

On Multidimensional Indices of Poverty

Martin Ravallion

The World Bank
Development Research Group
Director's office
February 2011



Abstract

There has been a growing interest in what have come to be termed “multidimensional indices of poverty.” Advocates for these new indices correctly point out that command over market goods is not all that matters to peoples’ well-being, and that other factors need to be considered when quantifying the extent of poverty and informing policy making for fighting poverty. However, the author argues that there are two poorly understood issues in assessing these indices. First, does one believe that any single index can ever be a sufficient statistic for poverty assessments? Second, when aggregation is called

for, should it be done in the space of “attainments,” using prices when appropriate, or that of “deprivations,” using weights set by the analyst? The paper argues that the goal for future poverty monitoring efforts should be to develop a credible set of multiple indices, spanning the dimensions of poverty most relevant to a specific setting, rather than a single multidimensional index. When weights are needed, they shouldn’t be set solely by an analyst measuring poverty. Rather, they should be, as much as possible, consistent with well-informed choices made by poor people.

This paper is a product of the Director’s office, Development Research Group. 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://econ.worldbank.org>. The author may be contacted at mravallion@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

On Multidimensional Indices of Poverty

Martin Ravallion¹

¹ The author is Director of the Development Research Group, World Bank, 1818 H Street NW, Washington DC, 20433, USA. For helpful discussions on this topic and comments on this paper the author is grateful to Sabina Alkire, Kathleen Beegle, Gabriel Demombynes, Quy-Toan Do, Jean-Yves Duclos, Francisco Ferreira, James Foster, Garance Genicot, Stephan Klasen, Peter Lanjouw, Michael Lokshin, Nora Lustig, Branko Milanovic, Mead Over, Dominique van de Walle, Roy Van der Weide and Hassan Zaman. These are the views of the author, and need not reflect those of the World Bank or any affiliated organization.

1. Introduction

There has been a growing interest in what have come to be termed “multidimensional indices of poverty.” There are many issues one might discuss related to these indices, including the choice of the functional form, the choice of poverty lines, and the robustness of the implied rankings to those choices.² However, these issues are essentially generic to all poverty measures (though with some more technical differences). The present discussion will focus instead on how multidimensional indices differ from more familiar approaches. This is a logical starting point for potential users trying to understand and apply these new measures; to assess their contribution we must first understand how they differ from standard measures.

Multidimensionality *per se* cannot be what distinguishes a multidimensional index of poverty (MIP). There is an obvious sense in which almost every poverty measure found in practice is “multidimensional.” Indeed, to my knowledge, the only truly one-dimensional indices are the rice-based measures once found in some countries in Asia, but no longer used.³ The main measures now found in practice use a composite measure of consumption or income with many components, relying heavily on market prices in aggregation.

Nor does the difference lie in the recognition of the fact that poverty is not just about low consumption of market commodities. It is widely agreed that there are also important non-market goods relevant to welfare, such as access to public services. Poverty is multidimensional. However, that does not imply that one needs a MIP. It is one thing to recognize that something is missing from a given measure, and needs to be considered, and quite another to create a single composite index. The more common approach is to collect multiple indicators of the various dimensions of poverty, invariably including an index of command over market goods, but also including indicators for health and education attainments and access to services. A well-known example is the United Nations’ *Millennium Development Goals*, which span multiple

² For example, in the case of poverty measurement, where there is almost always a degree of arbitrariness about the poverty line, best practice tests the robustness of poverty comparisons to the choices made, invoking the theory of stochastic dominance. For expositions in the standard “unidimensional” case see Atkinson (1987) and Ravallion (1994). Duclos et al. (2006) provide dominance tests for “multidimensional poverty.”

³ For example, the government of Vietnam measured poverty by rice consumption prior to the early 1990s. This was in part at least because of high inflation in the 1980s, and inadequate price indices. More standard multi-commodity poverty measures emerged soon after.

dimensions, but without forming a single composite index. At the country level, the World Bank's *Poverty Assessments* and the *Poverty Reduction Strategy Papers* of individual governments have typically drawn on multiple indicators (though naturally with varying emphasis), without forming a single composite index.⁴

This paper argues that the real differences between the recent measures that are called “multidimensional” and standard approaches lie elsewhere. The first difference is in whether one believes that a single index of poverty could ever be a sufficient statistic, or whether multiple indices are required, each measuring different things using the best data available for that task—presenting us with a “large and eclectic dashboard” (Stiglitz et al., 2009, p.62). A second difference is also evident on closer inspection, namely how the analyst chooses to collapse multiple dimensions into one, recognizing that some degree of aggregation will probably be called for even in the “dashboard” approach.

