Policy Research Working Paper 10008

Absolute and Relative Poverty Measurement

A Survey

Benoit Decerf

WORLD BANK GROUP
Development Economics
Development Research Group
April 2022

Policy Research Working Paper 1008

Abstract

This paper reviews the debate opposing the absolute and relative approaches to monetary poverty measurement. The arguments for combining both approaches into a single "overall" monetary poverty measure are introduced. The most salient proposals of hybrid poverty lines are presented. Then, the reasons why specific poverty indices may be required when a hybrid line is used are discussed. The class of hierarchical poverty indices is described, focusing in particular on the hierarchical headcount ratio.

This paper is a product of the Development Research Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The author may be contacted at bdecerf@worldbank.org.

Absolute and Relative Poverty Measurement: A Survey

Benoit Decerf[1]

World Bank

Keywords:
Poverty measurement, Absolute poverty, Relative poverty, Hybrid poverty line, Poverty index

JEL Codes: D63, I32

[1] World Bank, Development Research Group. Email : bdecerf@worlbank.org.

Acknowledgments : This survey will appear as a chapter in the forthcoming Research Handbook of Poverty and Inequality (Wagle, 2022). I am grateful to Udaya R. Wagle and to an anonymous referee for useful comments and suggestions. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author and should not be attributed in any manner to the World Bank, its affiliated organizations, or members of its Board of Executive Directors or the countries they represent.

1) Introduction

Since the second half of the XXth century, poverty measures have gained momentum as indicators of social progress (Ravallion, 2016). Until the end of the XXth century, poverty was mostly understood as insufficient command over economic resources and monitored using monetary poverty measures. Recently, the idea that poverty is a multidimensional phenomenon has gained momentum. In spite of this evolution, monetary poverty measures still play a central role in the monitoring of poverty, both at national and international levels. This is maybe not surprising given that private expenditures (possibilities) are a central component of individual well-being, especially for the poor. In this survey, we only consider monetary poverty measures, which we often refer to as "income" poverty measures. However, as it is widely recognized, income poverty measures ignore other key facets of well-being, such as education, health care or security. As a result, the conclusions derived from such measures should always be confronted with insights obtained from other indicators of well-being capturing these dimensions.

There are two main approaches to income poverty. These two approaches capture different facets of income poverty, but formally, they only differ by the type of poverty line they use. First, the absolute approach captures basic needs satisfaction, e.g., nutrition and clothing. An absolute poverty line has a fixed threshold in the space of consumption. That is, the real value associated to an absolute poverty threshold should be the same in all societies. Second, the relative approach captures a form of "social inclusion", whose cost in terms of consumption increases with standards of living. Therefore, a relative line increases with standards of living. Thus, the real value associated to a relative poverty threshold is larger in richer societies.

Longstanding debates took place between the proponents of each of the two approaches. Over time, a middle ground has gradually emerged, according to which both approaches capture an important facet of income poverty and should therefore be monitored simultaneously. Going further, Atkinson and Bourguignon (2001) and then Ravallion and Chen (2011) have argued that these two aspects should be monitored by an 'overall' poverty measure, which combines the absolute and relative aspects of income poverty. Indeed, if individual welfare depends on both own income and relative income, a welfare-consistent poverty measure should capture both own income and relative income.

The key question when designing an 'overall' poverty measure is the trade-off that the measure makes between own income and relative income. Welfare-consistency requires that this trade-off be related to the trade-off that the individuals themselves would make. The literature on 'overall' poverty measures has mostly focused on the implication of welfare-consistency for the poverty line. These efforts have led to the design of 'hybrid' poverty lines, which mix an absolute line with a relative line. However, recent progress has also been made on the implication of welfare-consistency for the poverty index. This is important because the poverty index governs the trade-off made by the poverty measure below the poverty line.

In this survey, we review the debate on absolute vs relative poverty measurement and we introduce the arguments for combining both kinds of approaches into a single measure. We present the most salient proposals of hybrid poverty lines. We then move into the selection of a poverty index, which, we argue, should be carefully selected considering that a hybrid poverty line is used. In this context, we describe the class of "hierarchical" poverty indices, their convenient properties, and the hierarchical head-count ratio. These recently developed poverty lines and poverty indices pave the way for the widespread use of overall poverty measures. This

expansion has already started. Following the recommendation of Atkinson (2016), the World Bank adopted in 2018 an overall measure in order to monitor the evolution of global income poverty.

2) The debate opposing absolute poverty and relative poverty

2.1 Framework

We introduce the minimal elements necessary to discuss and compare absolute poverty and relative poverty. Let y denote some (non-negative) real income and m denote the income standard in a given income distribution ($y_1,\ldots,y_n$), e.g. mean income or median income. Let z(m) denote the (non-negative) poverty threshold, also called poverty line, where z is the threshold function. An individual is considered poor when her income is smaller than the poverty threshold. An additive income poverty index is defined as

$$P = \frac{1}{n}\sum_{i=1}^{n} p(y_i, m) \qquad (1)$$

where p is the contribution function. The contribution of any non-poor individual is zero. FGT poverty indices (Foster-Greer-Thorbecke, 1984), which are pervasive in applications, define the contribution of poor individuals as

$$p(y_i, m) = \left(1 - \frac{y_i}{z(m)}\right)^{\alpha}$$

where the poverty aversion parameter α is such that $\alpha \geq 0$. The head-count ratio corresponds to $\alpha = 0$ and the poverty-gap ratio corresponds to $\alpha = 1$.

In this context, the only difference between an absolute measure and a relative measure is the definition of the threshold function z. Hence, for our purpose, we can safely ignore several tricky questions that are faced by anyone who wishes to construct a poverty measure. These questions relate to the welfare variable, the unit of observation, the unit of analysis, the equivalence scale and the poverty index.

