Journal of Public Economics 177 (2019) 104046

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Global poverty measurement when relative income matters

Martin Ravallion [a,*], Shaohua Chen [b]

[a] Department of Economics, Georgetown University, Washington, DC, USA
[b] Development Research Group, World Bank, Washington, DC, USA

ARTICLE INFO

Article history:
Received 8 May 2018
Received in revised form 12 May 2019
Accepted 16 July 2019
Available online 24 July 2019

JEL classifications:
I32
O10

Keywords:
Global poverty
Poverty lines
Relative income
Inequality

ABSTRACT

It is becoming well-recognized that relative measures are required for a complete picture of the extent of global poverty and its evolution. The standard assumption is that the national mean or median is the relevant comparison income for setting relative poverty lines. That assumption is questionable. The paper proposes a theoretical model of relative comparisons, which can be downward or upward in the distribution of income. On calibrating the approach to a new data-set on national poverty lines, the Gini-adjusted national mean emerges as the relevant comparison income. Calculations are provided of the implied global poverty measures since 1990. Both non-welfarist (capabilities-based) and welfarist interpretations are provided. Declining poverty incidence is indicated, but at a slower rate than the corresponding absolute measures suggest.

1. Introduction

It is now widely acknowledged (across multiple disciplines) that people are concerned about their relative incomes within some group of comparators.[1] In this light, how should we measure global poverty? Explicitly relative lines were proposed by Fuchs (1967), who suggested that the US poverty line 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 in the OECD and Eurostat, and is used by many national governments in the OECD (though 60% of the median is more common than 50%).[2] The UN's Sustainable Development Goals also include monitoring the share 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.[3]

Advocates of such relative 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), among others.

One concern with the Fuchs proposal and all its variants stems from the fact that the monetary line then has an elasticity of unity with respect to the median or mean. 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). By contrast, Kakwani (1986), Foster (1998) and Ravallion and Chen (2011) have proposed "weakly relative" lines that rise with the mean (or median), but with an elasticity less than unity.

An issue that has received little or no attention in the literature is what the comparison income should be at national level. The literature

* Corresponding author at: Department of Economics, Georgetown University, Washington, DC 20057, USA.
E-mail address: mr1185@georgetown.edu (M. Ravallion).

[1] Influential early contributions sociology and social psychology include Davis (1959) and Runciman (1966). In the context of understanding poverty see (inter alia) Abel-Smith and Townsend (1966), Townsend (1979) and Walker (2014). Contributions in economics include Duesenberry's (1949), Hirsch Drewnowski (1977), Frank (1985), Rayo and Becker (2007) and 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). Surveys of this literature can be found in Frey and Stutzer (2002), Senik (2005) and Clark et al. (2008).

[2] 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).

[3] 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).

https://doi.org/10.1016/j.jpubeco.2019.07.005
0047-2727/

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.[4] However, accepting that relative comparisons are welfare-relevant does not 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 coworkers) 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 relative poverty measurement and proposes new measures. We follow the existing literature in two key respects: First, to be judged "not poor" we say that a person needs to be neither absolutely poor nor relatively poor, as judged by the predicted poverty line in their country of residence. Following Atkinson and Bourguignon (2001), this can be interpreted as anchoring poverty measurement to two capabilities, in the sense of Sen (1985), namely a "subsistence" capability and a "social inclusion" capability. This is a non-welfarist interpretation of poverty, as being anchored to capabilities rather than "utility." However, we will note some implications of a welfarist interpretation. Second, we use national poverty lines to identify a schedule of relative global lines, also in keeping with the literature on global poverty measurement following Ravallion et al. (1991).[5]

The principal contribution of the paper concerns how the comparison income should be set in relative poverty measures. We question the long-standing assumption that the comparison income level in relativist comparisons at country level is the median or equally-weighted mean.[6] 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.[7] 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.[8] The paper provides a theoretical formulation that encompasses both upward and downward relative comparisons. This offers a new perspective on relative poverty.

The paper implements this approach on new data on national poverty lines assembled for this purpose, and survey-based distributions of consumption or income. Our data on national poverty lines suggest that the rank-weighted mean is the relevant comparison income, with lowest weight given to the richest. This implies that a Gini-discounted mean is called for in setting our relative poverty measures. We implement the new measures on a global basis, including countries at all levels of development. Our estimates draw on 1500 household surveys for 150 countries over 1990–2013.

After explaining our theoretical approach in the next section, we proceed to calibrate it to our new data set on national poverty lines (Section 3). In accounting for how national poverty lines vary across countries, we then show in Section 4 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 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. Section 5 provides our new estimates of global poverty measures. Some new insights also emerge, including that the rich world is making far less progress against poverty. Section 6 concludes.

2. Theoretical approach

We assume that an absolute line is given, but the problem remains of how to measure relative poverty. 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. This can be represented by a welfare function for household i in country j of the form:

$$u_{ij} = u\left(y_{ij}, y_{ij}/m_{ij}^{*}\right) \tag{1}$$

where y_{ij} is the individual's own consumption and m_{ij}^{*} is the individual's comparison income. The function $u(.)$ is assumed to be strictly and smoothly increasing in both arguments ($u_{y}' > 0$ and $u_{y/m}' > 0$ in obvious notation). Interpreting the distinction made by Davis (1959), 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."

A welfare function such as (1) can be 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 relative poverty line, z_j, is defined by:

$$u(z_j, z_j/m_j) = \overline{u}_j^{z} \tag{2}$$

This gives z_j as (implicitly) an increasing function of m_j for given $\overline{u}_j^{z}$, 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) $u_{ij} < \overline{u}_j^{z}$. 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 (2) will never be strongly relative given that the welfare function is strictly increasing in own consumption at given relative consumption. 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. More precisely, if (1) can be written as$u_{ij} = \tilde{u}(y_{ij}/m_j$

[4] 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.

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

[6] While our focus is on global poverty, it can also be noted that studies of the effects of relative income on subjective welfare have relied at times on equally-weighted means, as in Hagenaars and van Praag (1985) and Luttmer (2005).

