# Bounds on Welfare-Consistent Global Poverty Measures

Martin Ravallion

*Department of Economics*
*Georgetown University, Washington DC*

Shaohua Chen

*Development Research Group*
*World Bank, Washington DC*

**Abstract:** New measures of global poverty are presented that take seriously the idea of relative-income comparisons but also acknowledge a deep identification problem when the latent norms defining poverty vary systematically across countries. Welfare-consistent measures are shown to be bounded below by a fixed absolute line and above by weakly-relative lines derived from a theoretical model of relative-income comparisons calibrated to data on national poverty lines. Both bounds indicate falling global poverty incidence, but more slowly for the upper bound. Either way, the developing world has a higher poverty incidence but is making more progress against poverty than the developed world.

**Keywords:** Global poverty; poverty lines; relative income; inequality

**JEL classifications:** I32, O10

**Correspondence:** Martin Ravallion, Department of Economics, Georgetown University, Washington DC., 20057, USA. mr1185@georgetown.edu.

**Acknowledgments:** The authors thank Prem Sangraula and the World Bank's regional focal points for help in setting up the data base on national poverty lines and Qinghua Zhao for programming assistance. Helpful comments were received from Rebecca Blank, Benoît Decerf, Francisco Ferreira, Stephan Klasen, Dominique van de Walle and Michael Woolcock. These are the views of the authors and need not reflect those of their employers, including the World Bank.

## 1. Introduction

There is now ample support for the view that people are concerned about their relative incomes. Sociology and social psychology have long emphasized the relevance to defining poverty of concerns about shame, stigma and social exclusion.[1] Such "social effects" on welfare have also received attention in economics, including Duesenberry's (1949) model of how relative consumption influences savings, the arguments of Hirsch (1977) and Frank (1985) on how the evaluation of certain consumption goods depends on consumption relative to others, and the arguments and evidence that work effort is influenced by relative wages (Cohn, et al., 2014). The idea that welfare depends on relative income has also found support in survey data on subjective self-assessments of welfare, as in (for example) Luttmer (2005) and Knight et al. (2009).[2] And the idea has been invoked to explain the "Easterlin paradox" whereby average happiness appears not to rise much with economic growth (Easterlin, 1974; Clark et al., 2008). Economic theory has also provided a rationale for why relative income matters; for example, Rayo and Becker (2007) show that the welfare relevance of relative position can emerge as a response to the constraints faced in making choices (notably the difficulty in distinguishing close options and the boundedness of happiness). Furthermore, the literature suggests that poor people also care about relative incomes.[3]

In this light, how should we measure global poverty? An appealing guiding principle requires that poverty lines should be welfare-consistent, meaning that they are money metrics of some reasonable concept of welfare. As Sen (1983, p.168) puts it "…an absolute approach in the space of capabilities translates into a relative approach in the space of commodities."[4] Whether the absolute standard is an index of "utility" or an index of "capabilities" may be important for implementation, but the first-order issue is to demand welfare consistency in some defensible sense when measuring global poverty, i.e., those we judge to be equally well off are all either

[1] Important early contributions were made by Davis (1959) and Runciman (1966). In the context of understanding poverty see (inter alia) Abel-Smith and Townsend (1966), Townsend (1979) and Walker (2014).

[2] Surveys of this literature can be found in Frey and Stutzer (2002), Senik (2005) and Clark et al. (2008).

[3] Anthropologists have long described behaviors consistent with this idea; see, for example, Geertz (1976) and Fuller (1992). Rao (2001) describes the important of celebrations to social networks among poor people in rural India. Banerjee and Duflo (2008) document expenditures on celebrations and festivals by very poor people in surveys for a number of countries. Ravallion and Lokshin (2010) find that the poorest within a (very) poor country (Malawi) put low weight on relative position but this matters more to better-off strata. There is also evidence of adverse effects of relative position on health behaviors (Balsa et al., 2014). Smith et al (2012) provide a review of many studies showing behavioral responses to relative deprivation.

[4] Sen was commenting on the sociological approach to measuring poverty in Britain taken by Townsend (1979).

"poor" or "not-poor." The international poverty line for a given country can then be defined as the money needed to achieve a globally common level of welfare.[5]

If individual welfare depends on both own income and relative income then differences over time and space in the comparison group's level of living will require adjustments to any welfare-consistent monetary poverty line. In principle, the relative comparison might be upwards or downwards; in the former case, one is deemed to be relatively deprived if one is poorer than the average for some comparison group, while in the latter case one may be gratified in knowing that one is better off than that group. But the key point is that the income-poverty line becomes relative—specific to circumstances at each date and place.

This perspective immediately casts doubt on some of the prevailing stylized facts about poverty in the world. There is evidence that the incidence of absolute poverty—judged by a wide range of fixed real-income thresholds—is declining in the developing world, as shown in Chen and Ravallion (2004, 2010, 2013). Economic growth has played an important role, but it is less clear that this is also true when one takes account of relative income; indeed, the Easterlin paradox suggests otherwise. Is poverty also falling in growing economies when a welfare-consistent allowance is made for relative incomes? Similarly, it is widely believed that poverty is a much greater problem in the developing world than in today's rich world. Some observers have even been tempted to claim that there is really little difference between rich and poor countries in the personal experience of "poverty" once one takes account of the social effects on welfare.[6] Is that still true when one allows for relative poverty?

There are already measures in the literature that we might consider turning to in addressing these questions. Explicitly relative poverty lines appear to have been first proposed by Fuchs (1967) who argued that poverty lines for the US should be set at 50% of the current median. While not adopted officially in the US, a version of the Fuchs proposal has become the most common official method of measuring poverty in the OECD and Eurostat, and is used by many national governments in the OECD (though 60% of the median is more common than

[5] The definition of the poverty line as the point on the consumer's expenditure function corresponding to a reference level of utility needed to not be poor appears to have originated in Blackorby and Donaldson (1987). For further discussion see Ravallion (2016a, Chapter 5).

[6] For example, with reference to case studies (mainly using qualitative methods) in China, India, Norway, Pakistan, South Korea, Uganda and the United Kingdom, Walker (2014, p.14) claims that "…while material circumstances vary enormously across the case-study countries, poverty feels very similar in all settings; people simply cannot afford to live up to their own expectations and those of others."

50%).[7] The UN's Sustainable Development Goals also include monitoring the proportion of the population living below 50% of the median. Others have argued instead for using a fixed proportion of the mean rather than the median, and this too has been applied at country level, including in the UK.[8] Advocates of such relative poverty lines have often argued that the absolute lines do not keep up with evolving standards for defining poverty in growing economies. For example, Fuchs (1967, p.89) argued that ".. all so-called 'minimum' or 'subsistence' budgets are based on contemporary standards which will soon be out of date." Similar criticisms of the US official poverty lines have been made by Citro and Michael (1995) (in an expert committee report for the National Academy of Sciences) and Blank (2008), amongst others.

There is, however, a long-standing concern with the Fuchs proposal (and its variants as used by Eurostat and the OECD), stemming from the fact that the monetary line then has an elasticity of unity with respect to the median. This is dubbed a "strongly relative" poverty line by Ravallion and Chen (2011) who point out that (for a broad class of poverty measures) this violates an intuitively appealing axiom, namely that if all incomes increase (decrease) by the same proportion then an aggregate poverty measure must fall (rise); strongly relative measures turn out to have similar properties in practice to standard inequality measures.[9] This concern is probably the main reason why the Fuchs proposal has had very few followers in the developing world (or in the US). By contrast, what Ravallion and Chen call "weakly relative" lines also entail that the line rises with the mean or median, but with an elasticity less than unity.

A further issue, which has received little attention in the literature on poverty measurement, is what the comparison income should be. The literature on relative poverty has almost universally taken the comparison income to be either the (equally-weighted) mean or the median, although there has been some debate about which is better.[10] Accepting that relative

[7] Examples and discussions can be found in Fuchs (1967), Smeeding et al. (1990), Blackburn (1994), Atkinson (1998), Eurostat (2005), Nolan (2007) and OECD (2008, Chapter 5). In the context of developing countries, also see Atkinson and Bourguignon (2001) and Garroway and de Laiglesia (2012).

[8] See, for example, Drewnowski (1977), Duclos and Makdissi (2004), and de Mesnard (2007). The UK has used the mean in official poverty measures (Atkinson, 1998).

[9] A more formal discussion and evidence can be found in Ravallion (2016a, Chapter 8).

[10] Advocates of the median have argued that it is robust to measurement errors at the top and bottom while advocates of the mean have argued that using the same proportion of the median as the poverty line underestimates poverty (although there is no obvious reason why one would have to use the same proportion). A more sophisticated critique of the use of the median by de Mesnard (2007) points to some paradoxical theoretical results in poverty measurement that are avoided using the mean as the comparison income level.

comparisons are welfare-relevant does not, however, imply that the national average is the relevant comparator for global poverty measurement. Naturally there is heterogeneity in comparison groups. Research in sociology and social psychology has emphasized the role of comparisons with "similar others," also called "in-group members" as distinct from the "out-group" who are not relevant comparators (Davis, 1959). It is hardly obvious that the overall mean (or median) of the country of residence adequately characterizes the "in-group." Depending on how that group is specified (neighbors, friends, school cohort, or co-workers) one can clearly obtain quite complex formulations for a country-level measure of relative poverty. When measuring national poverty, the literature has subsumed this complexity into a single national metric. That is a seemingly reasonable simplification for the purpose of measuring poverty at the national level. But the key question is still begging: what is the relevant summary statistic for the national comparison income?

This paper revisits the conceptual basis of global poverty measurement and proposes new measures that unify the (very different) approaches taken in the past, notably between developed and developing countries. Our theoretical starting point is the assumption that welfare depends on both "own-income" and relative income, defined as the ratio of own-income to a country-specific comparison income. This provides a welfare-economic explanation for why we see higher real poverty lines in richer countries. We recognize, however, that there is a deep identification problem in using national lines to identify international relative lines, as has been done in the literature following Atkinson and Bourguignon (2001).[11] The problem is that the properties of the observed national poverty lines are consistent with two rival hypotheses, with very different implications for deriving international lines. It is one thing to believe that national lines reflect relative comparisons, but quite another to claim that they reveal the local costs of a globally common level of welfare (even when augmented to allow for measurement error and random idiosyncratic factors). That must be judged a strong assumption. The alternative interpretation is that richer countries adopt more generous reference welfare levels for defining poverty. This can generate higher lines in richer countries even without relative comparisons. Identification of a unique schedule of relative lines from cross-country variation in national lines is thus problematic, though this point has not been acknowledged in the literature.

[11] See Atkinson and Bourguignon (2001), Chen and Ravallion (2001, 2013), Ravallion and Chen (2011), and Jolliffe and Prydz (2017).

The paper makes three main contributions. The first is to formalize the aforementioned identification problem and so derive empirical bounds on the true global poverty measures so as to span the key parameter uncertainty.[12] The lower bound is an absolute line, fixed in real terms, while the upper bound is a schedule of weakly relative lines that rise with the country-specific comparison income consistently with national poverty lines. The welfare-consistent global poverty measure lies between these bounds, depending on how much the latent reference welfare level for defining poverty at the national level rises with the mean.

The second contribution concerns how the comparison income should be set, as required for the upper bound. Here our main point of departure from past work is that we take account of the bearing that inequality has on relativist comparisons. We question the long-standing assumption that the comparison income level in relativist comparisons at country level is the median or equally-weighted mean.[13] It is well recognized that the mean may be too heavily influenced by very high incomes, which are probably less relevant to the relativist comparisons that are likely to be made by most people, who know little about how rich the rich are. As Duesenberry (1949) recognized, it is probably not relative income that matters but relative (observable) consumption. Nor is the median a satisfactory fix. While concerns about measurement errors at the extremes are real, there is still ample information in the data, and it is far from obvious that such information should be entirely discounted.[14] We argue that a better approach is to postulate that, while the relativist comparison may put lower weight on richer people, it will never put zero weight on the rich, as is the case with the median.[15] We provide a simple theoretical formulation that encompasses bow upward and downward relative comparisons. This provides a new perspective on measuring relative poverty.

The third contribution is to provide new data on national poverty lines and survey-based distributions of consumption or income to implement the above ideas empirically. Our data on national poverty lines suggest that the rank-weighted mean is the relevant comparison income,

---

[12] Chen and Ravallion (2013) note in passing that one might interpret absolute and relative lines as lower and upper bounds but they do not discuss the identification problem that motivates this interpretation.
[13] While our focus here is on global poverty, it can also be noted that studies of the effects of relative income on subjective welfare have also relied at times on equally-weighted means, as in (for example) Hagenaars and van Praag (1985) and Luttmer (2005).
[14] The same point can be made about the use of a fixed proportion of any quantile corresponding to a fixed percentile. For example, Citro and Michael (1995) recommend using the 33rd percentile of the distribution of consumer spending on food, clothing, shelter and utilities. This idea was adopted in 2011 by the US Census Bureau's Supplementary Poverty Measure, which we return to.
[15] Note that the median is unresponsive to small changes in incomes sufficiently far above (or below) the median.

with lowest weight given to the richest. This implies that a Gini-discounted mean is called for in setting our upper bound. We implement the new measures for the lower and upper bounds on a global basis, including countries at all levels of development. Thus we provide globally-unified measures of poverty, in contrast to past work which has been bifurcated between "rich" and "poor" countries, with two distinct literatures. Our estimates draw on 1500 household surveys for 150 countries over 1990-2013.

The following section discusses our data on national poverty lines and some of their properties, as relevant to the rest of the paper. Section 3 reviews the measurement practices found in the literature. The paper's main contributions are found in Sections 4-6. Section 4 outlines our theoretical approach to measuring relative poverty. In accounting for how national poverty lines vary across countries, we then show in Section 5 that a weakly-relative poverty measure using a Gini-discounted mean dominates both strongly and weakly-relative measures using either the ordinary mean or the median. We find that higher inequality calls for a lower national comparison mean, but that a higher share of that mean should be passed onto the poverty line. The net effect is generally a higher national line than implied by standard (strongly) relative measures, most notably in poor countries. Finally, Section 6 provides our new estimates of global poverty measures. We find that the aforementioned stylized facts about global poverty—that it is falling and that poverty is a greater problem in the developing world—are robust to taking relative income seriously. Some new insights also emerge, including that the rich world is making far less progress against poverty. Section 7 concludes.

## 2. National poverty lines

National poverty lines have long provided the data used in setting global lines. In assessing poverty globally, the World Bank has 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 (World Bank, 1990; Ravallion et al., 1991). Two people with the same real consumption are treated the same way no matter where they live. Ravallion et al. (2009) compiled a sample of national lines, including 75 observations for developing countries. On this basis they set a line of $1.25 at 2005 PPP, which became the new international line for the World Bank. This was the mean poverty line of the poorest 15 countries in terms of consumption per capita. On allowing for the rates of price inflation in the set of national poverty lines used in deriving the $1.25 international line, Ferreira et al. (2016) updated the $1.25 line to $1.90 a day at 2011 PPP. While there has been some debate about the $1.90 line (see, for example, Klasen et al., 2016), it has since become widely accepted in the development community, as exemplified by its adoption in the UN's *Sustainable Development Goals*.

National lines have also been used to set international relative poverty lines (Atkinson and Bourguignon, 2001; Chen and Ravallion, 2001, 2011, 2013; Jolliffe and Prydz, 2017). While for a number of the OECD countries the national lines are directly proportional to the mean or median that is not true of most countries in the world. The methods of setting poverty lines vary, with numerous free parameters, including nutritional requirements, the composition of the food bundles and the allowances made for non-food spending. Through their parameterization at country level, national lines can be interpreted as social subjective lines that reflect prevailing concepts of what "poverty" means in each country.[16] It is then reasonable to expect that the variation in national lines across countries reflects differences in the comparison income.

