

Journal of Economic Inequality, forthcoming

Toward Better Global Poverty Measures

Martin Ravallion¹

Department of Economics, Georgetown University
Washington DC USA 20057
mr1185@georgetown.edu

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

JEL classifications: I32, E31, O10

¹ The author is grateful to Nora Lustig and Jacques Silber for encouraging him to write this paper. For helpful comments the author is grateful to Caitlin Brown, Shaohua Chen, Yuri Dikhanov, John Gibson, Stephan Klasen, Ranjan Ray and Dominique van de Walle.

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, including the quality of household survey data.² 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 used in global poverty measures. 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.

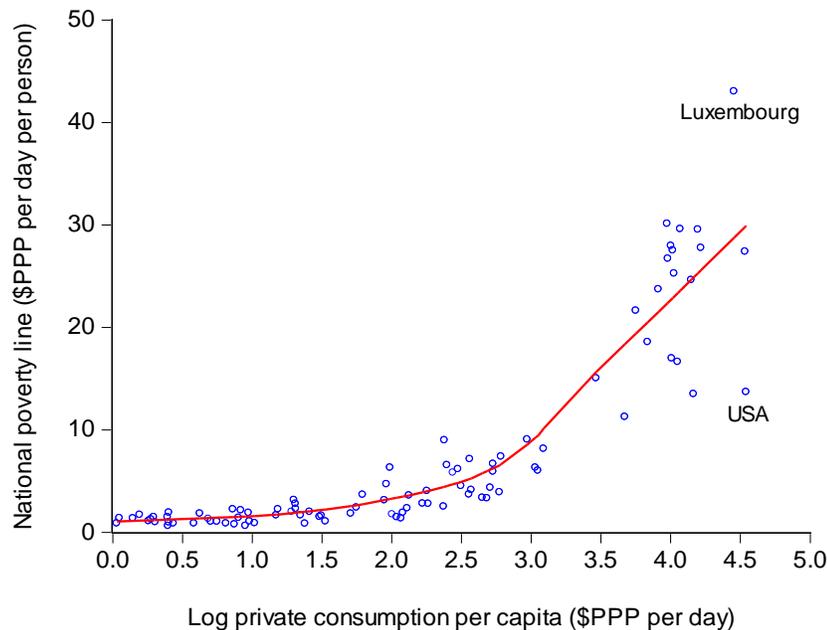
The paper discusses the three sets of problems in the above order, and suggests some possible solutions.

² Ravallion (2016, 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.

2 Taking social effects on welfare seriously

National poverty lines have long provided the data for setting global lines. Figure 1 plots the national lines for 95 countries against mean consumption. 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 \$s using the PPP exchange rates for consumption from the 2005 round of the International Comparison Program (ICP) (World Bank, 2008).⁴

Figure 1: Poverty lines across countries of the world



Source: Ravallion (2012).

The range in Figure 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 per day if one compares the empirical average line, conditional on mean

³ A 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 Figure 1. The pattern of rising lines with mean income, with an increasing slope, is still evident in this extended data set.

⁴ All 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.

consumption (as estimated by the locally smoothed scatter plot in Figure 1). 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. One sees that national lines tend to rise with mean consumption, although the relationship is clearly quite flat at low per capita consumption levels. The empirical elasticity rises from an average that is close to zero at low incomes up to unity.

In assessing poverty globally, it has long been 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 From 1990 the World Bank switched to an average line for a broader set of poor countries (World Bank, 1990; Ravallion et al., 1991). In the update by Ravallion et al. (2009) the line was set at \$1.25 at 2005 PPP, with a 95% confidence interval of (\$1.05, \$1.45).⁶ There has been huge progress over the last 30 years in reducing the incidence of global poverty by this standard (Chen and Ravallion, 2010a, 2013).

Of course, richer countries tend to have higher lines, as evident in Figure 1. No single line will ever be sufficient. 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 (Figure 1). But the key point is that whatever line is used, it is intended to be 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.

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.⁷

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

⁶ Ravallion et al. (2009) used Hansen's (2002) 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. (2015) derive the same line using different methods.

⁷ Early 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

In thinking about the implications of such social effects 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. When forming a global poverty 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).⁸ 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.⁹

(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).

⁸ Sen 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).

⁹ 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 US 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 US 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 US Census Bureau. While the US did not have an official poverty line 100 years ago, the most credible estimate at the time by Hunter (1904)

There have been various approaches to measuring relative poverty.¹⁰ 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. Most of the high-income countries represented in Figure 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 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."¹¹ 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

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).