[7] 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.

[8] Note that the median is unresponsive to small changes in incomes sufficiently far above (or below) the median.

) then the welfare-consistent poverty line takes the strongly-relative form $z_j = k_j m_j$ where $k_j \equiv \tilde{u}^{-1}(\overline{u}_j^z)$.

Also notice that relative poverty measures based on the above formulation need not be globally monotonic in y_{ij} for those deemed to be poor, as noted by Decerf (2017). 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 $u(y_{ij}^e, y_{ij}^e/\overline{m}) = u(y_{ij}, y_{ij}/m_j)$ for some globally constant reference mean $\overline{m}$. This is assured for a broad class of global poverty measures.[9]

The question is still begging: Are all income levels in society equally important in relative comparisons? The literature has said rather little about the appropriate comparison group in discussing relative poverty.[10] 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, although this has been ignored in the subsequent literature on poverty.

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 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.[11]

We propose an approach that encompasses both these "downward" and "upward" looking relativist comparisons. Our approach is 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 random draws of pairs of incomes in the country of residence in assessing how she is doing relative to others.[12] Naturally, she focuses more on the lower (upper) income within each pair if she makes downward (upward) comparisons. More generally one can imagine that she picks a comparison point somewhere in the (closed) interval between the two incomes, depending on whether she tends to look upward or downward. To formalize this idea, let $\varphi(y_{kj}, y_{lj})$ 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 $\varphi(y_{kj}, y_{lj})$ is a point somewhere in the closed interval $[\min(y_{kj}, y_{lj}), \max(y_{kj}, y_{lj})]$:

$$\varphi\left(y_{kj}, y_{lj}\right) \equiv (1-\delta)\ \min\left(y_{kj}, y_{lj}\right) + \delta\ \max\left(y_{kj}, y_{lj}\right) \quad \text{where } \delta \in [0, 1] \tag{3}$$

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:

$$m_j^* = \frac{1}{N_j^2}\sum_{k=1}^{N_j}\sum_{l=1}^{N_j}\varphi\left(y_{kj} y_{lj}\right) \tag{4}$$

On noting that $\sum\sum(y_{kj} + y_{lj}) = 2N_j^2 m_j$ and $\min(y_{kj} y_{lj}) = (y_{kj} + y_{lj} - |y_{kj} - y_{lj}|)/2$, with some algebraic manipulation we can re-write equation (4) as[13]:

$$m_j^* \cong \left[1-(1-2\delta)G_j\right] m_j \tag{5}$$

where G_j is the Gini coefficient for country j:

$$G_j \equiv \sum_{k=1}^{N_j}\sum_{l=1}^{N_j}\frac{|y_{kj} - y_{lj}|}{2N_j^2 m_j} \tag{6}$$

The comparison income in (5) encompasses both upward and downward relativism, depending on the parameter δ. 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 Sheshinski (1972) and Sen (1976), $(1 - G_j)m_j$, though in a different context, namely in measuring social welfare.[14] In that special case, one can also interpret the comparison income as a specific weighted mean. To see how, let incomes be ordered as $y_{1j} \geq y_{2j} \geq , \ldots, \geq y_{N_j}$. Then the comparison income can be re-written as[15]:

$$m_j^* = \frac{2}{N_j^2}\sum_{i=1}^{N_j} i y_{ij} \tag{7}$$

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.

[9] 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 $\overline{m}$) 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).

[10] 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).

[11] 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.

[12] 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.

[13] The approximation requires large N to be accurate.

[14] 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.

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

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

$$z_j = z\left(m_j^*, \overline{u}_j^z\right) \tag{8}$$

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

$$z'_m \equiv \frac{\partial z_j}{\partial m_j^*} = \frac{u'_{y/m} z_j / m_j^*}{m_j^* u'_y + u'_{y/m}} > 0; \quad z'_u \equiv \frac{\partial z_j}{\partial \overline{u}_j^z} = \frac{m_j^*}{m_j^* u'_y + u'_{y/m}} > 0 \tag{9}$$

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:

$$0 < \eta_j \equiv \frac{\partial \ln z_j}{\partial \ln m_j^*} = \frac{1}{1 + m_j^* MRS_j} < 1 \tag{10}$$

(where $MRS_j \equiv u_y'/u_{y/m}'$). We cannot predict how η_j varies with m_j^*; the necessary and sufficient condition for η_j to rise with 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[16]:

$$P_j = P\left[m_j / z\left(m_j^*, \overline{u}_j^z\right), L_j\right] \quad \text{with} \quad P'_{m/z} < 0 \tag{11}$$

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.) Holding $\overline{u}_j^z$ and L_j constant, we then have (taking log derivatives):

$$\frac{\partial \ln P_j}{\partial \ln m_j} = \frac{\partial \ln P_j}{\partial \ln (m_j / z_j)} \left(1 - \eta_j\right) < 0 \tag{12}$$

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. Thus our measure satisfies the weak relativity axiom of Ravallion and Chen (2011).

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 (11) can be re-written as[17]:

$$P_j = P\left[m_j / z\left(m_j^*, \overline{u}_j^z\right), G_j\right] \quad \text{with} \quad P'_G > 0 \tag{13}$$

Then we find that (holding m_j and $\overline{u}_j^z$ constant)[18]:

$$\left(\frac{\partial \ln P_j}{\partial \ln G_j}\right)_{dm=0} = \frac{\partial \ln P_j}{\partial \ln (m_j / z_j)} \cdot \frac{\eta_j (1 - 2\delta) G_j}{\left[1 - (1 - 2\delta) G_j\right]} + \left(\frac{\partial \ln P_j}{\partial \ln G_j}\right)_{dm/z=0} \tag{14}$$

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 also examine this issue empirically.

3. Data on 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). 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, 2010, 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.[19] 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).[20] 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).[21] The survey dates range from 2004 to 2012, with a median of 2011.