In elaborating these two differences I will illustrate the arguments using the most well-developed and broadly applied MIP to date, namely that developed by Alkire and Santos (2010a), which is a special case of the class of measures proposed by Alkire and Foster (2007). The following section discusses the Alkire-Santos index, and whether it can be considered sufficient for measuring poverty and informing policy making. Section 3 turns to the issue of how one can go about aggregating across multiple dimensions when a degree of aggregation is called for to reduce the dimensionality. Section 4 concludes.

2. Can we measure poverty adequately with any single index?

There are countless possibilities for forming composite indices by some form of essentially *ad hoc* aggregation—giving what I term elsewhere “mashup indices” (Ravallion, 2010a). The issue in this section is not how this can be done but whether it is useful to do so. One can easily imagine situations in which one would not want a mashup index. Imagine you go for your annual medical checkup. Your doctor does all the usual tests, but tells you that she will base her assessment solely on a single composite index—rescaling and averaging all the test

⁴ The Bank's Poverty Assessments (significant analytic reports covering virtually all developing countries) typically cover education, health and nutrition and access to infrastructure, in addition to income poverty.

results. You would be well advised to get a new doctor!⁵ Or imagine that a new car comes on the market that collapses all those dials on the dashboard into just one composite index, on which you are supposed to decide what to do (slow down or get fuel). You would surely not buy this car!

Essentially the same point applies to the task of prioritizing policies for fighting poverty in a given country (or other geographic area). We will naturally want to look at the country's attainments in various dimensions, rather than focusing on its performance with respect to a single composite index. Should we be focus on promoting job creation (say) or better health and education services? Such an approach does not deny that poverty is "multidimensional." Rather it says that forming a single (uni-dimensional) index may not be particularly useful for sound development policy making.

Consider now the MIP developed by Alkire and Santos (2010a) for the 2010 *Human Development Report* (UNDP, 2010). They choose 10 variables for their MIP under the same three headings—health, education and living standards—as the UNDP's *Human Development Index* (HDI). There are two variables for health (malnutrition, and child mortality), two for education (years of schooling and school enrolment), and six for deprivation in "living standards" (namely cooking with wood, charcoal or dung; not having a conventional toilet; lack of easy access to safe drinking water; no electricity; dirt, sand or dung flooring and not owning at least one of a radio, TV, telephone, bike or car). Poverty is measured separately in each of these 10 dimensions. The equally-weighted aggregate poverty measures for each of these three main headings are then weighted equally (one-third each) to form the composite index, also echoing the HDI. A household is identified as being poor if it is deprived across at least 30% of the weighted indicators. While the HDI uses aggregate country-level data, the Alkire-Santos MIP uses household-level data, which are then aggregated to the country level. Alkire and Santos construct their index for more than 100 countries.

⁵ In certain emergency situations (such as on the battle field), treatment decisions often require a prioritization of patients ("triage") and it appears that this is typically based on the probability of survival, which is a single index. But then one is not creating a "mashup index" since the variables and weights are entirely determined by their ability to predict that probability. There is nothing analogous to this probability in a MIP. As [Mead Over](#) points out in a blog comment: "In the physical examination example, where the situation is not life threatening, both the doctor and the patient presume that the patient's valuation of the information deserves priority, since presuming otherwise would be unnecessarily paternalistic."

The Alkire-Santos MIP is a special case of the theoretical measure proposed in an elegant formulation of the problem by Alkire and Foster (2007). This fact helps in understanding the Alkire-Santos MIP, but does not really get us very far since the theoretical formulation in Alkire and Foster (in common with other papers in this literature) takes virtually all the elements of the measure as given (determined outside the measure), notably the dimensions of poverty, the dimension-specific cutoffs, the weights on deprivations and the minimum number of deprivations needed to be deemed “poor.” As we will see, the devil is in these details.

What dimensions? A key step in implementing any multidimensional measure is to select a set of dimensions. There is, of course, ample scope for debate here. There is a (rather poorly-understood) issue about what dimensions are intrinsically, versus instrumentally, important. For example, we can probably all agree that “health” is valued intrinsically, independently of command over commodities. However, it is more contentious that education has such an intrinsic value—as implicitly assumed by the Alkire-Santos MIP—rather than being (very) important to income and (hence) command over commodities (and health too). And even if we agree that education is to be valued intrinsically, it is far from clear that “education poverty” should have the same weight as “health poverty.”

The data requirements of the Alkire-Foster index entail that relevant dimensions of poverty must invariably be left out in practice. Consider first the material goods entering the Alkire-Santos MIP. This is a rather narrow set of goods, leaving out a great many things that people do in fact consume. The consumption measure formed from a modern household budget or living-standards survey will aggregate (actual or imputed) expenditures on literally 100’s of consumption items (1,000 or more items in some surveys). Yet the Alkire-Santos MIP only identifies six factors for “living standards” (as described above). So their measure leaves out a great many of the multiple dimensions poverty—indeed, their MIP has far fewer dimensions of living standards than those included in a standard (“unidimensional”) consumption-based measure.