¹⁰ Examples include Townsend (1979), Mack and Lansley (1985), Atkinson (1998), Eurostat (2005), OECD (2008) and Garroway and de Laiglesia (2012).

¹¹ Clothing 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.

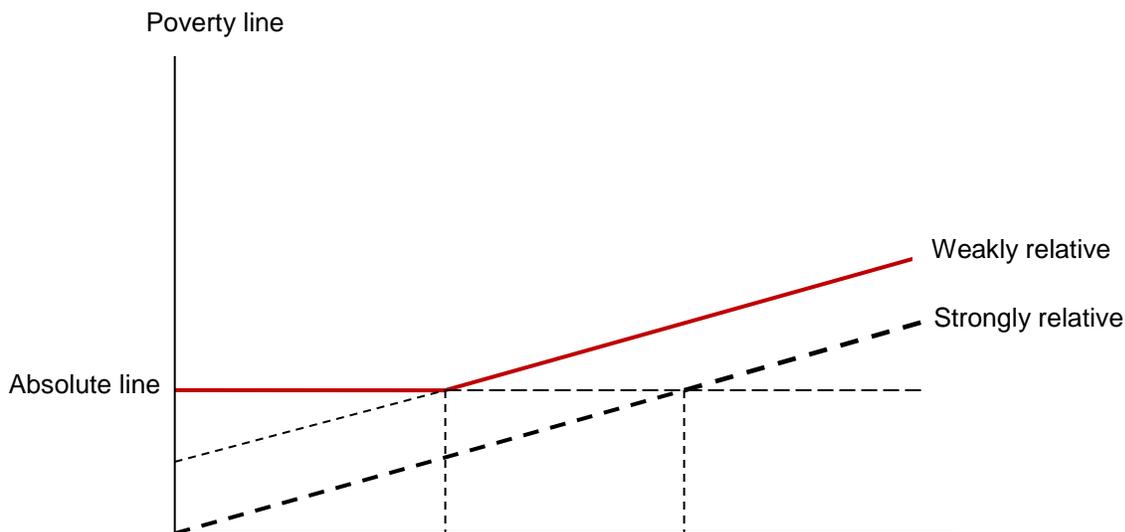
proportion of the mean, the resulting poverty measure depends solely on the distribution of relative incomes in the population, which can be loosely interpreted as “inequality.” If all income levels grow (or contract) at the same rate then the poverty measure will remain unchanged when the poverty line is set at a constant proportion of the mean or median. Seemingly perverse poverty comparisons have been found using strongly relative measures.¹²

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 (Figure 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 is implausible, and almost certainly understates the costs of social inclusion in poor countries.

These problems are avoided by the weakly relative line in Figure 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. 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 Figure 1). Chen and Ravallion find that the piece-wise linear functional form in Figure 2 gives a very good fit to the data on national lines. Ravallion and Chen (2013) provide global relative measures constructed on this basis.

¹² For example, the UNDP (2005, Box 3) (based on Nolan et al., 2005) shows how relative poverty measures for Ireland were rising despite higher absolute living standards for the poor. Similarly, Easton (2002) argues that relative poverty measures for New Zealand are deceptive in showing falling poverty despite lower absolute levels of living for the poor. 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.”

Figure 2: Strongly and weakly relative poverty lines



But are such weakly relative poverty lines welfare-consistent? This depends on why we see higher lines in richer countries, as in Figure 1. Social effects on welfare are one possible reason, but they are not the only possibility. The box outlines two models, both of which are consistent with the tendency for richer countries to have higher lines, but with very different implications for how we should measure global poverty.

If the positive gradient in poverty lines seen in Figure 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.

Box: Two theoretical models for interpreting Figure 1

Both models postulate that there are socially-specific standards for defining “poverty.” In one case the standards are set in the consumption space while in the other they are in the welfare space. Although both models can explain why national poverty lines are higher in richer countries (as in Figure 1), they have very different implications for the welfare-consistent measurement of global poverty.

Model 1: Relative deprivation as a welfare effect: There are two commodities (more can be added), consumed in amounts x_{1ij} and x_{2ij} for person i in country/date j . The derived utility depends on consumption relative to certain socially accepted thresholds (interpretable as “basic needs”) in each setting. The threshold for a given commodity rises with the overall mean, m_j , and also depends on the prices prevailing in j , denoted $p_j = (p_{1j}, p_{2j})$ (One might postulate other factors.) The thresholds for consumption are denoted $z_{1j} = z_1(m_j, p_j)$ and $z_{2j} = z_2(m_j, p_j)$ for goods 1 and 2 respectively.