Fig. 1 gives density functions for the poverty lines, survey means and medians. The skewness evident in Fig. 1 is as one would expect. The poverty lines are skewed further to the right than the medians, which are skewed further than the means. The range in national poverty

[16] 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).

[17] 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.

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

[19] 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).

[20] 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.

[21] 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.

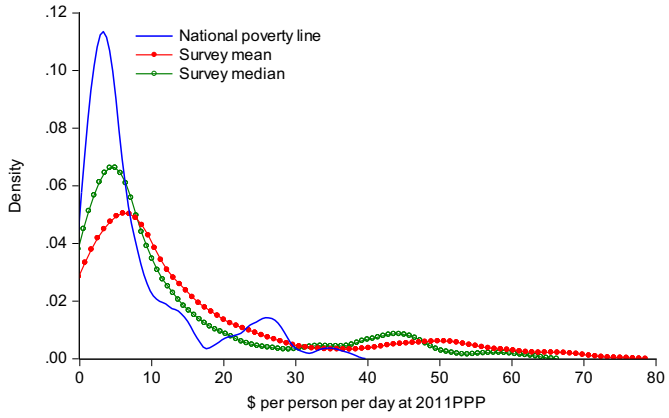

Fig. 1. Kernel density functions.

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 (Fig. 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. In past work the (equally-weighted) mean has been the main predictor.

Fig. 2(a) plots the data for the full sample (including OECD) against mean income.[22] Fig. 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 Fig. 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 consistently with past evidence (Ravallion et al., 1991; Ravallion, 2008). 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 the regression line is 0.485 (White s.e. = 0.020).[23] The overall elasticity (using a log-log regression) is 0.863 (s.e. = 0.027).[24]

It might be argued that the true relationship is not as strong as Fig. 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).[25]

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.[26] 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 Fig. 2.

Third, there may be omitted effects on national poverty lines correlated with mean income (and/or inequality). The political economy in specific countries may influence national poverty lines quite independently of the relativist comparisons in our theoretical model. It is of interest to see if the mean income effect persists when one allows for a country effect. This can be done by exploiting an alternative method of deriving national poverty lines, namely to find the lines that are implicit in data on the national 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*.[27] 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.

We have seen that national poverty lines are higher in richer countries and tend to rise over time with growth in the mean, and that this pattern is also found among the developing countries, most of which claim to use "absolute" lines. The specific source of this implicit relativism varies. In some cases the bundle of goods changes, as goods that were once thought of as luxuries come to be seen as essential for not being considered poor. Elsewhere (as noted above) national lines do not use an explicit bundle of goods but look instead at the expenditure

[22] 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).

[23] All standard errors of regression coefficients in this paper are corrected for a general form of heteroscedasticity using White's (1980) method.

[24] If we use the median instead of the mean, the slope is 0.564 (0.017) and the elasticity is 0.816 (0.026).

[25] Using the median instead for the non-OECD countries, the slope is 0.559 (0.042) and the elasticity is 0.740 (0.040).

[26] For a review of the methods used to set national poverty lines see Ravallion (2012).

[27] 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 $F_{it}^{-1}(H_{it})$.

or income level at which nutritional requirements are met in expectation. However, the relationship between nutritional intakes and income will of course change over time with changes in tastes and activity levels. The relationship we see between the poverty line and the mean can be thought of as the reduced form of a more complex relationship between the various free parameters of national lines and the overall mean.

Notice that the US is an outlier in Fig. 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.[28] 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 lower real lines—indeed, a (non-official) line that is roughly comparable with prevailing poverty lines today in the world's poorest countries.[29] 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 Fig. 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.[30] 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 Fig. 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 (Fig. 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 countries. Naturally then, as the mean rises, the ratio of the poverty line to the mean tends to fall, as can be seen in Fig. 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 (Fig. 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.

4. Calibration to national poverty lines

We begin with the following encompassing specification based on a linearization of Eq. (8) using (5) as the comparison income:

$$z_j = \alpha + \beta[1-(1-2\delta)G_j]m_j + \varepsilon_j \quad (j = 1, \ldots, n) \tag{15}$$

Here α, β, δ are parameters to be estimated and ε_j is an error term. Notice that if $\alpha = 0$ and $\delta = 0.5$ then (in expectation) we have the strongly-relative poverty lines 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.

We assume that $E(\varepsilon_j|m_j, G_j) = 0$ for identifying the comparison income from (1) at given reference welfare levels ($\overline{u}_j^z$). This may be considered a strong assumption once one notes that the error term implicitly includes $\overline{u}_j^z$ (Eq. 8). If countries with a higher mean tend to have more generous reference welfare levels for defining poverty then this will impart a bias to our estimates. However, this is not an issue if one does not demand welfare consistency. If instead one follows the capability-based approach of Atkinson and Bourguignon (2001) then one interprets the (predicted) national lines as reflecting the costs of social inclusion in different countries (with the absolute line interpreted as being required for the subsistence capability). By this interpretation, social inclusion requires that one lives above the reference level of welfare in the country of residence. Then we can interpret the gradient w.r.t. the comparison income as including any effect on that reference level of welfare. We will follow the Atkinson-Bourguignon capabilities-based interpretation, but note implications for a strict welfarist interpretation. We return to this point once we have estimated the regressions.

The estimates for Eq. (15) are found in column (1) of Table 1.[31] 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. 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:

$$z_j = \alpha + \beta(1-G_j)m_j + \varepsilon_j \tag{16}$$

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. Given the bearing that the functional form has on our subsequent results, we also performed the Ramsey RESET test for functional form; the linearity assumption in (16) cannot be rejected in any of the specifications.[32] 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.

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

[29] 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).

[30] For a survey of the literature see Ravallion (2016a, Part 2).

[31] 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 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, 2018).

[32] 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).