Nor does the index appear to span the relevant “non-income” dimensions. In a blog comment, [Duncan Green](#) criticized the Alkire-Santos MIP for leaving out “conflict, personal security, domestic and social violence, issues of power/empowerment” and “intra-household dynamics.”

Why is so much left out? In practice, the choice of dimensions for measuring poverty will naturally be constrained by the data. When following the Alkire-Foster approach, the options are constrained further by the fact that one must obtain all the data for the same sampled household. So they must all come from the same survey.⁶ Yet most surveys do not cover all the things one would like to know in a comprehensive assessment of poverty. This restricts the set of dimensions that can be included in the MIP. Nor can it be presumed that the dimensions that can be measured this way are representative of some subset of dimensions, within some seemingly reasonable taxonomy (such as “consumption poverty,” “health poverty” or “education poverty”). There will often be other data available on the selected dimensions, and other data on other relevant dimensions, but only from different surveys.

This aspect of the Alkire-Foster approach suggests that we will inevitably fall back in practice on the standard approach I described at the outset in which we use multiple indices rather than a single index. If one chooses not to form the composite at household level but to look instead at the separate dimensions of poverty then one is in a better position to span the relevant dimensions and to choose the best available data on each.

While this aspect of the Alkire-Foster methodology comes at a cost in terms of the coverage of the relevant dimensions of poverty, it can be acknowledged that it has an advantage too in that one can get some idea of the joint distribution of the multiple dimensions of poverty—to what extent the different dimensions of poverty that can be identified are shared by the same people. This adds something that cannot be easily identified when using multiple surveys (though simulation methods are sometimes used for that purpose). When a survey for a specific country does span multiple dimensions there can be much interest in exploring their joint distribution, though a MIP is not the only tool available for that purpose.⁷

Another data constraint also points to the need for multiple measures in practice, namely that the data we typically use in measuring poverty do not tell us much about consumption within the household. To use such data we need to make assumptions about intra-household

⁶ Unique identifiers can in principle link households across two or more surveys, but this is relatively rare in practice, especially in developing countries.

⁷ See, for example, the study by Lokshin and Ravallion (2005) of the joint distribution of wealth and power. Lokshin and Ravallion use standard statistical tools for this purpose, including correlation and regression methods and contingency tables and related statistics for ordinal categorical data.

distribution, such as the seemingly strong assumption of welfare equality within the household (Haddad and Kanbur, 1990). This data constraint points to the desirability of supplementary indicators of individual attainments. Data on child health and mortality has understandably been given high priority in looking for such indicators.

So why aggregate in the form of a MIP? Alkire and Santos (2010b, p.7) argue that their index "... goes beyond previous international measures of poverty to identify the poorest people and aspects in which they are deprived. Such information is vital to allocate resources where they are likely to be most effective." But is it the MIP or its components that matter for this purpose? Following Alkire and Foster (2007), the Alkire-Santos MIP has a neat decomposability; we can reverse the aggregation. This is useful, for only then will we have any idea how to go about addressing the poverty problem in that specific setting. But the question remains: why do we need the aggregation in the first place?

Consider the following stylized example of a policy problem. Suppose that there are two dimensions of welfare, "income" and "access to services." Assume that an "income-poor" but "services-rich" household attaches a high value to extra income but a low value to extra services, while the opposite holds for an "income-rich" but "services-poor" household.⁸ There are two policy instruments, a transfer payment and service provision. The economy is divided into geographic areas (which could be countries) and a given area gets either the service or the transfer. We then calculate a composite index like the Alkire-Santos MIP based on survey data on incomes and access to services. There is bound to be a positive correlation between average income and service provision, but (nonetheless) some places have high income poverty but adequate services, while others have low income poverty but poor services. The policy maker then decides whether each area gets the transfer or the service. Plainly, the policy maker should not be using the aggregate MIP for this purpose, for then some income-poor but service-rich households will get even better services, while some income-rich but service poor households will get the transfer. The total impact on (multidimensional) poverty would be lower if one based the allocation on the MIP rather than the separate poverty measures—one for incomes and one

⁸ Sufficient conditions are that there is declining marginal utility to both income and services and that the marginal utility of income (services) is non-decreasing in services (income).

for access to services. It is not the aggregate index that we need for this purpose but its components.

Once we recognize that not even the Alkire-Santos MIP can be considered a sufficient statistic for poverty assessments, and that it will need to be complemented by other measures for those things left out, we are essentially back to the standard approach of using multiple measures. Whether that still leaves a role for the MIP depends on whether we think that it captures an important subset of the relevant dimensions of poverty. Here the ambition to be “multidimensional” is in such marked tension with the need to obtain all dimensions for each surveyed household that one must question what role such an index can usefully play. Arguably it would be better to derive the best measure possible for each of a logically defensible set of grouped dimensions—such as “income poverty,” “health poverty” and “education poverty”.

