



OPHI WORKING PAPER NO. 43

Understandings and Misunderstandings of Multidimensional Poverty Measurement

Sabina Alkire* and James Foster^o

May 2011

Abstract

Multidimensional measures provide an alternative lens through which poverty may be viewed and understood. In recent work we have attempted to offer a practical approach to identifying the poor and measuring aggregate poverty (Alkire and Foster 2011). As this is quite a departure from traditional unidimensional and multidimensional poverty measurement – particularly with respect to the identification step – further elaboration may be warranted. In this paper we elucidate the strengths, limitations, and misunderstandings of multidimensional poverty measurement in order to clarify the debate and catalyse further research. We begin with general definitions of unidimensional and multidimensional methodologies for measuring poverty. We provide an intuitive description of our measurement approach, including a ‘dual cutoff’ identification step that views poverty as the state of being multiply deprived, and an aggregation step based on the traditional FGT measures. We briefly discuss five characteristics of our methodology that are easily overlooked or mistaken and conclude with some brief remarks on the way forward.

Keywords: poverty measurement, multidimensional poverty, deprivation, FGT measures, decomposability, joint distribution, axioms.

JEL classification: I3, I32, D63, O1

* The Oxford Poverty and Human Development Initiative (OPHI), Oxford Department of International Development, Queen Elizabeth House, University of Oxford.

^o The Elliott School of International Affairs, George Washington University, and the Oxford Poverty and Human Development Initiative (OPHI), Oxford Department of International Development, Queen Elizabeth House, University of Oxford.

The authors gratefully acknowledge support for this research from the United Kingdom Department for International Development (DFID), the Robertson Foundation, and the Institute of International Economic Policy (IIEP), George Washington University.

ISSN 2040-8188 ISBN 978-1-907194-27-6

Acknowledgements

The authors would like to thank Angus Deaton, Jean-Yves Duclos, Chrysanthi Hatzimasoura, Nora Lustig, Nicole Rippin, Maria Emma Santos, Shabana Singh, Gaston Yalonetzky, and the participants of the IIEP Development Lunch for helpful comments, and especially Martin Ravallion for helping to motivate the present piece.

The Oxford Poverty and Human Development Initiative (OPHI) is a research centre within the Oxford Department of International Development, Queen Elizabeth House, at the University of Oxford. Led by Sabina Alkire, OPHI aspires to build and advance a more systematic methodological and economic framework for reducing multidimensional poverty, grounded in people's experiences and values.

This publication is copyright, however it may be reproduced without fee for teaching or non-profit purposes, but not for resale. Formal permission is required for all such uses, and will normally be granted immediately. For copying in any other circumstances, or for re-use in other publications, or for translation or adaptation, prior written permission must be obtained from OPHI and may be subject to a fee.

Oxford Poverty & Human Development Initiative (OPHI)
Oxford Department of International Development
Queen Elizabeth House (QEH), University of Oxford
3 Mansfield Road, Oxford OX1 3TB, UK
Tel. +44 (0)1865 271915 Fax +44 (0)1865 281801
ophi@qeh.ox.ac.uk <http://ophi.qeh.ox.ac.uk/>

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by OPHI or the University of Oxford, nor by the sponsors, of any of the views expressed.

[T]he job of a 'measure' or an 'index' is to distill what is particularly relevant for our purpose, and then to focus specifically on that. ... The central issues in devising an index relate to systematic assessment of importance. Measurement has to be integrated with evaluation. This is not an easy task.

–Amartya Sen (1989)

1. Introduction

How we measure poverty can importantly influence how we come to understand it, how we analyze it, and how we create policies to influence it. For this reason, measurement methodologies can be of tremendous practical relevance.

In 2003, two signal articles were published. One, by François Bourguignon and Satya Chakravarty (2003), proposed a class of multidimensional poverty measures that extended the Foster Greer and Thorbecke (FGT) class of indices and discussed interrelationships among dimensions.¹ Another, by A. B. Atkinson (2003), linked the emerging axiomatic literature on multidimensional poverty measures to the 'counting' literature that had been implemented in Europe and urged that counting measures be connected more with welfare economics. In recent years, the literature on multidimensional poverty measurement has blossomed in a number of different directions. The 1997 *Human Development Report* and the 2000/1 *World Development Report* vividly introduced poverty as a multidimensional phenomenon, and the Millennium Declaration and MDGs have highlighted multiple dimensions of poverty since 2000. In the academic literature, new measurement methodologies are being created.² Interest in multidimensional poverty measurement is growing.