We have compiled a new data set of 145 national poverty lines. (A Statistical Annex is available describing the data sources.) This has entailed an extra 47 developing (non-OECD) countries on top of those used by Ravallion et al. (2009) as well as 24 OECD countries (not included in Ravallion et al., 2009).[17] For the developing countries, these are official national poverty lines or (when these could not be found) they are the lines set by the World Bank, as part of its analytic work at country level. For the US we have used the official poverty line. For the rest of the OECD countries we have used 60% of the per-capita median, though we also test sensitivity to using 50% of the median. Both the poverty lines and consumption levels are converted to per capita $US values using the PPP exchange rates for consumption from the 2011 ICP (World Bank, 2015).[18] The survey dates range from 2004 to 2012, with a median of 2011.

Figure 1 gives density functions for the poverty lines, survey means and medians. The skewness evident in Figure 1 is as one would expect. The poverty lines are skewed further to the

---

[16] The social subjective line is the level of income below which people in a specific social context tend to judge themselves as "poor" but above which they tend to see themselves as not poor. For further discussion and references see Ravallion (2016a, Chapter 4).

[17] Some countries also have national cut-off lines for means-tested social assistance. These are not strictly poverty lines so we chose not to include them.

[18] 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 2011 prices using the country's consumer price index (or the most appropriate index available), and then converted to PPP $'s using the 2011 PPP for consumption. When poverty lines are quoted as "per equivalent adult" (mainly OECD) we have re-scaled to "per capita" units by multiplying by the ratio of mean equivalent adults per household to mean household size.

right than the medians, which are skewed further than the means. The range in national poverty lines is large, from $0.69 to $36 per day. The overall mean line is $7.82 per person per day (s.e.=$0.68; n=146). (For the non-OECD countries the mean is $4.71 ($0.32; n=122).) The median is $4.38 and the mode is $3. While the World Bank's $1.90 line is well below the mode, it is clearly in a fairly dense part of the distribution (Figure 1). If we construct a band around the Bank's $1.90 line of (say) $1.80-$2.00 we find four countries (with their poverty lines): India ($1.82), Indonesia ($1.88), Ethiopia ($1.99) and Nepal ($2.00). The World Bank's international line is approximately Indonesia's line. China's national line is slightly above this group, at $2.29.

Recall that the Bank's $1.90 line is an update to 2011 prices of the $1.25 line proposed by Ravallion et al. (2009). In our new data set the mean poverty line of the poorest 15 countries in terms of the survey mean is $1.67 at 2011 PPP, slightly below the Bank's line. But one would not want to make too much of this difference. The $1.90 line is the mean for a somewhat larger group of countries, which could be considered justified by the fact that we have a larger data set of national lines than used by Ravallion et al. (2009). If one focuses instead on the poorest 25 (about the same proportion of the 122 non-OECD countries) then the mean national line is $1.91, almost exactly the Bank's 2011 line.

We do not, of course, have national lines for all country-year combinations; indeed, our 145 national lines account for only 10% of the number of estimates we will require of national poverty measures by date. So predicted values are needed to obtain a complete set of lines for our upper bound. In past work the (equally-weighted) mean has been the main predictor.

We will introduce a more general formulation of the comparison income later, but for the present descriptive purpose we also focus on the relationship between the national lines and the survey mean. Figure 2(a) plots the data for the full sample (including OECD).[19] Figure 2(b) gives the lines for the non-OECD countries but using instead a log scale for the mean to avoid the bunching up at low levels evident in Figure 2(a). Most countries are also identified. Of course there are comparability problems and measurement errors in the national lines. But the pattern is clear: national lines tend to rise with the overall mean. For example, while the mean for the poorest 15 countries is $1.67, for the richest 15 it is 20 times higher at $27 a day. The slope of

[19] These are mostly consumption means for developing countries, and mostly income means for OECD countries. However, this does not make any difference in the relationship (on adding a control variable for the type of survey).

the regression line is 0.485 (White s.e.=0.020).[20] The overall elasticity (using a log-log regression) is 0.863 (s.e.=0.027). (If we use the median instead of the mean, the slope is 0.564 (0.017) and the elasticity is 0.816 (0.026).) It is also notable that there is little sign of a "flat" segment at low means; the relationship is positive from the lower bound.

It might be argued that the true causal relationship is not as strong as Figure 2 suggests. Three concerns can be noted. First, the fact that some of the national lines are strongly relative lines is likely to be biasing the relationship. However, the relationship is still evident if one drops the OECD countries, though the slope falls slightly, to 0.454 (s.e.=0.039) while the elasticity falls to 0.773 (0.044). (Using the median instead for the non-OECD countries, the slope is 0.559 (0.042) and the elasticity is 0.740 (0.040).)

Second, a bias due to correlated measurement errors in the mean and poverty line might remain given that the national lines for developing countries are often calibrated to survey data (though the direction of bias is ambiguous in theory, noting that there is also the usual attenuation bias). For example, one method of setting national poverty lines identifies the poverty line as the total consumption expenditure level at which pre-determined food energy requirements are met in expectation.[21] Then, for fixed requirements, over (under) estimation of total expenditure will lead to an over (under) estimation of the poverty line. This is also likely using food Engel curves to set the non-food component of the poverty line. Acknowledging this concern, as a further check we used per capita private consumption expenditure (PCE) from the national accounts as the instrumental variable (IV) for the survey mean, under the assumption that the measurement errors in these two data sources are uncorrelated. That assumption can be questioned, although it should be noted that the national accounts in most developing countries are not calibrated to household surveys. (Consumption is generally derived as a residual after subtracting recorded sources of domestic absorption at the commodity level.) The IV estimate of the slope is 0.471 (0.026) for the full sample and 0.425 (0.043) for the non-OECD sub-sample. Using log PCE as the IV the estimated elasticity is 0.844 (0.030) and 0.744 (0.051) for the non-OECD sample. So (again) this does not suggest there is anything but a small bias in the relationship seen in the raw data in Figure 2.

[20] All standard errors of regression coefficients in this paper are corrected for a general form of heteroscedasticity using White's (1980) method.
[21] For a review of the methods used to set national poverty lines see Ravallion (2012).

Third, there may be omitted country effects correlated with mean income. An alternative method of deriving national poverty lines is to find the lines that are implicit in data on the poverty rate. Using fitted distributions, Jolliffe and Prydz (2016) estimate over 600 national poverty lines this way, as implicit in national poverty measures from the World Bank's *World Development Indicators* (such as World Bank, 2013).[22] The advantage of this method is that it generates multiple lines for each country, so we can add country fixed effects. Ravallion (2016b, Appendix) estimates the elasticity of the poverty line to the mean allowing for country effects and finds an OLS elasticity of 0.52 (s.e.=0.04; n=598). Without the country effects the elasticity is 0.74 (0.01; n=609). So the elasticity is lower when we allow country effects, but it remains positive and statistically significant. However, it should be noted that the aforementioned issue of correlated measurement errors is likely to be a greater concern for these implicit poverty lines, as argued by Ravallion (2016b).

Notice that the US is an outlier in Figure 2(a). The official poverty line for the US was $15.62 per person per day in 2011 (for two adults and two children). This is well below the line one would expect for a country with the US mean. Indeed, the US line is more typical of countries with about half the US mean (around the values expected in developing countries with the highest means). As noted, the US has been an exception to the otherwise common usage of strongly relative poverty measures in rich countries. Instead, the official US line (set by Orshansky, 1965) has only been adjusted for inflation over time, such that it has fallen relative to the mean and median. This has been a source of concern in the literature on poverty in the US, which has generally taken the view that the US line should have risen in real terms to better reflect rising overall living standards.[23] Proposed revisions to the official US line have met political resistance stemming from the fact that certain public spending allocations across programs and states depend in part on the official poverty rates (Blank, 2008). (Such political resistance to updating poverty lines is clearly not unique to the US.)

It remains that, over the longer term, poverty has clearly been relative in the US. While the official US poverty line has been held fixed in real terms since the mid-1960s, if one goes back to the literature on poverty measures for the US in the early 20th century one finds much

---

[22] Letting $F_{it}(.)$ denote the fitted cumulative distribution function for country *i* at date *t* and the observed headcount index as $H_{it}$, the implicit poverty line is $F_{it}^{-1}(H_{it})$.

[23] See the discussions in Citro and Michael (1995) and Blank (2008).

lower real lines—indeed, a (non-official) line that is roughly comparable with prevailing poverty lines today in the world's poorest countries.[24] A new "supplementary" poverty line was introduced by the US Census Bureau in 2011 that explicitly acknowledges the relevance of relative poverty in the US (Short, 2012). The next section will return to this new measure.

It is probably no surprise to readers that we see higher real lines in richer countries as evident in Figure 2. In identifying who is considered "poor" within its borders, a rich country tends to use a more generous allowance—just as one finds in survey data on individual perceptions of poverty.[25] The food bundles are almost always anchored to stipulated nutritional requirements, although these vary, with higher mean requirements in places and times with better nourished populations and often with higher activity levels. The food menus identified in practice for attaining given requirements also vary greatly, and are typically more generous (such as with larger allowances for protein and more diversified diets) in less poor places. Past research has also found that a large share of the mean-income gradient in national poverty lines is due to more generous allowances for non-food needs in richer countries (Ravallion et al., 2009). However, these observations can be interpreted in two very different ways: either a line with higher purchasing power is needed to attain the same level of welfare in a rich country as a poor one, or richer countries use a higher welfare threshold in defining poverty.

It is also notable that there is a positive intercept in Figure 2. This pattern seems intuitively plausible, as it is unlikely that the poverty lines used by countries could fall to zero in the limit as mean consumption falls to its lowest level. Using the non-OECD sample, the predicted poverty line based on a linear projection is $0.96 (s.e. = $0.25) for the country with the lowest mean, which is $0.76, for the Democratic Republic of the Congo (DRC). The DRC has an unusually low mean (Figure 2(b)). If one uses the country with the next lowest mean, Madagascar with a mean of $1.45, the predicted poverty line is $1.28.

So these data are more suggestive of weakly-relative lines, with an elasticity less than unity, but rising with the mean; using the linear projection for non-OECD countries, at the lowest mean consumption the elasticity is 0.36 (s.e.=0.12) while it approaches unity in high-income

[24] While the US did not have an official poverty line 100 years ago, the most credible and widely-cited 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, 2016a, Chapter 1). Kilpatrick (1973) found evidence that the mean subjective poverty line in the US (based on survey data) rose over time with average income with an elasticity of around 0.6. Also see the discussion in Blank (2008).
[25] For a survey of the literature see Ravallion (2016a, Part 2).

countries. Naturally then, as the mean rises, the ratio of the poverty line to the mean tends to fall, as can be seen in Figure 3 (using a log scale for the mean, to make the graph easier to read). The poverty lines tend to be roughly equal to the mean among the lowest-income countries (Figure 2(a)). Thus, for the poorest countries (lowest mean), a very high proportion of the population would live at or near the national line even with no inequality.

## 3. Relative poverty lines in past practice

As discussed in the introduction, our key guiding premise in formulating global poverty measures is that the required international comparisons of welfare must be anchored to a defensible and common concept of individual welfare. To the extent feasible with the data available, everyone's poverty status must be judged by a consistent welfare concept. We can dub this "welfare consistency."

From this perspective, all current practices are questionable. The welfare relevance of relative income implies that absolute lines in the income space do not correspond to a common level of welfare. 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. This has happened in the rich world over the last 100 years (including in the US as we have noted) and in recent times in growing developing countries including China, India, Indonesia and Vietnam.[26]

***Strongly-relative lines:*** The most common approach to measuring relative poverty is exemplified by the relative poverty measures which compare each household's observed income to a poverty line that is set at a constant proportion of the current median for the country of that household's residence. This poverty line can be written in the generic form:

$$z = k.y(p_z) \quad (1)$$

Here $z$ is the poverty line, $k$ is a constant, $y(.)$ is the quantile function (inverse of the cumulative distribution function, which is assumed to be continuous and monotonic increasing) and $p_z$ is a fixed percentile that defines the comparison group. In the case of the original Fuchs (1967) proposal, $k = p_z = 0.5$, although other parameter values have been used since, as noted.

[26] 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). Indonesia's official lines for a given year are anchored to the average consumption bundle of the 20% living above the previous year's line. Jolliffe and Prydz (2016) point to other examples of developing countries that have increased the real value of their national lines.

It is not clear why the quantile of any fixed percentile identifies a plausible comparison income. Why would incomes above or below this quantile not get a positive weight? The US is an interesting case. The new "Supplementary Poverty Measure" (SPM) produced by the US Census Bureau acknowledges past concerns that the US official poverty line has not been updated in real terms (Short, 2012). The SPM determines the poverty line by the quantile of the 33rd percentile of the distribution of a subset of consumption spending deemed to be "essential" (comprising food, clothing, shelter and utilities).[27] Thus the SPM sets $k$=1.2 and $p_z = 0.33$ in equation (1). However, it remains unclear why $y(p_z)$ is a plausible comparison income for any fixed $p_z$ (whether 0.5 or 0.33). In the case of the SPM it is also unclear why relative comparisons would only apply to "essential" goods. One can surely expect feelings of relative deprivation to respond as much to a lack of "non-essential" goods.

A further concern arises when the poverty line is set at a constant proportion of the mean or median, namely that the resulting poverty measure depends solely on the distribution of relative incomes in the population. 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.[28] Seemingly perverse poverty comparisons have been found using strongly relative measures.[29]

The relevance of strongly relative measures to developing countries is especially questionable. 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, which is somewhere around a survival level (Lindgren, 2015; Ravallion, 2016b). The value for the country with the lowest mean would be only $0.38 per day, which is almost certainly not enough for survival beyond a short time. Similarly, the Garroway and de Laiglesia (2012) relative line, set at 50% of the current median in each country, gives lines that are well below the lines typical of even low-income countries and even below likely biological minima (Ravallion, 2016b).

In short, strongly relative measures almost certainly understate the social inclusion needs of globally poor people and have a seemingly perverse implication for how these measures

---

[27] This follows the recommendation of a National Academy of Sciences Commission (Citro and Michael, 1995).
[28] Note that this property does not depend on whether the line is anchored to the mean or the median, given that the ratio of the median to the mean is constant in an inequality-neutral growth process. However, the choice between the mean and median can matter in other respects and objections to the use of the median have been identified by de Mesnard (2007) and Kampke (2010). We do not discuss these issues here.
[29] See, for example, the UNDP (2005, Box 3) and Easton (2002).

respond to economic growth and contraction. There is a quick fix for these problems, namely to add a positive intercept to (1). But this brings its own problems as we will see next.

***Weakly-relative lines:*** Kakwani (1986) proposed a poverty line of the form:

$$z = z_0 + \beta(m - z_0) \qquad (2)$$

where $z_0$ is the absolute line, which is taken to be given, *m* is the overall mean or median and $\beta$ is a parameter. If $0 < \beta < 1$ then the elasticity of the Kakwani poverty line w.r.t. *m* is positive but strictly less than unity; the limit is unity as *m* goes to infinity. Chakravarty et al. (2015) provide an axiomatic derivation for a line of the form in (2). Jolliffe and Prydz (2017) use a schedule of lines with the same form, which generalizes the Garroway and de Laiglesia (2012) proposal for developing countries to allow a positive intercept, thus making it weakly relative.

There are two concerns with (2). The first arises when we take it to data on national poverty, which are either absolute or strongly relative. Yet (2) is not a hybrid of absolute and relative lines. That would require an extra parameter, to deliver $z = (1-w)z_0 + wkm$ for $0 \leq w \leq 1$ and $0 < k < 1$, where $km$ is the strongly relative line with weight *w*. Working from this modification, we can fix $z_0$ exogenously and back out estimates of *w* and *k* from the data when valid solutions exist. Setting $z_0 = \$1.90$, the data in Figure 2(a) yield $\hat{w} = 0.791$ (s.e.=0.122) and $\hat{k} = 0.612\,(0.075)$. (One might also expect these parameters to vary; if one drops the OECD countries one finds that $\hat{w} = 0.675$ (s.e.=0.143) and $\hat{k} = 0.673(0.095)$.) However, valid solutions do not exist using the median as the comparison income, which would require values of $k > 1$.[30] So an internally consistent schedule of poverty lines linear in the median cannot be derived from these data when one uses the World Bank's absolute line of \$1.90 a day.