Clearly that task will still require some degree of aggregation. How should that be done? Here too there is an important, but poorly understood, distinction between the approach taken by recent MIPs and standard approaches to measuring poverty.

3. How should we aggregate?

One can distinguish two approaches to forming an aggregate poverty index. The first is to use prices (actual or imputed) to form a composite index for aggregate consumption, to be compared to a poverty line defined in the same space.⁹ Ideally this is not just consumption of market goods and services, but should include imputed values for non-market commodities. For market goods, either their market prices or appropriate shadow prices can be used. For non-market goods the missing “prices” will need to be assigned on *a priori* grounds or estimated. In practice, most poverty measures require imputations for missing prices, so this approach is a natural extension of prevailing practices. In principle we can broaden this approach to allow for non-commodity dimensions of welfare. The space defined by all primary dimensions of welfare (including commodities) can be called the “attainment space” (though the term “achievements” is also used in the literature), and the aggregation can be called “attainment aggregation.” The weights on attainments can be called “prices,” understood to include imputed prices.

⁹ Household consumption is typically normalized for differences in household size and the prices faced; equivalently the poverty line is so differentiated.

A simple example of a poverty measure using attainment-aggregation is the usual headcount index of poverty:¹⁰

$$P^A \equiv F_y(z) \tag{1}$$

where F_y is the distribution function for aggregate consumption y and z is the poverty line in that space. To keep things simple for expository purposes (including graphing), suppose that there are two attainments in amounts x_1 and x_2 , with prices p_1 and p_2 , so $y=p_1x_1+p_2x_2$.

The second approach measures poverty in each of the dimensions separately and then aggregates the dimension-specific “deprivations” into a composite index. Formal treatments of this approach can be found in Tsui (2002), Bourguignon and Chakravarty (2003), Duclos et al. (2006) and Alkire and Foster (2007). The Alkire-Santos MIP is an example. I shall call this “deprivation aggregation.”

To see more clearly how this second approach works, consider again the two continuous attainments, x_1 and x_2 , with distribution functions F_1 and F_2 respectively. Poverty lines, denoted z_1 and z_2 , are defined in each space and the weights on deprivations are w_1 and w_2 ($w_1+w_2=1$). Then a simple example of a poverty measure using deprivation-aggregation is the weighted incidence of poverty across the two dimensions:¹¹

$$P^D \equiv w_1F_1(z_1) + w_2F_2(z_2) \tag{2}$$

This is only one possible way of aggregating deprivations. Alternatively one can focus on the joint distribution, and ask what proportion of the population is poor in at least one of the two dimensions (Bourguignon and Chakravarty, 2003). Letting F_{12} denote the joint distribution function, the poverty measure is then $F_1(z_1) + F_2(z_2) - F_{12}(z_1, z_2)$. Alternatively, one might ask what proportion is poor in both dimensions, i.e., $F_{12}(z_1, z_2)$. The measure proposed by Alkire and Foster (2007) goes further in introducing an extra parameter, such that a household is


¹⁰ The headcount index has a number of well-known limitations, and there is a large literature on alternative measures; Atkinson (1987) lists a number of these. But the headcount index is all we need for the present expository purpose.

¹¹ Higher-order measures can also be defined, such as based on the Watts (1968) index (see Chakravarty et al., 2008) and the Foster, Greer and Thorbecke (1984) measure (Alkire and Foster, 2008). Tsui (2002) provides a general formulation of subgroup-consistent poverty measures based on deprivation aggregation. However, I will confine attention to the simple form in (2).

deemed to be poor if its weighted deprivation exceeds a critical value.¹² However, all these measures are essentially some weighted aggregation of deprivations, and (implicitly) a nonlinear function of the cut-offs z_1 and z_2 . This discussion will focus on the analytically convenient form in (2), though this simplification does not appear to come at much loss.

It is evident that these two approaches will not, in general, give the same measure, even when the poverty lines are consistent in that $z = p_1 z_1 + p_2 z_2$. This is clear from Figure 1. Attainment aggregation identifies as poor all those people whose consumption of the two goods is within the triangle with vertices, z/p_1 , 0 and z/p_2 ; instead, the deprivation approach identifies some subset of those with $x_1 < z_1$ or $x_2 < z_2$ (the two unbounded rectangles of width z_1 and z_2 in Figure 1). Without knowing the weights and data one cannot say which will give the larger count of who is poor. If deprivation-aggregation measure focuses on those who are poor in both dimensions ($x_1 < z_1$ and $x_2 < z_2$) then the “deprivation poor” will never outnumber the “attainment poor” ($y < p_1 z_1 + p_2 z_2$). However, that need not hold for other deprivation measures, including equation (2). For example, if the deprivation poor are defined as those who are poor in either dimension ($x_1 < z_1$ or $x_2 < z_2$) then they will never be fewer in number than the attainment poor.