We have simply called 'income' the welfare variable that is being studied. In practice, the exact definition of y requires making choices with no easy answer: should we consider income or consumption? If the former is selected, should we consider gross or net income? Should we include some imputed rent for home-owners? The unit of observation, i.e. the level at which income information is collected, is typically the household. In contrast, the unit of analysis, the level at which poverty is calculated, is typically the individual. Moving from household income to individual income requires taking a stand on the intra-household allocation of resources. The bulk of the literature assumes that income is equally shared, but a growing literature is trying to open the black box and estimate how resources are shared within the household (Cherchye et al, 2011; Browning et al, 2013; Dunbar et al, 2013; Lechene et al., 2019). Finally, an equivalence scale is used in order to compare households of different sizes and compositions. This tool accounts for the fact that larger households require lower expenditures per person in order not to be in poverty, because of the presence of public goods in the household, like home, transport, etc. It also accounts for the fact that some individuals, like children, may have smaller needs than adults.

As we explain below, the selection of the poverty index matters when combining absolute and relative poverty. We assume throughout that the poverty index is additive. The additive structure follows when requiring the measure to be decomposable in subgroups (Foster and

Shorrocks, 1991). Sen (1992) questioned the desirability of subgroup consistency, which may be less compelling when relative aspects matter. However, Decerf (2021) shows that a weak version of subgroup consistency provides the foundation for additive indices as in Eq. (1). For illustrative purposes, we assume in this section that the additive measure takes the FGT form. However, many other indices have been proposed (see Zheng (1991) for a survey).

2.2 Absolute poverty

The key specificity of absolute measures is that they are based on a poverty line representing consumption possibilities that are fixed across space and time. Hence, the real value represented by the poverty threshold does not change with the level of the income standard m. Thus, measure P is absolute if the threshold function is independent on m, i.e. for some $z_a>0$ we have $z(m)=z_a$ for all m.

The construction of an absolute threshold is a complex task that requires making a series of difficult choices. Ravallion (1998) describes a two-step procedure in order to pin down the absolute threshold. The first step specifies a reference level of welfare below which individuals would be considered poor. The second step identifies a money metric threshold corresponding to this level of welfare. This second step is typically achieved by selecting a bundle of goods whose consumption provides the desired level of welfare and computing the cost of this bundle at prevailing prices.

The dominant procedure for the selection of the bundle of goods is the cost of basic needs approach (Ravallion, 1993; Ravallion and Sen 1996; Wodon, 1997). The typical way of operationalizing the cost of basic needs approach is to select a food bundle representing the consumption habits of the poor. Then, an allowance for non-food goods such as clothes or housing is computed. One possibility is to derive this allowance from food Engel curves. For instance, if the poor spend one-half of their budget on food, then the poverty threshold is taken to be twice the cost of the food bundle.

In practice, designing an absolute measure is a highly complex task. In order to maintain the same level of welfare, the poverty threshold must be adapted when prices vary. When all prices increase in the same proportion, such adaptation is relatively straightforward. When relative prices change, substitution effects change the consumption habits of the poor. In that case, a trade-off may appear between specificity, i.e. linking the poverty threshold to the behavior of the poor in the given context, and welfare-consistency, i.e. setting a poverty threshold that leads to a constant level of welfare. The trade-off between specificity and welfare consistency follows from the fact that the consumption habits associated to the new price vector may provide the poor with a different level of welfare as the one they reached under the initial prices. When monitoring global poverty, the World Bank relies on purchasing power parity rates (PPPs) in order to compute poverty thresholds that can be compared across different countries. The question whether these PPP conversions do provide welfare-consistent poverty thresholds frequently leads to heated debate (Summers and Heston, 1991; Deaton and Heston, 2010; Ravallion, 2018), reflecting the difficulty to ensure welfare-consistency in practice. Recently, Allen (2017) and Moatsos (2016) suggest an alternative method in order to obtain comparable poverty thresholds. These authors start from fixed requirements for several basic needs (e.g. minimal number of calories and other nutrients or minimal surface of housing) and compute the minimal cost to fulfill all these requirements at prevailing prices. This method avoids the use of PPP conversion altogether, but may generate other problems. For instance, the kinds of food diets solving such linear optimization problem are typically not consistent with prevailing

tastes (Ravallion, 2020). That is, this alternative method does not perform well in terms of specificity.

This very short discussion is far from exhausting the questions revolving around the many choices necessary to build an absolute poverty line. A thorough discussion about the current state of the literature can be found in Ravallion (2016). Yet, we have illustrated that the particularity of an absolute line is to reflect the conviction that every person has certain basic needs or rights, irrespective of the society in which she lives. The absolute threshold aims at capturing the cost of minimally satisfying a few basic needs in the context considered.

In low-income countries, the poor spend most of their resources on food. In these countries, the cost of the reference food bundle represents a large fraction of the absolute threshold. Therefore, a person who remains below the line for an extended period of time experiences severe difficulties to remain adequately nourished, which may affect her long-term survival. For this reason, the absolute lines designed in low-income countries capture a functioning that could be called 'subsistence'.

2.3 Relative poverty

The key specificity of relative measures is that they are based on a poverty line whose threshold is computed as a constant fraction of the income standard, i.e. for some $b \in (0,1)$ we have $z_r(m)=bm$. Hence, the real value represented by the relative poverty threshold changes with the level of the income standard m. The poverty threshold captures the income necessary to achieve what is considered a minimal living standard in the society considered. Below this threshold, individuals may experience difficulties to engage in the everyday life of their society, i.e. they are at risk of social exclusion.

Relative poverty measures find their foundation in the theory of relative deprivation of Runciman (1966) and in the work of Townsend (1979). In a nutshell, the idea behind relative poverty measures is that relative economic position is the key determinant of welfare and not the absolute economic position (Ravallion, 2008). A large literature has defended this viewpoint (Easterlin, 1995; van de Stadt et al. 1985; Alpizar et al. 2005). There is in particular a large body of evidence suggesting that relative income is a key determinant of subjective well-being (Clark and Oswald, 1996; Luttmer, 2005; Perez-Truglia, 2020). Thus, individual welfare is influenced by the income of some reference group.