A second concern arises when $m < z_0$, as the poverty line implied by (2) is then lower than $z_0$, which is a logical contradiction. And we cannot rule out the possibility that $m < z_0$. Indeed, we will see in section 5 that fitting the linear schedule in (1) to our data on national lines yields predicted lines for the poorest countries that are well below $z_0 = \$1.90$ (at 2011 PPP).

---

[30] Using the median one obtains an unconstrained $\hat{k} = 1.008\,(0.144)$ on the full sample and $\hat{k} = 1.021\,(0.175)$ on the non-OECD sample.

The problem is more common using the median; while our data indicate that only three countries have a survey mean less than \$1.90 a day, 18 countries have a median less than \$1.90.[31]

There are other examples of weakly-relative measures in the literature. Foster's (1998) "hybrid line" is 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 (as in Figure 2(a)). 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 rising above some point. However, it has the undesirable feature that the relative component goes to zero at zero mean. This almost certainly understates the costs of social inclusion in poor countries.

A schedule of weakly-relative lines that avoids the aforementioned problems is the piece-wise linear form:

$$z = \max(z_0, \ \alpha + \beta m) \tag{3}$$

where $\alpha \geq 0$. By construction, the hybrid line can never be below the absolute line. This is the form used by Ravallion and Chen (2011, 2013).[32] Following Atkinson and Bourguignon (2001) this hybrid line can be thought of as combining a capability for "subsistence" with country-specific requirements for a "social inclusion". Our formulation in (3) is a generalization of the Atkinson-Bourguignon proposal; the generalization is to add a parameter, $\alpha$, which can be interpreted as the lower bound to social-inclusion needs; the Atkinson-Bourguignon lines are the special case with $\alpha = 0$. As we will see below, this generalization is crucial to welfare-economic interpretation of global poverty measures.

Neither the strongly nor weakly-relative measures described above are globally monotonic in own income, meaning that when comparing any two people (wherever they may live) the one with the higher income cannot have higher measured poverty. While monotonicity holds within countries, it need not hold between them. So it is possible that a person who is absolutely poor is deemed less poor than someone who is only relatively poor, as noted by Decerf (2017). However, as we will show in the following section, as soon as one is explicit about the welfare-theoretic foundation of the poverty measure this concern vanishes.

[31] Using the Jolliffe-Prydz schedule based on (2) but using the median we find in our data set that 15 countries have a poverty line less than \$1.90.

[32] Other examples can be found in Budlender et al. (2015) and Jolliffe and Prydz (2017) (who also consider the piece-wise linear form as an option to the linear form in (2)).

All the approaches above—in a literature spanning 50 years—either explicitly or implicitly treat the mean or median as the comparison income in setting the relative line. As we will argue next, this is questionable if one thinks further about the nature of such comparisons.

## 4. Welfare-consistent relative poverty lines

The essential idea of "relative comparison" is that individual economic welfare depends (at least in part) on how the individual is doing relative to a set of comparators in society. Welfare-consistent weakly relative lines at the national level can then be rationalized by postulating an individual welfare function for household *i* in country *j* of the form:

$$u_{ij} = u(y_{ij}, y_{ij} / m_{ij}^{*}) \tag{4}$$

where $y_{ij}$ is the individual's own consumption and $m_{ij}^{*}$ is the individual's <u>comparison income</u>. The welfare function is assumed to be strictly and smoothly increasing in both arguments ( $u'_y > 0$ and $u'_{y/m} > 0$ in obvious notation). If $y_{ij} < m_{ij}^{*}$ then person *i* can be said to experience "relative deprivation," while if $y_{ij} \geq m_{ij}^{*}$ then she experiences "relative gratification" (interpreting the distinction made by Davis, 1959).

A welfare function such as (4) can be readily used to motivate relative poverty measures. In the literature, $m_{ij}^{*}$ is assumed to be either the mean or median consumption or income for the country and date of residence. We relax this assumption shortly, but it is of interest to briefly work through its implications. If we take the comparison income to be the mean ($m_j$), such that $m_{ij}^{*} = m_j$ for all *i*, then the welfare-consistent international poverty line, $z_j$, is defined by:

$$u(z_j, z_j / m_j) = \bar{u}_j^{z} \tag{5}$$

This gives $z_j$ as (implicitly) an increasing function of $m_j$ for given $\bar{u}_j^{z}$, which is the fixed level of welfare to not be deemed poor in country *j*. It is clear then that $y_{ij} < z_j$ implies (and is implied by) $u_{ij} < \bar{u}_j^{z}$. A globally-welfare consistent poverty line can be defined as one based on a constant welfare level, $u^{z}$ for all $j$.

Notice that the poverty line defined by (5) will never be strongly relative given that the welfare function is strictly increasing in own consumption at given relative consumption, which seems a very reasonable assumption. The implicit welfare-consistent line will rise with the mean, with a positive elasticity less than unity. Strongly relative lines only emerge as the limiting case in which $u'_y$ goes to zero, such that welfare depends solely on relative income. (Note that if (4) can be written as $u_{ij} = \tilde{u}(y_{ij} / m_j)$ then the welfare-consistent poverty line takes the form $z_j = k_j m_j$ where $k_j \equiv \tilde{u}^{-1}(\bar{u}_j^z)$.)

Also notice that globally-consistent poverty measures based on the above formulation need not be globally monotonic in $y_{ij}$ for those deemed to be poor, as noted in the previous section. This is a moot point, however, given that $y_{ij}$ is not a valid money-metric of welfare when relative income matters. A more appealing property in this context is monotonicity in the individual equivalent income, $y_{ij}^e$, defined implicitly by $u(y_{ij}^e, y_{ij}^e/\bar{m}) = u(y_{ij}, y_{ij}/m_j)$ for some globally constant reference mean $\bar{m}$.[33] This is assured for a broad class of global poverty measures.[34]

The question is still begging: Are all income levels in society equally important in relative comparisons? The literature in economics has said rather little about the appropriate comparison group in discussing relative poverty.[35] The assumption of an (equally-weighted) mean or the median is almost universal in this literature. When forming the comparators for deciding whether a person is relatively deprived one might not want to put equal weight on the richest stratum as the poor or middle class. Indeed, Duesenberry's (1949) original formulation of the relative-income hypothesis postulated an un-equally weighted mean.

When we allow the weights to vary by level of income, the extent of inequality can influence the level of the reference income used for relative comparisons. Suppose that the poor and middle class are the more relevant comparators for most people. Higher inequality suggests that this reference group is relatively poorer, implying a higher relative income at given own income. The use of the median as the reference is one response to the concern that the rich get too high a weight in the mean. However (as noted), while we might agree that the rich are less

---

[33] If we also impose homotheticity, such that $u(y_{ij}, y_{ij}/m_j)$ is linear in $y_{ij}$, then $y_{ij}^e/z = y_{ij}/z_j$ (the "welfare ratio" in Blackorby and Donaldson, 1987). We will not require this property for our analysis.

[34] This is the class of measures whereby individual poverty can be defined as $p(y_{ij}^e)$ with $p = 0$ for $y_{ij}^e \geq z$ (where $z$ is the poverty line corresponding to $\bar{m}$) and monotonically decreasing when $y_{ij}^e < z$ ; the global measure is a population-weighted aggregate of this individual measure. This holds for the entire class of additive measures characterized by Atkinson (1987).

[35] As noted by Chen (2015). The comparison group has received somewhat more attention in the literature on subjective welfare following Clark and Oswald (1996); also see the survey in Clark et al. (2008).

relevant comparators, it surely cannot be plausible that they are irrelevant as comparators. Against this view, it might be argued that relativist comparisons tend to be more "upward looking"—that the comparators for the poor are the middle class, and for the latter, the rich. Then the argument reverses, with higher inequality requiring a higher poverty line.[36]

We propose an approach that encompasses both "downward" and "upward" looking relativist comparisons, motivated by the following thought experiment. In keeping with the fact that we are measuring poverty at the country level, we follow the literature in postulating a common comparison income within a given country. (In principle our approach could be applied at a more disaggregated sub-national level, but that is not the present application.) To allow for either downward or upward comparisons, one can imagine a person making a random draw of a pair of incomes in the country of residence, so as to assess how she is doing relative to others.[37] Naturally, she focuses more on the lower (upper) income if she makes downward (upward) comparisons. More generally one can imagine that she picks a comparisons point somewhere in the (closed) interval between the two incomes, depending on whether the observer tends to look upward or downward. To formalize this idea, let $\phi(y_{kj}, y_{lj})$ denote the contribution of the (*k*, *l*) pair drawn in country *j* to the assessment of the comparison mean for that country. We assume that $\phi(y_{kj}, y_{lj})$ is a point somewhere in the closed interval $[\min(y_{kj}, y_{lj}), \max(y_{kj}, y_{lj})]$:

$$\phi(y_{kj}, y_{lj}) \equiv (1-\delta)\min(y_{kj}, y_{lj}) + \delta \max(y_{kj}, y_{lj}) \text{ where } \delta \in [0,1] \tag{6}$$

The thought experiment is repeated for multiple pairs. With a large sample, in a population of size $N_j$, one will end up with an unbiased estimate of the comparison mean:

$$m_j^* = \frac{1}{N_j^2} \sum_{k=1}^{N_j} \sum_{l=1}^{N_j} \phi(y_{kj} y_{lj}) \tag{7}$$

With some algebraic manipulation we can re-write this as:[38]

---

[36] Note that this is a separate issue to the point noted in the introduction that high inequality may yield a direct disutility, thus requiring a higher monetary poverty line in high-inequality settings to assure welfare consistency.

[37] This corresponds to one of the assumptions made by relative deprivation theory in sociology, namely that social comparisons are random in the relevant population (Davis, 1959). Alternatively, one might imagine taking random draws of single income levels within the population for the purpose of assessing a person's relative position. However, given that social comparisons can either look upwards or downwards, rather than draw a single income it would be more informative to imagine drawing a pair to help assess one's relative position.

[38] The derivation uses the fact that $\sum\sum(y_{kj} + y_{lj}) = 2N_j^2 m_j$ and $\min(y_{kj} y_{lj}) = (y_{kj} + y_{lj} - |y_{kj} - y_{lj}|)/2$. The approximation requires large $N_j$ to be accurate.

$$m_j^* \cong [1-(1-2\delta)G_j]m_j \tag{8}$$

where $G_j$ is the Gini coefficient for country *j*:

$$G_j \equiv \sum_{k=1}^{N_j}\sum_{l=1}^{N_j}\frac{|y_{kj}-y_{lj}|}{2N_j^2 m_j} \tag{9}$$

The comparison income in (8) encompasses both upward and downward relativism, depending on the parameter $\delta$. We will say that relative comparisons tend to be downward looking if $\delta < 0.5$ and upward looking if $\delta > 0.5$. If $\delta = 0.5$ then we have the current practice in the literature of treating the overall mean as the comparison income.

A limiting case is of interest both theoretically and (as we will see) empirically, namely the case of downward comparisons such that a person's own income is assessed against the smaller of the two sampled incomes in the thought experiment, i.e., $\delta = 0$. Then we can recognize $m_j^*$ as the distribution-corrected mean proposed by Sen (1976), $(1-G_j)m_j$, though in a different context, namely in measuring social welfare.[39] In that special case, one can also interpret the comparison income as a specific weighted mean. To see how, let incomes be ordered as $y_{1j} \geq y_{2j} \geq,...,\geq y_{N_j}$. Then the comparison income can be re-written as:[40]

$$m_j^* = \frac{2}{N_j^2}\sum_{i=1}^{N_j} i y_{ij} \tag{10}$$

Unlike the median (or some other quantile for any fixed $p_z$), all income levels are deemed relevant to the relativist comparisons made against $m_j^*$, but the weights attached to those incomes fall with the rank in the distribution, starting from the poorest.

We do not assume that $\delta = 0$, but test this empirically. The upper income in each sampled pair may still hold a clue, and at the opposite extreme of $\delta = 1$ the comparison income becomes $(1+G_j)m_j$. Then the appropriate rank-weighted mean puts highest weight on the richest, and lowest on the poorest.

---

[39] Sen (1976) derives this measure (sometimes known as the Sen-evaluation function)from a set of axioms, the key axiom being rank-order weights on incomes in the social welfare function. Yitzhaki (1979) derives the Sen evaluation function in a different way, as mean income less the extent of the aggregate relative deprivation.

[40] This uses the fact that $\sum_{i=1}^{n} i = [n(1+n)]/2$. Thus the weights add to unity.

We now have a clearer theoretical rationale for the comparison income $m_j^*$. The globally welfare-consistent poverty line can be written as:

$$z_j = z(m_j^*, \bar{u}_j^z) \tag{11}$$

(recalling that $\bar{u}_j^z$ is the reference level of welfare above which one is not considered poor). The function is increasing in both arguments:

$$z'_m \equiv \frac{\partial z_j}{\partial m_j^*} = \frac{u'_{y/m} z_j / m_j^*}{m_j^* u'_y + u'_{y/m}} > 0\,;\; z'_u \equiv \frac{\partial z_j}{\partial \bar{u}_j^z} = \frac{m_j^*}{m_j^* u'_y + u'_{y/m}} > 0\,.$$

The corresponding poverty measures are weakly relative in that the elasticity of $z_j$ with respect to $m_j^*$ (equal to its elasticity w.r.t. $m_j$) is less than unity; specifically:

$$0 < \eta_j \equiv \frac{\partial \ln z_j}{\partial \ln m_j^*} = \frac{1}{1 + m_j^* MRS_j} < 1 \tag{12}$$

(where $MRS_j \equiv u'_y / u'_{y/m}$). We cannot predict how $\eta_j$ will vary with $m_j^*$; the necessary and sufficient condition for $\eta_j$ to rise with higher $m_j^*$ is that the elasticity of $MRS_j$ w.r.t. $m_j^*$ is less than -1.

It is of interest to see how poverty measures based on these poverty lines respond to changes in the mean (aggregate growth or contraction) and inequality. We can write the poverty measure in the generic form:[41]

$$P_j = P[m_j / z(m_j^*, \bar{u}_j^z), L_j] \text{ with } P'_{m/z} < 0 \tag{13}$$

where $L_j$ is a vector of parameters fully describing the Lorenz curve. When $L_j$ is fixed we can say that changes in the mean are distribution-neutral. (Of course, we are also interested in what happens with growth when $L_j$ varies, which we examine empirically in Section 6.) Holding $\bar{u}_j^z$ and $L_j$ constant, we then have (taking log derivatives):

$$\frac{\partial \ln P_j}{\partial \ln m_j} = \frac{\partial \ln P_j}{\partial \ln(m_j / z_j)}(1 - \eta_j) < 0 \tag{14}$$

[41] This holds for all poverty measures that are homogeneous of degree zero in the mean and the poverty line. This is a broad class encompassing the class of additive measures described in Atkinson (1987).

In other words, as long as the poverty measure is weakly relative ($\eta_j < 1$), distribution-neutral growth in the mean will reduce that measure.

On the other hand, the partial effect of an increase in the Gini index is theoretically ambiguous. This will depend in part on precisely how the Lorenz curve shifts and there are infinitely many possibilities. However, it is of interest to consider situations in which (13) can be re-written as: [42]

$$P_j = P[m_j / z(m_j^*, \bar{u}_j^z), G_j] \text{ with } P'_G > 0 \tag{15}$$

Then we find that (holding $m_j$ and $\bar{u}_j^z$ constant):[43]

$$\left(\frac{\partial \ln P_j}{\partial \ln G_j}\right)_{dm=0} = \frac{\partial \ln P_j}{\partial \ln(m_j / z_j)} \cdot \frac{\eta_j (1-2\delta) G_j}{[1-(1-2\delta)G_j]} + \left(\frac{\partial \ln P_j}{\partial \ln G_j}\right)_{dm/z=0} \tag{16}$$

The sign of this expression is ambiguous. If $\delta = 0.5$ or relative comparisons are upward-looking ($\delta > 0.5$) then poverty will increase with a higher Gini index (holding the mean constant). Clearly this is also possible, but is not assured, with downward comparisons. We will examine this issue empirically in Section 6.