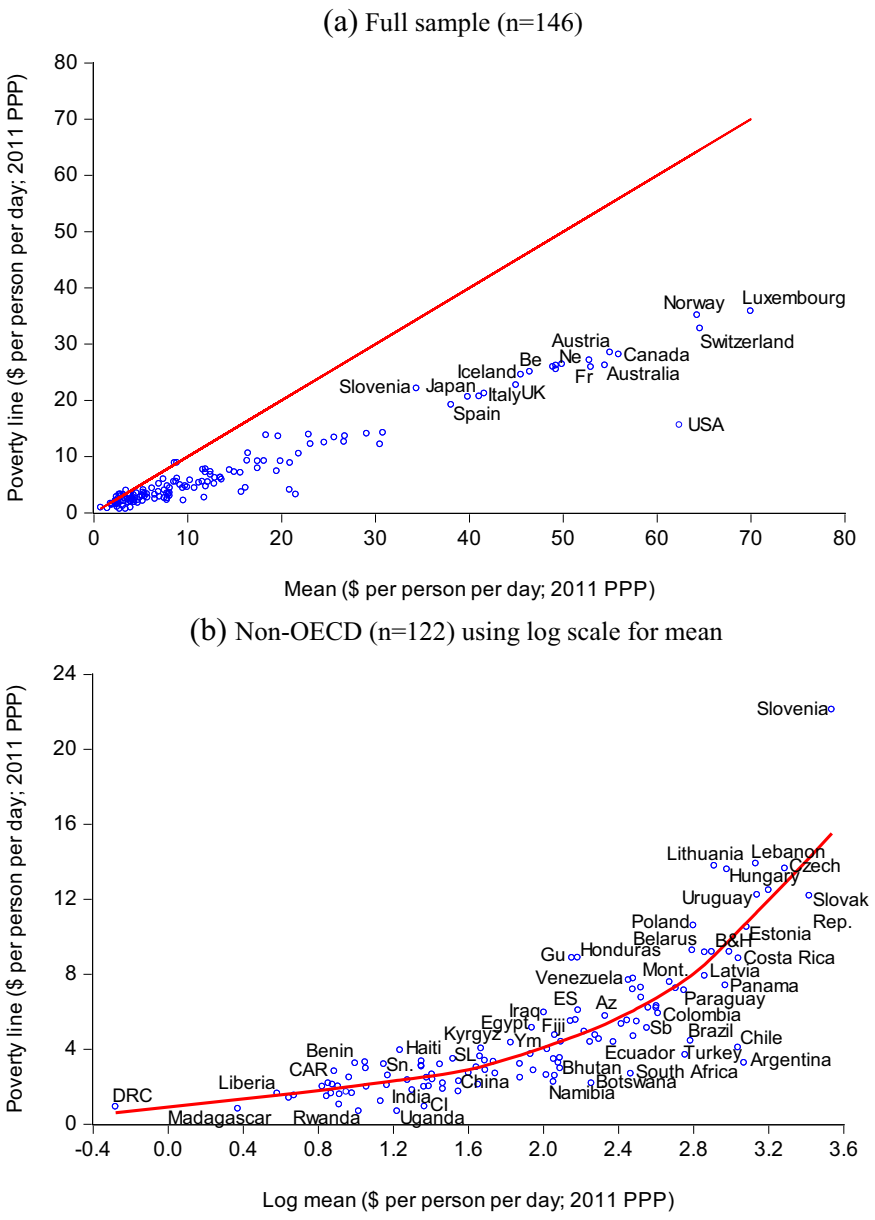

Fig. 2. Poverty lines across countries Note: Empirical non-parametric regression (locally weighted polynomial). Az: Azerbaijan; Be: Belgium; B&H: Bosnia and Herzegovina; CAR: Central African Republic; CI: Cote d'Ivoire; ES: El Salvador; Fr: France; Gu: Guatemala; Mont.: Montenegro; Ne: Netherlands; Sb: Serbia; Sn.: Senegal; SL: Sierra Leone; UK: United Kingdom; USA: United States of America; Ym: Yemen.

So our data suggest that high-inequality countries tend to have lower national poverty lines at given mean income. This is consistent with our theoretical argument made about downward relative-income comparisons in Section 2.

We make six remarks about our preferred estimates based on (16). First, recall that correlated latent heterogeneity in the reference welfare levels ($\overline{u}_j^z$) can lead to a bias in our estimates of the parameters of the welfare-consistent relative lines. To see this more clearly, note that our theoretical model implies that the error term in (16) include a term $\overline{u}_j^z$. Then it is readily verified that $\hat{\beta}_{OLS}$ converges in large samples to $\beta + \phi\pi$ where $\phi > 0$ is the regression coefficient of ε_j on $\overline{u}_j^z$ and $\pi \geq 0$ is the regression coefficient of $\overline{u}_j^z$ on m_j^*. Thus $\hat{\beta}_{OLS} \geq \beta$. (Clearly our IV estimate above 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. Our estimates can then be interpreted as an upper bound to the welfare-consistent relative

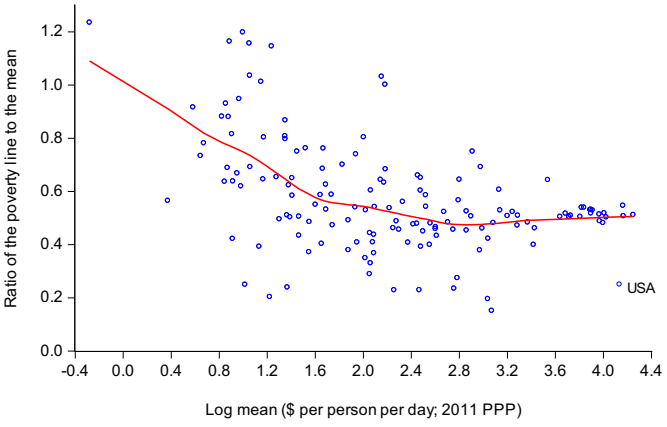

Fig. 3. Ratio of the national poverty line to the current survey mean.

Table 1
Regressions for national poverty lines.

