sociological approach to the measurement of poverty: A reply to Professor Peter Townsend. *Oxford Econ. Pap.* **37**, 669–676 (1985)
- Silver, M.: The hedonic country product dummy method and quality adjustments for purchasing power parity calculations. IMF Working Paper WP/09/271, International Monetary Fund, Washington, DC (2009)
- Smith, A.: *An Inquiry into the Nature and Causes of the Wealth of Nations*. Electronic Classic Edition, Pennsylvania State University (1776)
- Summers, R., Heston, A.: The Penn World Table (Mark 5): An extended set of international comparisons, 1950–1988. *Q. J. Econ.* **106**, 327–368 (1991)
- Townsend, P.: The meaning of poverty. *Br. J. Sociol.* **13**(3), 210–227 (1962)
- Townsend, P.: *Poverty in the United Kingdom: a Survey of Household Resources and Standards of Living*. Penguin Books, Harmondsworth (1979)

- Townsend, P.: A sociological approach to the measurement of poverty: a rejoinder to Professor Amartya Sen. *Oxf. Econ. Pap.* **37**(4), 659–68 (1985)
- United Nations: The Millennium Development Goals Report. United Nations Secretariat, New York (2011)
- United Nations Development Program (UNDP): Human Development Report. UNDP, New York (2005)
- World Bank: World Development Report: Poverty. Oxford University Press, Oxford (1990)
- World Bank: Global Purchasing Power Parities and Real Expenditures. 2005 International Comparison Program, World Bank, Washington, DC (2008)
- World Bank: Purchasing Power Parities and Real Expenditures of World Economies. World Bank, Washington, DC (2014)