## 5. Identification using national poverty lines

In using data on national poverty lines to calibrate welfare-consistent international lines one requires a strong identifying assumption, namely that the national lines represent the local costs of a common global level of welfare needed to not be considered poor, i.e., that $\bar{u}_j^z = \bar{u}^z$ for all *j* in equation (11). While it is plausible that national poverty lines reflect relative comparisons, it is hardly evident that the implicit reference level of welfare can be treated as globally constant. Of course, one can allow for idiosyncratic factors, including measurement errors, which can be swept into a regression error term. But it is a strong assumption that these omitted country-specific factors are orthogonal to the variables of interest, notably the mean. More plausibly, the latent national welfare norm is higher in richer countries. (One would then see a pattern such as in Figure 2 even if there is no welfare effect of relative comparisons.) If the gradient in national poverty lines in Figure 2 is entirely due to the effects of differences in the mean on the reference

[42] Kakwani (1993) characterizes a class of proportional shifts in Lorenz curves that satisfy this property, although this is only one possibility consistent with our assumption.

[43] Note that $P'_G$ holds $m_j / z_j$ constant while the following partial log derivative allows $z_j$ to vary.

welfare level then this would eliminate the need for relative poverty measures, given that on ethical grounds one would want to impose a common welfare standard globally. This would only require a fixed real line, as is common practice for global poverty measurement. However, given that (to some extent) the gradient in national poverty lines also reflects the welfare relevance of relative income, a fixed real line will not be welfare consistent.

This identification problem leads us to propose lower and upper bounds to the true welfare-consistent global poverty measure. The true measure lies somewhere between a lower bound defined by a global absolute line, fixed in terms of real income, and an upper bound of a schedule of relative lines. Note, however, that we are not starting from a piece-wise linear poverty line (equation (3)) as the hybrid of absolute and relative lines (as in Atkinson and Bourguignon, 2001). Rather we are arguing for upper and lower bounds as a response to a deep identification problem in inferring welfare-consistent international poverty lines from data on national lines. Since the upper bound cannot be below the lower one, a piece-wise structure is generally required.

The lower bound will be set by the World Bank's $1.90 a day line, while the upper bound will be determined by the data on national lines in Section 2. In identifying the upper bound we begin with the following encompassing specification based on a linearization of equation (11) using (8) as the comparison income:

$$z_j = \alpha + \beta[1-(1-2\delta)G_j]m_j + \varepsilon_j \quad (j = 1,..,n) \tag{17}$$

Here $\alpha, \beta, \delta$ are parameters to be estimated and $\varepsilon_j$ is an error term that includes $\bar{u}_j^z$. As discussed, for identifying the upper bound we assume that $E(\varepsilon_j | m_j, G_j) = 0$. Notice that if $\alpha = 0$ and $\delta = 0.5$ then (in expectation) we have the strongly-relative poverty lines found in the literature that are set at a constant proportion of the mean. On the other hand if $\alpha > 0$ (again with $\delta = 0.5$) then we have a schedule of weakly-relative lines using the equally-weighted mean as the comparison income. However, what do the data suggest?

The estimates for equation (17) are found in column (1) of Table 1.[44] We can reject both null hypotheses, $\delta = 0.5$ and $\delta = 1$, but we cannot reject the null that $\delta = 0$ at the 5% level.

[44] We tested sensitivity to excluding the 24 OECD countries given that their means are so much higher and their poverty lines are mostly set as a constant proportion of the mean or median. However, this made little difference to

Recall that $\delta = 0$ implies that the comparison mean is rank-weighted, with lowest weight on the richest. On imposing this data-consistent restriction we are drawn to our preferred model:

$$z_j = \alpha + \beta(1 - G_j)m_j + \varepsilon_j \qquad (18)$$

The OLS estimate of this equation for the full sample is given in column (3) of Table 1 while column (4) gives the estimate for the non-OECD countries. We also provide in columns (5) and (6) an IV estimate, using PCE per capita as the IV for the mean (as discussed in Section 2). This makes little difference.

Equation (18) is to be interpreted as an upper bound given that our OLS estimate of $\beta$ will be biased upwards by any latent tendency for countries with a higher mean income to have a higher reference level of welfare to not be deemed poor. To see this more clearly, let the error term in (18) include a term $\bar{u}_j^z$ which varies across countries according to $m_j^*$. Then it is readily verified that $\hat{\beta}_{OLS}$ converges in large samples to $\beta + \varphi\pi$ where $\varphi > 0$ is the regression coefficient of $\varepsilon_j$ on $\bar{u}_j^z$ and $\pi \geq 0$ is the regression coefficient of $\bar{u}_j^z$ on $m_j^*$. Thus $\hat{\beta}_{OLS} \geq \beta$. (Clearly our IV estimate does not address this concern.) Thus the gradient w.r.t. the mean derived from the national lines will overestimate the value required for welfare consistency. The upper bound assumes no bias due to latent heterogeneity in country-specific reference welfare levels for defining poverty ($\pi = 0$). At the other extreme, the lower bound assumes maximum bias, so that $\beta = 0$, giving a standard absolute line.

Focusing on the non-OECD group, across the range of the data, the share of the mean that is passed onto the poverty line varies from 0.25 (at the highest Gini of 0.63) to 0.58 (at the lowest of 0.17). The elasticity of the expected value of the poverty line to $m_j^*$ implied by (18) is $1/[1 + \alpha/(\beta m_j^*)]$ and the elasticity rises with $m_j^*$ (the elasticity of $MRS_j$ to $m_j^*$ is -2). Figure 4 plots the relationship (for non-OECD). The elasticity of the poverty line w.r.t. the Gini-discounted mean varies from about 0.4 (at the lowest $m_j^*$) to 0.9. The elasticity is close to unity for the better off developing countries (in terms of $m_j^*$). Given the bearing that the functional form has on these

---

the results so we kept the full sample. Using the full sample also has the attraction that we end up with a truly global poverty line (as argued in Ravallion and Chen, 2013).

results, we also performed the Ramsey RESET test for functional form; the linearity assumption in (18) cannot be rejected in any of the specifications.[45]

The US remains an outlier even when one uses our Gini-discounted mean. If one adds a dummy variable for the US to the regression in Column (5) the coefficient is -$11.95 per person per day (s.e.=$0.22), implying that the expected poverty line in the US (given its value of $m_j^*$) is $27.57 rather than the actual value of $15.62 (per capita for a family of two adults and two children).[46] The estimated parameters do not change much when one includes a dummy variable for the US; we get $\hat{\beta}$ = 0.722 (s.e.=0.007) and $\hat{\alpha}$ = 0.764 (0.135) ($R^2$=0.970; n=145). Inverting the relationship, the expected value of $m_j^*$ for a country with the US poverty line is $20.58 (s.e.=$0.20), which is 55% of the US value ($37.14). The country with the closest $m_j^*$ to this is the Czech Republic $(m^* = \$19.70;\ z = \$13.64)$.

It will be recalled that some of the literature on relative poverty in the OECD countries has used the median rather than the mean. Of course, these two variables are highly correlated; r=0.995 between the mean and median and r=0.999 between the Gini-discounted mean and the median. Nonetheless, one can still separate their effects and when one does it is clear that our distribution-corrected mean is a stronger predictor than the median. Adding the median (*M*) to (18) we obtain the following regression:

$$z_j = \underset{(0.183)}{0.866} + \underset{(0.631)}{0.877}\,(1 - G_j)\,m_j - \underset{(0.519)}{0.139}\,M_j + \hat{\varepsilon}_j \quad R^2=0.956,\ \text{SEE}=1.737,\ n=145 \qquad (19)$$

The coefficient on the median is small, with the wrong sign, while the coefficient on the Gini-discounted mean changes little, although (naturally) its standard error rises when *M* is included. The Gini-discounted mean yields a lower standard error of the regression than the median on its own, although the median predicts slightly better than the (equally-weighted) mean.[47]

We make two further remarks related to other poverty lines in the literature. First, if we calibrate strongly relative lines that are directly proportional to the (equally-weighted) mean then the best fit is to use a constant of almost exactly 0.5 (0.497; s.e.=0.008), i.e., to set the line at

[45] Using the full sample the RESET (2) gave F(2,140)=0.172 (prob.=0.842); using the non-OECD sample the test result is 0.349 (prob.=0.706).
[46] The new Supplementary Poverty Line for the US (Short, 2012) is $17.60 per person per day (for two adults and two children, owning their home with a mortgage). So this cuts $2 a day off the gap, which still leaves $10 a day.
[47] The SEE using the median is 1.772 versus 1.723 using the Gini-discounted mean and 2.080 for the ordinary mean.

about half the mean. Figure 5, panel (a), compares this strongly relative line with the weakly relative lines implied by the analysis above; there is little difference for the OECD so we focus on the non-OECD sub-sample. We see that the strongly relative lines are considerably lower among the countries with a low mean; for the poorest 50 countries in terms of the mean the gap is about \$0.70 a day, which is likely to be sizeable for poverty measures.[48] For 81% of the (non-OECD) countries the weakly relative line is above the strongly relative line; in other words, the fact that the share of the mean is higher for the weakly relative line generally outweighs the fact that we discount the mean for the extent of inequality. Panel (b) of Figure 5 compares two sets of weakly relative lines, one with our distributional correction and the other based solely on the (equally-weighted) mean. Given that the Gini coefficient is only weakly correlated with the mean ($r$=-0.15), there is less sign of a systematic difference between the two sets of relative lines, although a persistent gap is still evident among low income countries.

Second, in the light of these findings, consider again the World Bank's \$1.90 a day absolute line in 2011 prices. This is higher than the expected line in the poorest countries based on our preferred specification. Using our parameter estimate for the full sample (column (5) of Table 1) the expected value of the poverty line is \$1.17 for the country with the lowest value of $m_j^*$ ($m_{\min}^* = \min(m_j^*, j=1,..,n)$), namely the DRC where $m^*$=\$0.44. However, there is naturally a degree of uncertainty in any estimate of $m_{\min}^*$. As already noted, the DRC value is an outlier. The next smallest value is \$0.86 and with the next 7 the values are found in the interval \$0.86-\$1.14. The mean of the poorest 10 is exactly \$1.00, for which the predicted poverty line is \$1.56, also well below the World Bank's \$1.90 line. The latter can still be interpreted as the average poverty line corresponding to a set of the poorest countries (in the spirit of Ravallion et al., 1991); \$1.90 lies between the mean line for the poorest 23 and 24 countries ranked by $m_j^*$ (with means of \$1.86 and \$1.92 respectively). Given that a degree of consensus has developed around the \$1.90 line it is the natural choice for the lower bound.

In summary, the comparison income that is most consistent with the data on national poverty lines is the rank-weighted income in (10). We find that 70% of an increase in the Gini-discounted mean is passed onto the poverty line. Statistically, this model outperforms past methods of measuring relative poverty when judged against national poverty lines.

---

[48] To be more precise the mean gap for the 46 countries with a mean less than \$6 per day is \$0.69 (s.e.=\$0.20).

## 6. Implications for global poverty measures

Building on the results so far, we now provide lower and upper bounds to the true welfare-consistent schedule of poverty lines. To recap, our lower bound ($z_j^L$) is the World Bank's $1.90 absolute line at 2011 PPP while the upper bound is given by our schedule of weakly-relative lines based on (16). With some rounding off, the upper bound is:

$$z_j^U \equiv \$1.90 + \max(0.7m_j^* - \$1.00,\ 0) \qquad (20)$$

(Notice that this approximation is equally valid whether or not one includes the OECD.) We find 13 countries (11% of the non-OECD sample) for which $1.90 is binding (i.e., $m_j^* < \$1.43$). The elasticity of the upper bound to $m_j^*$ goes from zero to a value close to unity (0.97) in the country with the highest $m_j^*$ (Norway). The mean (and median) elasticity in the OECD countries is 0.96. So the upper bound is very close to being a strongly relative measure in the OECD.

In our empirical implementation we will confine attention to the headcount index of poverty (the proportion of the population living below the poverty line), so we will not need to calculate a distribution of equivalent incomes. All we need to know is whether income is above or below the relevant poverty line. Table 2 gives the headcount index for both the lower and upper bounds over 1990-2013. The regional classification in the upper panel is for all low and middle-income countries and a few high-income countries that are eligible for World Bank loans or recently graduated to the Bank's "high-income" category. We also give a separate row for all high-income countries (excluding those elsewhere classified to avoid double-counting). Countries stay in their base-year category over time. For developing countries we use *PovcalNet* while for other countries we use the *Luxembourg Income Study* (LIS).[49] For about two-thirds of countries we use a comprehensive measure of consumption, which we take to be the preferred indicator when there is a choice, while we use disposable income for the rest; the latter is more common in Latin America and the OECD, and LIS uses incomes. We use over 1500 household surveys spanning 150 countries. Estimates are done at 3-yearly intervals up to 2011, then annual to 2013. Estimates for survey years are interpolated to line up with these reference years (Chen and Ravallion, 2004).

[49] The assumptions made in constructing the *PovcalNet* data set are described in Chen and Ravallion (2010, 2013). While LIS is clearly the best available option for developed countries it has some inadequacies for the purpose of measuring global poverty as discussed in Ravallion (2015).

We find that the percentage of the world's population living below both bounds has fallen over time (Figure 6(a)). The trend rate of decline for the upper bound is 0.7 percentage points per year (a regression coefficient on the year of -0.688; s.e.=0.028). The corresponding trend for the lower bound is one percentage point per year (-1.055; s.e.=0.043); if this is maintained then the poverty rate for the lower bound will reach zero by 2025, although it may well be the case that it is harder to maintain the same trend rate of decline as the poverty rate gets closer to zero.[50] So the gap between the poverty rates for the two bounds has risen over time, with a trend rate of increase of about 0.4 percentage points per year (0.367; s.e.=0.025).

Table 3 gives the corresponding counts of the numbers of people below each line. The number living below the lower bound is falling at 47 million per year (a regression coefficient on time of -47.437; s.e.=4.472) while the trend rate of decline in the number of people living below the upper bound is 11 million per year (-10.971; s.e.=3.109) (Figure 6(b)). The number of people between the two bounds is rising at a rate of 36 million per year (36.467; s.e.=2.107). As can be seen in Figure 7, the falling global count of the poor by the lower bound has come with a similar increase in the numbers of people in the developing world who are not poor by this measure but live below the upper bound. The regression coefficient of the number of people in developing countries living below the upper bound on the number living below the lower bound is 0.213 (s.e.=0.045; n=11), which is significantly positive but also significantly less than unity (prob.<0.00005). By interpretation, slightly less than 80% of those who rise above the absolute lower bound end up living between the bounds—no longer poor by the global absolute line but still poor by standards typical of the country they live in.

There are some marked regional differences. As found in past work, the poverty rate for the lower (absolute) line has fallen sharply in East Asia over this period, but there have also been recent (since around 2000) reductions at a comparable pace in South Asia and Sub-Saharan Africa (SSA). In 1990, East Asia had the highest poverty rate of about 60% but this had fallen to under 4% by 2013, while by 2013 SSA has the highest poverty rate (by far) of 41%.

The poverty rate for the lower bound has been under 1% in the high-income countries, and has shown little sign of a trend (fluctuating in a range from 0.6 to 0.9%). Measurement issues are also a concern for these countries, given the presence of negative incomes in the

[50] For further discussion on this point see Ravallion (2013).

primary data; if consumption data were more readily available for these countries we suspect that the poverty rate for the lower bound would be zero or close to it.