	(1)	(2)	(3)	(4)	(5)	(6)
	Full sample	Non-OECD	Full sample	Non-OECD	Full sample	Non-OECD
	OLS	OLS	OLS	OLS	IV	IV
Intercept (α)	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 (β)	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 (δ)	−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; OLS: Ordinary Least Squares; IV: Instrumental variable.
*** 1%.

poverty lines. At the other extreme, the lower bound assumes maximum bias, so that $\beta = 0$, giving a standard absolute line. As noted, this is not a concern for the non-welfarist interpretation.

The second remark concerns the implied elasticity of the poverty line w.r.t. the comparison mean, which (as noted in the introduction) is a key parameter choice in poverty measurement. Focusing on the non-OECD group, across the range of the data, we find that 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 (16) is $1/[1 + \alpha/(\beta m_j^*)]$and the elasticity rises with m_j^* (the elasticity of MRS_j to m_j^* is −2). Fig. 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^*).

A third remark concerns the US. This country 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).[33] 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$).

A fourth remark relates to the (widespread) use of the median as the comparison income. 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 (16) we obtain the following regression:

$$z_j = \underset{(0.183)}{0.866} + \underset{(0.631)}{0.877}(1-G_j)m_j - \underset{(0.519)}{0.139}\,M_j + \hat{\varepsilon}_j \quad R^2 = 0.956, \text{SEE} = 1.737, \text{n} = 145 \tag{17}$$

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.[34]

Fifth, it is of interest to note that 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 about half the mean. Fig. 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.[35] 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 Fig. 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.

A final remark concerns 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, \ldots, 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 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

[33] 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.

[34] The SEE using the median is 1.772 versus 1.723 using the Gini-discounted mean and 2.080 for the ordinary mean.

[35] 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).

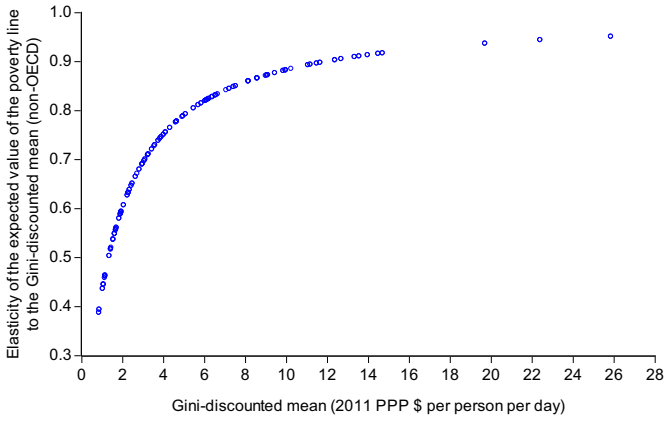

Fig. 4. Elasticity of the poverty line to the Gini-discounted mean.

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 an absolute line.

5. Implications for global poverty measures

Recall that to be judged "not poor" we say that a person needs to be neither absolutely poor nor relatively poor, as judged by the predicted poverty line in their country of residence. Thus, the absolute line of $1.90 a day is taken to override the relative line if the latter is smaller. With some rounding off, our schedule of poverty lines based on the results of the previous section is:

$$z_j = 1.90 + \max[0.7(1-G_j)m_j - 1.00, 0] \quad (18)$$

(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 our relative lines 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 our relative lines are very close to being a strongly relative measure in the OECD.

Recall that a welfarist interpretation would identify the absolute line ($1.90) as the lower bound to the welfare-consistent line while Eq. (18) is interpreted as the upper bound. To accommodate both interpretations, we will provide the absolute line as well as our hybrid lines based on (18), incorporating relative poverty.

(a) Weakly relative poverty lines with distribution adjustmentversus strongly relativelines

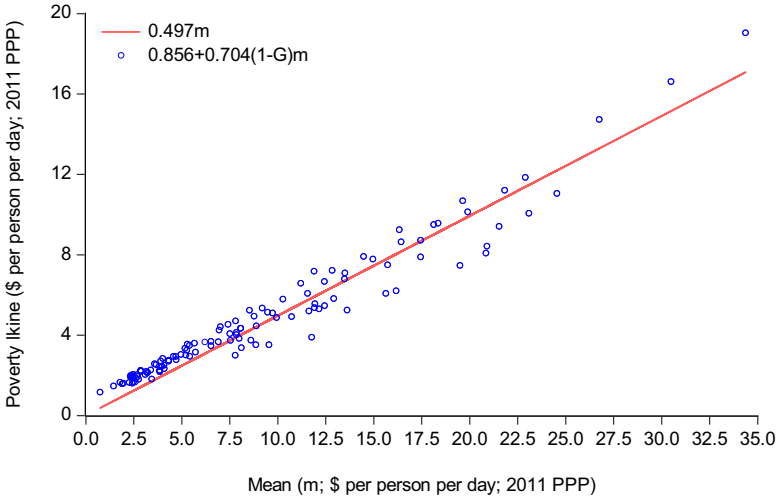

(b) Weakly relative poverty lines with and without distributional adjustment

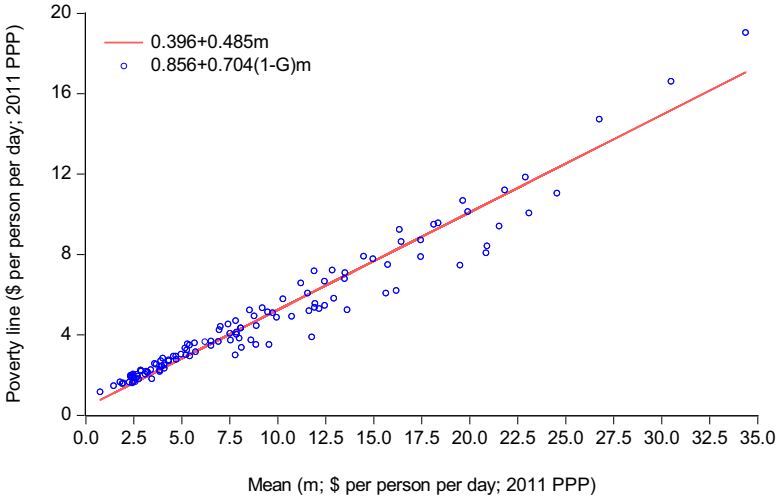

Fig. 5. Comparison of relative lines (non-OECD) Note: The distributional adjustment uses the Gini-discounted mean (see text).