Utility is $u_{ij} = u(x_{1ij}/z_{1j}, x_{2ij}/z_{2j})$, which is maximized subject to the usual budget constraint:

$y_{ij} = p_{1j}x_{1ij} + p_{2j}x_{2ij}$. The poverty line is the cost of the threshold bundle, $z_j = p_{1j}z_{1j} + p_{2j}z_{2j}$, which rises with m_j . Globally, everyone living at the poverty line can afford the relevant threshold bundle, assuming a common utility level, $\bar{u}^z \equiv u(1,1)$. Thus the lines are welfare consistent.

Model 2: A socially determined welfare threshold: In this model there are no social effects on welfare, so the utility of person i in country/date j depends on own consumption, as $u_{ij} = u(x_{1ij}, x_{2ij})$ (with the function u having standard properties), which is again maximized subject to $y_{ij} = p_{1j}x_{1ij} + p_{2j}x_{2ij}$.

As usual, the maximum attainable utility is given by the indirect utility function $v(y_{ij}, p_j)$. The other key difference with Model 1 is that the reference level of utility needed to not be considered poor is now taken to be a rising function of the mean; let this be $u_j^z = f(m_j, p_j)$. The solution of

$u_j^z = v(z_j, p_j)$ for z_j is an increasing function of m_j (as well as p_j).

Where does u_j^z come from in Model 2? We can posit a “first-best” distribution for a given mean. This can be characterized as maximizing some weighted aggregate of utilities, with the weights reflecting the government’s social preferences. (With diminishing marginal utility of income the optimum equalizes the weighted marginal utilities across people.) One can then define u_j^z as the lowest utility found in this optimal plan, with an implied z_j . The first-best is not, however, implemented given other constraints (notably on information and administrative capabilities), so the current actual distribution has incomes below z_j . In other words, the prevailing national poverty line is the minimum income in the government’s ideal distribution of income for that country and time, given the total income. Naturally, when the latter rises, the socially ideal minimum also rises.

As noted in the box, 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 is 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 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. The lower bound of the optimum (and indeed all its quantiles) will almost certainly rise with the overall mean. Thus we can expect higher monetary 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 not welfare-consistent.

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 assumed to reflect the costs of attaining a common level of welfare.¹³ The true welfare-consistent measure is somewhere between the two bounds.

3 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 per person less than this line.¹⁴ 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*.¹⁵ 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).¹⁶

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

¹³ 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 Figure 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).

¹⁴ The counts may or not be equally weighted; more sophisticated "distribution-sensitive" poverty measures give higher weight to gains to the poorest.

¹⁵ How well they do this in practice is another matter; on these schemes see Ravallion (2014) and Dutta et al. (2014).

¹⁶ 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.

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).¹⁷ 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 Figure 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 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.

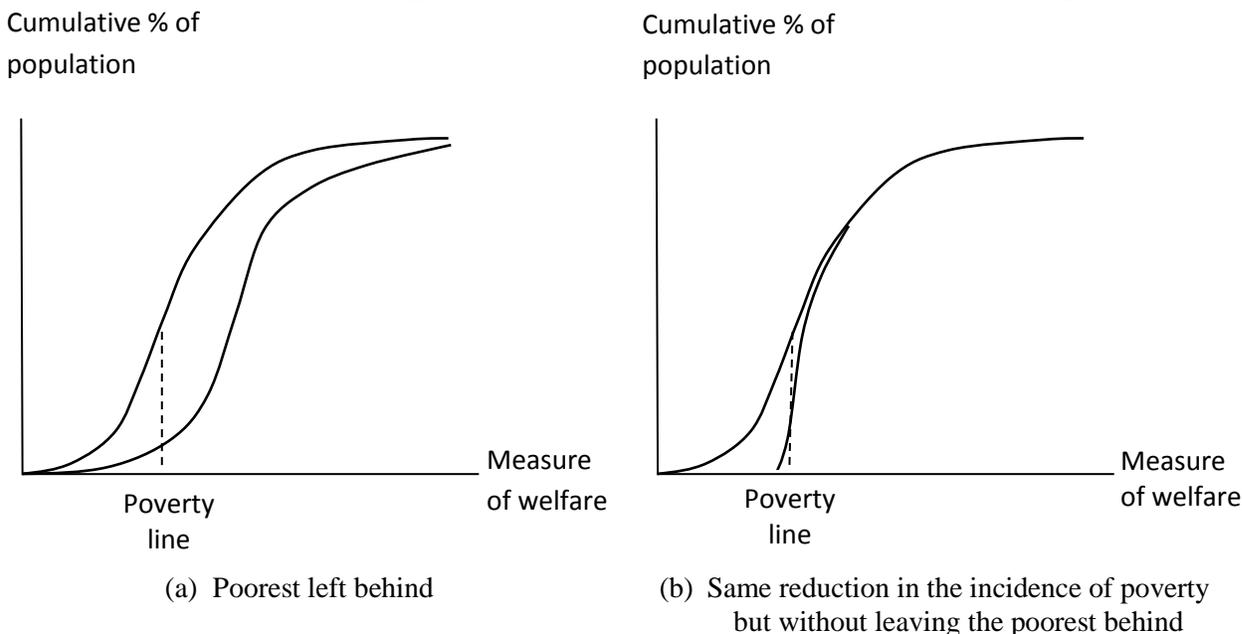
This neglect of explicit attention to the poorest is due, at least in part, to the difficulties in identifying the floor. The lowest observed level of living in a survey may differ greatly from that household’s normal level of living. For example, food consumption in a survey typically relates to the 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, 2015a). Success in assuring that no person is left behind can be monitored from existing data sources under certain assumptions. The approach recognizes that