In a time of considerable creative ferment, it can be helpful to clarify the potentials of different measurement methodologies, and to distinguish misunderstandings from limitations. This paper does so for a multidimensional methodology we first developed in 2007 and which has been implemented at the national and international levels. A practical aim of Alkire and Foster (2007, 2011) was to construct poverty measurement methods that could be used with discrete and qualitative data (for example, functionings like literacy or physical security) as well as continuous and cardinal data (as income and consumption are commonly viewed). A theoretical aim was to re-examine the identification step (addressing the question 'who is poor?'), which poses a much greater challenge when there are multiple dimensions. The result is a methodology for measuring poverty in the sense of Sen (1976) that first identifies who is poor, then aggregates to obtain overall measures of poverty that reflect the multiple deprivations experienced by the poor.

Our methodology is perhaps best seen as a general framework for measuring multidimensional poverty since many key decisions are left to the user. These include the selection of dimensions, dimensional cutoffs (to determine when a person is deprived in a dimension), dimensional weights (to indicate the relative importance of the different deprivations), and a poverty cutoff (to determine when a person has enough deprivations to be considered to be poor).³ We note that this flexibility makes it particularly useful for measurement efforts at the country level where these decisions can fit the purpose of the

¹ Foster, Greer and Thorbecke (1984).

² The literature is growing rapidly; some key papers in the axiomatic tradition include Anand and Sen (1997), Brandolini and D'Alessio (1998), Chakravarty et al (1998), Tsui (2002), Atkinson (2003), Bourguignon and Chakravarty (2003), Deutsch and Silber (2005), Duclos, Sahn and Younger (2006), Chakravarty and D'Ambrosio (2006), Kakwani and Silber (2008a, b) and Thorbecke (2008). For recent critiques of multidimensional measurement, see Ravallion (2010, 2011).

³ The preliminary step of selecting dimensional variables is closely linked to the selection of dimensional weights. For example, choosing one out of several possible variables is tantamount to assigning that dimension full weight and the remaining dimensions zero weight.

measure and can embody normative judgements regarding what it means to be poor. The method delivers an aggregate poverty measure that reflects the prevalence of poverty and the joint distribution of deprivations. Useful partial indices are reported that reveal the intuition and layers of information embedded in the summary measure. For example, our most basic index is the product of two partial indices that measure prevalence and intensity. The overall measure can be additively decomposed by population subgroups and (after the poor have been identified) by dimension. The methodology satisfies a set of basic axioms for multidimensional poverty measurement.

The reason to delineate clearly the value-added of any new set of tools – such as this methodology – is twofold. First, if the value-added is well understood, and if misunderstandings are clarified, then researchers may be able to sharpen the contributions of the new technique by developing complementary methodologies, by improving applications, or by tackling standing questions. Second, a clear articulation of potential applications can enable others to mine the new tools for their full range of empirical insights.

The paper proceeds as follows. Section 2 presents unidimensional approach to poverty measurement while Section 3 describes multidimensional methods. Section 4 provides an intuitive description of our measurement approach and the insights that it generates. Section 5 discusses five characteristics of our methodology that are easily overlooked or mistaken, while Section 6 concludes with some brief remarks on the way forward.

2. Unidimensional Methods

In his seminal 1976 paper ‘Poverty: An Ordinal Approach to Measurement’, Amartya Sen described two main steps that poverty measurement must address:

- (i) identifying the poor among the total population, and
- (ii) constructing a numerical measure of poverty.

Sen’s two-step procedure of identification and aggregation has become the standard conceptual framework for poverty measurement, and we follow his approach in our discussion of unidimensional and multidimensional methods.

