

# Annual Review of Economics On Measuring Global Poverty

# Martin Ravallion

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



#### www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Annu. Rev. Econ. 2020. 12:167-88

First published as a Review in Advance on June 10, 2020

The Annual Review of Economics is online at economics.annualreviews.org

https://doi.org/10.1146/annurev-economics-081919-022924

Copyright © 2020 by Annual Reviews. All rights reserved

JEL codes: I32, O10

# Keywords

poverty, prices, poverty lines, functionings, nutrition, relative income

#### Abstract

This article critically assesses prevailing measures of global poverty. A welfarist interpretation of global poverty lines is augmented by the idea of normative functionings, the cost of which varies across countries. In this light, current absolute measures are seen to ignore important social effects on welfare, while popular, strongly relative measures ignore absolute levels of living. It is argued that a new hybrid measure is called for, combining absolute and weakly relative measures consistent with how national lines vary across countries. Illustrative calculations indicate that we are seeing a falling incidence of poverty globally over the past 30 years. This is mainly due to lower absolute poverty counts in the developing world. While fewer people are poor by the global absolute standard, more are poor by the country-specific relative standard. The vast bulk of poverty, both absolute and relative, is now found in the developing world.

# **1. INTRODUCTION**

Poverty measures have long been important indicators of social progress. Over 100 years ago, Arthur Bowley (1915, p. 213), the inaugural Professor of Statistics at the London School of Economics, wrote,

There is, perhaps, no better test of the progress of a nation than that which shows what proportion are in poverty; and for watching the progress the exact standard selected as critical is not of great importance, if it is kept rigidly unchanged from time to time.

The use of poverty measures in assessing social progress gained momentum in the second half of the twentieth century, in both rich and poor countries (Ravallion 2016).

Today we find two main approaches to measuring poverty and monitoring progress in reducing it. The first focuses on absolute measures that strive to use poverty lines with constant real value, in keeping with Bowley's (1915) advice. For example, this is essentially what the official poverty measures for the United States strive to do. It is also how the World Bank measures global poverty, aiming to apply a rigidly unchanged real line across countries as well as over time. The second approach uses relative measures for which the poverty line varies in real terms, being set at a constant proportion of the current mean or median—an approach that emerged in the 1960s and became popular in western Europe in the late twentieth century. There has been much debate on the choice between absolute versus relative measures.

This article provides a critical overview of the economic foundations of both approaches and asks, Does either make sense? The article's answer is no. It is argued that a new approach is needed for measuring and monitoring global poverty going forward.

Poverty is taken here to be an objective economic deprivation—low economic welfare, or standard of living.<sup>1</sup> A poverty line is a money metric of welfare, and the international poverty line is the money needed in a specific country and on a particular date to achieve a level of economic welfare fixed across countries (to measure global poverty) and over time (to monitor progress).

How is economic welfare measured? The quality of household surveys is important. While survey data have improved, there are continuing challenges, such as selective compliance in surveys. The assumptions made in measuring household consumption or income can also matter, as do those made about how household income is shared within the household. Price indices are also important. While these can all be challenging issues, they are not intrinsic to poverty measurement (but apply more broadly to the measurement of real incomes and their distribution, including in policy evaluation). So they are passed over here.<sup>2</sup>

A potentially contentious issue that is fundamental to poverty measurement is whether an individual's relative income in the country of residence matters. Here, existing poverty measures tend to opt for one of two very different assumptions, corresponding to the absolute and relative measures above: (*a*) that relative income does not matter to economic welfare or (*b*) that relative income is all that matters.

Neither is plausible. When applied globally, the fixed real line advocated by Bowley (1915) cannot capture relative economic deprivation at the country level or the need for higher outlays for economic well-being in richer countries. However, it is no less obvious that the absolute standard of living, at a given relative income, also matters—thereby ruling out measures in which the poverty line is set at a constant proportion of the mean or median.

<sup>&</sup>lt;sup>1</sup>The term income poverty is sometimes used, although other variables generally enter the calculations. <sup>2</sup>Fuller coverage of these issues, with emphasis on their implications for poverty measurement, can be found in Ravallion (2016) and Atkinson (2019).

While the principle of welfare consistency has value in assessing methodological choices, it does not say anything about the reference level of welfare to not be deemed poor. Practitioners have sought guidance from other sources, including officially stipulated nutritional intakes. Borrowing from Sen's (1983) capabilities approach, nutritional status can be thought of as a key functioning relevant to a person's welfare (or capabilities). However, nutritional status is only one such functioning, and others clearly matter (as Sen recognizes). Once we allow for social inclusion as a welfare-relevant functioning and study how poverty lines vary in practice, a clearer picture emerges of global poverty, without requiring either of the assumptions above.

The article begins with a discussion of the theoretical rationales for absolute and relative poverty measures. Section 3 provides an overview of existing national poverty lines. Sections 4 and 5 review current practices for absolute and relative poverty measures, respectively, and relevant evidence from the literature. Illustrative calculations are also provided of the implied global poverty measures. Section 6 concludes.

# 2. THEORETICAL STARTING POINTS

Two theoretical frameworks can be used in thinking about global poverty measurement—namely, the welfarist approach and the capabilities approach. This article draws on both. This section starts with a welfarist model, according to which someone is said to be poor if her attained level of economic welfare is below some critical level. As an ethical starting point, poverty comparisons are taken to be absolute in some agreed welfare space, although they could be either absolute or relative in the consumption or income space. The issue is then how we think about economic welfare. Here, some useful insights can be found in the capabilities approach.

To encompass both absolute and relative measures, let us assume that the welfare of an individual living in a household with consumption or income y, facing prices  $\mathbf{p}$  (a vector) and with personal characteristics  $\mathbf{x}$  (including household and environmental attributes), can be represented by a function of the form  $v(y, y/m, \mathbf{p}, \mathbf{x})$ , where m is the mean (or median) income of the country of residence. The function v is assumed to be strictly and smoothly increasing in y and nondecreasing in y/m. Absolute poverty measures correspond to the case in which v(.) is invariant to y/m, while relative measures allow v(.) to be strictly increasing in y/m.

The welfare-consistent international poverty line, z, is then defined implicitly by

$$v\left(z,\frac{z}{m},\mathbf{p},\mathbf{x}\right)=\bar{u}.$$
 1.

Here,  $\bar{u}$  is the level of welfare to not be deemed poor. Welfare consistency in global poverty measurement demands that  $\bar{u}$  is fixed across all countries. Under these assumptions, we can see that  $y \le z$  is equivalent to  $v(.) \le \bar{u}$ , noting that the solution of Equation 1 for z can be written as

$$z = z(m, \mathbf{p}, \mathbf{x}, \bar{u}).$$
 2.

If we can identify the function z(.) by observing national poverty lines, then we can retrieve the key features of the underlying welfare function. (Section 3 returns to this issue.) Even without any data, we can immediately notice some implications for the debate on absolute versus relative measures. It is readily verified that the solution for z in Equation 2 will rise with the mean, with a positive elasticity less than unity, making this a schedule of weakly relative lines (as defined by Ravallion & Chen 2011).<sup>3</sup> Strongly relative lines only emerge as the limiting case in which v(.) is invariant to y (at a given y/m) but strictly increasing in y/m. Then  $z = k(\mathbf{p}, \mathbf{x}, \bar{u})m$ .

<sup>&</sup>lt;sup>3</sup>On implicitly differentiating Equation 1 with respect to *m* and solving, one finds that  $0 < \frac{\partial lnz}{\partial lnm} = \frac{1}{1+mMRS} < 1$ , where  $MRS \equiv v_y/v_{y/m}(>0)$ .

This formulation begs the question of what the reference welfare level is for not being deemed poor ( $\bar{u}$  in Equation 1). One might say that this is arbitrary and only require internal consistency given that choice. But this is not very satisfactory, since the choice often matters to the measures and the conclusions drawn, including policy implications. Setting reference welfare levels in measurement can be recognized as a longstanding issue in applied economics. Any price index found in practice has some implicit welfare anchor, and the index value will (in general) vary as the reference welfare level varies. (For example, national consumer price indices are typically anchored to consumption bundles somewhere around the mean or median of the distribution of income.) The measurement challenge is there and cannot be ignored.

The concept of functionings found in the capabilities approach helps us think about  $\bar{u}$  in Equation 1. Economists often think of welfare as a function of commodities consumed—the utility function representing preferences. As has long been recognized, such a utility function cannot plausibly be treated as being independent of personal characteristics. However, it is more believable that welfare is a stable, interpersonally comparable function of what a person can be and do—her functionings.<sup>4</sup> Two functionings have been prominent in the measurement of poverty—namely, nutritional status and social inclusion. The idea of functionings provides an extra structural layer to the standard model of consumer choice—a layer that helps address the identification problem in making interpersonal comparisons for measuring poverty when welfare-relevant personal characteristics vary.