¹⁷ 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 (2016, Ch. 3).

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.

Figure 3: Same reduction in the poverty count but different implications for the poorest



To make this approach operational, Ravallion (2015a) assumes that the probability of any observed consumption being the floor falls linearly up to an upper bound. 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.¹⁸ 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, 2015a). The world’s poorest have not been (literally) “left behind,” but they have gained disappointingly little over the last 30 years.

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.

¹⁸ This is proved in Ravallion (2015a). 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.

4 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.¹⁹

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.²⁰ 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. The ICP does not provide sub-national PPPs.

The PPPs have entered the calculation of the World Bank's global poverty measures in two ways.²¹ 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

¹⁹ See 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.

²⁰ For example, the Asian Development Bank in Manila runs the ICP for all of Asia, including South and West Asia.

²¹ This follows from the approach of basing international lines on national lines proposed by Ravallion et al. (1991).

using the best available price index for that country.²² 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. Also, 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.²³

PPP revisions and global poverty measures: The ICP's estimates from the 2011 round were released in World Bank (2014). 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 new PPPs suggest far less poverty judged relative to a poverty line with constant US purchasing power; adjusting the \$1.25 a day line in 2005 prices for US inflation gives a 2011 line of \$1.44. Using this line, by one estimate, the new PPPs imply almost half the global poverty rate for 2011 as the old PPPs (Dykstra et al., 2014). A debate ensued about the new ICP and its implications for the global economic landscape. The calculations of the impact of the new PPPs are sensitive to the level of the poverty line, as shown by Edward and Sumner (2015).

Fixing the US purchasing power of the international line over time is hard to defend given the generally higher inflation rates in developing countries.²⁴ Although there are some methodological differences, Jolliffe and Prydz (2015) and Klasen et al. (2015) have calculated international lines for 2011 that are more consistent with the methods used by Ravallion et al. (2009) to set the \$1.25 line. Depending on the assumptions made and the data used (including on national lines), the international line so obtained is around \$1.70-\$1.90. Then one finds much less change in the aggregate poverty count (somewhere between a small downward revision and

²² PPP 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.

²³ Chen and Ravallion (2010a) show that the effect of the 2005 ICP revisions on their own was to bring down the poverty count; the fact that the poverty count rose was due instead to the new and better sample of national poverty lines (a larger and more globally representative set of national lines than had been used by Ravallion et al. 1991).

²⁴ Thus, 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.

a small upward revision), although there are some large differences for individual countries. It is clear that the initial assessments suggesting a substantially lower 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. 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, such 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, 2015b, Majumder et al., 2015).²⁵ 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). The puzzles are not just for India. There has been a general downward drift in the PLIs in the 2011 ICP, undoing the ICP's prior trend in the opposite direction (Ravallion, 2015b). 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

²⁵ The 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.

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, 2015b). While the literature offers some clues, the source of the Asia drift in the PLIs for 2011 remains unclear.

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.²⁶

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

²⁶ The 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.²⁷ 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.