Unidimensional methods can be applied when a well-defined single-dimensional resource variable, such as income, has been selected as the basis for poverty evaluation. This variable is typically assumed to be cardinal; however, in some cases the variable may only have ordinal significance (i.e., the direction of change is discernable, but not its magnitude). Identification in the unidimensional environment typically proceeds by setting a poverty line corresponding to a minimum level below which one is considered poor. Aggregation is usually achieved through the use of a numerical poverty measure that determines the overall level of poverty in a distribution given the poverty line.

To facilitate our discussion, we begin with three examples of unidimensional poverty measures drawn from the FGT class. Consider the following vectors constructed from a distribution y of a given resource (say income), and a poverty line z . The *deprivation vector* g^0 replaces each income below the poverty line with one and replaces nonpoor incomes with zero. The *normalized gap vector* g^1 replaces each poor income y_i with the normalized gap $(z-y_i)/z$ and assigns zero to the rest. The *squared gap vector* g^2 replaces each poor income y_i with the squared normalized gap $[(z-y_i)/z]^2$ and assigns zero to the rest. Each of these vectors is censored in that every nonpoor person is assigned a value of zero.

The headcount ratio is simply $P_0 = \mu(g^0)$, or the mean of the deprivation vector; it indicates the prevalence of poverty. The poverty gap measure is $P_1 = \mu(g^1)$; it measures the average depth of poverty across the society as a whole. The squared gap, or distribution sensitive FGT, measure is $P_2 = \mu(g^2)$; it emphasizes the conditions of the poorest of the poor. All three can be applied to cardinal variables; only the headcount ratio can also be used with an ordinal variable. The measures satisfy an array of axioms, including a subgroup decomposability property that views overall poverty as a population share weighted average of subgroup poverty levels. The headcount ratio violates an intuitive monotonicity axiom (if a poor person's resource level falls, poverty should rise); the poverty gap measure satisfies monotonicity, but not the transfer principle (poverty should fall if two poor resource levels are brought closer together by a progressive transfer between them); the final FGT index satisfies monotonicity and the transfer principle.

Unidimensional methods require a single dimensional variable and a single cutoff, but place no a priori restrictions on how the resource variable has been constructed. It could be a single resource variable, such as income, added up across all sources. It could be total expenditure added up across different categories reported in an expenditure survey, or perhaps drawn from consumption surveys that require respondents to recall quantities and prices. Of course, the interpretation of the variable and its cutoff level is very different if total income or total expenditure is used, with the former reflecting 'what could be' and the latter reflecting 'what is' (Atkinson, 1989). However, the underlying principle of aggregation is the same: adding up monetary values to obtain a total resource level that can be compared to a monetary cutoff.⁴

Unidimensional tools might also be applied to other aggregate variables, such as those obtained by combining fundamentally distinct components that are not measured in the same units and may have no natural or observable means of conversion into a common variable. One method is to use a composite indicator, which aggregates across several component variables by multiplying each by a factor and adding up; another example might be a utility function that aggregates components in a nonlinear fashion. However, there are several challenges with such a strategy. The entire enterprise depends on the validity of the aggregate variable in representing the actual resources or achievements of people and the actual tradeoffs among component variables, and yet evidence in this regard may be sparse, lending little guidance to the selection of weights or of a proper functional form. One or more components may be only ordinally meaningful, which rules out the usual forms of aggregation across component variables (since every cardinal representation is admissible for an ordinal variable). Or the components may well represent distinct resources and categories of needs, which cannot be merged or freely traded for one another. Applying unidimensional methods to an aggregate composite indicator suggests that the presence or the extent of shortfalls in component variables are of no particular concern, and thus do not independently affect whether a person is poor, or the overall level of poverty. For these reasons it may make sense to explore alternatives that can complement unidimensional methods.

3. Multidimensional Methods

Suppose, then, that we have data on achievements in several dimensions, distributed across a population. Following Sen (1976) we ask: Who is poor and how should overall poverty be measured in this setting?

⁴ These methods may also be used with non-aggregated cardinal variables (such as years of schooling), ordinal variables (such as self-reported health), and categorical variables (such as modes of access to drinking water), although the latter two categories will restrict the poverty measures that can be used.

As noted above, if the underlying concept of poverty admits a natural way of aggregating the various dimensions into an overall variable, then a unidimensional methodology can be used. In this approach, the poor are identified on the basis of a single cutoff, and overall poverty is evaluated using a unidimensional measure such as a member of the FGT class. The many dimensions are merged into one and viewed through a unidimensional lens.