The theory of relative deprivation postulates that the income of the reference group exerts a negative externality on individual welfare. This negative externality is consistent with many empirical results, although one should be aware that some theories and empirical works show that the externality could in fact be positive. For instance, richer communities are associated with better local public goods, which have a positive impact on welfare (Bowles and Gintis, 2002). Yet, there seems to be a broad consensus in the relative poverty literature that the negative externalities dominate the positive externalities.

The construction of a relative line requires making a series of non-obvious choices. First, one must select the relevant reference group. When computing national poverty, it is natural to consider all individuals living in the country. This choice becomes somewhat more difficult when computing poverty in a set of countries like the European Union (Brandolini and Rosolia, 2019). Second, one must select the statistic capturing the income standard. This statistic is the benchmark against which the individual's relative income is defined. In practice, most relative

lines are based on mean income or median income. There are debates about which of these two should be used. Median income is less sensitive than mean income to incomes in the top of the distribution, and has thus a lower variance in random population samples (Atkinson et al., 2002). However, relative measures based on median income may exhibit strange behavior when inequality changes (de Mesnard, 2007; Kämpke, 2010). Finally, one must select the proportion of the income standard defining the poverty threshold. Typical values range from 0.4 to 0.6 (see for instance Garroway and Laiglesia, 2012).

2.4 Absolute poverty versus relative poverty

Assessing poverty with an absolute measure may yield very different conclusions than doing this with a relative measure. When comparing these two kinds of measures, Notten and De Neubourg (2011) show that they typically yield different levels of poverty. Moreover, they show that the sign of the evolution of poverty is often reversed when switching from one kind of measure to the other. Reversals also frequently happen when comparing two different countries at a given point in time.

Unsurprisingly, there have been conceptual debates among defenders of each of these two approaches. The intellectual exchange between Sen (1993, 1995) and Townsend (1995) is a notable example. The latter argues that poverty was essentially a relative phenomenon, whereas the former defends the view that relative aspects must 'take the back seat' when individual resources are too low. Interestingly, the capability theory proposed by Sen allows reconciling to some extent the two views. According to this theory, poverty is absolute in the space of functionings, but the resources necessary to meet a fixed level of functionings may depend on the context. In particular, the cost of participating in the everyday activities of a society may increase with its income standard. As a result, an absolute threshold in the space of functionings may translate into a relative threshold in the space of income.

In practice, the national poverty measures of most developing countries are absolute, whereas those of most OECD countries are relative. There are exceptions, notably the United States that use an absolute measure, which was established building on the work of Orshansky (1965).

Even when a country has made its choice between the two approaches, the controversies are not necessarily over. In the long run, countries using absolute measures tend to revise their poverty threshold upward as their income standards increase. For instance, developing countries like China, India or Indonesia increased their absolute threshold over time (Ravallion, 2012; Jolliffe and Prydz, 2021). Also, some developing countries set their national poverty line using methodologies that may incorporate some elements of relativity. For instance, the food poverty line in the Arab Republic of Egypt is computed as the cost of the average food bundle consumed by households in the bottom 40% of the distribution (World Bank, 2011). When standards of living increase, households adapt their consumption patterns towards more expensive calories, which implicitly raises the real value of the food poverty line. Among developed countries, the US absolute threshold has been criticized (Citro and Michael, 1995), leading to the introduction of a second 'supplemental poverty measure' (SPM), which incorporates similar elements of relativity in its design. In the short run, countries using relative measures may observe some counterintuitive evolutions of their measure during periods of strong economic growth (Whelan and Maître, 2010) or periods of rapid economic decline (Halleröd and Heikkilä, 2002).

When performing cross-country comparisons, each approach might be criticized on the basis that it does not take into account the essence of the other approach. On the one hand, relative

lines tend to zero when the income standard tends to zero. As a result, relative thresholds in very low-income countries are so low that some question that they would even allow for biological survival (Ravallion, 2020). On the other hand, any given absolute line may be highly relevant in some countries, but might seem completely irrelevant in others. For instance, any individual living in the rich world on an amount corresponding to the extreme poverty threshold of the World Bank (Ferreira et al, 2016) would be unable to meaningfully participate in her society, and few would consider that such individual is not poor.

From these debates, several authors concluded that absolute measures and relative measures capture different facets of poverty, and that it would make sense to consider both simultaneously (Bourguignon, 1999). Following two measures of poverty, one absolute and one relative, would provide insights both on the satisfaction of basic needs and on social inclusion.

3) Measures combining absolute and relative poverty

3.1 The case for 'overall' poverty measures

Using simultaneously two poverty measures is not the same as using a single measure combining absolute and relative poverty. Two main arguments have been used to justify the use of such an 'overall' measure of income poverty.

First, Atkinson and Bourguignon (2001) argue that such measure is necessary for the monitoring of global poverty. Taking a global perspective requires a unified framework that considers the functionings sustaining the poverty measurement practices in low-income countries and high-income countries. The former rely on absolute measures capturing subsistence, whereas the latter rely on relative measures capturing social inclusion. Only such an overall measure is able to capture simultaneously these two functionings.

Second, Ravallion and Chen (2011) argue that a welfare perspective requires an overall measure because both own income and relative income are important determinants of individual welfare. Formally, the level of individual welfare for own income y and relative income y/m is w(y,y/m), where w is strictly increasing in y and weakly increasing in y/m. The formulation of individual welfare as w(y,y/m) is particularly convincing as it encompasses the absolute and relative approaches in a coherent way. Also, Ravallion (2020) shows how to augment the welfarist perspective by incorporating the idea of normative functionings considered by Atkinson and Bourguignon (2001).