Thus, economic welfare is now represented by a primal welfare function  $u(\mathbf{f})$ , which is taken to be scalar-valued and strictly increasing in the vector of functionings  $\mathbf{f}$ . It is assumed that  $u(\mathbf{f})$  is stable and interpersonally comparable. While there can be welfare-relevant differences in personal characteristics, as represented here by a vector  $\mathbf{x}$ , these are assumed to matter to economic welfare only insofar as they alter functionings, which are taken to be the primitives of economic welfare for the purpose of assessing poverty. This can be made fully consistent with the prior welfare function,  $v(y, y/m, \mathbf{p}, \mathbf{x})$ , given that we can imagine that functionings are generated by higher income, but they also depend on relative income as well as prices and personal characteristics [i.e.,  $\mathbf{f} = \mathbf{f}(y, y/m, \mathbf{p}, \mathbf{x})$ ].

The task of anchoring the reference level of economic welfare for deciding who is poor requires setting a vector of fixed normative functionings  $\mathbf{f}^*$ . For example, the normative function for nutrition may be to reach stipulated nutritional requirements for good health and normal activities, while the normative functionings for social inclusion may be socially acceptable clothing and housing. Thus, we have an answer, in theory at least, as to what the reference level of welfare should be in Equations 1 and 2—namely,  $\bar{u} = u(\mathbf{f}^*)$ . There is still a problem for practice in determining  $\mathbf{f}^*$  and the function u(.), reflecting the trade-offs between functionings.

The above formulation is somewhat more demanding than the capabilities approach, which only asks if the normative functions are attainable for someone with  $y, m, \mathbf{p}$ , and  $\mathbf{x}$ . Sen (1983, 1985) argues that poverty should be assessed by capabilities, defined as the set of all attainable functionings—interpretable as the person's opportunities rather than actual outcomes. In that formulation, there is no explicit welfare function defined on functionings. But this greater generality comes at a cost. Trade-offs can be expected—more of one functioning will tend to mean less of another. The interpersonal comparisons required for deciding if one person is poorer than another will often call for a judgement on the properties of some welfare function of the form  $u(\mathbf{f})$ .

A number of further remarks can be made, of relevance to the subsequent discussion:

<sup>&</sup>lt;sup>4</sup>Sen (1987) discusses the relationship between economic welfare (standard of living) and functionings.

- 1. This is clearly a broader conceptualization of what poverty means than one based solely on monetary income or expenditure. Here, income is simply a convenient metric for representing an underlying multidimensional welfare function. Using a money metric of welfare does not, of course, mean one thinks that people only care about income. Nonetheless, the measures found in practice may often be seen as incomplete, pointing to the need for complementary measures capturing those things left out; common examples include access to nonmarket goods and intrahousehold inequalities. Alternatively, one can think about this dashboard of measures as a multidimensional implementation of the capabilities approach, with economic welfare as one of the dimensions.<sup>5</sup>
- 2. This can be thought of as an approach to measuring absolute poverty, but it is absolute in the space of welfare. Relative income is taken to enter the welfare function directly, but we can also imagine that the vector **x** includes aspects of the environment, as relevant to welfare. Nutritional adequacy depends on both food consumption (quantities and composition) and the local social environment, insofar as this influences how diets are evaluated, or the local health environment, which matters to nutrient absorption. Similarly, social inclusion can be expected to depend on how personal consumption compares to the average income in the place of residence, either through perceived relative deprivation or risk sharing.
- 3. The above formulation has not equated economic welfare with the maximand of choice over consumption. We can readily imagine a deeper model in which functionings are related to quantities of goods consumed relative to social norms, with quantities chosen to maximize economic welfare. The function z(.) in Equation 2 is then the consumer's expenditure function. This provides a rationale for the poverty line as the cost of a poverty bundle of goods—namely, the vector of utility-compensated demands,  $\mathbf{q}[\mathbf{p}, m, \mathbf{x}, u(\mathbf{f}^*)]$ , along the indifference surface corresponding to  $u(\mathbf{f}^*)$ :<sup>6</sup>

$$z = \mathbf{pq} \left[ \mathbf{p}, m, \mathbf{x}, u(\mathbf{f}^*) \right].$$
 3.

- 4. However, it may be considered a strong assumption that  $u(\mathbf{f})$  is the choice maximand. We might evaluate a person's economic welfare by  $u(\mathbf{f})$  without assuming that the person maximizes  $u(\mathbf{f})$ . Instead, we might postulate a subjective welfare function,  $U[u(\mathbf{f}), m, \mathbf{x}]$ . Some elements of  $\mathbf{x}$ , such as personality traits, can matter to U but not u. (For example, one can be "poor but happy.") Consumption choices will only maximize  $u(\mathbf{f})$  if  $U[u(\mathbf{f}), m, \mathbf{x}]$  for an affine transform of U(.)] is additively separable between  $u(\mathbf{f})$  and  $m, \mathbf{x}$ . Without such separability, the choices made and attained functionings will also depend on the properties of the subjective welfare function, U. An implication is that we cannot in general infer a unique economic welfare function for assessing poverty status by observing only how consumption choices and attained functionings vary with prices, incomes, and characteristics.<sup>7</sup> An external judgement is required.
- 5. Even if the normative functionings are identical between different countries, the poverty lines required to reach them vary with  $\mathbf{p}$ , m, and  $\mathbf{x}$ , interpretable as differences in the cost of attaining  $\mathbf{f}^*$ . As noted, richer countries face higher costs, ceteris paribus.

<sup>&</sup>lt;sup>5</sup>A further issue is whether one keeps these other dimensions separate or aggregates them into a composite index. For further discussion of this topic, see Ravallion (2011).

<sup>&</sup>lt;sup>6</sup>This requires some extra technical assumptions, although familiar ones from consumer theory.

<sup>&</sup>lt;sup>7</sup>Given data on ( $\mathbf{p}, \mathbf{q}, y, \mathbf{x}$ ), suppose that we can integrate back to some indirect utility function  $v(y, \mathbf{p}, \mathbf{x})$  under standard assumptions. Then, any function  $V[v(y, \mathbf{p}, \mathbf{x}), \mathbf{x}]$  will be consistent with the same data. This relates to a longstanding identification problem in welfare measurement, such as in setting equivalence scales (Browning 1992).

6. Calibrating the poverty line to only one (or some subset) of the functionings relevant to welfare will not in general be welfare consistent. In particular, suppose that one finds a nutrition-based poverty line,  $z_n$ , such that nutritional adequacy  $(f_n^*)$  is just reached, i.e. (in obvious notation),

$$f_n\left(z_n,\frac{z_n}{m},\mathbf{p},\mathbf{x}\right) = f_n^*.$$

This will not yield  $z_n = z$  unless one assumes that economic welfare depends only on nutritional status. Section 4 returns to this point.

Prevailing practice can be interpreted as an approximation—often a rough approximation, as we see below—to the theoretical ideal described above.<sup>8</sup>

#### **3. NATIONAL POVERTY LINES**

National poverty lines have been used as data for setting international poverty lines.<sup>9</sup> So we should understand how those lines are set and see how they vary across countries.

Stipulated nutritional requirements (such as based on WHO 1985) have long been used in setting poverty lines at the country level. In various ways, virtually all the national lines found in developing countries are so anchored (Ravallion 2012).<sup>10</sup> (Most of the rich countries use strongly relative lines, which Section 5 returns to.) Certain generic health risks are known to rise when intakes fall persistently below these levels, although specific individuals can still live healthy and normal lives. These requirements are averages, embodying assumptions about the desired growth paths (for children) and activity levels.

A common method of setting national lines is to identify a bundle of foods that attain the stipulated nutritional requirements and then price that bundle locally. An allowance for nonfood spending is invariably included, often anchored to the food Engel curve.<sup>11</sup> For example, the official poverty line for the United States was set at three times the cost of a 1962 Economy Food Plan (Orshansky 1965); the three times reflects the assumed food share of one third. (The line is updated over time using a national consumer price index.) Another common approach is to estimate the total (food + nonfood) consumption expenditure or income level at which nutritional requirements are met on average in the specific setting. (This can be interpreted as inverting a stochastic version of Equation 4.)

Nonetheless, there is an inescapable normative element in all such poverty lines. Nutritional requirements depend on the level of physical activity assumed, which is a matter for judgement, as is the choice of food and nonfood needs deemed to be required for attaining any given set of nutritional requirements. Different judgements can be expected and defended, and those judgements will undoubtedly vary with the setting—the relevant place and time, or the purpose of the measures.<sup>12</sup>

<sup>&</sup>lt;sup>8</sup>The focus here is on global poverty measures. Attempts to implement a less rough approximation in countryspecific research include Ravallion & van de Walle (1991) and Dimri & Maniquet (2019).