In our empirical implementation we confine attention to the headcount index (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. Countries stay in their base-year regional category over time. For developing countries we use *PovcalNet* while for other countries we use the *Luxembourg Income Study* (LIS).[36] 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).

We find that the percentage of the world's population living below our global relative lines based on (16) has fallen over time (Fig. 6(a)). The trend rate of decline is 0.7 percentage points per year (a regression coefficient on the year of −0.688; s.e. = 0.028). The trend rate of decline in the number of people living below our relative lines is 11 million per year (−10.971; s.e. = 3.109) (Fig. 6(b)). Fig. 6(a) and (b) also give the series for the corresponding absolute line, which (as expected) implies a steeper decline in the poverty measures. The numbers of people who are poor by our proposed new measures but not absolutely poor by the World Bank's $1.90 line has tended to rise over time (Fig. 6(b)).

There are some marked regional differences, as summarized in Fig. 7.[37] Sub-Saharan Africa (SSA) is the region with the highest poverty rate, except in 1990, when East Asia had the highest rate. By 2013, we find that about half of SSA's population falls below our new relative lines, as compared to 30% in East Asia, which is also the region with the steepest decline in poverty. While there was little progress in most regions in the 1990s, this changed in the new millennium. 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. 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.

For high income countries, the poverty rate in 2013 is 19% (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 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 is higher in the developing world—roughly double the rate found in high-income countries. By 2013, a higher share of the population of the developing world was relatively poor but not absolutely poor than for the rich world. In all, 91% of poverty by our measure is in the developing world, where there has also been a marked re-alignment of poverty, notably between East Asia, South Asia and Africa, as can be seen in Fig. 7.

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 3, 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. Table 2 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; this is equivalent to a country-fixed effects regression in the log levels.[38] 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 poverty measure is theoretically ambiguous.

We see that higher inequality is associated with higher poverty rates and a higher mean comes with lower poverty. 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. 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.

6. Conclusions

We have proposed a new approach to measuring relative poverty globally. 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 theoretical basis for the choice. The past use of either the mean or median for relativist comparisons in the literature on poverty is 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 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 that has been documented before in the literature is not found for our new series of weakly relative lines. 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

[36] 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).

[37] The working paper version, Ravallion and Chen (2018, Tables 2 and 3), provide detailed estimates by region over time.

[38] Note that these are descriptive regressions showing the partial correlations, rather than estimates of a causal model. Also note that (like any regression) these are averages; we cannot rule out the possibility that a specific country with no change in the mean and rising inequality may still see a falling poverty measure. However, we found no case in which the mean fell only slightly (by no >1%) and the Gini index rose yet the poverty rate fell.

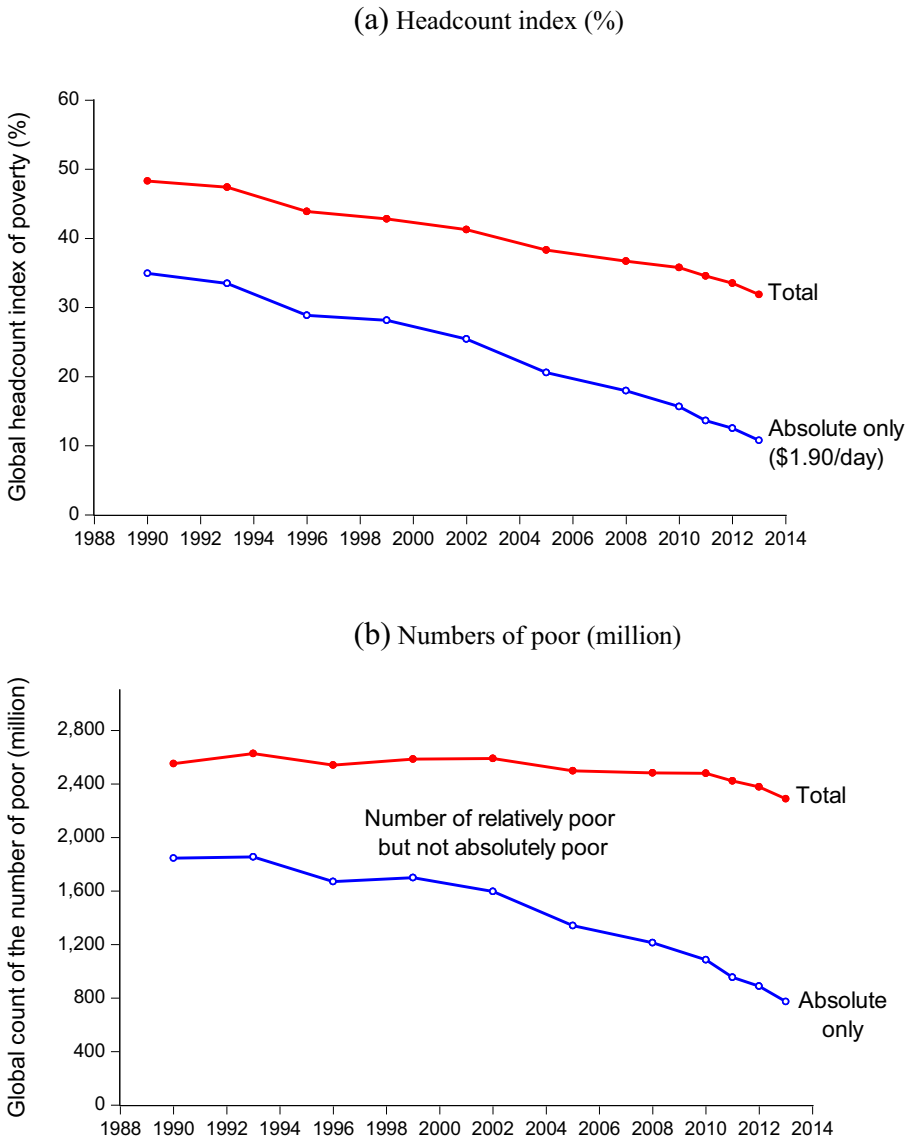

Fig. 6. Global poverty measures.