The welfare perspective also reveals the reason why the use of a dashboard made of two separate poverty measures, one absolute and one relative, may be deemed unsatisfactory. Indeed, the dashboard approach yields a partial ranking of income distributions because such approach must remain silent when the two measures "disagree", i.e. yield opposing rankings of two distributions. Yet, for any definition of individual welfare, there exist pairs of distributions for which the absolute measure P^A disagrees with the relative measure P^R even if the welfare of all individuals is larger for one of the two distributions. That is, we could have two income distributions $(y_1,\ldots,y_n)$ and $(x_1,\ldots,x_n)$ for which $P^A(y_1,\ldots,y_n)>P^A(x_1,\ldots,x_n)$ but $P^R(y_1,\ldots,y_n)<P^R(x_1,\ldots,x_n)$ even if $w(y_i,y_i/m)> w(x_i,x_i/m)$ for all i (or the other way around). Hence, the dashboard approach may remain silent even when the comparison is totally unambiguous.

This insight is graphically illustrated in Figure 1, where we compare an individual 1 earning y^1 in a society with income standard m^1 with individual 2 earning y^2 in a society with income standard m^2 with $m^2 > m^1$. Individual 1 is absolutely poor but not relatively poor while individual 2 is relatively poor but not absolutely poor. Assume that all other individuals are non-poor, so the comparison of the two societies only depends on the comparison of individuals 1 and 2. The plain lines in figures 1.(a) and 1.(b) are iso-poverty curves, i.e. sets of bundles (y,m) that lead to a constant contribution p(y,m) for the respective poverty measures. The higher the iso-poverty curve, the lower the poverty contribution. Figure 1.(a) reveals graphically how an absolute measure compares individuals living in societies with different income standards. Figure 1.(b) does the same for a relative measure. The absolute measure finds more poverty in society 1 but the relative measure finds more poverty in society 2. That is, the two measure disagree and the dashboard approach remains silent. Figure 1.(c) shows the iso-welfare curves for some individual welfare function w. According to this welfare function, the welfare of individual 2 is larger than the welfare of individual 1, because the higher income of individual 2 more than compensates her smaller relative income. A welfare-consistent 'overall' poverty measure should thus conclude that poverty is smaller in society 2.

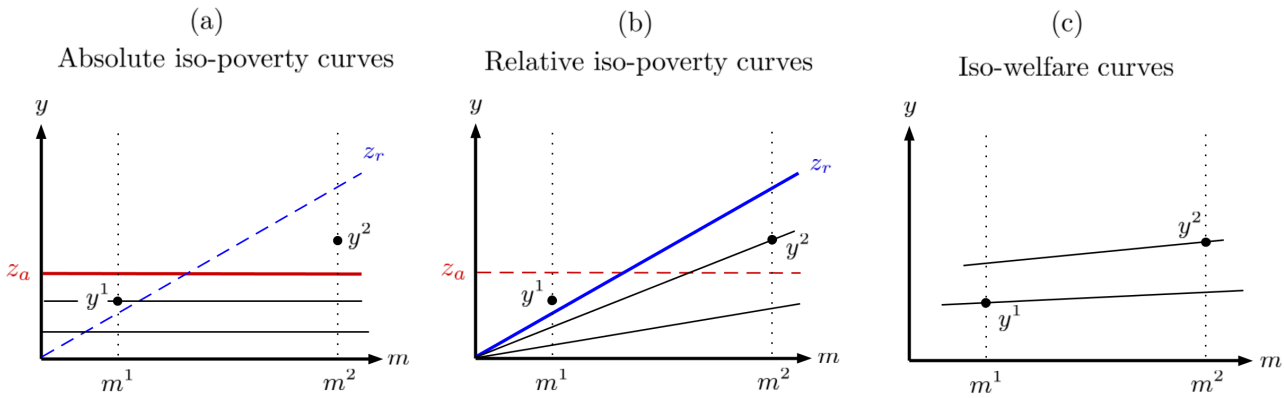

Figure 1: Contrasting the implicit comparisons of 'bundles' (y^1,m^1) and (y^2,m^2) according to an absolute measure (a), a relative measure (b) and some welfare function w (c).

Of course, identifying distributions for which the welfare of all individuals is increased requires knowing the welfare function w, which is a strong requirement. Some may thus argue in favor of the dashboard approach, on the grounds that two measures would present the whole information and leave the observer decides for herself how to trade-off the absolute and relative aspects. Unfortunately, this view implies aggregating each aspect across all individuals into P^A and P^R before weighing P^A and P^R. This runs counter to the view that a sound welfare analysis requires to first aggregate different dimensions at the individual level and then aggregate across individuals (Fleurbaey, 2009). Also, as we show below in Figure 2, a mild assumption on the welfare function may already help settle some cases for which the dashboard approach remains silent.

Together, there are strong arguments for the use of an overall poverty measure. Such a measure is called for because both own income and relative incomes are key determinants of individual welfare. It is also called for in contexts for which policy makers care for both subsistence and social inclusion.

The central question when designing a poverty measure that accounts for both own income and relative income is the trade-off that the measure makes between these two variables. For a welfare-consistent measure, this trade-off is encapsulated in the welfare function w. This trade-

off is the key for the comparison of two individuals living in societies with different income standards. In words, how does the measure compare Juan who lives in Colombia with Kader who lives in Bangladesh?

3.2 Design of the 'hybrid' poverty line

The literature on 'overall' poverty measures mostly focuses on the design of 'hybrid' poverty lines, which account for both the absolute and relative facets of income poverty. We review here the most salient proposals. For the sake of short notation, we sometimes denote an hybrid line by z, instead of z(m), thereby ignoring its dependence on the level of income standard m.

The early proposal by Foster (1998) defines the hybrid line z^F from an absolute line z_a and a relative line z_r as

$$z^F = z_r^{\rho} z_a^{1-\rho},$$

where $\rho \in (0,1)$ is the elasticity of the poverty line with respect to the income standard. As this elasticity parameter is constant, a 1-percent increase in the income standard m always leads to a ρ-percent increase in the poverty threshold. Foster hopes that the adoption of such hybrid line would help steer the debate about whether the poverty line should be absolute or relative towards the question 'to what extent the poverty line should be relative'? Foster does not tie his proposal to a definition of individual welfare w.