Figure 1: Attainment versus Deprivation Aggregation



¹² This extra parameter is also taken as given. Alkire and Foster discuss how it might be chosen and how one can test robustness to that choice.

The academic literature to date on multidimensional poverty has focused more on the deprivation-aggregation approach. This is not surprising since the attainment approach is a natural generalization of prevailing approaches to poverty measurement using price-weighted commodity aggregation, and so there is not much more to say about it in an academic paper; taking the attainment-aggregation approach, the remaining challenges are all in applications, notably in estimating missing prices.

However, there have also been theoretical arguments against aggregation using prices in the attainment space, and these arguments appear to have had influence. The following discussion takes a closer look at these arguments.

In the literature on the deprivation-aggregation approach, the weights on deprivations are taken to be known and explicit, while the weights on attainments are implicit. The papers in this literature provide some elegant mathematical formulations of their measures, but provide little guidance on where the weights attached to deprivations come from.¹³ In practice, the weights on deprivations are set by the analyst, with no obvious reason to suppose that they would be accepted by those one is trying to help by measuring poverty: policy makers and, of course, poor people.

The literature has also been close to silent about the tradeoffs between attainments that are built into a MIP. Interest in those tradeoffs does not rest on any view that the poverty measure should be seen as some policy maximand. Rather the interest stems from the need to understand the properties of the index. The tradeoff is given by the marginal rate of substitution (MRS), which is simply the marginal weight on one attainment relative to that on another.¹⁴ Neither the marginal weights on attainments nor the implied MRSs have received much attention in the literature on MIPs.

¹³ References are made to “participatory processes” in setting the weights (see, for example, Alkire and Foster, 2007), though this is not how the weights were set in any of the examples to date, including Alike and Santos (2010a).

¹⁴ While the term “MRS” is more commonly used in reference to the tradeoff built into an individual utility function, it can be defined for any suitably analytic function of two or more variables. Note also that, similarly to the usage in utility theory, defining the MRS does not imply that the function concerned is a maximand in an optimization problem. The MRS is simply an interesting property of the function, telling us its relative weights at each data point.

There is more than one way to define the MRS for a poverty measure—depending on how the changes in x_1 and x_2 are distributed. A natural option is to focus on the tradeoff facing someone at the poverty line. For the attainment-aggregation approach, it is plain that the poverty measure’s MRS—the increment to z_1 needed to compensate for less of z_2 keeping the poverty measure constant—is simply the relative price, p_1/p_2 .¹⁵ As long as the poverty bundle is consistent with the choices made by someone living at the poverty line, the poverty orderings based on this approach will be consistent with consumer welfare, in the sense that if someone living at the poverty line becomes worse (better) off then measured poverty rises (falls). Under these conditions, the poverty line is the point on the consumer’s expenditure function (inverse of the indirect utility function) corresponding to the poverty level of utility. Then any exogenous welfare-reducing (increasing) change—such as due to a change in relative prices, or any other shift parameter of the individual utility function—will be poverty increasing (decreasing) for all standard poverty measures. Thus welfare consistency is assured with appropriate calibration.

Compare this with the deprivation-aggregation approach. Again focus on the tradeoff between z_1 and z_2 . This is equivalent to the MRS when every value taken by one of the two dimensions is increased by a constant amount that is sufficient to compensate for a one unit decrease in the other dimension, where “compensation” is defined as holding the value of the aggregate poverty measure constant. The MRS of the deprivation-aggregation measure in (2) is:

$$MRS^D \equiv \frac{w_1 f_1(z_1)}{w_2 f_2(z_2)} \quad (3)$$

where the density functions for the two attainments are denoted f_1 and f_2 . There is no good reason to think this would necessarily, or ever, coincide with the relative price, p_1/p_2 . The MRS for a poverty measure based on the proportion of people for who are poor in at least one of the two dimensions is even less transparent, and is given by:

$$MRS^D \equiv \frac{f_1(z_1) - \int^{z_2} f_{12}(z_1, x_2) dx_2}{f_2(z_2) - \int^{z_1} f_{12}(x_1, z_2) dx_1} \quad (4)$$

Again, it would be a fluke if this equaled p_1/p_2 . The discontinuities in the Alkire-Foster index defy any similar (analytic) representation of its MRS. While the tradeoffs between attainments in

¹⁵ Note that the poverty measure is a function of z_1 and z_2 not x_1 and x_2 , so the MRS is defined in terms of derivatives with respect to z_1 and z_2 .

empirical implementations (such as the Alkire-Santos index) would be difficult to calculate, there is no reason to suppose that they would accord with relative prices, even for market goods.