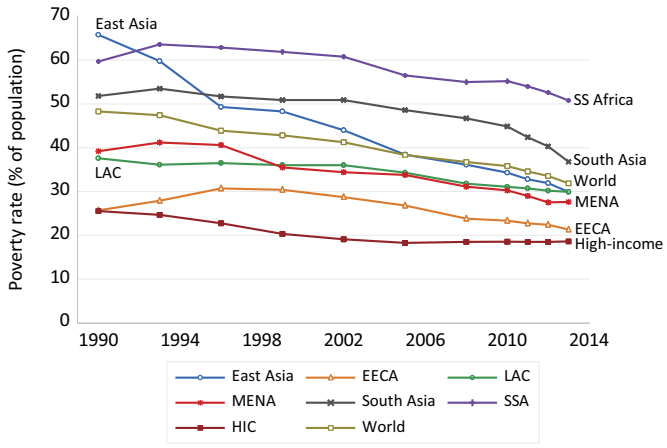

Fig. 7. Regional poverty rates Note: EECA: Eastern Europe and Central Asia; LAC: Latin America and the Caribbean; MENA: Middle-East and North Africa; SSA: Sub-Saharan Africa; HIC: High income countries not elsewhere classified.

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 measures is primarily about the distribution of the gains from economic growth.

Whether one focuses on absolute poverty or our relative poverty, the incidence of poverty is appreciably higher in the developing world. Over 90% of the poor by our proposed lines are found in the developing world. However, the developing world has been making greater progress over time against poverty, judged by either absolute or relative measures. 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 poverty measures are responsive to both the mean and inequality, although the relative measure responds less elastically.

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

Table 2
Cross-country regressions for changes in poverty measures.

	Growth rate in the headcount index				
	Absolute poverty lines ($1.90 a day)		New poverty lines ($z_j = \$1.90 + \max[0.7(1 - G_j)m_j - \$1.00, 0]$)		
Growth rate in the survey mean ($g(m)$)	−2.239***	−2.394***	−0.434***	−0.453***	−1.664***
	(0.280)	(0.283)	(0.050)	(0.048)	(0.193)
$g(m) \times$ 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***	n.a.	0.427***	0.460***
		(0.461)		(0.130)	(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; Growth rates are annualized log differences; $g_i(x_{it}) \equiv \ln(x_{it}/x_{it-\tau})/\tau_i$.
*** 1%.

received from Rebecca Blank, Benoît Decerf, Francisco Ferreira, Stephan Klasen, Dominique van de Walle, Michael Woolcock, seminar participants at Georgetown University and the Paris School of Economics, and the journal's editor and referees. Financial support for this research was provided by the World Bank. These are the views of the authors and need not reflect those of their employers, including the World Bank.