<sup>&</sup>lt;sup>9</sup>See Ravallion et al. (1991), Atkinson & Bourguignon (2001), Ravallion et al. (2009), Ravallion & Chen (2011, 2019), Jolliffe & Prydz (2017), and Atkinson (2019).

<sup>&</sup>lt;sup>10</sup>The term developing countries can be defined in more than one way, but here it can be interpreted as the Part 2 member countries of the World Bank.

<sup>&</sup>lt;sup>11</sup>Further discussion of the methods used in practice in setting poverty lines can be found in Ravallion (2012; 2016, chapter 4) and Atkinson (2019, chapter 2).

<sup>&</sup>lt;sup>12</sup>For example, in determining eligibility, some antipoverty programs in the United States use a multiple of the official line.

**Figure 1** shows a recent compilation of national lines, converted from each local currency unit (LCU) to purchasing power parity (PPP) using the results of the 2011 International Comparison Program (ICP). (Section 4 returns to PPPs.) These are drawn from national or World Bank efforts to construct poverty lines appropriate to each country. The original lines do not of course use PPPs but are set in terms of local prices and local perceptions of what poverty means, typically anchored to both the set minima for nutritional intakes and prevailing diets. The national lines are set either by country-governments (the national statistics office, almost always) or by the World Bank in its country-level analytic work, typically in consultation with the government of the country concerned.

We see that richer countries tend to have higher lines. The fit in **Figure 1***a* is very close to linear, but it is not homogeneous. There is a positive intercept, which is unsurprising; one would not expect the poverty line to fall to zero in the lower limit as the mean falls. For the sample of non–Organization for Economic Co-operation and Development (OECD) countries, the predicted poverty line for the country with the lowest mean is \$0.96 a day [standard error (SE) = 0.25].<sup>13</sup> So the cross-county data in **Figure 1** are suggestive of weakly relative lines with an elasticity less than unity, but rising with the mean, from 0.36 (SE = 0.12) at the lowest mean income to near unity in high-income countries.

**Figure 1** uses cross-sectional data, so the pattern may be driven by latent country effects in national lines. While those lines are rarely revised quickly—there is (understandable) political resistance—they have risen over time with sustained gains in overall living standards. This has happened in the rich world over the past 100 years. For example, around the turn of the twentieth century, the most widely used poverty line in the United States was little more than \$1 per person per day in 2005 prices, while it is closer to \$15 a day now (Ravallion 2016). In recent times, we have also seen rising real poverty lines over time in growing developing countries, including China, India, Indonesia, and Vietnam.<sup>14</sup> A more convincing test of whether national lines for developing countries behave as weakly relative lines can be performed using the Jolliffe & Prydz (2017) data set of implicit national poverty lines, which has sufficient observations over time to permit panel-data analysis.<sup>15</sup> Then one finds that the strong positive relationship between national poverty lines and mean consumption holds even if one includes country fixed effects. Regressing the log poverty line (at PPP) on the log survey mean and including the 102 feasible country effects, one finds that the average elasticity is 0.52 (SE = 0.04; n = 598).<sup>16</sup>

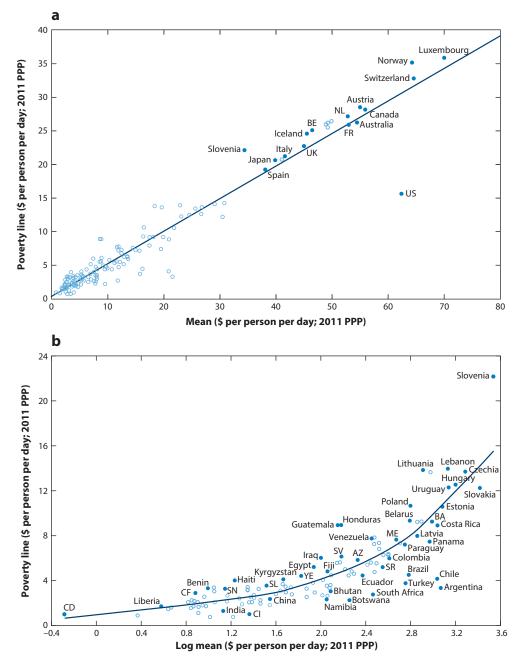
Thus, the national poverty lines in developing countries are implicitly weakly relative lines over time, with an average elasticity around 0.5—significantly positive but significantly less than unity.

One can interpret this pattern in terms of the model in Section 2. The national lines in **Figure 1** can be thought of as the local cost of normative functionings. The latter can vary across countries, but (probably more importantly) so too can the costs of attaining those functionings. In particular, the cost of social inclusion is almost certainly higher in richer countries, reflecting both goods

<sup>&</sup>lt;sup>13</sup>This is the Democratic Republic of the Congo, which has an unusually low mean. If one uses the country with the next lowest mean, Madagascar, the predicted poverty line is \$1.28.

<sup>&</sup>lt;sup>14</sup>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 & Prydz (2017) point to other examples of developing countries that have increased the real value of their national lines.

<sup>&</sup>lt;sup>15</sup>Jolliffe & Prydz (2017) estimate 609 poverty lines for 118 countries, as implicit in national poverty measures from the World Bank's *World Development Indicators*. 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  $F_{it}^{-1}(H_{it})$ . <sup>16</sup>Using the log of gross national income instead of the survey mean, it is 0.43 (SE = 0.01; *n* = 595).



#### Figure 1

Poverty lines across countries in US dollars using purchasing power parity (PPP) exchange rates. (*a*) Full sample (n = 146). (*b*) Non–Organization for Economic Co-operation and Development (OECD) (n = 122) using log scale for mean. Linear regression line in panel *a*. Empirical nonparametric regression function (locally weighted polynomial) in panel *b*. Dark blue filled circles are labeled with names or country codes; unfilled light blue circles are additional unlabeled countries. Abbreviations: AZ, Azerbaijan; BA, Bosnia and Herzegovina; BE, Belgium; CD, Democratic Republic of the Congo; CF, Central African Republic; CI, Cote d'Ivoire; FR, France; ME, Montenegro; NL, Netherlands; SL, Sierra Leone; SN, Senegal; SR, Serbia; SV, El Salvador; UK, United Kingdom; US, United States of America; YE, Yemen. Data from Ravallion & Chen (2019).

prices (allowing for nontraded goods) and the demands deemed necessary to assure that the normative functionings are attained.

However, there is nothing to guarantee that the national lines correspond to normative functionings all ethical observers would consider appropriate. In making global poverty comparisons, it would be unwise to focus on the lowest observed national lines. Some averaging is clearly called for, as are tests for robustness to the use of higher lines. The need for averaging arises from the variability in the methods and normative standards used to define poverty lines among countries at a given mean income. There is almost certainly some random measurement error in the observed national lines, also pointing to the need for averaging.

In using national poverty lines to infer a welfare-consistent schedule of global lines, we should also recognize an identification problem. Higher national lines could reflect either higher costs of attaining a given level of economic welfare or higher reference levels of welfare for deciding who is poor (higher  $\bar{u}$  in Equations 1 and 2). Absent a resolution of this identification problem, it can be argued that absolute and (weakly) relative poverty measures should be viewed as lower and upper bounds, respectively, on an unknown true welfare-consistent poverty measure (Ravallion & Chen 2011, 2019; Ravallion 2016).

This identification problem also arises when normative functionings are the anchor for the reference level of welfare in setting poverty lines (Section 2), to the extent that richer countries adopt more generous normative functionings. However, this identification problem need not be a concern if one takes a nonwelfarist approach that respects whatever normative functionings are deemed relevant in a given society without demanding welfare consistency. This appears to be the approach that Atkinson & Bourguignon (2001) have in mind (and is discussed further in Atkinson 2019). Note, however, that still does not justify strongly relative lines, as discussed above.

#### 4. ABSOLUTE POVERTY MEASURES

In the bulk of its global poverty work, the World Bank has insisted that the global poverty line should have constant purchasing power across countries. For measuring the global poverty rate, this is equivalent to using real household consumption (or disposable income when consumption is not available) per person as the welfare metric. This section reviews the Bank's method and a recently proposed alternative.

#### **4.1. Purchasing Power Parities**

In the context of measuring global poverty, the most widely used price index is based on the PPP rate derived from the ICP's price surveys at the country level.<sup>17</sup> The normalized PPPs based on those prices are essentially multicountry versions of a Fisher price index. Converting LCUs at the PPP rate (instead of the market or official exchange rate) is believed to better reflect the prices actually faced in each country. Official exchange rates cannot be relied on for this purpose, since many goods are not traded internationally, and they tend to be cheaper in poorer countries where wage rates are lower. So market exchange rates are thought to exaggerate the extent of global poverty by overstating the cost-of-living in poor countries.