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.²⁸ 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.²⁹ 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. Figure 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.³⁰ 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 2011.³¹

²⁷ 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.

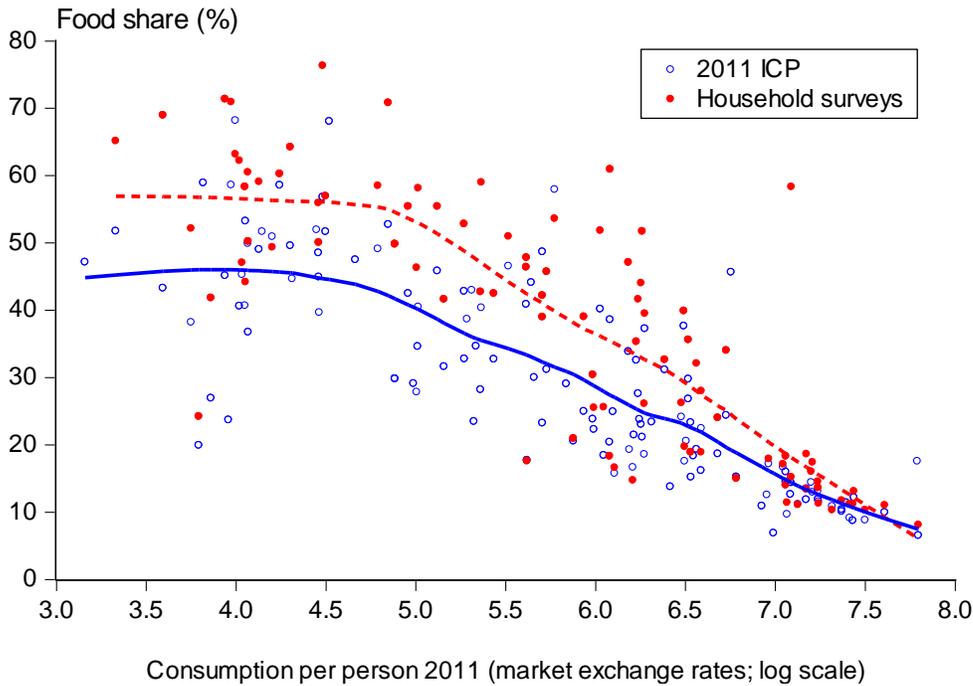
²⁸ For example, there was an urban bias in the sampling for the 2005 PPP for China, which required a correction using supplementary data on the cost-of-living differentials between urban and rural areas, as described in Chen and Ravallion (2010b). Failing to correct for this bias entailed a large over-estimation of the global poverty count.

²⁹ 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.)

³⁰ 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 downloadable data files.

³¹ The calculations were done by World Bank staff mainly using the micro data; see Ravallion and Chen (2015).

Figure 4: Two Engel curves



Source: Ravallion and Chen (2015).

We see that there is a sizeable gap between the two Engel curves in Figure 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%).

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 “PPPs 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 essentially what Kakwani and Son (2015) propose with regard to the 2011 ICP. They update the 2005 line of \$1.25 in local currency for inflation in that country to derive the 2011 line, and then they back out the 2011 line in \$'s using the PPPs from the 2011 ICP. If one stays with the \$1.25 a day line at 2005 PPP, then the value of the international line for 2011 will vary across countries. Kakwani and Son report that their “equivalent poverty line” for 2011 varies across the Asia region from \$1.02 a day in Tajikistan to \$2.36 in Fiji, even though the line is constant using 2005 PPPs. Given what we have learnt about why PPPs change independently of domestic inflation rates, it is no surprise that the equivalent poverty lines for 2011 calculated by Kakwani and Son vary so much across countries even with zero variance in 2005.

The Kakwani-Son proposal amounts to ignoring the new ICP results for the purpose of measuring global poverty. By their method, the poverty measures for each country are exactly what one gets using the 2005 PPPs and updating the poverty line in local currency using the CPI; in other words, one would have the same poverty measures found in *PovcalNet* using the 2005 ICP for the PPP conversions. Of course, if one maintains that the PPPs from the 2011 ICP round are better for setting the international line then one should not follow the Kakwani-Son proposal.

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.³² Essentially their PPPs become functions of country-specific parameters of the Engel curves relating budget shares to incomes. Against 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 differences.³³ The key assumption of a geographically stable Engle 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

³² Antecedents using empirical Engel curves to estimate price indices include Hamilton (2001) and Almås (2012).