Atkinson and Bourguignon (2001) and Atkinson (2019) propose defining the hybrid line z^{AB} as

$$z^{AB} = \max(z_r, z_a).$$

Under their capability perspective, an individual is poor if she fails either of the two functionings under consideration, i.e. subsistence as captured by z_a and economic inclusion as captured by z_r. The particularity of this kind of hybrid line is that its threshold is absolute in low-income countries, and then relative in higher income countries. This kind of hybrid line is often referred to as a 'societal' poverty line in the literature.

Ravallion (2008) and Ravallion and Chen (2011) are the first to take a welfare perspective. Starting from a welfare function w(y,y/m), these authors make the point that the threshold function z should be "welfare-consistent", which formally means that the level of welfare w(z(m), z(m)/m) should not depend on the value of m, where z(m) is the hybrid line in a society whose level of income standard is m. Given that individual welfare is positively affected by relative income, the poverty line z(m) should increase with the income standard m. They also show that the relative line entering the definition of the societal line cannot correspond to a constant proportion of the income standard because individual welfare is strictly increasing in own income when holding relative income constant. In their terminology, a *weakly* relative line is defined as

$$z_{wr} = c + bm,$$

where the slope parameter $b \in (0,1)$ and the intercept $c \geq 0$. They call *strongly* relative any line z_r that corresponds to the definition of z_{wr} but for which $c = 0$. They propose a Weak Relativity Axiom (WRA) requiring that the poverty measure should strictly decrease when all incomes are multiplied by a constant larger than one. The motivation for the WRA being that the welfare of all individuals is strictly larger after such equi-proportionate growth. When the index is the head-count ratio, they show that such axiom implies for the intercept $c > 0$. As a result, the 'societal' poverty line should be

$$z^{RC} = \max(z_{wr}, z_a)$$

with $c > 0$. They note that the elasticity of their poverty line z^{RC} is not constant but tends to increase towards one as the level of income standard m tends to infinity, a property that matches their empirical data.

Chakravarty et al. (2016) suggest to axiomatize the hybrid poverty line. They impose two axioms on the individual welfare function w that jointly imply that w should be linear in the income standard m. The route that they outline seems promising, even if their own attempt is not exempt of limitations. In particular, their axioms are not ordinal requirements. That is, their axioms do not directly constrain the underlying preference over own income and relative income, but they rather constrain the welfare function that represents this preference. As a result, this preference may have undesirable features. For instance, one of the two welfare functions that they characterize is $w^C(y,y/m)=(q-a)y+am$, where $q>0$ and $a<0$. One of its undesirable feature is that individual welfare can be larger for zero income in some society than for strictly positive income in another society with larger income standard, i.e. $w^C(0,m) > w^C(y,m')$ for some $y>0$ and $m<m'$. Another is that this welfare function may be reduced when own income and the income standard are multiplied by the same constant larger than one, i.e. $w^C(ky,km) < w^C(y,m)$ for some $y>0$ and $m>0$ and any $k>1$, which happens as soon as $w^C(y,m) < 0$. The welfare functions that they derive imply the use of the following hybrid poverty line

$$z^C = sz_a + (1-s)m,$$

where $s\in(0,1)$. Formally, line z^C is a weakly relative line.

Ravallion and Chen (2019) provide the theoretical foundation for the use of a different statistic than mean or median income in order to define the income standard m. Their model points to the use of a distribution-corrected mean.

In order to set the parameters values for their societal poverty lines, many authors regress observed national poverty thresholds on the income standard (Atkinson and Bourguignon, 2001; Ravallion and Chen, 2011; Jolliffe and Prydz, 2021). In most global poverty applications, z_a is taken to be the extreme poverty threshold of the World Bank. In particular, the societal poverty line z^{JP} regressed by Jolliffe and Prydz (2021) has been adopted by the World Bank, which now reports global poverty figures on the basis of z^{JP}. The poverty line z^{JP} is a calibration of z^{RC} with parameters $c = 1, b = 0.5, z_a = 1.9$ (in 2011 PPP$) where m is captured by median income.

3.3 Selection of an 'overall' poverty index

The literature on poverty indices based on two poverty lines, one absolute and one relative, is much thinner. When constructing an poverty measure based on a hybrid line, authors typically take the poverty index to be the head-count ratio. This index is the fraction of individuals below the poverty line, which has the nice feature that it is very easy to interpret. Also, the properties of this index are known (Zheng, 1994). Yet, an often overlooked issue is that the properties of classical poverty indices, like FGT indices, have been characterized under the assumption that the poverty line is absolute. In other words, their properties under a hybrid line are unknown. In particular, there is no guarantee that the trade-off these indices make between own income and relative income is related to individual welfare w. If this is not the case, then 'overall' poverty measures require poverty indices that are specifically designed for their purpose.

Atkinson and Bourguignon (2001) are the first to recognize the issue and they propose a parametric family of poverty indices that aggregate income gaps with respect to the two poverty

lines z_a and z_r. They do not study the properties of their indices, but apply three of these indices in order to assess global poverty. Their 'overall' measures P^{AB} are based on a societal line z^{AB}, where z_a is the extreme poverty threshold of the World Bank, z_r is a strongly relative line and three different indices defined by Eq. (1) for $\alpha = 0$, $\alpha = 1$ and $\alpha = 2$. Anderson and Esposito (2014) propose another family of poverty indices along similar lines.