References

Abel-Smith, Brian, Townsend, Peter, 1966. The Poor and the Poorest: A New Analysis of the Ministry of Labour's Family Expenditure Surveys of 1953–54 and 1960. Bell, London.
Atkinson, Anthony B., 1987. On the measurement of poverty. Econometrica 55, 749–764.
Atkinson, Anthony B., 1998. Poverty in Europe. Blackwell, Oxford.
Atkinson, Anthony B., Bourguignon, Francois, 2001. Poverty and inclusion from a world perspective. In: Stiglitz, J.E., Muet, P.-A. (Eds.), Governance, Equity and Global Markets. Oxford University Press, Oxford.
Blackburn, McKinley, 1994. International comparisons of poverty. Am. Econ. Rev. 84 (2), 371–374.
Blank, Rebecca M., 2008. How to improve poverty measurement in the United States. J. Policy Anal. Manag. 27 (2), 233–254.
Chen, Shaohua, Ravallion, Martin, 2001. How did the world's poor fare in the 1990s? Rev. Income Wealth 47 (3), 283–300.
Chen, Shaohua, Ravallion, Martin, 2004. How have the World's poorest fared since the early 1980s? World Bank Res. Obs. 19, 141–170.
Chen, Shaohua, Ravallion, Martin, 2010. The developing world is poorer than we thought, but no less successful in the fight against poverty. Q. J. Econ. 125 (4), 1577–1625.
Chen, Shaohua, Ravallion, Martin, 2013. More relatively poor people in a less absolutely poor world. Rev. Income Wealth 59 (1), 1–28.
Chen, Xi, 2015. Relative deprivation and individual well-being. IZA World Labor 140.
Citro, Constance, Michael, Robert (Eds.), 1995. Measuring Poverty: A New Approach. National Academy Press, Washington, DC.
Clark, Andrew, Oswald, Andrew, 1996. Satisfaction and comparison income. J. Public Econ. 61, 359–381.
Clark, Andrew, Frijters, Paul, Shields, Michael, 2008. Relative income, happiness and utility: an explanation for the Easterlin paradox and other puzzles. J. Econ. Lit. 46 (1), 95–144.
Cohn, Alain, Fehr, Ernst, Herrmann, Benedikt, Schneider, Frédéric, 2014. Social comparison and effort provision: evidence from a field experiment. J. Eur. Econ. Assoc. 12 (4), 877–898.
Davis, James, 1959. A formal interpretation of the theory of relative deprivation. Sociometry 22, 280–296.
Decerf, Benoit, 2017. Why not consider that being absolutely poor is worse than being only relatively poor? J. Public Econ. 152, 79–92.
Drewnowski, Jan, 1977. Poverty: its meaning and measurement. Dev. Chang. 8, 183–208.
Duclos, Jean-Yves, Makdissi, Paul, 2004. Restricted and unrestricted dominance for welfare inequality and poverty orderings. J. Public Econ. Theory 6, 145–164.
Duesenberry, James S., 1949. Income, Saving and the Theory of Consumer Behavior. Harvard University Press, Cambridge, Mass.
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. J. Econ. Inequal. 14, 141–172.
Foster, James, 1998. Absolute versus relative poverty. Am. Econ. Rev. Pap. Proc. 88 (2), 335–341.
Frank, Robert H., 1985. Choosing the Right Pond: Human Behavior and the Quest for Status. Oxford University Press, New York.
Frey, Bruno, Stutzer, Alois, 2002. What can economists learn from happiness research? J. Econ. Lit. XL, 402–435.
Fuchs, Victor, 1967. Redefining poverty and redistributing income. Public Interest 8, 88–95.
Garroway, Christopher, de Laiglesia, J.R., 2012. On the relevance of relative poverty for developing countries. OECD Development Centre Working Paper 314. OECD, Paris.
Hagenaars, Aldi J.M., van Praag, Bernard M.S., 1985. A synthesis of poverty line definitions. Rev. Income Wealth 31, 139–154.
Hunter, Robert, 1904. Poverty. MacMillan Company, London.
Jolliffe, Dean, Prydz, Espen, 2016. Estimating international poverty lines from comparable national thresholds. J. Econ. Inequal. 14 (2), 185–198.
Jolliffe, Dean, Prydz, Espen, 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 University Press, Cambridge.
Kakwani, Nanak, 1993. Poverty and economic growth with application to Côte D'Ivoire. Rev. Income Wealth 39, 121–139.
Kilpatrick, Robert, 1973. The income elasticity of the poverty line. Rev. Econ. Stat. 55, 327–332.
Klasen, Stephan, Krivobokova, Tatyana, Greb, Friederike, Lahoti, Rahul, Pasaribu, Syamsul, Wiesenfarth, Manuel, 2016. International poverty measurement: which way now? J. Econ. Inequal. 14 (199-125).
Knight, John, Song, Lina, Gunatilaka, Ramani, 2009. Subjective well-being and its determinants in rural China. China Econ. Rev. 20 (4), 635–649.
Luttmer, Erzo, 2005. Neighbors as negatives: relative earnings and well-being. Q. J. Econ. 120, 963–1002.
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.
Nolan, Brian, 2007. A Comparative Perspective on the Development of Poverty and Exclusion in European Societies. International Policy Analysis, Bonn.
Organization for Economic Co-Operation and Development (OECD), 2008. Growing Unequal? Income Distribution and Poverty in OECD Countries. OECD, Paris.
Orshansky, Mollie, 1965. Counting the poor: another look at the poverty profile. Soc. Secur. Bull. 28, 3–29.
Ravallion, Martin, 2008. On the welfarist rationale for relative poverty lines. In: Basu, Kaushik, Kanbur, Ravi (Eds.), Social Welfare, Moral Philosophy and Development: Essays in Honour of Amartya Sen's Seventy Fifth Birthday. Oxford University Press, Oxford.
Ravallion, Martin, 2012. Poverty lines across the world. In: Philip, N. (Ed.), The Oxford Handbook of the Economics of Poverty. Oxford University Press, Jefferson, Oxford.
Ravallion, Martin, 2015. The Luxembourg income study. J. Econ. Inequal. 13 (4), 527–547.
Ravallion, Martin, 2016a. The Economics of Poverty: History, Measurement and Policy. Oxford University Press, New York.
Ravallion, Martin, 2016b. Toward better global poverty measures. J. Econ. Inequal. 14, 227–248.
Ravallion, Martin, Chen, Shaohua, 2011. Weakly relative poverty. Rev. Econ. Stat. 93 (4), 1251–1261.
Ravallion, Martin, Chen, Shaohua, 2018. Welfare-Consistent Global Poverty Measures When Relative Incomes Matter (NBER Working Paper 23739).
Ravallion, Martin, Datt, Gaurav, van de Walle, Dominique, 1991. Quantifying absolute poverty in the developing world. Rev. Income Wealth 37, 345–361.
Ravallion, Martin, Chen, Shaohua, Sangraula, Prem, 2009. Dollar a day revisited. World Bank Econ. Rev. 23 (2), 163–184.
Rayo, Luis, Becker, Gary S., 2007. Evolutionary efficiency and happiness. J. Polit. Econ. 115 (2), 302–337.
Runciman, W.G., 1966. Relative Deprivation and Social Justice. Routledge and Kegan Paul, London.
Sen, Amartya, 1976. Real national income. Rev. Econ. Stud. 43 (1), 19–39.
Sen, Amartya, 1985. Commodities and Capabilities. North-Holland, Amsterdam.
Senik, Claudia, 2005. Income distribution and well-being: what can we learn from subjective data? J. Econ. Surv. 19 (1), 43–63.

Sheshinski, Eytan, 1972. Relation between a social welfare function and the Gini Index of inequality. J. Econ. Theory 4 (1), 98–100.
Short, Kathleen, 2012, "The research supplemental poverty measure: 2011." Current Population Reports P60–244. U.S. Census Bureau.
Smeeding, T.M., Rainwater, L., Rein, M., Hauser, R., Schaber, G., 1990. Income poverty in seven countries: initial estimates from the LIS database. In: Smeeding, T.M., O'Higgins, M., Rainwater, L. (Eds.), Poverty, Inequality and Income Distribution in Comparative Perspective. Urban Institute Press, Washington, DC.
Townsend, Peter, 1979. Poverty in the United Kingdom: A Survey of Household Resources and Standards of Living. Penguin Books, Harmonsworth.
Walker, Robert, 2014. The Shame of Poverty. Oxford University Press, Oxford.
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 University, Oxford.
World Bank, 2015. Purchasing Power Parities and the Real Size of World Economies: A Comprehensive Report of the 2011 International Comparison Program. World Bank, Washington, DC.
Yitzhaki, Shlomo, 1979. Relative deprivation and the Gini coefficient. Q. J. Econ. 93 (2), 321–324.