There are numerous issues about how PPPs are calculated.<sup>18</sup> One concern is that prevailing PPPs are designed for comparing national accounts aggregates across countries, not for measuring

<sup>&</sup>lt;sup>17</sup>The ICP is a huge, global statistical effort involving the statistics offices (in 2011) of 145 national governments, the regional development banks, and Eurostat and led by the World Bank under the auspices of the United Nations Statistical Commission. The ICP does not currently allow spatial variation in PPPs within countries, although national poverty measures often allow for cost-of-living differences within countries. <sup>18</sup>See (inter alia) Summers & Heston (1991), Deaton (2010), Deaton & Heston (2010), and Ravallion (2018).

poverty. PPPs for the poor have been estimated by Deaton & Dupriez (2011), who constructed a set of PPPs that accord with the consumption patterns of people living near the international poverty line, based on household surveys. As it turned out, there was not much difference between poverty measures based on the Deaton–Dupriez PPPs for the poor and the standard PPPs for the 2005 ICP round (Chen & Ravallion 2010). A similar conclusion is reached by Dikhanov et al. (2017) using 2011 ICP data for Africa. This finding is not because poor people have similar consumption patterns to the averages in national accounts, but rather because the required reweighting has a similar structure across countries.

Another concern is whether the ICP's methods deal adequately with housing, which is an example of a comparison-resistant good, for which comparable prices cannot be readily observed across countries (Deaton & Heston 2010). (Another example is government services.) Too low a weight on housing rental in rich countries could arise from the use of cross-country average weights on commodity groups in PPPs. Thus, when measuring absolute global poverty, a case might be made for setting a higher line at PPP in richer countries to properly reflect latent costof-living differences. (This is not currently done by the Bank, which now includes rich countries in PovcalNet but applies the same PPP line to them.)

The PPPs change with each new ICP round, owing to changes in methodology and new data. Different methods have been tried to address the problems, such as for comparison-resistant services. This has generated some confusion. For example, echoing Deaton (2010), Allen (2017, p. 3691) claims that the World Bank's methods imply "that the number of poor in India. . . increased markedly despite India's economic growth—a perverse result indeed!" However, Allen is not actually referring to how India's poverty rate has changed over time; the Bank's methods have long indicated falling poverty measures in India with economic growth (see, for example, Datt & Ravallion 2011). Rather Allen's comment refers to the comparison of two sets of estimates using old and new PPPs (for different ICP base years, with differences in methods). The Bank follows standard practice in doing the PPP conversion only at the ICP base year; the price adjustments over time use the best available local price indices.

The changes in PPPs have been a mystery to many users, often associated with the ICP's decentralized implementation and not helped by restrictions on public access to the complete ICP microdata on prices. Ravallion (2018) documents excess sensitivity of PPP changes to market exchange rates, suggesting that the PPPs may put higher weight on internationally traded goods than do domestic deflators. This may not be surprising, since traded goods are more easily compared across countries for the price surveys. But there are other puzzles. For example, the last set of PPPs, from the 2011 ICP, indicated less poverty in the Asia region than prior ICP rounds. The reason is unclear. Based on the available documentation, Ravallion (2018) suggests that the most plausible explanation is that the ICP implementation for Asia (by the Asian Development Bank) did a better job of covering rural areas, where prices tend to be lower than in urban areas. However, the urban bias in the ICP's price surveys remains in much of the rest of the world. So the partial methodological improvement in Asia creates a concern in global poverty comparisons.

Motivated by such concerns, Atkinson (2019) recommends that the Bank's global measures should not be updated in the light of each new ICP round (at least until 2030, the United Nations' target date for eliminating \$1.90 a day poverty). Some observers have recommended abandoning PPPs for this purpose and have proposed alternatives; one recent example is discussed in detail below in this section. Another option is to use a moving average of PPPs.

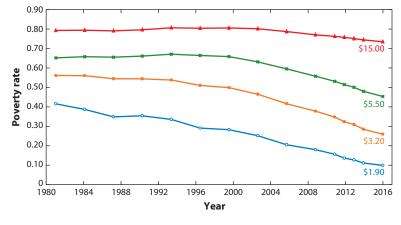
#### 4.2. Global Poverty Lines

The research papers produced to underpin the Bank's more high-profile reports have used multiple poverty lines (at PPP) and tested the robustness of key qualitative claims (notably whether poverty is falling) to the choice of poverty line. For example, there are many such lines in Chen & Ravallion (2010); indeed, they test and accept first-order dominance over 30 years for a wide range of possible lines up to the US official line, which was around \$15 per person per day in 2010 prices (for a family of four with two children).

One might stop there and declare the job done, noting that the ordinal poverty comparison is then robust to both poverty lines and measures within a broad class (Atkinson 1987). However, as noted, many users of poverty measures want a single line, or possibly two at most. The World Bank (1990) and Ravallion et al. (1991) turned to data on national poverty lines and proposed that the focus in setting absolute international lines should be on the national lines (at PPP) found in poor countries. The national lines found in the poorest countries are understood to be frugal specifications for what is needed to not be deemed poor globally.

Based on an expanded and improved data set on national lines and various averaging methods, Ravallion et al. (2009) set the international line at \$1.25 a day using the 2005 ICP. The United Nations' Sustainable Development Goals aim (among many other goals) to assure that nobody lives below this line by 2030.<sup>19</sup> Using domestic price indices and the 2011 ICP, the \$1.25 line was updated to \$1.90 a day by Ferreira et al. (2016). Let us call this the benchmark line. This is understood to be a low line; anything less than this would be hard to defend as one would be using a line that is lower than the average found in the poorest countries. Higher lines can be justified and used. While users of the World Bank's PovcalNet can enter any desired line, the landing page has built-in options for \$3.20 and \$5.50 lines, which are more representative of poverty lines found in middle-income countries, in addition to the default \$1.90 benchmark. The international poverty lines are converted to LCUs in 2011 at PPP and then adjusted for local inflation to obtain poverty lines for each survey date in each country.

Figure 2 plots the incidence of poverty in the world as a whole from 1981 to 2015 using all three World Bank lines (\$1.90, \$3.50, and \$5.50), based on the 2011 ICP, and using all the data available in PovcalNet at the time of writing (July 2019). For comparison, the figure also gives



#### Figure 2

Global absolute poverty rates for various poverty lines. The figure gives the percentage of the world's population below each poverty line in US dollars using purchasing power parity exchange rates. Data from PovcalNet (http://iresearch.worldbank.org/PovcalNet/povDuplicateWB.aspx).

<sup>&</sup>lt;sup>19</sup>This goal is based on the calculations in Ravallion (2013), but with one important difference: the 2013 paper outlined a scenario that would get to a poverty rate for the developing world as a whole of 3%, not 0%, by 2030. That would still involve lifting about one billion people out of poverty. Section 6 returns to this issue.

the series for \$15 a day, which is about the official poverty line for the United States in 2011 (for a family of four with two children). Over the period as a whole, a decline in the global poverty rate is evident, although it is notable that this accelerated in the new millennium. The \$1.90 poverty rate fell from 42.1% in 1981 to 28.6% in 1999 (0.7% points per year) and to 9.9% in 2015 (1.1% points per year). Little or no progress was made for the higher lines prior to 2000, but progress is evident since. (Of course, in the limit, as the line rises, the rate of change goes to zero.) The proportion of the world's population living below the US poverty line rose slightly between 1981 and 1999 (from 80.5% to 81.8%) but then fell to 74.6% in 2015.

#### 4.3. Measures Based on the Minimum Cost of Nutritional Adequacy

Given the aforementioned concerns about the Bank's PPPs, an alternative method that does not use those PPPs is of obvious interest. New measures of global absolute poverty have been proposed by Allen (2017), who uses linear programming (LP) to estimate country-specific least-cost diets for attaining globally fixed nutritional requirements, which he then values at local prices, and adds spending on a fixed bundle of his selected nonfood goods (including an explicit allowance for housing). The main difference is in how the poverty lines are calculated. Allen uses ICP prices but essentially weights them differently, depending on the solutions to the LP problem.