The upshot of these observations is that aggregation across deprivations cannot in general yield poverty measures that are consistent with the welfare of someone living at the poverty line. This is because deprivation aggregation essentially ignores all implications for welfare measurement of consumer choice in a market economy. While those implications need not be decisive in welfare measurement, it is clearly worrying if the implicit tradeoff between any two market goods built into a poverty measure differs markedly from the tradeoff facing someone at the poverty line. When calibrated correctly, an attainment-aggregation measure guarantees that poor people would accept the tradeoffs built into the poverty measure. But there is no obvious calibration method for which this holds using deprivation aggregation. For example, everyone may agree that they are better off in situation A than B (such as before and after a change in prices or some other welfare shock), yet a measure using deprivation aggregation can show higher poverty in A, given that it does not reflect the tradeoffs that consumers have chosen. Similarly, the approach will identify some people as poor because they are lacking in one or more things that they can afford, but have no interest in acquiring, such as due to differences in demographics or relative prices. Poverty comparisons between people and over time (which are never easy) could be especially problematic.

These observations may not carry much weight with advocates of deprivation-aggregation, since they reject prices as weights (Tsui, 2002; Bourguignon and Chakravarty, 2003; Alkire and Foster, 2007). What are their arguments?

Three (related) critiques of using prices for aggregation can be identified in the literature. In the first, Alkire and Foster (2007, p.7) argue that the attainment approach entails a “loss of information on dimension-specific shortfalls.” It is true that attainment aggregation does not use the information on how far x_i falls short of z_i for each dimension i and each household. However, such dimension-specific poverty lines are not typically data, but must be assigned. One can equally well defend attainment aggregation on the grounds that it does not require this extra task.

Second, critics of using prices to aggregate in the attainment space argue that this practice avoids the problem of measuring “multidimensional poverty” by turning it into a more familiar

“one-dimensional” poverty measure; see, for example, Bourguignon and Chakravarty (2003) and Alkire and Foster (2007). However, as is plain from the above discussion, both approaches collapse the multiple dimensions into one; they just do it in different spaces. The real issue is how one does this aggregation, and whether one accepts theoretical restrictions implied by consistency with consumer welfare. So this critique brings us back to the point already discussed.

Third, advocates of deprivation-aggregation criticize attainment aggregation on the grounds that prices are missing or deemed unreliable (Tsui, 2002; Bourguignon and Chakravarty, 2003; Alkire and Foster, 2007). This is the more important issue, which deserves closer scrutiny. Of course, by either approach, weights must be assigned, and switching the space in which they are assigned cannot on its own address any of the concerns about using market prices as weights. One can agree that market prices do not accord with shadow prices in general without preferring to aggregate in the deprivation space, which rejects the use of all prices; this holds across dimensions of poverty that relate directly to market goods as well as non-market goods. It is one thing to recognize that not all goods are market goods, or that there are market distortions, and quite another to ignore market prices when they are data.

The case for deprivation aggregation cannot rest solely on the deficiencies of market prices as a means of valuation. Even if p_1 and p_2 were the true shadow prices, there can be no presumption that $P^D = P^A$; this would be a fluke. Furthermore, suppose that for some z_1 and z_2 , the weights in the deprivation-aggregation measure were chosen to deliver an MRS between attainments of p_1/p_2 . Then for any (non-uniform) distributions, the two measures would deviate for any changes in z_1 or z_2 , as this would change the densities determining the MRS in the deprivation-aggregation measure.

It has been argued that setting even initially arbitrary weights on deprivations should be viewed as the start of a public debate on what weights are appropriate (Alkire and Foster, 2007). Public opinion might be considered an important clue to setting weights on deprivations, or shadow prices on attainments. Setting the weights in some initially *ad hoc* way might then be thought of as the first step in an “iterative public debate” about what the weights should be.

Stimulating such a debate could well be a valuable contribution. However, there is little sign as yet that this has led to new weights in past mashup indices. Consider, for example, the

oldest of the mashup indices still in use, the HDI. Its weights were set 20 years ago, with equal weight to the (scaled) sub-indices for health, education and GDP.¹⁶ Equality of the weights was, of course, an arbitrary judgment, and it might have been hoped that the weights would evolve in the light of the subsequent public debate. But that did not happen. The weights on the three components of the HDI (health, education and income) have not changed in 20 years, and it is hard to believe that the HDI got it right first go.¹⁷