³³ In what appears to be the only test of the Engel curve method to date Gibson et al. (2014) found that it performed poorly in Vietnam when compared to reliably-observed geographic price relatives. Further tests are needed.

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).³⁴ 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.

5 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 longstanding 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).

³⁴ Both PPPs are set to unity for India in the MRS calculations.

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 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 Purchasing Power Parity 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 longstanding 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, adjustments will 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.

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.

References

- Ahluwalia, Montek S., Nicholas G. Carter, and Hollis B. Chenery, 1979, "Growth and Poverty in Developing Countries," *Journal of Development Economics* 6:299-341.
- Almås, Ingrid, 2012, "International Income Inequality: Measuring PPP Bias by Estimating Engel Curves for Food," *American Economic Review* 102(2): 1093-1117.
- Atkinson, Anthony, 1998, *Poverty in Europe*. Oxford: Blackwell Press.
- Atkinson, Anthony, and Francois Bourguignon, 2001, "Poverty and Inclusion from a World Perspective," in J. E. Stiglitz and P-A. Muet (eds.), *Governance, Equity and Global Markets*, Oxford: Oxford University Press.
- Balassa, Bela, 1964, "The Purchasing Power Parity Doctrine: A Reappraisal," *Journal of Political Economy*, 72: 584-596.
- Banerjee, Abhijit and Esther Duflo, 2007, "The Economic Lives of the Poor," *Journal of Economic Perspectives* 21(2): 141-167.
- Bhagwati, Jagdish, 1984, "Why are Services Cheaper in Poor Countries?" *Economic Journal*, 94: 279-286.
- Blank, Rebecca M., 2008, "How to Improve Poverty Measurement in the United States," *Journal of Policy Analysis and Management* 27(2): 233-254.
- Chandy, Lawrence, and Homi Kharas, 2014, "[What Do New Price Data Mean For the Goal Of Ending Extreme Poverty?](#)" Brookings, May 5.
- Chen, Shaohua and Martin Ravallion, 2010a, "The Developing World is Poorer than we Thought, but no Less Successful in the Fight Against Poverty," *Quarterly Journal of Economics*, 125(4), pp1577-1625.
- _____ and _____, 2010b, "China is Poorer than we Thought, but no Less Successful in the Fight Against Poverty" in Sudhir Anand, Paul Segal, and Joseph Stiglitz (ed), *Debates on the Measurement of Poverty*, Oxford: Oxford University Press.
- _____ and _____, 2013, "More Relatively Poor People in a less Absolutely Poor World," *Review of Income and Wealth* 59(1): 1-28.
- Citro, Constance and Robert Michael, 1995, *Measuring Poverty: A New Approach*. Washington, DC: National Academy Press.

- Deaton, Angus, 2010, "Price Indices, Inequality and the Measurement of World Poverty," *American Economic Review* 100(1): 5-34.
- Deaton, Angus, and Bettina Aten, 2014, "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.
- Deaton, Angus and Olivier Dupriez, 2011, "Purchasing Power Parity Exchange Rates for the Global Poor," *American Economic Journal: Applied Economics* 3(2): 137-166.
- Deaton, Angus, and Alan Heston, 2010, "Understanding PPPs and PPP-Based National Accounts," *American Economic Journal: Macroeconomics* 2(4): 1-35.
- Duesenberry, James S., 1949, *Income, Saving and the Theory of Consumer Behavior*, Cambridge, Mass.: Harvard University Press.
- Dutta, Puja, Rinku Murgai, Martin Ravallion and Dominique van de Walle, 2014, *Right-to-Work? Assessing India's Employment Guarantee Scheme in Bihar*. Washington DC: World Bank.
- Dykstra, Sarah, Charles Kenny, and Justin Sandefur, 2014, "[Global Absolute Poverty fell by Almost Half on Tuesday](#)," Center for Global Development, May 2.
- Easton, Brian, 2002, "Beware the Median," *Social Policy Research Center Newsletter* 82: 6-7.
- Edward, Peter and Andy Sumner, 2015, "New Estimates of Global Poverty and Inequality: How Much Difference do Price Data Really Make?" ECINEQ Working Paper 2015-365.
- Eurostat, 2005, "Income Poverty and Social Exclusion in the EU25," *Statistics in Focus 03 2005*, Office of Official Publications of the European Communities, Luxembourg.
- Fleurbaey, Marc, and François Maniquet, 2011, *A Theory of Fairness and Social Welfare*, Cambridge: Cambridge University Press.
- Foster, James, 1998, "Absolute versus Relative Poverty," *American Economic Review, Papers and Proceedings* 88(2): 335-341.
- Foster, James, J. Greer, and Erik Thorbecke, 1984, "A Class of Decomposable Poverty Measures," *Econometrica* 52: 761-765.
- Frank, Robert H., 1985, *Choosing the Right Pond: Human Behavior and the Quest for Status*. New York: Oxford University Press.
- Friedman, Jonathan, 1990, "Being in the World: Globalization and Localization," *Theory, Culture and Society* 7(2): 311-328.