Decerf (2017) argues that 'overall' measures P^{AB} may make counter-intuitive poverty comparisons. The reason being that these measures violate the intuition that extremely poor individuals in low-income countries should be considered poorer than individuals who are only-relatively poor in middle-income countries. Indeed, the former are deprived from a minimal level of subsistence whereas the latter are not, because their income is above the extreme poverty threshold. He proposes a family of '*hierarchical'* indices that systematically consider that an absolutely poor is poorer (i.e., has a larger contribution) than an only-relatively poor

$$p^H(y_i, m) = \begin{cases} 1 - \lambda \dfrac{y_i}{m} & when\ y_i < z_a \\ (1-\lambda) - (1-\lambda)\dfrac{y_i - z_a}{z - z_a} & when\ z_a < y_i < z \end{cases}$$

where z_a is the absolute threshold, z is the 'societal' line and $\lambda \in [0,1]$ is the parameter that tunes the priority given to an absolutely poor over an only-relatively poor. The priority given to the absolutely poor is larger for larger values of λ (Decerf and Ferrando, 2021). However, the sensitivity to errors in the measurement of the absolutely poor individual's income is also larger for larger values of λ.

The fact that hierarchical indices implicitly consider any absolutely poor to be poorer than any only-relatively poor is graphically illustrated in Figure 2. In that figure, we compare an individual 3 earning y^3 in a society with income standard m^3 with an individual 4 earning y^4 in a society with income standard m^4 with $m^4 > m^3$. Individual 3 is absolutely poor while individual 4 is only-relatively poor. Assume that all other individuals are non-poor, so the comparison of the two societies only depend on the comparison of individuals 3 and 4. The iso-poverty curves in Figure 2.(a) reveal graphically how an Atkinson-Bourguignon measure P^{AB} compares individuals living in societies with different income standards. In a nutshell, its inter-personal comparisons correspond to those of an absolute measure in countries with income standards below m* and they correspond to those of a relative measure in countries with income standards above m*. Crucially, measure P^{AB} implicitly considers that the only-relatively poor individual 4 is poorer than the absolutely poor individual 3. This is unlike the hierarchical measure P^H, whose iso-poverty curves are flat below the absolute threshold z_a as illustrated in Figure 2.(b). As a result, measure P^H implicitly considers that the only-relatively poor individual 4 is *less* poor than the absolutely poor individual 3. The comparisons of measure P^H corresponds to those of an absolute measure below the absolute threshold z_a, and it is only above z_a that relative considerations come into play. This example constitutes another illustration of the limitation of using a dashboard of two poverty measures, one absolute and one relative, instead of an 'overall' measure. Indeed, the absolute measure and the relative measure would disagree when comparing individuals 3 and 4, at least if these measures are gap-sensitive. As a result, the dashboard approach remains silent even if individual 4 is unambiguously less poor than individual 3 if one endorses the view that the absolutely poor are poorer than the only-relatively poor.

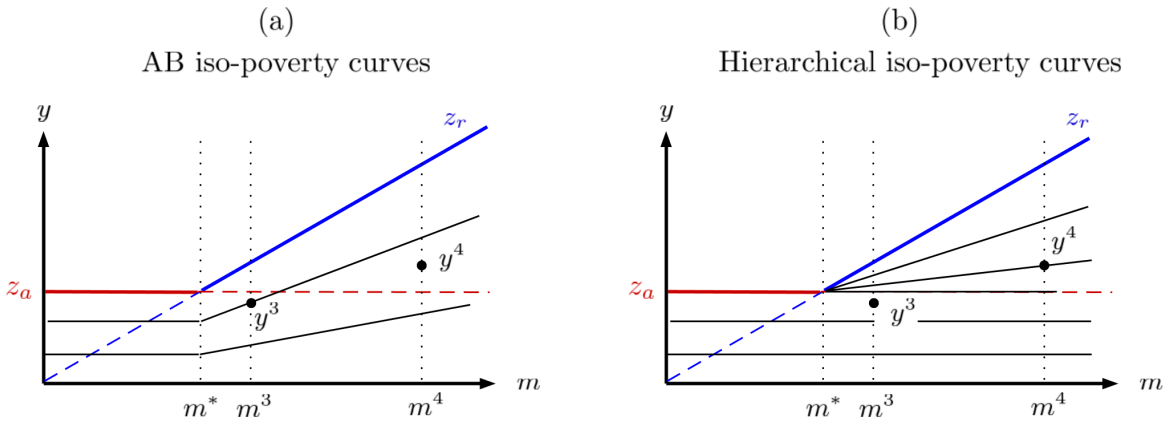

Figure 2: Contrasting the implicit comparisons of 'bundles' (y^3,m^3) and (y^4,m^4) according to an Atkinson-Bourguignon measure P^{AB} (a), and to a hierarchical measure P^H (b).

Decerf (2021) shows that, when considering a societal line z^{AB} or z^{RC}, a set of properties à la Foster and Shorrocks (1991) characterizes the family of poverty indices that satisfy the 'hierarchical' assumption, namely that absolutely poor individuals are poorer than only relatively poor individuals. The particularity of this characterization result is that it does not assume that the poverty line is absolute.

Decerf and Ferrando (2021) assess the evolution of global poverty based on the family of hierarchical indices p^H. First, they conduct a survey whose results suggest that 'hierarchical' inter-personal comparisons collect broad support in the US, the UK and South Africa. Second, they show that the overall poverty trend is sometimes independent on the (arbitrary) value given to the priority parameter λ even for some cases for which absolute measures P^A follow an opposite trend as relative measures P^R. That is, when endorsing hierarchical inter-personal comparisons, it is sometimes possible to unambiguously settle a disagreement between absolute and relative poverty. This possibility is illustrated in Figure 2, where all the hierarchical measures P^H conclude that there is more poverty in society 3. Third, they show that the selection of a poverty index may have an impact at least as large on the evolution of overall poverty than the design of a societal poverty line.

3.4 Joint selection of the poverty line and the poverty index

Recently, Decerf et al. (2022) propose a theoretical framework that allows making progress on two limitations of the literature. First, the literature assumes the existence of a common welfare function w that makes the same trade-off between own income and relative income for all individuals. Second, the literature has only studied the implications of welfare-consistency on the poverty line. Decerf et al. (2022) study the implications for the poverty line and for the poverty index of the "escaping-poverty" property, i.e., the requirement that the poverty measure is reduced when a poor individual escapes poverty.