Setting initial weights and revising them in the light of subsequent debate would also point to the need to know the tradeoffs in the most relevant space for understanding what the weights really mean. Arguably, the fact that MIPs have assigned weights in the deprivation space rather than the attainment space does not make it easy for the debate to proceed on a well-informed basis. I would conjecture that most people will find it easier to attach a monetary value (or market-good equivalent) to a non-market good than to assess what tradeoff is acceptable between the corresponding two poverty measures. If the welfare dimension for which prices are missing is (say) health status, it is surely easier for people to judge how much money they would pay for better health than what tradeoff they would accept between poverty in income space and poverty in health space. Indeed, given the opaqueness about the tradeoffs in the primary attainment dimensions built into MIPs such as the Alkire-Santos index, it can be argued that users (including policy makers) may end up tacitly accepting, and acting upon, tradeoffs that they would find objectionable when revealed. This can hardly be helpful in advancing open public debate about the weights. The weights need to be transparent. Here again it is far from clear that the deprivation space is superior to the attainment space for this purpose.

A seemingly natural hybrid can be proposed that blends attainment and deprivation aggregation. Suppose that there is a third attainment dimension, x_3 , for a non-market good for which the analyst is unwilling to assign a shadow price. We might then measure poverty by:

$$P^H \equiv (w_1 + w_2)F_y(p_1z_1 + p_1z_1) + w_3F_3(z_3) \quad (5)$$

¹⁶ The weights on the HDI's primary dimensions have varied over time due to (often seemingly arbitrary) changes in the bounds used for scaling the indices. However, the weights on the HDI's core dimensions have never been explicitly identified or discussed in the HDRs. See Ravallion (2010b).

¹⁷ In switching to a geometric mean in the 2010 HDR, the weights on the three achievement variables changed, though their logs are still equally weighted; see Ravallion (2010b) for further discussion.

(where $w_1 + w_2 + w_3 = 1$). However, this hybrid might not entail quite as much progress as one thinks at first sight. While there is no explicit “price” for dimension 3, it still has an implicit price and MRS (given by $w_3 f_3(z_3) / [(w_1 + w_2) f_y(p_1 z_1 + p_1 z_1)]$). The hybrid approach does not avoid the need to set “prices”—it only makes them less transparent.

If one is unwilling to set a shadow price on the non-market good, then it can be argued that neither of these approaches—attainment aggregation versus deprivation aggregation—should be used. The fact that poverty is “multidimensional” does not in itself imply that we need either approach. For example, we might instead measure consumption poverty with the best data available on command over market goods, while also using the best data on other dimensions of poverty that are left out, as appropriate to the country context. This is common practice. Unambiguous country rankings may be hard to justify, although still possible under certain conditions (Atkinson and Bourguignon, 1982; Duclos et al., 2006). But, thankfully, the inability to determine the “price” need not delay many of the policy-related applications of poverty measurement, as discussed in section 2.

4. Conclusions

Advocates of the new class of “multidimensional indices of poverty” are right to point out that command over market goods is not all that matters, and that other dimensions of welfare need to be considered in our efforts to fight poverty. Far less convincing is their claim that a single composite index can be defended with rigor or even help much in those efforts. Recognizing that poverty is not just about lack of household command over market goods does not imply that one needs to collapse the multiple dimensions into one (uni-dimensional) index.

It is not credible to contend that any single index could capture all that matters in all settings. No consensus exists on what dimensions to include and how they should be weighted to form the composite index. We can all agree that reducing child mortality is a hugely important development goal, but how can one contend—as the Alkire-Santos index does, for example—that avoiding the death of a child is equivalent to alleviating the combined deprivations of having a dirt floor, cooking with wood, and not having a radio, TV, telephone, bike or car? Or that attaining these material conditions is equivalent to an extra year of schooling or to not have any malnourished family member? These are difficult judgments to make. Thankfully resolving all

such tradeoffs is not commonly needed in policy making at the country level. But when one faces a tradeoff, because a policy spans more than one dimension, those with a stake in the outcomes will almost certainly be in a better position to do determine what weights to apply than the analyst calibrating a measure of poverty. The specific country and policy context should determine what tradeoff is considered appropriate. This will often be a political decision, though hopefully a well informed one.

So we may well be asking too much of a single measure of “poverty” to have it include things like child mortality, schooling, violence or empowerment, on top of material living standards. It is one thing to agree that consumption of market commodities is an incomplete metric of welfare—and that for the purpose of assessing poverty one needs to also account for indicators of non-market goods and services—and quite another to say that a single “poverty” measure should embrace all these things.

This suggests that we may need to focus our efforts and resources on developing the best possible distinct measures of the various dimensions of poverty deemed relevant to a given setting—aiming for a credible set of “multiple indices” rather than a single “multidimensional index.” Granted, that route forgoes the option of identifying a joint distribution, but that is hardly a large loss in practice when that one’s characterization of the joint distribution is so severely data-constrained (notably in finding a single household survey that covers all the dimensions that matter). And, of course, taking the multiple-index route in the context of poverty monitoring does not preclude researching the joint distributions in specific data sets, when feasible.