- Friedman, Milton, 1957, *A Theory of the Consumption Function*, Princeton NJ: Princeton University Press.
- Fuller, C.J., 1992, *The Camphor Flame. Popular Hinduism and Society in India*, Princeton: Princeton University Press.
- Garroway, C. and de Laiglesia, J.R., 2012, “On the Relevance of Relative Poverty for Developing Countries,” OECD Development Centre Working Paper 314, OECs, Paris.
- Geertz, Clifford, 1976, *The Religion of Java*. Chicago: University of Chicago Press.
- Gibson, John, Trinh Le and Bonggeun Kim, 2014, “Prices, Engel Curves and Time-Space Deflation: Impacts on Poverty and Inequality in Vietnam,” University of Waikato, New Zealand.
- Hamilton, Bruce, 2001, “Using Engel’s Law to Estimate CPI Bias,” *American Economic Review* 91(3): 619-630.
- Hammond, Peter, 1976, “Equity, Arrow’s Conditions and Rawls’ Difference Principle,” *Econometrica* 44: 793-804.
- Hansen, Bruce E., 2000, “Sample Splitting and Threshold Estimation,” *Econometrica* 68: 575-603.
- Hill, Robert, 2004, “Constructing Price Indexes across Space and Time: The Case of the European Union,” *American Economic Review* 94(5): 1379-1410.
- Hirsch, Fred, 1977, *The Social Limits to Growth*. London: Routledge & Kegan Paul.
- Hunter, Robert, 1904, *Poverty*, London: MacMillan Company.
- Inklaar, Robert, 2013, “Price Levels and Economic Growth: Making Sense of Revisions to Data on Real Incomes—A Comment,” *Review of Income and Wealth*, 59(4): 614-22.
- Inklaar, Robert, and Prasada Rao, 2014, “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.
- Johnson, David and Timothy Smeeding, 2012, “A Consumer’s Guide to Interpreting Various US Poverty Measures,” Fast Focus 14, Institute for Research on Poverty, University of Wisconsin.
- Jolliffe, Dean, and Espen Prydz, 2015, “Global Poverty Goals and Prices. How Purchasing Power Parity Matters.” Policy Research Working Paper 7256, World Bank.

- Kakwani, Nanak, and Hyun H. Son, 2015, "Global Poverty Estimates Based on 2011 Purchasing Power Parity: Where Should the New Poverty Line be Drawn?," *Journal of Economic Inequality*, forthcoming.
- Klasen, Stephan, Tatyana Krivobokova, Friederike Greb, Rahul Lahoti, Syamsul Pasaribu, and Manuel Wiesenfarth, 2015, "International Poverty Measurement: Which Way Now?" *Journal of Economic Inequality*, forthcoming.
- Lindgren, Mattias, 2015, "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.
- Luttmer, Erzo, 2005, "Neighbors as Negatives: Relative Earnings and Well-Being," *Quarterly Journal of Economics* 120: 963-1002.
- Mack, Joanna, and Stewart Lansley, 1985, *Poor Britain*. London: Allen and Unwin.
- Majumder, Amita, Ranjan Ray, and Sattwik Santra, 2015, "Preferences, Purchasing Power Parity, and Inequality Analytical Framework, Propositions, and Empirical Evidence," Policy Research Working Paper 7395, World Bank.
- Nolan, Brian, Tereas Munzi and Tim Smeeding, 2005, "Two Views of Irish Poverty Trends," Background Paper to *Human Development Report 2005*, New York: UNDP.
- Organization for Economic Co-Operation and Development (OECD), 2008, *Growing Unequal? Income Distribution and Poverty in OECD Countries*, Paris, OECD.
- Rao, Vijayendra, 2001, "Poverty and Public Celebrations in Rural India," *Annals of the American Academy of Political and Social Science* 573(1): 85-104.
- Ravallion, Martin, 2008, "On the Welfarist Rationale for Relative Poverty Lines," in Kaushik Basu and Ravi Kanbur (eds) *Social Welfare, Moral Philosophy and Development: Essays in Honour of Amartya Sen's Seventy Fifth Birthday*, Oxford: Oxford University Press.
- _____, 2010, "Understanding PPPs and PPP-Based National Accounts: A Comment," *American Economic Journal: Macroeconomics* 2(4): 46-52.
- _____, 2012, "Poverty Lines across the World," in P.N. Jefferson (ed), *Oxford Handbook of the Economics of Poverty*, Oxford: Oxford University Press.
- _____, 2013a, "Price Levels and Economic Growth," *Review of Income and Wealth* 59(4): 593-613.
- _____, 2013b, "Reply to Inklaar," *Review of Income and Wealth* 59(4): 623-628.