Their results provide support for the use of 'societal' poverty lines together with 'hierarchical' poverty indices. Classical indices like FGT indices are not hierarchical. Their results imply that using the head-count ratio is less attractive when constructing an 'overall' poverty measure. In particular, in the presence of heterogenous preferences, the head-count ratio need not be reduced when a poor individual escapes poverty. Indeed, when preferences are heterogenous, the head-count ratio no longer corresponds to the fraction of poor individuals. The reason being

that some individuals may be non-poor, i.e., be better-off than at the poverty bundle, even though they earn an income smaller than the societal line. For this reason they call the societal line a "deprivation" line, rather than a "poverty" line. They suggest to use instead a modification of the head-count ratio that encapsulates the main features of their theory. This index can be written as

$$HH = H_A + \omega\, H_R \quad \text{where the weight} \quad \omega = \frac{z-\hat{y}_R}{z-z_a} \quad \text{and where}$$

H_A is the fraction of absolutely deprived individuals (all individuals i for whom $y_i<z_a$), H_R is the fraction of only-relatively deprived individuals (all individuals i for whom $z_a \leq y_i<z$) and $\hat{y}_R$ is the average income among the only-relatively deprived individuals, which implies that the weight ω takes a value smaller than 1. They call this index the hierarchical head-count ratio. This index corresponds to the hierarchical index p^H with parameter value $\lambda=0$.

4 Concluding remarks

The debate between absolute and relative poverty measures has been long-lasting because both sides have good arguments. As Bourguignon (1999) observed, these two approaches capture different facets of poverty, namely basic needs and a form of social inclusion, and it makes sense to monitor both simultaneously. Ravallion and Chen (2011) provide a compelling argument that an 'overall' measure combining both aspects might be superior to following two poverty measures separately, one absolute and one relative. They note that individual welfare depends on both own income and relative income. As a result, a welfare-consistent income poverty measure should aggregate the two aspects at the individual level. Moreover, such measures should make a trade-off between own income and relative income that is related to individual preferences. They show that this implies the use of a 'societal' poverty line based on a weakly relative line. As noted by Atkinson and Bourguignon (2001), an 'overall' measure should not only use a specific 'societal' poverty line, but also a specific poverty index. Decerf et al. (2022) show that the 'overall' measure is reduced when a poor individual escapes poverty only if the poverty index is 'hierarchical'. They suggest using a particular hierarchical index that is a simple modification of the head-count ratio. Together, the recent developments in the literature on 'overall' poverty measures pave the way for applications and for their gradual integration in the monitoring of income poverty all over the world.

5 References

Allen, R. C. (2017). Absolute poverty: When necessity displaces desire. *American Economic Review*, *107*(12), 3690-3721.

Alpizar, F., Carlsson, F., & Johansson-Stenman, O. (2005). How much do we care about absolute versus relative income and consumption?. *Journal of Economic Behavior & Organization*, *56*(3), 405-421.

Anderson, E., & Esposito, L. (2014). On the joint evaluation of absolute and relative deprivation. *The Journal of Economic Inequality*, *12*(3), 411-428.

Atkinson, A. (2016). *Monitoring Global Poverty: Report of the commission on global poverty*. The World Bank, Washigton, DC.

Atkinson, A. B. (2019). *Measuring poverty around the world*. Princeton University Press.

Atkinson, A. B., & Bourguignon, F. (2001). Poverty and inclusion from a world perspective.

Atkinson, T., Cantillon, B., Marlier, E., & Nolan, B. (2002). *Social indicators: The EU and social inclusion*. Oup Oxford.

Blanco, M., Engelmann, D., & Normann, H. T. (2011). A within-subject analysis of other-regarding preferences. *Games and Economic Behavior*, *72*(2), 321-338.

Bourguignon, F. (1999). Absolute poverty, relative deprivation and social exclusion. In *Villa Borsig Workshop Series* (pp. 1-4).

Bowles, S., & Gintis, H. (2002). Social capital and community governance. *The economic journal*, *112*(483), F419-F436.

Brandolini, A., & Rosolia, A. (2019). The distribution of well-being among Europeans. *Bank of Italy Occasional Paper*, (496).

Browning, M., Chiappori, P. A., & Lewbel, A. (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *Review of Economic Studies*, *80*(4), 1267-1303.

Chakravarty, S. R., Chattopadhyay, N., Deutsch, J., Nissanov, Z., & Silber, J. (2016, November). Reference groups and the poverty line: An axiomatic approach with an empirical illustration. In *Inequality after the 20th Century: Papers from the Sixth ECINEQ Meeting*. Emerald Group Publishing Limited.

Cherchye, L., De Rock, B., & Vermeulen, F. (2011). The revealed preference approach to collective consumption behaviour: Testing and sharing rule recovery. *The Review of Economic Studies*, *78*(1), 176-198.

Citro, C., Michael, R. (1995). Measuring Poverty : A New Approach. Washington, DC : Natl. Acad.

Clark, A. E., & Oswald, A. J. (1996). Satisfaction and comparison income. *Journal of public economics*, *61*(3), 359-381.

Deaton, A., & Heston, A. (2010). Understanding PPPs and PPP-based national accounts. *American Economic Journal: Macroeconomics*, *2*(4), 1-35.

Decerf, B. (2017). Why not consider that being absolutely poor is worse than being only relatively poor?. *Journal of Public Economics*, *152*, 79-92.

Decerf, B. (2021). Combining absolute and relative poverty: income poverty measurement with two poverty lines. *Social Choice and Welfare*, *56*(2), 325-362.

Decerf, B., & Ferrando, M. (2021). Unambiguous trends combining absolute and relative income poverty: new results and global application. *The World Bank Economic Review*.