Some degree of aggregation will almost certainly be needed even within the multiple-index approach. This raises an important but poorly understood issue concerning the space we choose to aggregate multiple dimensions. We must decide whether this is to be done in the space of attainments—the variables we use in assessing individual welfare for the purpose of deciding who is poor and how much—or in the space of deprivations—the assessments we make about whether one is poor in a given attainment space. While arguments can be made for and against each approach, I have questioned the main arguments made in favor of deprivation aggregation, as has become popular in the recent “multidimensional indices of poverty.”

The role played by prices lies at the heart of the matter. It is widely agreed that prices can be missing for some goods and deceptive for others. There are continuing challenges facing applied economists in addressing these problems. However, it is one thing to recognize that markets and prices are missing or imperfect, and quite another to ignore them in welfare and poverty measurement. There is a peculiar inconsistency in the literature on multidimensional indices of poverty whereby prices are regarded as an unreliable guide to the tradeoffs, and are largely ignored, while the actual weights being assumed in lieu of prices are not made explicit in the same space as prices. We have no basis for believing that the weights being used are any better than market prices, when available. Nor do we have any basis for believing that the weights bear any resemblance to defensible shadow prices. Aggregating under such conditions risks stifling, rather than promoting, open debate about what tradeoffs are in fact acceptable, when such tradeoffs need to be set.

References

- Alkire, Sabina and James Foster, 2007, "Counting and Multidimensional Poverty Measurement," Oxford Poverty and Human Development Initiative, Working Paper 7, University of Oxford.
- Alkire, Sabina and Maria Emma Santos, 2010a, "Acute Multidimensional Poverty: A New Index for Developing Countries," Oxford Poverty and Human Development Initiative, Working Paper 38, University of Oxford.
- _____ and _____, 2010b, "Multidimensional Poverty Index," Research Brief, Oxford Poverty and Human Development Initiative, University of Oxford.
- Alkire, Sabina, Maria Emma Santos, Suman Seth and Gaston Yalonetzky, 2010, "Is the Multidimensional Poverty Index Robust to Different Weights?" mimeo, Oxford Poverty and Human Development Initiative, University of Oxford.
- Atkinson, Anthony B., 1987, "On the Measurement of Poverty," *Econometrica* 55: 749-64.
- Atkinson, Anthony B., and Francois Bourguignon, 1982, "The Comparison of Multi Dimensional Distributions of Economic Status," *Review of Economic Studies* 49: 183-201.
- Bourguignon, Francois and Satya Chakravarty, 2003, "The Measurement of Multidimensional Poverty," *Journal of Economic Inequality* 1: 25-49.
- Chakravarty, Satya, J. Deutsch and Jacques Silber, 2008, "On the Watts Multidimensional Poverty Index and its Decomposition," *World Development* 36(6): 1067-78.
- Duclos, Jean-Yves, David Sahn and Stephen Younger, 2006, "Robust Multidimensional Poverty Comparisons," *Economic Journal* 116: 943-68.
- Foster, James, J. Greer, and Erik Thorbecke, 1984, "A Class of Decomposable Poverty Measures," *Econometrica* 52: 761-765.
- Haddad, Lawrence and Ravi Kanbur, 1990, "How Serious is the Neglect of Intra-Household Inequality?," *Economic Journal* 100: 866-881.
- Lokshin, Michael and Martin Ravallion, 2005, "Rich and Powerful? Subjective Power and Welfare in Russia," *Journal of Economic Behavior and Organization* 56: 141-172.
- Ravallion, Martin, 1994, *Poverty Comparisons*, Chur, Switzerland: Harwood Academic Press.
- _____, 2008, "Poverty Lines," in *The New Palgrave Dictionary of Economics*, 2nd Edition, Larry Blume and Steven Durlauf (eds) London: Palgrave Macmillan.

- _____, 2010a, “[Mashup Indices of Development](#),” Policy Research Working Paper 5432, World Bank, Washington DC.
- _____, 2010b, “[Troubling Tradeoffs in the Human Development Index](#).” Policy Research Working Paper 5484, World Bank, Washington DC.
- Stiglitz, Joseph, Amartya Sen and J.P. Fitoussi JP, 2009, *Report by the Commission on the Measurement of Economic Performance and Social Progress*, www.stiglitz-sen-fitoussi.fr.
- Tsui, Kai-Yuen, 2002, “Multidimensional Poverty Indices,” *Social Choice and Welfare* 19: 69-93.
- United Nations Development Programme (2010) *Human Development Report*, New York: Palgrave Macmillan for the UNDP.
- Watts, Harold W., 1968, “An Economic Definition of Poverty,” in Daniel P. Moynihan (ed.), *On Understanding Poverty*, New York, Basic Books.