- _____, 2014a, “An Emerging New Form of Social Protection in 21st Century China,” in *Oxford Companion to the Economics of China*, edited by Shenggen Fan, Ravi Kanbur, Shang-jin Wei, Xiaobo Zhang, Oxford: Oxford University Press.
- _____, 2014b, “[India’s Puzzling New PPP](#),” Center for Global Development Blog Post.
- _____, 2015a, “Are the World’s Poorest Being Left Behind?” ECINEQ Working Paper 369.
- _____, 2015b, “An Exploration of the International Comparison Program’s New Global Economic Landscape,” NBER Working Paper 20338.
- _____, 2016, *The Economics of Poverty. History, Measurement and Policy*. New York and Oxford: Oxford University Press,
- Ravallion, Martin and Shaohua Chen, 2011, “Weakly Relative Poverty,” *Review of Economics and Statistics* 93(4): 1251-1261.
- _____ and _____, 2013, “A Proposal for Truly Global Poverty Measures,” *Global Policy*, 4(3): 258-265.
- _____ and _____, 2015, “[Rising Food Prices in Poor Countries: A New Clue to those Puzzling PPP Revisions](#),” Center for Global Development Blog Post.
- Ravallion, Martin, Gaurav Datt, and Dominique van de Walle, 1991, “Quantifying Absolute Poverty in the Developing World,” *Review of Income and Wealth* 37: 345-361.
- Ravallion, Martin, Shaohua Chen, and Prem Sangraula, 2009, “Dollar a Day Revisited,” *World Bank Economic Review* 23(2):163-184.
- Rawls, John, 1971, *A Theory of Justice*, Cambridge MA: Harvard University Press.
- Runciman, W.G., 1966, *Relative Deprivation and Social Justice*, London: Routledge and Kegan Paul.
- Samuelson, Paul, 1964, “Theoretical Notes on Trade Problems,” *Review of Economics and Statistics*, 46: 145-154.
- Sen, Amartya K., 1983, “Poor, Relatively Speaking,” *Oxford Economic Papers* 35(2): 153-69.
- _____, 1985, “A Sociological Approach to the Measurement of Poverty: A Reply to Professor Peter Townsend,” *Oxford Economic Papers* 37: 669-676.

- Silver, Mark, 2009, "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.
- Smith, Adam, 1976, *An Inquiry into the Nature And Causes of the Wealth of Nations*, London: Methuen & Co., Ltd.
- Summers, Robert, and Alan Heston, 1991, "The Penn World Table (Mark 5): An Extended Set of International Comparisons, 1950-1988," *Quarterly Journal of Economics*, 106: 327-368.
- Townsend, Peter, 1962, "The Meaning of Poverty," *British Journal of Sociology* 13(3): 210-227.
- _____, 1979, *Poverty in the United Kingdom: A Survey of Household Resources and Standards of Living*. Harmondsworth: Penguin Books.
- _____, 1985, "A Sociological Approach to the Measurement of Poverty: A Rejoinder to Professor Amartya Sen," *Oxford Economic Papers* 37(4): 659-68.
- United Nations, 2011, *The Millennium Development Goals Report*, New York: United Nations Secretariat.
- United Nations Development Program (UNDP), 2005, *Human Development Report*, New York: UNDP.
- World Bank, 1990, *World Development Report: Poverty*, Oxford: Oxford University.
- _____, 2008, *Global Purchasing Power Parities and Real Expenditures. 2005 International Comparison Program*, World Bank, Washington DC.
- _____, 2013, *World Development Indicators*, World Bank, Washington DC.
- _____, 2014, *Purchasing Power Parities and Real Expenditures of World Economies*, World Bank, Washington DC.