When we turn to the upper bound, SSA remains the region with the highest poverty rate but the gap with other regions is (naturally) attenuated. We find that about half of SSA's population falls below our upper bound in 2013, as compared to 30% in East Asia, which is also the region with the steepest decline in poverty judged by the upper bound. The poverty rate rose for the upper bound in SSA and South Asia in the 1990s, but has been falling steadily since the late 1990s. The relative poverty rate in SSA has fallen by a substantial 10% points since 2000 and it has fallen by 18% points in South Asia since 1993. Nonetheless, over the period as a whole, there has been an increase in the number of relatively poor in both SSA and South Asia, though this has levelled off in SSA in the last few years of the series while South Asia has seen falling numbers of relatively poor since 2008. There has also been a more-or-less steady decline in the relative poverty rate for Latin American and the Caribbean and for Eastern Europe and Central Asia since 2000 (with rising poverty in the economic transition period of the 1990s). The relative poverty rate has fallen steadily in the Middle East and North Africa (MENA) though the numbers of relatively poor in MENA have shown no clear trend in either direction.

For high income countries, the poverty rate in 2013 is 19% for the upper line. (For the US the relative poverty rate is 22% in 2013 and has been around 21-22% since the 1980s.) While the high-income countries as a whole have seen a decline in the relative poverty rate since 1990 it is rather small, and there has been little progress against poverty in these countries (as a whole) since 2000, with the (relative) poverty rate staying around 18-19%. It should be recalled that our measure for the upper bound is close to being strongly relative for the high-income countries, so this lack of progress by our measure reflects a lack of progress in improving the relative distribution of income from the point of view of the bottom quintile.

The incidence of poverty judged by the upper bound is higher in the developing world—roughly double the rate found in high-income countries (Table 2). Also, by 2013, a greater proportion of the population of the developing world was relatively poor but not absolutely poor (living between the bounds) than for the rich world. By numbers of people, the bulk (91%) of poverty judged by the upper bound, as well as (of course) the lower bound, is in the developing world (Table 3). Within the developing world, there has been a marked re-alignment of relative poverty, notably between East Asia, South Asia and Africa, as can be seen in Figure 8.

It is of interest to see how our measures at country level respond to differences in both the survey mean and inequality. To investigate this we study the data for the earliest and latest survey rounds for each country. (The median date of the first survey is 1993 while the median for the second is 2012.) As noted in Section 4, given that our upper bound uses weakly-relative measures, the effect of a change in the mean holding the Lorenz curve constant is unambiguous. However, that is not so when one also allows the Lorenz curve to shift with the mean. For both the earliest and latest survey rounds, Figure 9 plots the relationship between the estimates for each bound and the respective means, all in logs. We find that a higher mean is associated with lower poverty measures. This echoes the past finding in the literature that growth tends to be distribution-neutral on average (Ravallion and Chen, 1997; Ferreira and Ravallion, 2009). As one would expect, there is a stronger relationship for the lower bound, but it tends to flatten out at high mean countries, with little difference in the implied elasticities at high means.

To isolate the partial effects (holding either the mean or inequality constant), Table 4 provides regressions of the growth rate in the headcount index on the growth rate in the real survey mean and the growth rate in the Gini index.[51] The regression coefficients can be interpreted as elasticities. It will be recalled from Section 4 that the effect of a change in inequality (at a given mean) on our upper-bound poverty measure is theoretically ambiguous.

We see that higher inequality is associated with higher poverty rates for both bounds while a higher mean comes with lower poverty, as predicted. This also holds when we drop the Gini index. Also note that the elasticities w.r.t. the mean are similar with and without the Gini index, implying that growth rates have little or no correlation (either way) with changes in inequality.[52] While the elasticities are statistically significant in both cases they are appreciably lower for the upper-bound measures, which are also less responsive (in terms of elasticities) to inequality.

For the upper bound there is also a significant interaction effect between the growth rate in the mean and the (log) initial value of the mean, implying a lower growth elasticity at higher means; the elasticity reaches zero at a point close to the highest mean. However, no such interaction effect was in evidence for the lower bound.[53]

[51] This is equivalent to a country-fixed effects regression in the levels. Note that these are descriptive regressions showing the partial correlations, rather than estimates of a causal model.

[52] This was found by Ravallion and Chen (1997) and has been verified often in the literature; for a review see Ferreira and Ravallion (2009).

[53] The t-ratio on the interaction effect for the lower bound is -1.24. The regression is not reported for this case.

## 7. Conclusions

A strong case can be made that welfare-relevant concerns about relative incomes need to be built into prevailing assessments of global poverty and how it is changing. National poverty lines have provided the main data for setting global poverty measures. However, there is a poorly-recognized but deep identification problem in using data on national lines to infer welfare-consistent international lines. The World Bank's longstanding international line has been anchored to the national lines (at purchasing power parity) found in the poorest countries. The line is then held fixed across all countries. This approach assumes that relative income does not matter to welfare. A large body of research since the World Bank's approach was introduced in 1990 suggests this is an overly strong assumption, and that a fixed real line will not be welfare-consistent in any reasonable sense. People at the same level of welfare are not being treated equally.

One might argue instead for using national lines for setting an international schedule of relative lines, and a strand of the literature has followed this approach. However, this makes an equally strong assumption, namely that one can statistically ignore the latent heterogeneity across countries in the "social norms" that define the critical welfare level to not be considered poor. We do not know whether higher lines in richer countries reflect higher expenditures (in real terms) needed to attain the same level of welfare globally or higher reference levels of welfare.

Recognizing the uncertainty on this point leads us to propose lower and upper bounds to welfare-consistent global poverty measures. The natural choice for the lower bound is the World Bank's $1.90 a day international line, at 2011 purchasing power parity (corresponding to the $1.25 a day line at 2005 PPP). In setting the upper bound we have proposed a new approach to measuring relative poverty. We follow other recent work in proposing that the relative line should only be "weakly relative," meaning that its elasticity with respect to the mean should be less than unity. However, we have taken a new approach to the question of what the comparison mean should be. We have argued that past debates about whether one should use the mean, median or something else in measuring relative poverty have not made clear the welfare-economic basis for the choice. We have interpreted relative poverty lines in a welfare-consistent way allowing for relative comparisons. The past use of either the mean or median for relativist comparisons in the literature on poverty is then questionable. For example, it does not seem plausible that such comparisons would give the rich either equal weight (as in the mean) or zero weight (the median).

On starting from a model that encompasses a wide range of options, we have found that a rank-weighted mean is more consistent with the data we have assembled on national poverty lines. This implies that one should use a Gini-discounted mean in setting international relative poverty lines. When judged by the fit to national poverty lines, our method provides a marked improvement over past approaches using either the mean or the median. Our results imply that relative-poverty lines should be adjusted downwards in high-inequality countries, but that does not imply less poverty since we also find that a higher share of the mean should be reflected in the poverty line. When compared to half the mean, our proposed relative poverty lines are higher on average and higher for 71% of the countries in our full sample of national poverty lines (77% of the non-OECD sample). Indeed, our relative line is also higher for a majority (60%) of those countries with a Gini coefficient above the median.

The considerable pace of progress against absolute global poverty in the income space that has been documented before in the literature is not found for the weakly relative lines that form our upper bound. But progress is still evident. Our continuously relative measure (updating the line over time as well as across space) shows a declining poverty rate globally. This also holds for all regions, including the high-income countries, though the pace of progress against poverty has been noticeably less in those countries as a whole, and progress against relative poverty in the high-income countries has stalled since the Great Recession. Consistent with the data, our relative poverty lines have a higher elasticity with respect to the mean in richer countries, approaching unity in the richest. So the evident lack of progress against poverty in the rich world implied by our upper bound measures is primarily about the distribution of the gains from economic growth.

Whether one focuses on absolute poverty (our lower bound) or relative poverty (upper bound), the incidence of poverty is appreciably higher in the developing world. Over 90% of the poor by our upper line are found in the developing world, which is home to virtually all of those who are poor by our lower line. However, the developing world has been making greater progress over time against poverty, judged by either bound. Side-by-side with the falling numbers of absolutely poor in the developing world, we find that there have been rising numbers of people who are still poor by the standards typical of the country they live in. Both the lower and upper-bound poverty measures are responsive to both the mean and inequality, although the upper bound measure responds less elastically.

## References


Abel-Smith, Brian, and Peter Townsend, 1966, *The Poor and the Poorest: A New Analysis of the Ministry of Labour's Family Expenditure Surveys of 1953-54 and 1960.* London: Bell.

Atkinson, Anthony B., 1987, "On the Measurement of Poverty," *Econometrica* 55: 749-764.

__________________, 1998, *Poverty in Europe*, Oxford: Blackwell.

Atkinson, Anthony B., 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.

Balsa, Ana, Michael French, and Tracy Regan, 2014, "Relative Deprivation and Risky Behaviors," *Journal of Human Resources* 49(2): 446–471.

Banerjee, Abhijit, and Esther Duflo, 2008, "What is Middle Class about the Middle Classes Around the World?", *Journal of Economic Perspectives* 22(2): 3-28.

Blackburn, McKinley, 1994, "International Comparisons of Poverty", *American Economic Review* 84(2): 371-374.

Blackorby, Charles, and Donald Donaldson, 1987, "Welfare Ratios and Distributionally Sensitive Cost-Benefit Analysis," *Journal of Public Economics* 34: 265-290.

Blank, Rebecca M., 2008, "How to Improve Poverty Measurement in the United States," *Journal of Policy Analysis and Management* 27(2): 233-254.

Budlender, Joshua, Murray Leibbrandt and Ingrid Woolard, 2015, "Weakly Relative Poverty Lines in Developing Countries: The case of South Africa," Paper presented at the 2015 Conference of the Economic Society of South Africa.

Chakravarty, Satya, Nachiketa Chattopadhyay, Zoya Nissanov and Jacques Silber, 2015, "Reference Groups and the Poverty Line: An Axiomatic Approach with an Empirical Illustration," WIDER Working Paper 2015/002.

Chen, Xi, 2015, "Relative Deprivation and Individual Well-Being," *IZA World of Labor* 140.

Chen, Shaohua, and Martin Ravallion, 2001, "How Did the World's Poor Fare in the 1990s?", *Review of Income and Wealth* 47(3): 283-300.

____________, and ______________, 2004,"How Have the World's Poorest Fared Since the Early 1980s?" *World Bank Research Observer* 19: 141-70.

____________, and ______________, 2010, "The Developing World is Poorer than we Thought, but no Less Successful in the Fight against Poverty," *Quarterly Journal of Economics* 125(4): 1577-1625.

____________, 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. (eds.) *Measuring Poverty: A New Approach.* Washington, DC: National Academy Press.

Clark, Andrew, Paul Frijters, and Michael Shields, 2008, "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," *Journal of Economic Literature* 46(1): 95-144.

Clark, Andrew, and Andrew Oswald, 1996, "Satisfaction and Comparison Income," *Journal of Public Economics* 61: 359-381.

Cohn, Alain, Ernst Fehr, Benedikt Herrmann, and Frédéric Schneider, 2014, "Social Comparison And Effort Provision: Evidence From A Field Experiment," *Journal of the European Economic Association* 12(4): 877–898.

Davis, James, 1959, "A Formal Interpretation of the Theory of Relative Deprivation," *Sociometry* 22: 280-296.

de Mesnard, Louis, 2007, "Poverty Reduction: The Paradox of the Endogenous Poverty Line," LEG - Document de travail - Economie 2007-05, LEG, Laboratoire d'Economie et de Gestion, CNRS UMR 5118, Université de Bourgogne.

Decerf, Benoit, 2017, "Why Not Consider that Being Absolutely Poor is Worse than Being only Relatively Poor?" *Journal of Public Economics* 152: 79-92.

Drewnowski, Jan, 1977, "Poverty: Its Meaning and Measurement," *Development and Change* 8: 183-208.

Duclos Jean-Yves and Paul Makdissi, 2004, "Restricted and Unrestricted Dominance for Welfare Inequality and Poverty Orderings," *Journal of Public Economic Theory* 6: 145-164.

Duesenberry, James S., 1949, *Income, Saving and the Theory of Consumer Behavior,* Cambridge, Mass.: Harvard University Press.

Easterlin, Richard A., 1974, "Does Economic Growth Improve the Human Lot? Some Empirical

Evidence," in P.A. David and W.R. Melvin (eds) *Nations and Households in Economic Growth*, Palo Alto: Stanford University Press.

Easton, Brian, 2002, "Beware the Median," *Social Policy Research Center Newsletter* 82: 6-7.

Ferreira, F., Chen, S., Dabalen, A., Dikhanov, Y., Hamadeh, N., Joliffe, D. Narayan, A., Prydz, E., Revenga, A., Sangraula, P., Serajuddin U., Yoshida, N., 2016, "A Global Count of the Extreme poor in 2012: Data Issues, Methodology and Initial Results," *Journal of Economic Inequality* 14: 141-172.

Ferreira, Francisco, and Martin Ravallion, 2009, "Poverty and Inequality: The Global Context," in *The Oxford Handbook of Economic Inequality*, edited by Wiemer Salverda, Brian Nolan and Tim Smeeding, Oxford: Oxford University Press.

Foster, James, 1998, "Absolute versus Relative Poverty," *American Economic Review, Papers and Proceedings* 88(2): 335-341.

Frank, Robert H., 1985, *Choosing the Right Pond: Human Behavior and the Quest for Status*. New York: Oxford University Press.

Frey, Bruno, and Stutzer, Alois, 2002, "What Can Economists Learn from Happiness Research?," *Journal of Economic Literature*. XL: 402–435.

Fuchs, Victor, 1967, "Redefining Poverty and Redistributing Income," *The Public Interest* 8: 88-95.

Fuller, C.J., 1992, *The Camphor Flame. Popular Hinduism and Society in India* Princeton: Princeton University Press.

Garroway, Christopher, and J.R. de Laiglesia, 2012, "On the Relevance of Relative Poverty for Developing Countries," OECD Development Centre Working Paper 314, OECD, Paris.

Geertz, Clifford, 1976, *The Religion of Java.* Chicago and London: University of Chicago Press.

Hagenaars, Aldi J.M., and Bernard M.S. van Praag, 1985, "A Synthesis of Poverty Line Definitions," *Review of Income and Wealth,* 31:139-154.

Hirsch, F., 1977, *The Social Limits to Growth*. London: Routledge & Kegan Paul.

Hunter, Robert, 1904, *Poverty*, London: MacMillan Company.

Jolliffe, Dean, and Espen Prydz, 2016, "Estimating international poverty lines from comparable national thresholds," *Journal of Economic Inequality* 14(2): 185-198.

___________, and __________, 2017, “Global Societal Poverty: A Relative and Relevant Measure,” Paper Presented at the Poverty Goals and Prices. How Purchasing Power Parity Matters.” Policy Research Working Paper 8073, World Bank.

Kakwani, Nanak, 1986, *Analyzing Redistribution Policies: A Study using Australian Data.* Cambridge: Cambridge University Press.

______________, 1993, “Poverty and Economic Growth with Application to Côte D'Ivoire,” *Review of Income and Wealth* 39, 121-139.

Kampke, Thomas, 2010, “The use of Mean Values vs. Medians in Inequality Analysis,” *Journal of Economic and Social Measurement* 35: 43–62.

Kilpatrick, Robert, 1973, “The Income Elasticity of the Poverty Line,” *Review of Economics and Statistics* 55: 327-332.

Klasen, Stephan, Tatyana Krivobokova, Friederike Greb, Rahul Lahoti, Syamsul Pasaribu, and Manuel Wiesenfarth, 2016, “International Poverty Measurement: Which Way Now?” *Journal of Economic Inequality* 14: 199-125.

Knight, John, Lina Song, Ramani Gunatilaka, 2009, “Subjective Well-Being and Its Determinants in Rural China,” *China Economic Review* 20(4): 635-49.