Allen (2017) provides an alternative approach to the World Bank's global poverty measures. His method avoids some concerns about how PPPs are constructed and used in global poverty measurement but raises new issues. Allen is resurrecting an earlier approach. For the United States in 1940, Stigler (1945) calculated the bundle that minimized the cost of attaining predetermined nutritional requirements. This can be represented as an LP problem (although the simplex method had not yet been available when Stigler made his estimates). However, Stigler found that the implied diets were unlikely to be socially acceptable—the poverty lines derived this way were not considered behaviorally plausible even for poor Americans in the early 1940s.<sup>20</sup> This was later confirmed by Smith (1959), who found that very few people (in Michigan in the mid-1950s) actually consumed anything like the LP solutions. Least-cost diets have been found to have too little diversity, which is also valued by nutritionists (Masters et al. 2018, using data from Africa).

In the light of the results of Stigler (1945) and Smith (1959), the least-cost method was subsequently rejected in the bulk of the literature. For example (in reference to Stigler's least-cost diets), Sen (1981, p. 27) writes, "Such minimum cost diets are typically very inexpensive, but exceedingly dull and very often regarded as quite unacceptable." In the United States, the official line "...was not designed to be a minimum cost food plan but a palatable food plan" meeting recommended dietary allowances (Hanson 2008, p. 573). The World Bank and other researchers working on poverty, including in poor countries, have not used the least-cost method in setting poverty lines.<sup>21</sup> None of the 75 national lines for developing countries used by Ravallion et al. (2009) to locate the Bank's \$1.25 line in 2005 ICP prices used this method.<sup>22</sup> Instead, modern methods identify a food bundle consistent with prevailing tastes in each setting, respecting the influence of local food habits as well as recommended nutritional intakes.<sup>23</sup>

<sup>&</sup>lt;sup>20</sup>To quote Stigler (1945, p. 313): "It would be the height of absurdity to practice extreme economy at the dinner table in order to have an excess of housing or recreation or leisure."

<sup>&</sup>lt;sup>21</sup>The method has been used at times as a lower-bound benchmark to the cost of adequate diets (as in Masters et al. 2018 and Hirvonen et al. 2019).

<sup>&</sup>lt;sup>22</sup>I also checked this with three World Bank specialists on the national poverty lines, who confirmed this claim.
<sup>23</sup>This is also evident in the fact that the consumption patterns of migrants do not adapt quickly to the new set of relative prices in the destination; see Atkin (2016) using data on Indian migrants.

What then is Allen's case for resurrecting Stigler's least-cost method? If we are willing to equate nutritional intakes (relative to the stipulated "requirements") with welfare then Allen's proposal is close to the welfarist model described in Section 2. However, that assumption is hard to defend. Even if one were to think that nutritional status is the sole determinant of human welfare, nutrient absorption is lower in less healthy environments.<sup>24</sup> Thus, quantitative intakes need not be a particularly good proxy for nutritional status, such as is indicated by anthropometric data. Putting this issue aside, we can also question whether nutritional status is an adequate welfare indicator. As noted in Section 2, it can be deceptive to think about the relevant anchoring functionings in too partial a way when measuring poverty. When welfare depends on both nutritional status (relative to requirements) and social inclusion, ignoring the latter can lead to welfare-inconsistent poverty measures. Credible measures require that we allow for the functioning of social inclusion, as both a factor influencing the food consumption bundle relevant to attaining nutritional requirements in a specific social context and as an independent determinant of welfare. This point is not new in the history of thought on poverty, but its salience is undervalued by least-cost, nutrition-based absolute measures.

Allen (2017) grants that his solution is not reasonable for people living in rich countries, but he claims that it is fine for poor people in poor countries, for whom "necessity displaces desire" (in the title of Allen's paper) and so that "linear programming is much more germane to poor people" (Allen 2017, p. 3695). Against this claim, there is ample evidence from the spending behavior of poor people globally that they care about more than their nutritional intakes.<sup>25</sup> Allen's own results indicate that his LP predictions are high on food grains and fats and low on meat, fish, vegetables, and fruits relative to actual consumption in developing countries in the 1960s (Ravallion 2017). There is clearly less variety in his LP solutions, consistent with the findings of Stigler and Smith for the United States. The reason is obvious: Both food and nonfood choices are clearly influenced by other factors, including social roles of consumption and connectivity in local communities. Leastcost, nutritionally adequate diets may well be just as socially unacceptable in poor countries today as Stigler found them to be in 1940s America.

Allen (2017) claims that he gets a higher poverty count than for the Bank's \$1.90 a day line. Ostensibly, this is surprising, given that one would expect his method to generate even more frugal poverty lines than the \$1.90 line. However, a careful reading of Allen's paper suggests otherwise, especially when it is read alongside the working paper version (Allen 2016) that gives a rather different interpretation of his own numbers, emphasizing a broad agreement with the Bank's line. One can always raise the nutrient requirements and so raise the poverty line. Nutrient specifications that appear to be more consistent with practice yield an overall poverty count that is appreciably lower than the Bank's.<sup>26</sup> Urban bias in the ICP prices Allen uses is also likely to give higher poverty measures than for the national lines in poor countries used to anchor the World Bank's international line.

Does this matter to the comparisons over time? Allen does not provide his poverty measures over time. Ravallion (2017, table 2) compares estimates using Allen's lines with the Bank's for a

<sup>&</sup>lt;sup>24</sup>See, for example, Duh & Spears (2016) using data for India.

<sup>&</sup>lt;sup>25</sup>Examples can be found in Banerjee & Duflo (2008) on spending patterns, Rao (2001) on celebrations in India, and Milanovic (2008) on qat consumption in Yemen. Section 5 gives further examples.

<sup>&</sup>lt;sup>26</sup>As noted by Ferreira (2017), in an earlier working paper, Allen (2016) had focused instead on a nutrient specification that turned out to give him exactly \$1.90 a day on average. Indeed, Allen identifies this as a key finding of the working paper, arguing that his approach "...provides a clear rationale for why \$1.90 per day is a good standard" (Allen 2016, p. 1). The story changed with the final published version, where Allen focuses instead on a specification that gives a higher poverty count than for a uniform \$1.90 a day. No justification is given for this choice. The published version heralds the higher poverty count and provides no poverty measures for any of his other specifications.

common set of countries (as used by Allen). The population-weighted poverty rate using Allen's lines for 1990 is lower than for \$1.90 a day (45% versus 50%, respectively). The decline in the poverty rate over a 20-year period is 33.6% for Allen's lines against 33.8% for the Bank's. The choice makes almost no difference over time.

Given that there are strengths and weaknesses to both the Bank's method and Allen's, it is reassuring that the levels of global poverty (for at least some of Allen's nutrient specifications) and the trends over time are so similar, despite the methodological differences.

The next section critically reviews another strand of the literature in which the functioning of social inclusion has been given more (explicit) emphasis—with more dramatic implications for the picture of global poverty.

# 5. RELATIVE POVERTY MEASURES

The theoretical model in Section 2 postulated that welfare depends in part on relative income, suggesting that a higher real income may be needed to attain the same level of welfare in a richer country (as implied by Equation 1).<sup>27</sup> This receives support from various strands of the literature. The relevance of concerns about shame, stigma, relative deprivation, and social exclusion has long been emphasized in the literature in sociology and social psychology.<sup>28</sup> 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 (Luttmer 2005, Knight et al. 2009).<sup>29</sup> 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).

It might be conjectured that these concerns are less pertinent in poorer places. However, that is questionable. Anthropologists have long described behaviors consistent with the idea that social effects matter to poor people; see, for example, Geertz (1976) and Fuller (1992). Rao (2001) describes the importance of celebrations to social networks among poor people in rural India. Banerjee & Duflo (2008) document expenditures on celebrations and festivals by very poor people in surveys for a number of countries. There are many potential reasons why we observe such behavior among poor people. One possibility is a direct relative comparison. Another is more indirect: Such behaviors can stem from insurance motives in settings with repeated interaction (as argued by Ravallion 2008). The key point here is that the incomes of others around you matter, even when you are poor.<sup>30</sup>

In carrying this idea to measurement practice, we face two sources of uncertainty that have been somewhat neglected in the literature. First, saying that people (including poor people) care about relative income does not imply that it is relevant to the concept of economic welfare by which we judge one person to be poorer than another. That is a judgement we must make about what constitutes economic welfare, recognizing that this need not accord with the maximand of choice (Section 2).

<sup>28</sup>Including Davis (1959), Runciman (1966), Townsend (1979), and Walker (2014).

<sup>&</sup>lt;sup>27</sup>Recall that higher lines at PPP may be needed in rich countries to address the concerns about comparisonresistant goods (Section 4). Higher allowances for the costs of housing, in particular, may be called for. This is a problem for measuring absolute poverty, which is not the topic of this section.

<sup>&</sup>lt;sup>29</sup>Surveys of this literature can be found in Frey & Stutzer (2002), Senik (2005), and Clark et al. (2008).

<sup>&</sup>lt;sup>30</sup>Smith et al. (2012) provide a review of many studies showing behavioral responses to relative deprivation.