Decerf, B., & Ferrando, M., & Quinn, N. (2022). *Global Income Poverty Measurement with Preference Heterogeneity : Theory and Application. CentER Discussion Paper 2022-007.*

Dunbar, G. R., Lewbel, A., & Pendakur, K. (2013). Children's resources in collective households: identification, estimation, and an application to child poverty in Malawi. *American Economic Review*, *103*(1), 438-71.

de Mesnard, L. (2007). Poverty reduction: the paradox of the endogenous poverty line. *LEG, Economy Series, Working Paper*, (2007-05).

Easterlin, R. A. (1995). Will raising the incomes of all increase the happiness of all?. *Journal of Economic Behavior & Organization*, *27*(1), 35-47.

Ferreira, F. H., Chen, S., Dabalen, A., Dikhanov, Y., Hamadeh, N., Jolliffe, D., ... & Yoshida, N. (2016). A global count of the extreme poor in 2012: data issues, methodology and initial results. *The Journal of Economic Inequality*, *14*(2), 141-172.

Fleurbaey, M. (2009). Beyond GDP: The quest for a measure of social welfare. *Journal of Economic literature*, *47*(4), 1029-75.

Foster, J., Greer, J., & Thorbecke, E. (1984). A class of decomposable poverty measures. *Econometrica: journal of the econometric society*, 761-766.

Foster, J. E., & Shorrocks, A. F. (1991). Subgroup consistent poverty indices. *Econometrica: Journal of the Econometric Society*, 687-709.

Garroway, C., & De Laiglesia, J. R. (2012). On the relevance of relative poverty for developing countries.

Halleröd, B., & Heikkilä, M. (2002). Poverty and social exclusion in the Nordic countries. In *Nordic social policy* (pp. 203-232). Routledge.

Jolliffe, D., & Prydz, E. B. (2021). Societal poverty: A relative and relevant measure. *The World Bank Economic Review*, *35*(1), 180-206.

Kämpke, T. (2010). The use of mean values vs. medians in inequality analysis. *Journal of economic and social measurement*, *35*(1-2), 43-62.

Lechene, V., Pendakur, K., & Wolf, A. (2019). *OLS estimation of the intra-household distribution of consumption* (No. W19/19). IFS Working Papers.

Luttmer, E. F. (2005). Neighbors as negatives: Relative earnings and well-being. *The Quarterly journal of economics*, *120*(3), 963-1002.

Moatsos, M. (2016). Global absolute poverty: behind the veil of dollars. *Journal of Globalization and Development*, *7*(2).

Notten, G., & De Neubourg, C. (2011). Monitoring absolute and relative poverty:"Not enough" is not the same as "much less". *Review of Income and Wealth*, *57*(2), 247-269.

Perez-Truglia, R. (2020). The effects of income transparency on well-being: Evidence from a natural experiment. *American Economic Review*, *110*(4), 1019-54.

Orshansky, M. (1965). Counting the poor: Another look at the poverty profile. *Soc. Sec. Bull.*, *28*, 3.

Ravallion, M. (2008). *On the welfarist rationale for relative poverty lines* (Vol. 4486). World Bank Publications.

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). *Poverty comparisons*. Taylor & Francis.

Ravallion, M. (2016). *The economics of poverty: History, measurement, and policy*. Oxford University Press.

Ravallion, M. (2018). An exploration of the changes in the international comparison program's global economic landscape. *World development*, *105*, 201-216.

Ravallion, M. (2020). On measuring global poverty. *Annual Review of Economics*, *12*, 167-188.

Ravallion, M., & Chen, S. (2011). Weakly relative poverty. *Review of economics and statistics*, *93*(4), 1251-1261.

Ravallion, M., & Chen, S. (2019). Global poverty measurement when relative income matters. *Journal of public economics*, *177*, 104046.

Ravallion, M., & Sen, B. (1996). When method matters: Monitoring poverty in Bangladesh. *Economic development and cultural change*, *44*(4), 761-792.

Runciman, W. G. (1966). *Relative deprivation and social justice: A study of attitudes to social inequality in twentieth-century England* (Vol. 13). Routledge & Kegan Paul.

Sen, A. (1983). Poor, relatively speaking. *Oxford economic papers*, *35*(2), 153-169.

Sen, A. (1985). A sociological approach to the measurement of poverty: a reply to Professor Peter Townsend. *Oxford Economic Papers*, *37*(4), 669-676.

Sen, A. (1992). *Inequality reexamined*. Oxford University Press.

Summers, R., & Heston, A. (1991). The Penn World Table (Mark 5): an expanded set of international comparisons, 1950–1988. *The Quarterly Journal of Economics*, *106*(2), 327-368.

Townsend, P. (1979). Poverty in the United Kingdom. Harmondsworth, Middlesex: Penguin Books

Townsend, P. (1985). A sociological approach to the measurement of poverty—a rejoinder to Professor Amartya Sen. *Oxford economic papers*, *37*(4), 659-668.

Van de Stadt, H., Kapteyn, A., & Van de Geer, S. (1985). The relativity of utility: Evidence from panel data. *The review of Economics and Statistics*, 179-187.

Wagle (2022). Research Handbook of Poverty and Inequality, Edward Elgar. Forthcoming.

Whelan, C. T., & Maitre, B. (2010). Poverty in Ireland in comparative European perspective. *Social Indicators Research*, *95*(1), 91.

Wodon, Q. T. (1997). Food energy intake and cost of basic needs: Measuring poverty in Bangladesh. *The Journal of Development Studies*, *34*(2), 66-101.

World Bank (2011) Arab Republic of Egypt. Poverty in Egypt 2009-09. Withstanding the Global Economice crisis. Report No. 60249-EG

Zheng, B. (1994). Can a poverty index be both relative and absolute?. *Econometrica: Journal of the Econometric Society*, 1453-1458.

Zheng, B. (1997). Aggregate poverty measures. *Journal of economic surveys*, *11*(2), 123-162.