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.

Milanovic, Branko, and John E. Roemer, 2016, “Interaction of Global and National Income Inequalities,” *Journal of Globalization and Development* 7(1): 109-115.

Nolan, Brian, 2007, *A Comparative Perspective on the Development of Poverty and Exclusion in European Societies*, Bonn: International Policy Analysis.

Organization for Economic Co-Operation and Development (OECD), 2008, *Growing Unequal? Income Distribution and Poverty in OECD Countries*, Paris, OECD.

Orshansky, Mollie, 1965, “Counting the Poor: Another Look at the Poverty Profile,” *Social Security Bulletin* 28: 3-29.

Pradhan, Menno, and Martin Ravallion, 2000, “Measuring Poverty Using Qualitative Perceptions of Consumption Adequacy,” *Review of Economics and Statistics* 82(3): 462-471.

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.

______________, 2012, "Poverty Lines Across the World" in *The Oxford Handbook of the Economics of Poverty*, edited by Philip N. Jefferson, Oxford: Oxford University Press.

______________, 2013, "How Long Will It Take To Lift One Billion People Out Of Poverty?" *World Bank Research Observer* 28 (2): 139-158.

______________, 2015, "The Luxembourg Income Study," *Journal of Economic Inequality* 13(4): 527-547.

______________, 2016a, *The Economics of Poverty: History, Measurement and Policy*. New York: Oxford University Press.

______________, 2016b, "Toward Better Global Poverty Measures," *Journal of Economic Inequality* 14: 227-248.

Ravallion, Martin, and Shaohua Chen, 1997, "What Can New Survey Data Tell Us About Recent Changes in Poverty and Distribution?", *World Bank Economic Review* 11(2): 357-382.

______________, and ____________, 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.

Ravallion, Martin, Shaohua Chen and Prem Sangraula, 2009, "Dollar a Day Revisited," *World Bank Economic Review* 23(2):163-184.

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, and Michael Lokshin, 2010, "Who Cares About Relative Deprivation?" *Journal of Economic Behavior and Organization*. 73(2): 171-185.

Rayo, Luis, and Gary S. Becker, 2007, "Evolutionary Efficiency and Happiness," *Journal of Political Economy* 115(2): 302-337.

Runciman, W.G., 1966, *Relative Deprivation and Social Justice*, London: Routledge and Kegan Paul.

Saunders, Peter and Timothy Smeeding, 2002, "Beware the Mean!" *Social Policy Research Center Newsletter* 81(1), 4-5.

Sen, Amartya, 1973, *On Economic Inequality*. Oxford: Clarendon Press.

___________, 1976, "Real National Income," *Review of Economic Studies* 43(1): 19-39.

___________, 1983, "Poor, Relatively Speaking," *Oxford Economic Papers,* 35(2): 153-69.

Senik, Claudia, 2005, "Income Distribution and Well-Being: What can we learn from Subjective Data?" *Journal of Economic Surveys* 19(1): 43-63.

Short, Kathleen, 2012, "The Research Supplemental Poverty Measure: 2011." *Current Population Reports P60-244*. U.S. Census Bureau.

Smeeding, T.M., L. Rainwater, M. Rein, R. Hauser and G. Schaber, 1990, "Income Poverty in Seven Countries: Initial Estimates from the LIS Database," in T.M. Smeeding, M. O'Higgins and L. Rainwater, eds., *Poverty, Inequality and Income Distribution in Comparative Perspective*, Washington, DC: Urban Institute Press.

Smith, Heather, Thomas Pettigrew, Gina Pippin, and Silvana Bialosiewicz, 2012, "Relative Deprivation: A Theoretical and Meta-analytic Critique," *Personality and Social Psychology Review* 16(3): 203–232.

Townsend, Peter, 1979, *Poverty in the United Kingdom: A Survey of Household Resources and Standards of Living.* Harmonsworth: Penguin Books.

United Nations Development Program (UNDP), 2005, *Human Development Report,* New York: UNDP.

Walker, Robert, 2014, *The Shame of Poverty*, Oxford: Oxford University Press.

White, Halbert, 1980, "A Heteroskedasticity-consistent Covariance Matrix Estimator and a Direct Test for Heteroscedasticity," *Econometrica* 48(4): 817–838.

World Bank, 1990, *World Development Report: Poverty*, Oxford: Oxford University.

__________, 2015, *Purchasing Power Parities and the Real Size of World Economies: A Comprehensive Report of the 2011 International Comparison Program*. Washington, DC: World Bank.

Yitzhaki, Shlomo, 1979, "Relative Deprivation and the Gini Coefficient," *Quarterly Journal of Economics* 93(2): 321–24.

**Figure 1: Kernel density functions**

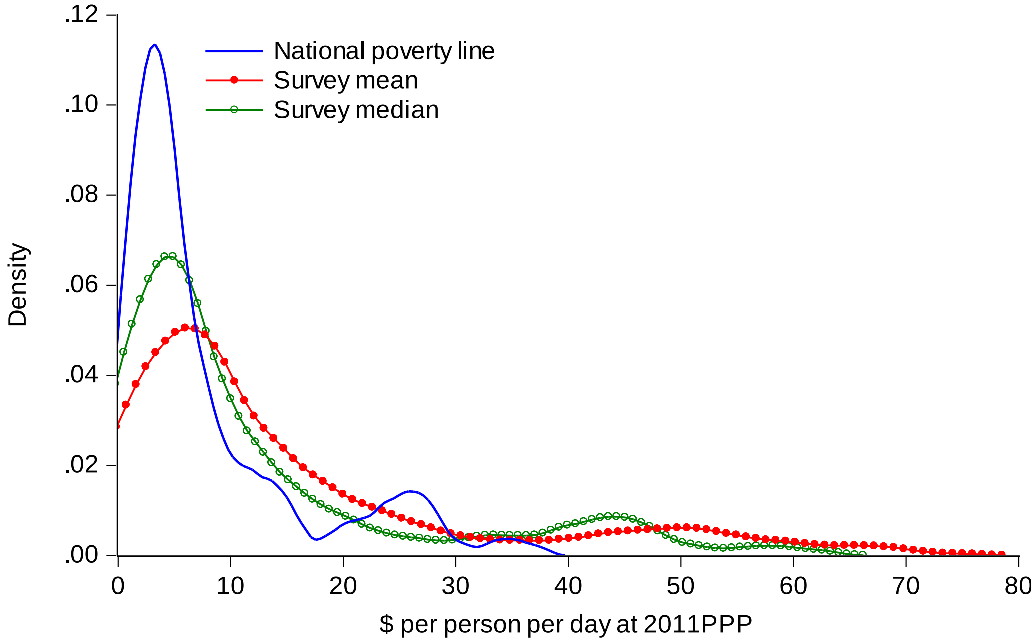

**Figure 2: Poverty lines across countries**

(a) Full sample (n=146)

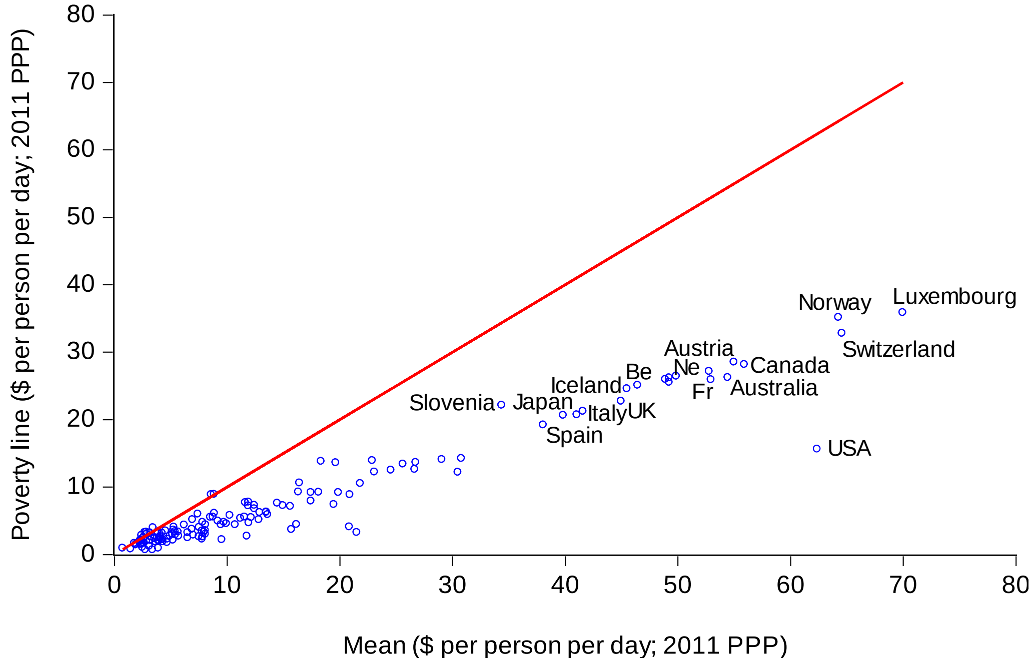


(b) Non-OECD (n=122) using log scale for mean

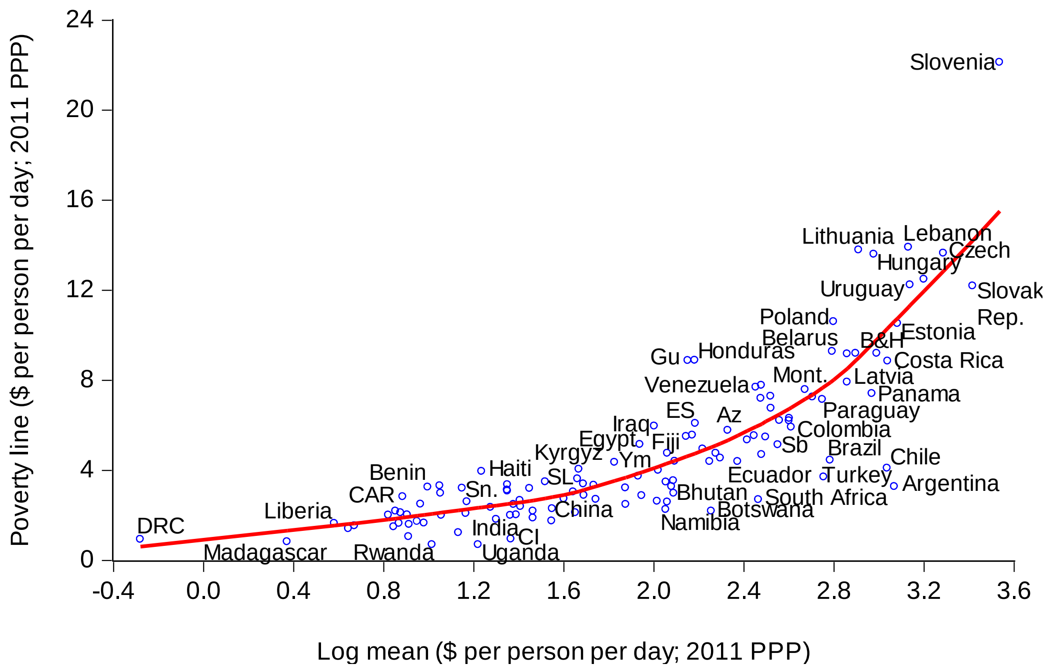


Note: Empirical non-parametric regression (locally weighted polynomial). Az: Azerbaijan; Be: Belgium; B&H: Bosnia and Herzegovina; CAR: Central African Republic; CI: Cote d'Ivoire; ES: El Salvador; Fr: France; Gu: Guatemala; Mont.: Montenegro; Ne: Netherlands; Sb: Serbia; Sn.: Senegal; SL: Sierra Leone; UK: United Kingdom; USA: United States of America; Ym: Yemen.

**Figure 3: Ratio of the national poverty line to the current survey mean**

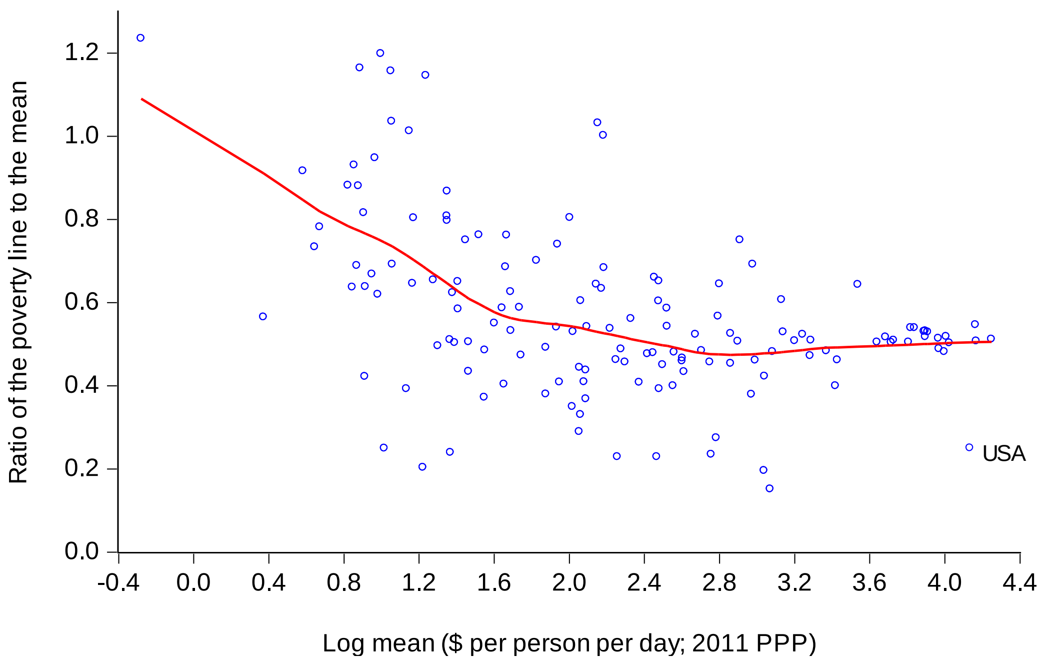


**Figure 4: Elasticity of the poverty line to the Gini-discounted mean**

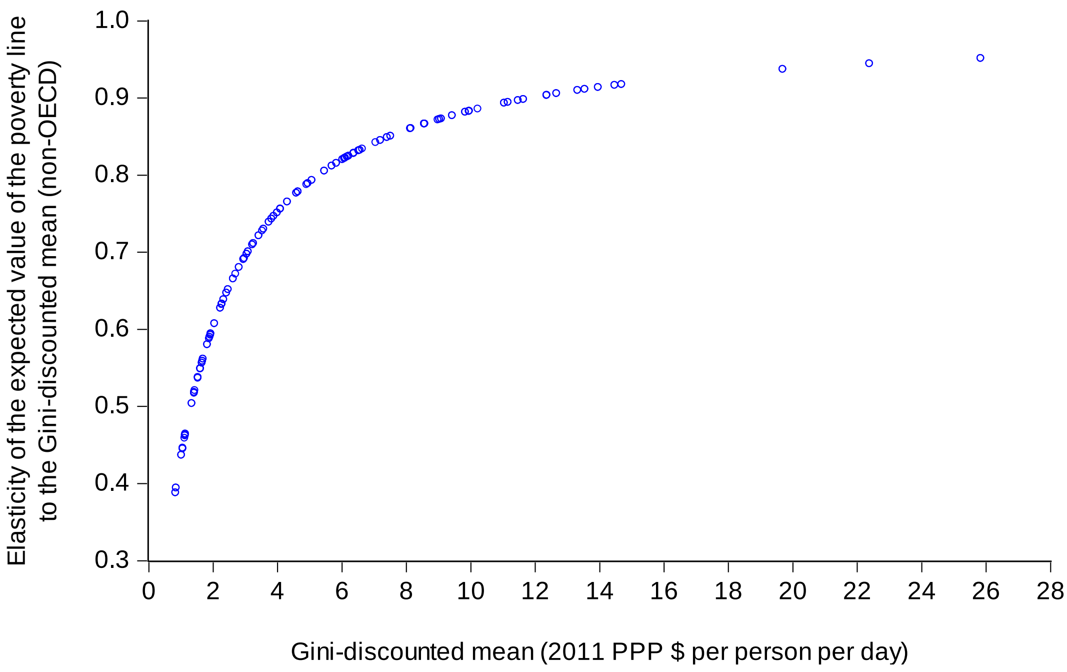

**Figure 5: Comparison of relative lines (non-OECD)**

(a) Weakly relative poverty lines with distribution adjustment versus strongly relative lines