Second, what is the relevant comparison group and what is the relevant statistic about that group? In the context of measuring global poverty, it is natural to treat the comparison group as fellow citizens nationally, although in reality it may be more local, or even global. But is the relevant statistic the mean or median, as commonly assumed, or something else?

#### 5.1. Strongly Relative Lines

The most common approach to measuring relative poverty compares each household's 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 relative poverty line  $(z^R)$  can be written in the generic form

$$z^{\mathrm{R}} = k.y\left(\pi_{z}\right).$$

Here, k is a constant, y(.) is the quantile function (inverse of the CDF), and  $\pi_z$  is a fixed percentile that defines the comparison group. In the case of the original Fuchs (1967) proposal,  $k = \pi_z = 0.5$ , although other parameter values have been used since. This method has been popular among some statistical agencies (notably Eurostat and the OECD) and some researchers; for example, in work for the OECD, Garroway & de Laiglesia (2012) estimate such relative poverty measures for developing countries using  $\pi_z = 0.5$  and k = 0.4, 0.5, 0.6.

The point made in Section 2 that a poverty line set at a constant proportion of the mean cannot in general be globally welfare consistent also applies when the comparison income is  $y(\pi_z)$ . Other concerns also loom large. 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 United States is an interesting case. The new supplemental 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). (As seen in **Figure 1**, the official line is well below what one would expect given average income.) The SPM uses the quantile of the thirty-third percentile of the distribution of a subset of consumption spending deemed to be essential (comprising food, clothing, shelter, and utilities), following Citro & Michael (1995). (Thus, the SPM sets k =1.2 and  $\pi_z = 0.33$  in Equation 5.) However, it remains unclear why  $y(\pi_z)$  is a plausible comparison income for any fixed  $\pi_z$ . In the case of the SPM, it is also unclear why relative comparisons would apply only to essential goods. Indeed, one might expect feelings of relative deprivation to respond as much to a lack of nonessential goods.

Ravallion & Chen (2019) provide a theoretical formulation of the comparison income that encompasses both upward and downward relative comparisons. Instead of the ordinary mean or median, the model points to a distribution-corrected mean, the properties of which depend on whether people tend to look up or down (in terms of incomes) when they assess how they are doing relative to others. The discussion returns to the comparison income once some other issues are addressed. For now, one can treat the simple mean in the country of residence as the comparison income.

When the poverty line is set at a constant proportion of the mean (or median), there is a further concern 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.<sup>31</sup> Seemingly perverse poverty comparisons have been found using strongly relative

<sup>&</sup>lt;sup>31</sup>Note that this property does not depend on whether the line is anchored to the mean or the median; the ratio of the median to the mean is constant in an inequality-neutral growth process. However, objections to the use of the median have been identified by de Mesnard (2007) and Kampke (2010).

measures.<sup>32</sup> As we have seen, national poverty lines in developing countries have an average elasticity to the mean of about 0.5—appreciably (and significantly) less than unity.

Strongly relative lines are especially questionable in poor countries. 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 only \$0.64 a day (2005 ICP prices), which is somewhere around a survival level that Lindgren (2015) estimates to be \$0.67 a day (also in 2005 ICP prices). The value for the country with the lowest mean would be only \$0.38 per day, which is unlikely to be enough for survival beyond a short time. Similarly, the Garroway & de Laiglesia (2012) measures assume lines that are well below those typical of even low-income countries and even below likely biological minima.

In short, strongly relative measures almost certainly understate the nutritional and social inclusion needs of globally poor people and have a seemingly perverse implication for how these measures respond to economic growth and contraction. While strongly relative measures have been more popular in rich countries, they are hard to accept elsewhere and hence globally.

# 5.2. Weakly Relative Lines

The literature has suggested some possible solutions to these deficiencies of strongly relative measures. In passing, Kakwani (1986) suggests

$$z^{\mathrm{R}} = z^{\mathrm{A}} + \beta(m - z^{\mathrm{A}}), \qquad 6.$$

where  $z^A$  (>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 poverty line with respect to *m* is positive but less than unity (making it weakly relative); the limit of the elasticity is unity as *m* goes to infinity. Chakravarty et al. (2015) provide an axiomatic derivation for a line of the form in Equation 6. Jolliffe & Prydz (2017) and the World Bank (2018) use a schedule of lines with essentially the same form, which generalizes the Garroway & de Laiglesia (2012) proposal for developing countries to allow a positive intercept, thus making it weakly relative.

An alternative approach is found in Foster (1998), who proposes using the geometric mean of an absolute line and a strongly relative line. This accords nicely with the above estimated elasticity of 0.5 based on national lines with country fixed effects (Section 3). However, that is an average elasticity for developing countries. As we have seen, the elasticity tends to rise as mean income rises (Section 3). The fact that the Kakwani (1986) proposal allows the elasticity to vary and to go toward unity at high incomes is more appealing when we look for an encompassing schedule of global lines consistent with the data on national lines. So the following discussion will take Equation 4 as the starting point.

A problem arises in Equation 6 when  $m < z^A$ , as the implied line is then lower than  $z^A$ . We cannot rule out  $m < z^A$ ; indeed, the data used by Ravallion & Chen (2019) indicate that this is the case for 11% of non-OECD countries. A schedule of hybrid absolute plus relative (A + R) lines that avoids all the aforementioned problems is the piece-wise linear form:

$$z^{A+R} = max\left(z^{A}, \alpha + \beta m\right) = z^{A} + max\left(\alpha + \beta m - z^{A}, 0\right) \ (\alpha \ge 0).$$

$$7.$$

(Note that the A + R line can never be below the absolute line.) This is the formula used by Ravallion & Chen (2011) and Chen & Ravallion (2013). An antecedent is found in Atkinson & Bourguignon (2001) (and Atkinson 2019). However, there is an important difference. The Atkinson–Bourguignon lines are the special case of Equation 7 in which one sets  $\alpha = 0$  [i.e., they

<sup>&</sup>lt;sup>32</sup>See, for example, UNDP (2005, box 3) and Easton (2002).

are strongly relative above some critical level of income  $(z^A/\beta)$ ]. Then the aforementioned objections to strongly relative lines return. What Equation 7 gives us instead is a straightforward generalization of the Atkinson & Bourguignon (2001) proposal by adding a parameter,  $\alpha$ , that can be interpreted as the lower bound to social inclusion needs.

#### 5.3. Empirical Implementation

We saw in Section 4 that the World Bank's absolute lines have been set according to the national lines found in poor countries. International relative poverty lines have also been anchored to national lines, but now the focus is on how they vary with average income across countries, as discussed in Section 3. That is the approach followed here, in keeping with a strand of the literature.<sup>33</sup>

However, it is acknowledged that there is uncertainty about whether any differences in the latent reference levels of welfare are statistically ignorable (as noted in Section 3). Richer countries may tend to use more generous reference welfare levels for defining poverty. Then the true welfare-consistent poverty measure will be bounded below by  $z^A$  and above by  $z^{A+R}$  (Ravallion & Chen 2011, 2019). This is less of a concern if one follows the approach of Atkinson & Bourguignon (2001), for 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 with respect to the comparison income as including any effect on that reference level of welfare.

For the absolute lines ( $z^A$  in Equation 7), this article uses the World Bank's \$1.90 a day line in 2011 prices. Following past literature, the schedule of weakly relative lines is calibrated to national lines. The data on national poverty lines suggest that the rank-weighted mean is the relevant comparison income, with lowest weight given to the richest (Ravallion & Chen 2019). This implies that a Gini-discounted mean is called for [i.e.,  $m_j^* = (1 - G_j)m_j$ ], where  $G_j$  is the Gini index for country *j*. On calibrating to the data set of national lines in **Figure 1**, one obtains the following schedule of A + R lines:

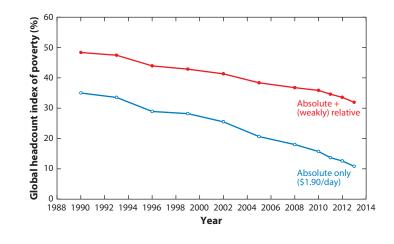
$$z_{i}^{A+R} = max [\$1.90, \$0.90 + 0.7(1 - G_{i})m_{i}] = \$1.90 + max [0.7(1 - G_{i})m_{i} - \$1.00, 0].$$
 8.

Thus, a person is not poor globally if she is neither absolutely poor (relative to  $z^A$ ) nor poor by the expected standard for the country she lives in.