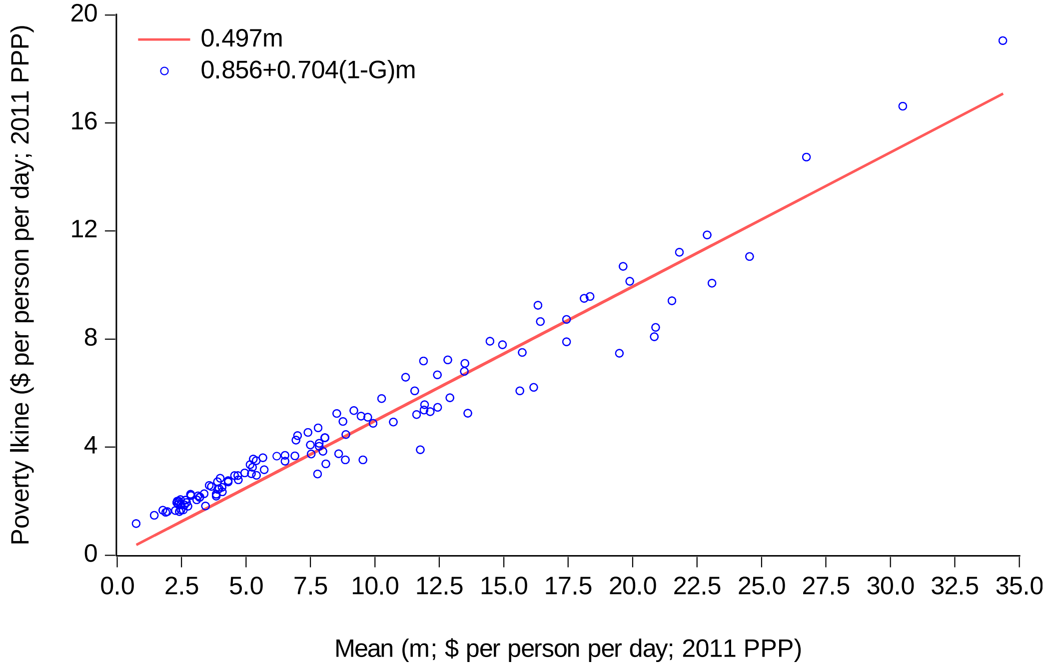


(b) Weakly relative poverty lines with and without distributional adjustment

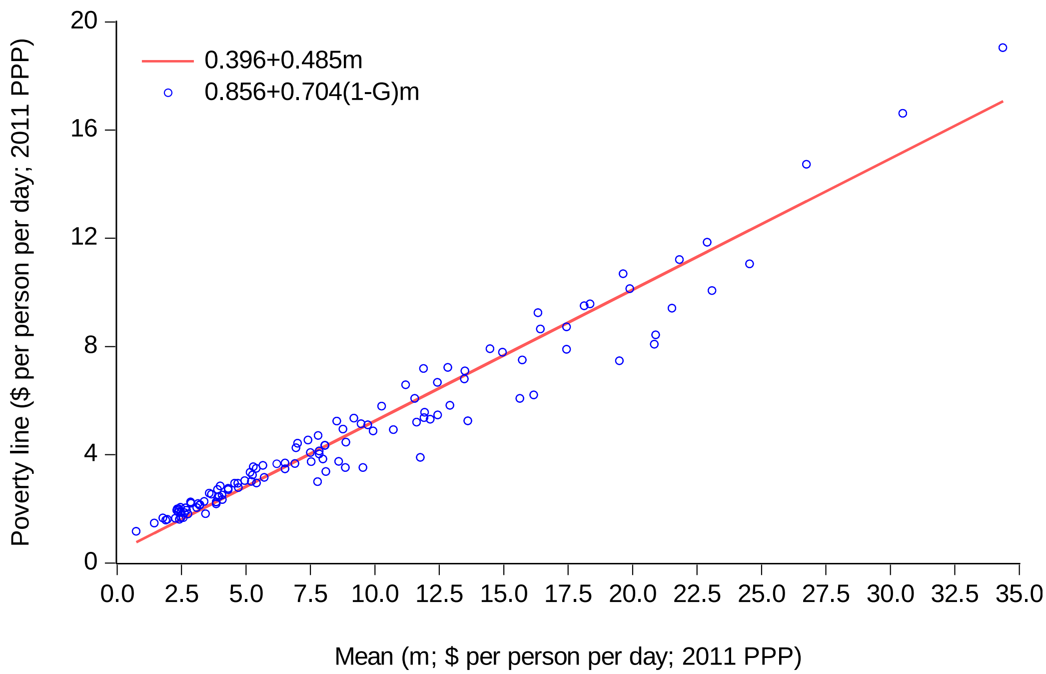


Note: The distributional adjustment uses the Gini-discounted mean (see text).

**Figure 6: Global poverty measures for lower and upper bounds**

(a) Headcount index (%)

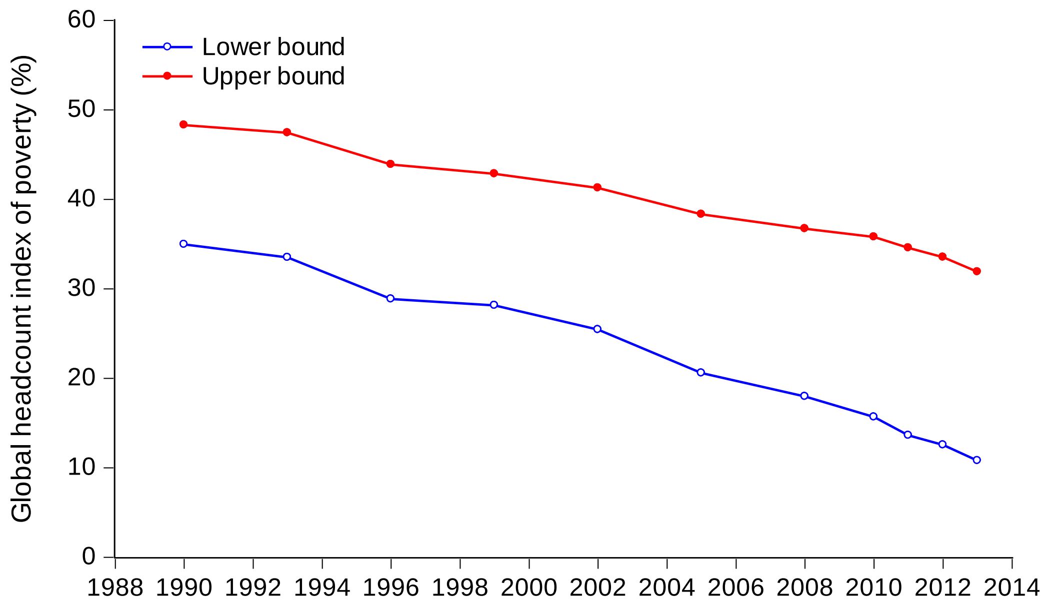


(b) Numbers of poor (million)

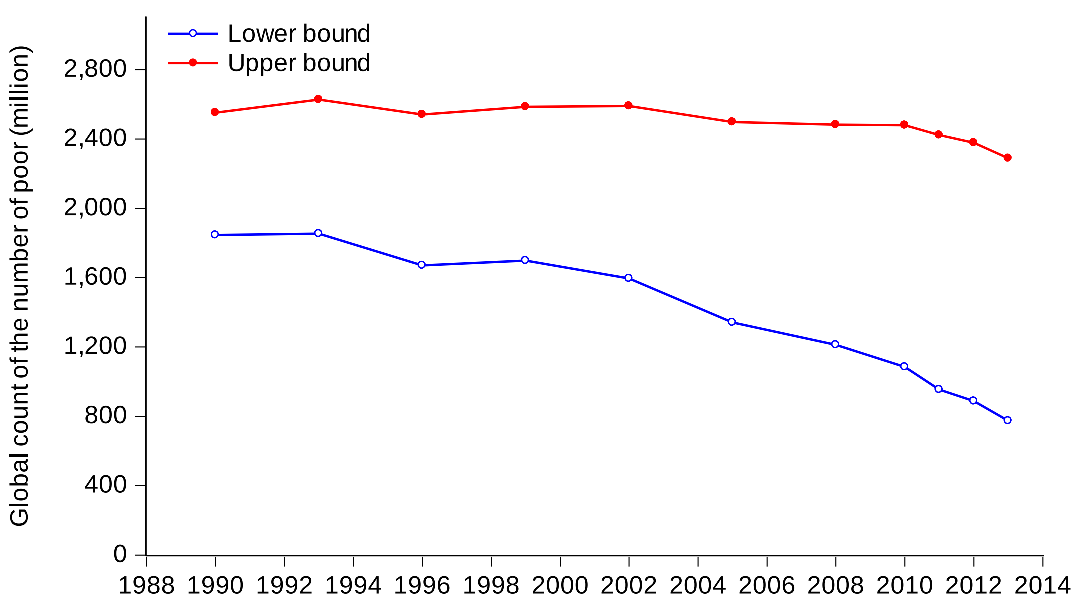

**Figure 7: Components of global poverty count for the upper bound line**

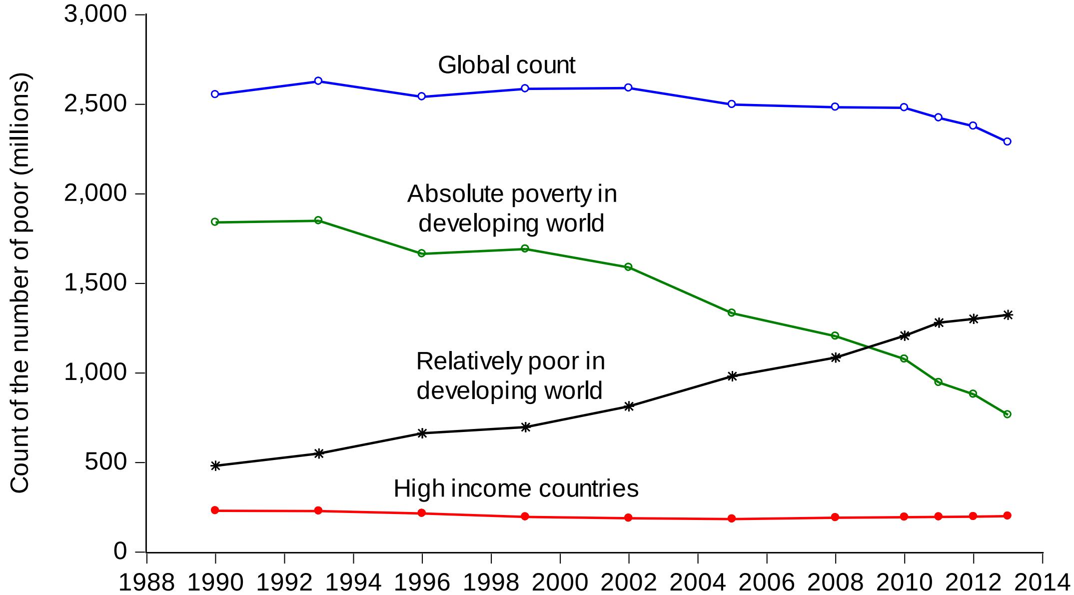


Note: The count of the "absolutely poor in developing world" is the number of people living below the lower bound, while the count of "relatively poor in developing world" is the number between the two bounds,

**Figure 8: Changing composition of global poverty by region**

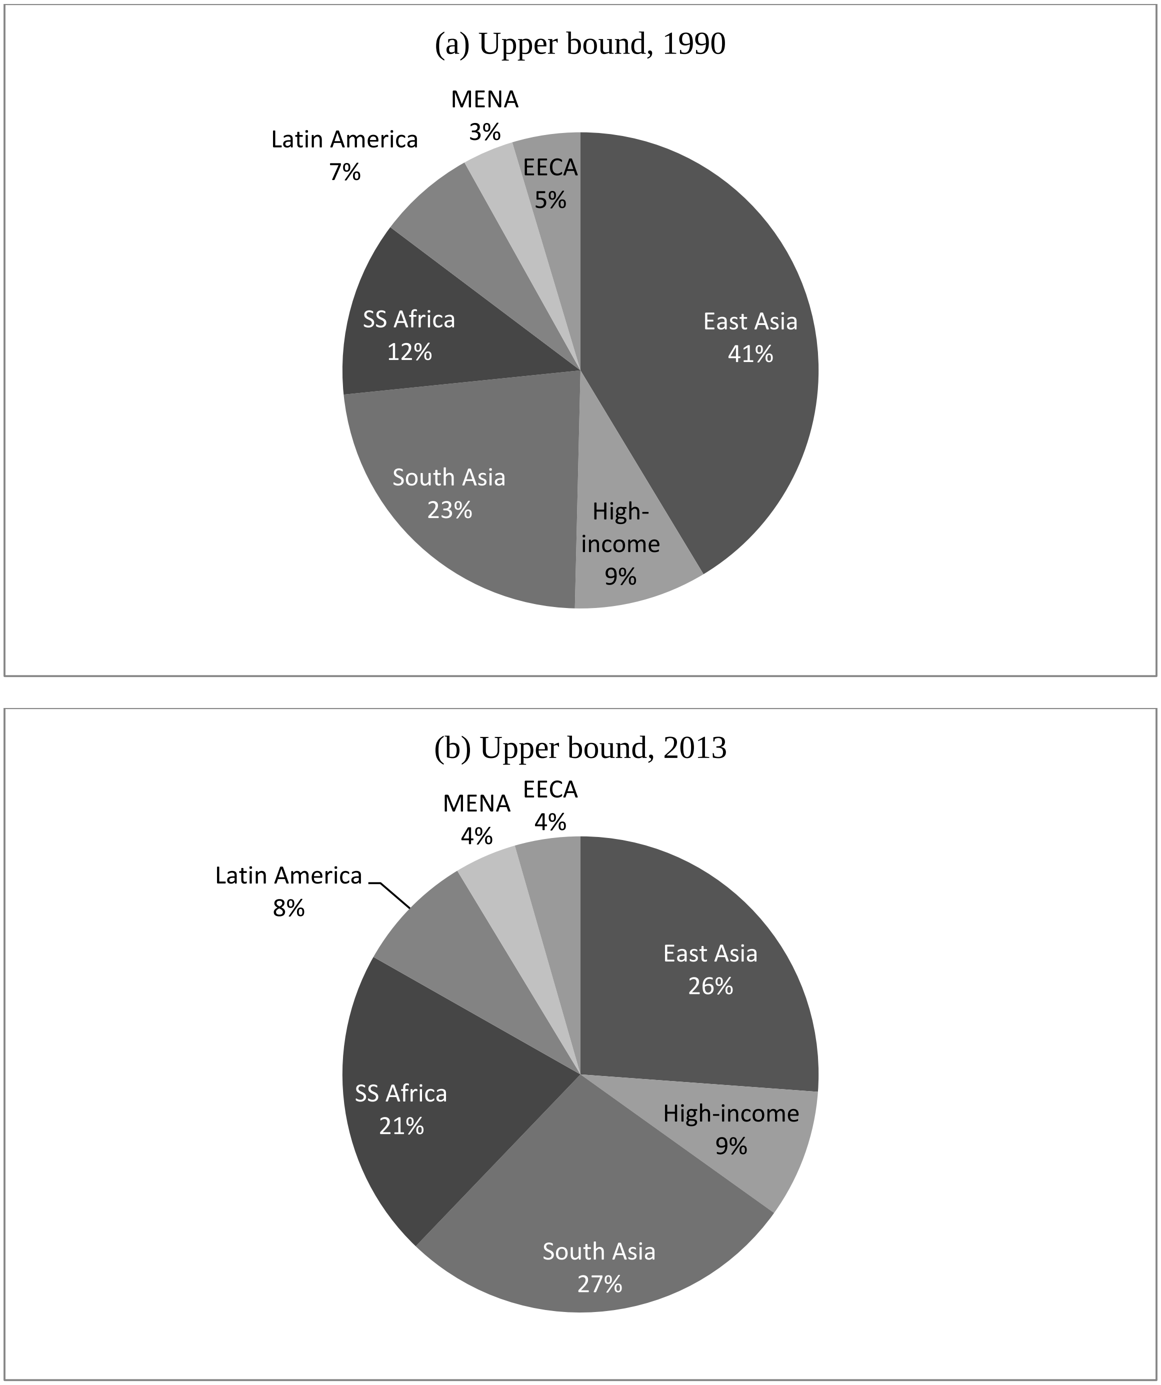


Notes: EECA: Eastern Europe and Central Asia; MENA: Middle-East and North Africa; SS: Sub-Saharan; "Latin America" includes the Caribbean.

**Figure 9: Descriptive relationship between the poverty measure and the overall mean**

(a) Earliest survey

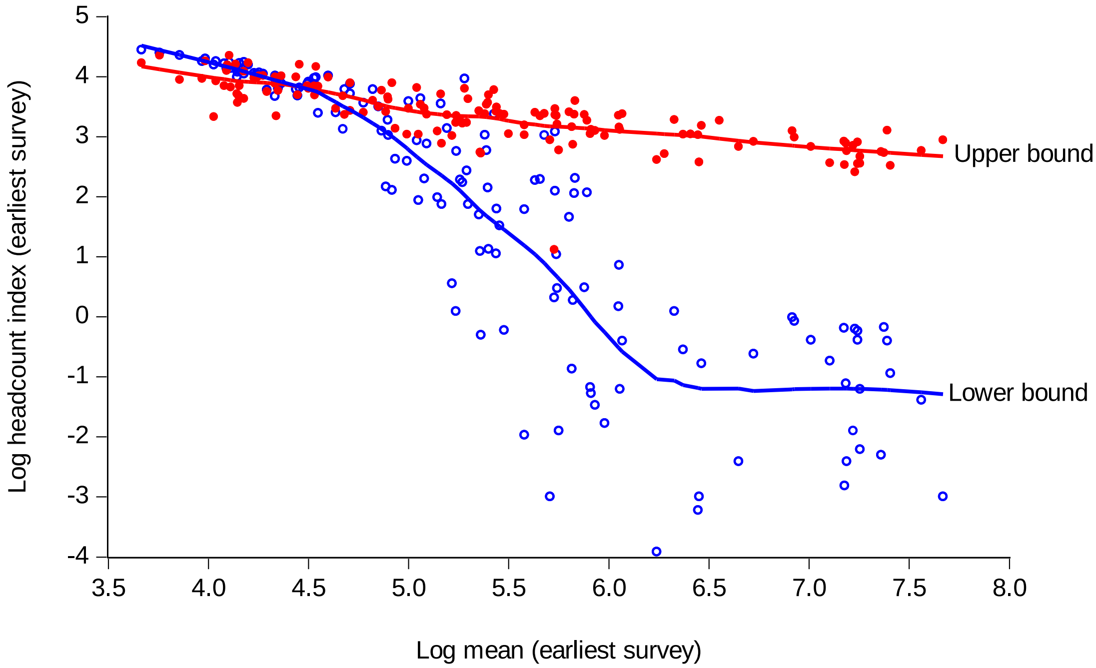


(b) Latest survey

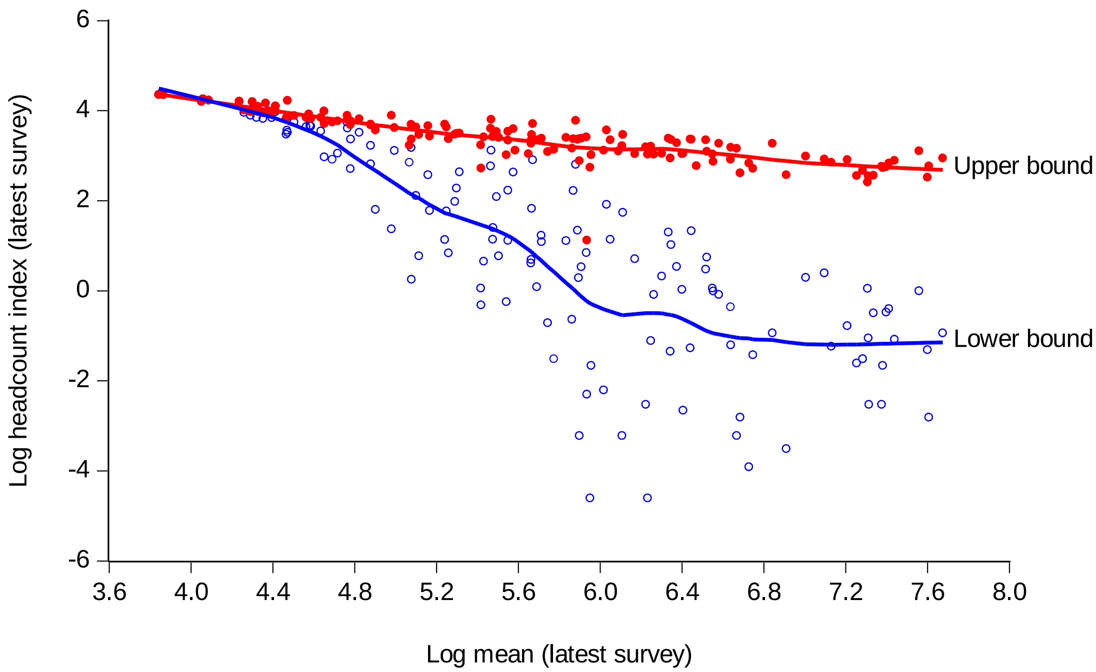

## Table 1: Regressions for national poverty lines

| | (1) Full sample OLS | (2) Non-OECD OLS | (3) Full sample OLS | (4) Non-OECD OLS | (5) Full sample IV | (6) Non-OECD IV |
|---|---|---|---|---|---|---|
| Intercept ( $\alpha$ ) | 1.072*** | 0.891*** | 0.856*** | 0.887*** | 1.001*** | 1.102*** |
| | (0.313) | (0.217) | (0.163) | (0.233) | (0.234) | (0.054) |
| Gini-discounted mean ( $\beta$ ) | 0.781*** | 0.697*** | 0.704*** | 0.695*** | 0.687*** | 0.650*** |
| | (0.106) | (0.085) | (0.018) | (0.050) | (0.029) | (0.054) |
| Weight on higher income in any pair ( $\delta$ ) | -0.115 | -0.002 | n.a. | n.a. | n.a. | n.a. |
| | (0.111) | (0.097) | | | | |
| $R^2$ | 0.958 | 0.813 | 0.956 | 0.813 | 0.956 | 0.804 |
| SEE | 1.698 | 1.540 | 1.723 | 1.534 | 1.733 | 1.536 |
| N | 145 | 121 | 145 | 121 | 143 | 119 |

Notes: White standard errors in parentheses; ***: 1%; **: 5%; *: 10%; OLS: Ordinary Least Squares; IV: Instrumental variable.

**Table 2: Headcount indices of poverty for upper and lower bounds by region and year**

| Region | Poverty line bound | 1990 | 1993 | 1996 | 1999 | 2002 | 2005 | 2008 | 2011 | 2012 | 2013 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| East Asia + Pacific | Upper | 65.8 | 59.8 | 49.3 | 48.3 | 44.0 | 38.4 | 36.1 | 32.8 | 31.9 | 29.9 |
| | Lower | 60.2 | 52.4 | 39.4 | 37.2 | 29.0 | 18.4 | 14.9 | 8.4 | 7.1 | 3.5 |
| EECA | Upper | 25.7 | 27.9 | 30.7 | 30.4 | 28.7 | 26.8 | 23.8 | 22.7 | 22.4 | 21.3 |
| | Lower | 1.9 | 5.5 | 7.3 | 8.0 | 6.3 | 5.0 | 3.1 | 2.6 | 2.4 | 2.2 |
| Latin America | Upper | 37.6 | 36.1 | 36.5 | 36.0 | 36.0 | 34.3 | 31.7 | 30.7 | 30.2 | 29.9 |
| | Lower | 15.8 | 14.4 | 14.2 | 13.9 | 13.0 | 10.8 | 7.1 | 6.0 | 5.6 | 5.4 |
| MENA | Upper | 39.2 | 41.2 | 40.6 | 35.5 | 34.4 | 33.8 | 31.1 | 29.0 | 27.5 | 27.6 |
| | Lower | 6.0 | 5.6 | 4.8 | 3.8 | n.a. | 3.0 | 2.1 | n.a. | n.a. | n.a. |
| South Asia | Upper | 51.8 | 53.5 | 51.7 | 50.9 | 50.9 | 48.6 | 46.7 | 42.4 | 40.3 | 36.8 |
| | Lower | 44.6 | 44.8 | 40.3 | n.a. | 38.5 | 33.6 | 29.4 | 19.9 | 17.5 | 15.1 |
| SS Africa | Upper | 59.7 | 63.6 | 62.9 | 61.9 | 60.8 | 56.5 | 55.0 | 54.0 | 52.6 | 50.8 |
| | Lower | 54.3 | 58.4 | 57.7 | 57.1 | 55.6 | 50.0 | 47.0 | 44.1 | 42.6 | 41.0 |
| World (excl. HIC) | Upper | 53.0 | 52.0 | 48.0 | 47.1 | 45.4 | 42.0 | 40.0 | 37.4 | 36.2 | 34.2 |
| | Lower | 42.0 | 40.1 | 34.4 | 33.4 | 30.0 | 24.2 | 21.1 | 15.9 | 14.6 | 12.6 |
| HIC | Upper | 25.5 | 24.7 | 22.8 | 20.3 | 19.1 | 18.3 | 18.5 | 18.5 | 18.5 | 18.6 |
| | Lower | 0.6 | 0.6 | 0.7 | 0.7 | 0.8 | 0.9 | 0.8 | 0.8 | 0.7 | 0.7 |
| World | Upper | 48.3 | 47.4 | 43.9 | 42.8 | 41.3 | 38.3 | 36.7 | 34.6 | 33.5 | 31.9 |
| | Lower | 35.0 | 33.5 | 28.9 | 28.2 | 25.4 | 20.6 | 18.0 | 13.6 | 12.5 | 10.8 |

Notes: EECA: Eastern Europe and Central Asia; MENA: Middle-East and North Africa; HIC: high-income countries (excluding those eligible for World Bank loans, which are assigned to their respective regions); SS: Sub-Saharan; "Latin America" includes the Caribbean; "n.a." indicates that the survey coverage rate is considered too low (under 50%); estimates are still included in the aggregates. Extrapolations/interpolations are done to line up survey years at country level with the reference years given in the table.

**Table 3: Number of poor in million for upper and lower bounds by region and year**

| | Poverty line bound | 1990 | 1993 | 1996 | 1999 | 2002 | 2005 | 2008 | 2011 | 2012 | 2013 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| East Asia + Pacific | Upper | 1055.5 | 1000.8 | 856.4 | 867.9 | 813.8 | 727.6 | 699.6 | 648.9 | 635.0 | 598.8 |
| | Lower | 965.9 | 876.8 | 683.8 | 669.0 | 535.1 | 349.2 | 288.2 | 166.9 | 141.8 | 71.0 |
| EECA | Upper | 118.1 | 129.9 | 143.2 | 142.2 | 133.7 | 125.0 | 111.5 | 107.7 | 106.5 | 101.8 |
| | Lower | 8.9 | 25.4 | 34.1 | 37.3 | 29.3 | 23.2 | 14.6 | 12.5 | 11.5 | 10.3 |
| Latin America | Upper | 169.1 | 171.0 | 181.6 | 187.7 | 195.4 | 193.8 | 186.1 | 186.8 | 185.5 | 186.0 |
| | Lower | 71.2 | 68.3 | 70.7 | 72.2 | 70.6 | 60.8 | 41.9 | 36.4 | 34.1 | 33.6 |
| MENA | Upper | 88.7 | 99.8 | 104.7 | 97.1 | 99.3 | 102.7 | 99.5 | 98.0 | 94.5 | 96.4 |
| | Lower | 13.7 | 13.6 | 12.4 | 10.5 | n.a. | 9.2 | 6.7 | n.a. | n.a. | n.a. |
| South Asia | Upper | 587.0 | 646.1 | 663.4 | 692.0 | 731.1 | 733.7 | 739.5 | 700.9 | 674.7 | 624.4 |
| | Lower | 505.0 | 541.5 | 517.0 | n.a. | 552.4 | 508.3 | 464.7 | 327.9 | 293.3 | 256.2 |
| SS Africa | Upper | 303.5 | 351.7 | 377.0 | 402.6 | 428.2 | 430.9 | 455.3 | 484.9 | 485.2 | 481.7 |
| | Lower | 276.1 | 323.1 | 346.1 | 371.3 | 391.3 | 381.5 | 389.1 | 395.7 | 393.1 | 388.7 |
| World (excl. HIC) | Upper | 2321.9 | 2399.0 | 2326.3 | 2389.4 | 2401.7 | 2313.7 | 2291.1 | 2226.9 | 2181.2 | 2089.3 |
| | Lower | 1840.5 | 1848.5 | 1663.9 | 1691.9 | 1588.1 | 1332.4 | 1205.0 | 946.3 | 880.3 | 766.0 |
| HIC | Upper | 230.2 | 227.9 | 215.2 | 196.3 | 188.6 | 184.5 | 191.5 | 195.9 | 196.7 | 199.6 |
| | Lower | 5.8 | 5.9 | 6.7 | 7.2 | 7.7 | 8.6 | 7.8 | 8.5 | 7.9 | 7.6 |
| World | Upper | 2552.3 | 2627.1 | 2541.2 | 2586.0 | 2590.2 | 2498.1 | 2483.0 | 2423.0 | 2377.8 | 2288.5 |
| | Lower | 1846.2 | 1854.4 | 1670.6 | 1699.1 | 1595.8 | 1341.0 | 1212.8 | 954.8 | 888.2 | 773.6 |

Notes: See Table 2.

**Table 4: Descriptive cross-country regressions for changes in poverty measures**

| | Growth rate in the headcount index | | | | |
|---|---|---|---|---|---|
| | Lower bound poverty lines ($g(F(z^L))$) | | Upper bound poverty lines ($g(F(z^U))$) | | |
| Growth rate in the survey mean ($g(m)$) | -2.239*** (0.280) | -2.394*** (0.283) | -0.434*** (0.050) | -0.453*** (0.048) | -1.664*** (0.193) |
| $g(m)$ x log initial mean | n.a | n.a | n.a | n.a | 0.240*** (0.037) |
| Growth rate in the Gini index ($g(G)$) | n.a. | 2.478*** (0.461) | n.a. | 0.427*** (0.130) | 0.460*** (0.112) |
| $R^2$ | 0.337 | 0.455 | 0.316 | 0.404 | 0.499 |
| SEE | 0.079 | 0.072 | 0.016 | 0.015 | 0.014 |
| N | 136 | 136 | 144 | 144 | 144 |

Notes: White standard errors in parentheses; ***: 1%; **: 5%; *: 10%. Growth rates are annualized log differences; $g_i(x_{it}) \equiv \ln(x_{it} / x_{it-\tau}) / \tau_i$ .