**Figure 3** implements the measures on a global basis, drawing on 1,500 household surveys for 150 countries from 1990 to 2013. For both the absolute and A + R measures, the percentage of the world's population living in poverty has fallen over time. The trend rate of decline for the A + R measures is 0.7 percentage points per year (a regression coefficient on the year of -0.688; SE = 0.028). The corresponding trend for the absolute measures is one percentage point per year (-1.055; SE = 0.043). If this is maintained, then the poverty rate for the \$1.90 line will reach zero by 2025. However, as we see below, when one focuses instead on the view from the bottom, it appears very unlikely that the world will maintain the same trend rate of decline as the poverty rate gets closer to zero.

The fall in the global poverty rate as judged by the hybrid A + R lines is due to falling absolute poverty counts. Indeed, the proportion who are relatively poor but not absolutely poor—the gap between the poverty rates for the absolute and A + R lines—has risen over time, with a trend rate of increase of about 0.4 percentage points per year (0.367; SE = 0.025). In 1990, 1.85 billion people (35% of the world's population) lived below the \$1.90 line, and a further 700 million (13%) lived in

<sup>&</sup>lt;sup>33</sup>Including Chen & Ravallion (2001, 2013), Atkinson & Bourguignon (2001), Ravallion & Chen (2011), Jolliffe & Prydz (2017), and Atkinson (2019).

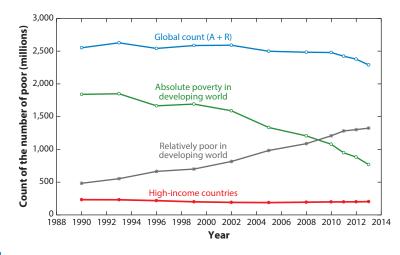


#### Figure 3

Global absolute and relative poverty measures. The absolute only line gives the percentage of the world's population below \$1.90 a day using purchasing power parity exchange rates. The absolute and (weakly) relative line gives the percentage of the world's population living below the larger of the \$1.90 absolute line and a weakly relative line that rises with mean income. Figure adapted from Ravallion & Chen (2019).

relative but not absolute poverty (i.e., they were poor by typical standards of the country they lived in but not globally poor by the \$1.90 standard). By 2013, the count for the absolute standard had fallen to 770 million (11% of the world's population), while that for the A + R had fallen by much less, to 2.3 billion (32%). The count of those who are not poor by the \$1.90 line but still poor by a line typical of their country of residence has more than doubled, from 0.7 billion to 1.5 billion.

**Figure 4** provides the global count of the number of people living below the A + R lines. The count of the absolutely poor in developing world is the number of people living below the



#### Figure 4

Components of global poverty count for the A + R lines (given by the larger of the \$1.90 absolute line and a weakly relative line that rises with mean income). The global count is of those living below the  $z^{A + R}$  lines; the count of the absolutely poor in developing world is the number of people living below  $z^{A}$ , while the count of relatively poor in developing world is the number between the two lines. Authors' calculations. Abbreviations: A, absolute; R, relative.

\$1.90 line, while the count of relatively poor in developing world is the number between that line and the A + R lines in the developing world. The counts for high-income countries are for the A + R lines and are almost entirely of those living in relative poverty.

We see that the falling global count of the poor by the \$1.90 line 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 A + R lines. Slightly less than 80% of those who rise above the absolute line end up living between the two—no longer poor by the global absolute line but still poor by standards typical of the country they live in.

Whether one focuses on absolute-only poverty or A + R poverty, the incidence of poverty is appreciably higher in the developing world than in the advanced countries (as a whole). Over 90% of the poor by the A + R line are found in the developing world, which is home to virtually all those who are poor by the lower line. Side by side with the falling numbers of absolutely poor in the developing world, we find rising numbers of people who are still poor by the standards typical of the country they live in.

#### 6. CONCLUSIONS

This article has argued that global poverty measures should be anchored to a common concept of economic welfare based on two key functionings—namely, nutritional status and social inclusion. International poverty lines are interpreted as money metrics of that concept of welfare. Nutritional status alone cannot be considered a sufficient statistic, including for poor people. The minimum cost of a given nutritional status is a questionable guide when (as is evidently the case) people do not maximize their caloric intakes subject to their income and the prices faced. Concerns about relative deprivation and social inclusion also have a legitimate place in poverty analysis, including in poor countries.

Existing approaches to setting international lines for measuring absolute poverty are plainly inadequate if economic welfare depends (in part) on relative income in the society in which they live, or there are higher costs of social inclusion in richer countries—costs that are unlikely to be captured fully by the usual PPP deflators. Current approaches to measuring relative poverty are also inadequate under the assumption that economic welfare depends on own income at a given relative income. Assuming that both own income and relative income matter, the elasticity of the poverty line to the mean should be positive but less than unity, which rules out both the absolute and (strongly) relative approaches found in practice. A hybrid approach combining absolute and weakly relative measures is called for to reflect both subsistence and social inclusion. By the proposed approach, a person is poor if she is either below the common global standard or living below the poverty line one would expect given the average income in the country of residence. This gives us truly global poverty measures—that span countries at all levels of development.

This article has provided illustrative calculations. Progress in reducing global poverty is evident for both the absolute and the hybrid A + R lines, although with rising counts of those who are relatively poor but no longer absolutely poor. There are very few people in the rich world, and even in many middle-income countries, who are poor by absolute standards typical of the poorest countries. Nonetheless, they are still poor by the standards of what poverty means in their own country.

#### **DISCLOSURE STATEMENT**

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

#### ACKNOWLEDGMENTS

For comments on an earlier version of this article, the author thanks Harold Alderman, Yuri Dikhanov, Francisco Ferreira, Madhulika Khanna, Christoph Lakner, Mattias Lindgren, Will Masters, Milan Thomas, Dominique van de Walle, and the journal's anonymous reviewers.

# LITERATURE CITED

- Allen R. 2016. Absolute poverty: when necessity displaces desire. Discuss. Pap. Econ. Soc. Hist. 141, Oxf. Univ., Oxford, UK
- Allen R. 2017. Absolute poverty: when necessity displaces desire. Am. Econ. Rev. 107(12):3690-721
- Atkin D. 2016. The calorie costs of culture: evidence from India migrants. Am. Econ. Rev. 106(4):1144-81
- Atkinson AB. 1987. On the measurement of poverty. Econometrica 55:749-64
- Atkinson AB. 2019. Measuring Poverty Around the World. Princeton, NJ: Princeton Univ. Press
- Atkinson AB, Bourguignon F. 2001. Poverty and inclusion from a world perspective. In *Governance, Equity and Global Markets*, ed. JE Stiglitz, P-A Muet, pp. 151–64. Oxford, UK: Oxford Univ. Press
- Banerjee A, Duflo E. 2008. What is middle class about the middle classes around the world? J. Econ. Perspect. 22(2):3–28
- Bowley AL. 1915. The Nature and Purpose of the Measurement of Social Phenomena. London: P.S. King Sons
- Browning M. 1992. Children and household economic behavior. J. Econ. Lit. 30:1434-75
- Chakravarty S, Chattopadhyay N, Nissanov Z, Silber J. 2015. *Reference groups and the poverty line: an axiomatic approach with an empirical illustration*. WIDER Work. Pap. 2015/002, UNU-WIDER, Helsinki, Finl.
- Chen S, Ravallion M. 2001. How did the world's poor fare in the 1990s? Rev. Income Wealth 47(3):283-300
- Chen S, Ravallion M. 2010. The developing world is poorer than we thought, but no less successful in the fight against poverty. Q. J. Econ. 125(4):1577–625
- Chen S, Ravallion M. 2013. More relatively poor people in a less absolutely poor world. *Rev. Income Wealtb* 59(1):1–28
- Citro C, Michael R. 1995. Measuring Poverty: A New Approach. Washington, DC: Natl. Acad.
- Clark A, Frijters P, Shields M. 2008. Relative income, happiness and utility: an explanation for the Easterlin paradox and other puzzles. *J. Econ. Lit.* 46(1):95–144
- Cohn A, Fehr E, Herrmann B, Schneider F. 2014. Social comparison and effort provision: evidence from a field experiment. *J. Eur. Econ. Assoc.* 12(4):877–98
- Datt G, Ravallion M. 2011. Has India's economic growth become more pro-poor in the wake of economic reforms? World Bank Econ. Rev. 25(2):157–89
- Davis J. 1959. A formal interpretation of the theory of relative deprivation. Sociometry 22:280–96
- Deaton A. 2010. Price indices, inequality and the measurement of poverty. Am. Econ. Rev. 100(1):5-34
- Deaton A, Dupriez O. 2011. Purchasing power parity exchange rates for the global poor. Am. Econ. 3.: Appl. Econ. 3(2):137–66
- Deaton A, Heston A. 2010. Understanding PPPs and PPP-based national accounts. Am. Econ. J.: Macroeconomics 2(4):1–35
- de Mesnard L. 2007. Poverty reduction: the paradox of the endogenous poverty line. LEG Econ. Ser. Work. Pap. 2007–05, Univ. Bourgogne, Dijon, Fr.
- Dikhanov Y, Hamadeh N, Vigil-Oliver W, Degefu TB, Song I. 2017. Poverty-specific purchasing power parities in Africa. Policy Res. Work. Pap. 8150, World Bank, Washington, DC
- Dimri A, Maniquet F. 2019. Income poverty measurement in India: defining group-specific poverty lines or taking preferences into account? J. Econ. Inequal. https://doi.org/10.1007/s10888-019-09434-6
- Duesenberry JS. 1949. Income, Saving and the Theory of Consumer Behavior. Cambridge, MA: Harvard Univ. Press
- Duh J, Spears D. 2016. Health and hunger: disease, energy needs and the Indian calorie consumption puzzle. *Econ. J.* 127:2378–409
- Easterlin RA. 1974. Does economic growth improve the human lot? Some empirical evidence. In *Nations and Households in Economic Growth*, ed. PA David, WR Melvin, pp. 89–125. Palo Alto, CA: Stanford Univ. Press

Easton B. 2002. Beware the median. Soc. Policy Res. Cent. Newsl. 82:6-7

- Ferreira F. 2017. Global poverty today, the 1908 winter in St. Petersburg, and 'controversy bias.' World Bank/Let's Talk Development, Dec. 11. https://blogs.worldbank.org/developmenttalk/globalpoverty-today-1908-winter-st-petersburg-and-controversy-bias
- Ferreira F, Chen S, Dabalen A, Dikhanov Y, Hamadeh N, et al. 2016. A global count of the extreme poor in 2012: data issues, methodology and initial results. *7. Econ. Inequal.* 14:141–72
- Foster J. 1998. Absolute versus relative poverty. In *Papers and Proceedings of the Hundred and Tenth Annual Meeting of the American Economic Association*, Vol. 88, No. 2, ed. JD Baldwin, RL Oaxaca, pp. 335–41. Chicago: Am. Econ. Assoc.
- Frank RH. 1985. Choosing the Right Pond: Human Behavior and the Quest for Status. New York: Oxford Univ. Press
- Frey B, Stutzer A. 2002. What can economists learn from happiness research? J. Econ. Lit. 40:402-35
- Fuchs V. 1967. Redefining poverty and redistributing income. Public Interest 8:88-95
- Fuller CJ. 1992. The Campbor Flame. Popular Hinduism and Society in India. Princeton, NJ: Princeton Univ. Press
- Garroway C, de Laiglesia JR. 2012. On the relevance of relative poverty for developing countries. OECD Dev. Cent. Work. Pap. 314, OECD, Paris
- Geertz C. 1976. The Religion of Java. Chicago/London: Univ. Chicago Press
- Hanson K. 2008. Mollie Orshansky's strategy to poverty measurement as a relationship between household food expenditures and economy food plan. *Appl. Econ. Perspect. Policy* 30(3):572–80
- Hirsch F. 1977. The Social Limits to Growth. London: Routledge Kegan Paul
- Hirvonen K, Bai Y, Headey D, Masters WA. 2019. Affordability of the EAT-*Lancet* reference diet: a global analysis. *Lancet* 8(1):E59–66
- Jolliffe D, Prydz E. 2017. *Global societal poverty: a relative and relevant measure*. Policy Res. Work. Pap. 8073, World Bank, Washington, DC
- Kakwani N. 1986. Analyzing Redistribution Policies: A Study Using Australian Data. Cambridge, UK: Cambridge Univ. Press
- Kampke T. 2010. The use of mean values versus medians in inequality analysis. J. Econ. Soc. Meas. 35:43-62
- Knight J, Song L, Gunatilaka R. 2009. Subjective well-being and its determinants in rural China. China Econ. Rev. 20(4):635–49
- Lindgren M. 2015. The elusive quest for the subsistence line: How much does the cost of survival vary between populations? Comp. Inst. Anal. Work. Pap. 2015:1, Lund Univ., Lund, Sweden
- Luttmer E. 2005. Neighbors as negatives: relative earnings and well-being. Q. 7. Econ. 120:963-1002
- Masters WA, Bai Y, Herforth A, Sarpong DB, Mishili F, et al. 2018. Measuring the affordability of nutritious diets in Africa: price indexes for diet diversity and the cost of nutrient adequacy. Am. J. Agric. Econ. 100(5):1285–301
- Milanovic B. 2008. Qat expenditures in Yemen and Djibouti: an empirical analysis. J. Afr. Econ. 17(5):661– 87
- Orshansky M. 1965. Counting the poor: another look at the poverty profile. Soc. Secur. Bull. 28:3-29
- Rao V. 2001. Poverty and public celebrations in rural India. Ann. Am. Acad. Political Soc. Sci. 573(1):85-104
- Ravallion M. 2008. On the welfarist rationale for relative poverty lines. In Arguments for a Better World: Essays in Honor of Amartya Sen, Vol. I: Ethics, Welfare, and Measurement, ed. K Basu, R Kanbur, pp. 375–96. Oxford, UK: Oxford Univ. Press
- Ravallion M. 2011. On multidimensional indices of poverty. 7. Econ. Inequal. 9(2):235-48
- Ravallion M. 2012. Poverty lines across the world. In *The Oxford Handbook of the Economics of Poverty*, ed. PN Jefferson, pp. 75–104. Oxford, UK: Oxford Univ. Press
- Ravallion M. 2013. How long will it take to lift one billion people out of poverty? World Bank Res. Obs. 28(2):139-58
- Ravallion M. 2016. The Economics of Poverty: History, Measurement, and Policy. New York: Oxford Univ. Press Ravallion M. 2017. An interesting step backwards in measuring global poverty. Economics & Poverty, Dec.
  - 3. https://examplewordpresscom61323.files.wordpress.com/2017/12/comments-robert-allen-absolute-poverty-aer.pdf

- Ravallion M. 2018. An exploration of the changes in the International Comparison Program's global economic landscape. *World Dev.* 105:201–16
- Ravallion M, Chen S. 2011. Weakly relative poverty. Rev. Econ. Stat. 93(4):1251-61
- Ravallion M, Chen S. 2019. Global poverty measurement when relative income matters. J. Public Econ. 177:104046
- Ravallion M, Chen S, Sangraula P. 2009. Dollar a day revisited. World Bank Econ. Rev. 23(2):163-84
- Ravallion M, Datt G, van de Walle D. 1991. Quantifying absolute poverty in the developing world. *Rev. Income Wealth* 37:345–61
- Ravallion M, van de Walle D. 1991. The impact on poverty of food pricing reforms: a welfare analysis for Indonesia. *7. Policy Model*. 13:281–99
- Runciman WG. 1966. Relative Deprivation and Social Justice. London: Routledge Kegan Paul
- Sen A. 1981. Poverty and Famines: An Essay on Entitlement and Deprivation. Oxford, UK: Oxford Univ. Press
- Sen A. 1983. Poor, relatively speaking. Oxf. Econ. Pap. 35(2):153-69
- Sen A. 1985. Commodities and Capabilities. Amsterdam: North-Holland
- Sen A. 1987. The Standard of Living. Cambridge, UK: Cambridge Univ. Press
- Senik C. 2005. Income distribution and well-being: What can we learn from subjective data? *J. Econ. Surv.* 19(1):43–63
- Short K. 2012. The research supplemental poverty measure: 2011. Curr. Popul. Rep. P60-244, US Census Bur., Washington, DC
- Smith H, Pettigrew T, Pippin G, Bialosiewicz S. 2012. Relative deprivation: a theoretical and meta-analytic critique. Pers. Soc. Psychol. Rev. 16(3):203–32
- Smith V. 1959. Linear programming models for the determination of palatable human diets. J. Farm Econ. 41(2):272–83
- Stigler G. 1945. The cost of subsistence. J. Farm Econ. 27(2):303-14
- Summers R, Heston A. 1991. The Penn World Table (Mark 5): an extended set of international comparisons, 1950–1988. Q. J. Econ. 106:327–68
- Townsend P. 1979. Poverty in the United Kingdom: A Survey of Household Resources and Standards of Living. Harmondsworth, UK: Penguin Books
- UNDP (United Nations Dev. Program). 2005. Human Development Report. New York: UNDP
- Walker R. 2014. The Shame of Poverty. Oxford, UK: Oxford Univ. Press
- World Bank. 1990. World Development Report: Poverty. New York: Oxford Univ. Press
- World Bank. 2018. Piecing Together the Poverty Puzzle. Washington, DC: World Bank
- WHO (World Health Organ.). 1985. Energy and Protein Requirements. Tech. Rep. Ser. 724, WHO, Geneva