What if an aggregate variable cannot be plausibly constructed, and instead there are several important distinct dimensions? How should we identify the poor and measure poverty in this case? Bourguignon and Chakravarty (2003) propose the use of dimension-specific lines – which are called deprivation cutoffs in Alkire and Foster (2007) – as the basis for determining who is deprived and in which dimension. They then posit the existence of an identification function, which determines whether a person is deprived enough to be called poor, and a poverty measure, which evaluates how much poverty there is overall. Axioms analogous to the ones used in the unidimensional case ensure that the measure properly reflects poverty and that it can be decomposed by subgroup. The axioms also ensure that the poverty measure is consistent with the identification function.

Much of the research in this area has been concerned with finding an appropriate poverty measure, rather than devising new methods of identifying the poor.⁵ Two benchmark identification approaches are discussed by Atkinson (2003): the union and intersection approaches. Under union identification, a person who is deprived in any dimension is considered poor; under intersection identification, only persons who are deprived in all dimensions are considered poor. Both approaches are easy to understand and have useful characteristics, such as being able to be applied to ordinal variables. However, they can be particularly ineffective at separating the poor from the nonpoor. In a recent study by Alkire and Seth (2009) that uses ten dimensions to identify the poor in India, the union approach identifies 97 per cent of the population as poor, whereas the intersection approach identifies one-tenth of 1 per cent. Such a range of values is common in many studies.

Bourguignon and Chakravarty's (2003) discussion of identification concerns general forms of identification functions rather than specific examples, and it is clear from the context that tradeoffs are being made between continuous dimensional variables. But this leads the discussion back to the original question of whether a coherent aggregate variable can be constructed from the individual dimensions. If the answer is no, as postulated above, then it may be somewhat difficult to justify the aggregation needed for a general identification function. If yes, then there may be good reason to explore a unidimensional method.

One important omission in this literature is a proper discussion of the axiomatic structure for identification functions (or, more generally, for overall methodologies) that could help guide the construction of new identification techniques. Indeed, too little attention has been paid to developing practical alternatives to the union, intersection, and unidimensional identification approaches. This is a key motivation behind Alkire and Foster (2007, 2011).

Before moving on to a description of our methodology, we would like to mention an alternative approach that might be used in this context. Recall that the multidimensional methodologies discussed above draw upon a matrix in which each row contains the vector of achievement levels associated with a given person and each column contains the vector of people's achievement levels in a given dimension. This general structure allows the measurement method to take advantage of the information contained

⁵ See for example Tsui (2002), Chakravarty and D'Ambrosio (2006), Chakravarty and Silber (2008), Bossert, Chakravarty and D'Ambrosio (2009), and Rippin (2010).

in the joint distribution⁶ of achievements. In contrast, we can define a restricted type of methodology that applies deprivation cutoffs within *each* dimension and ignores all information on links *across* dimensions. A *marginal method* is one that assigns the same level of poverty to any two matrices that generate the same marginal distributions. For example, consider a method that employs a union identification approach and aggregates using some increasing function of the unidimensional headcount ratios (one for each dimension). The resulting poverty level would surely reflect the prevalence of deprivation in society. However, it would not be able to say whether the deprivations are spread evenly across the population or whether they are concentrated in an underclass of multiply deprived persons. We elaborate on this distinction in section 5.1.

One useful feature of a marginal method is that it can be estimated using dimensional data from different sources, where the underlying columns of dimensional achievements are not linked and may even have different population sizes. Indeed, this approach to measurement can be extended to include cases where the columns are drawn from *different* reference groups within a population, such as children and adults. The Human Poverty Index developed by Anand and Sen (1997) is an example of this form of extended marginal measure, since it aggregates dimensional deprivation indicators from different populations (namely, adult illiteracy, the probability at birth of not surviving until the age of 40, the percentage of households lacking drinking water, and the percentage of malnourished children according to weight for age). However, while multiply sourced marginal measures can provide useful estimates of poverty or aggregate deprivation, they are based on anonymous or unlinked data and are unable to *identify* who is multidimensionally poor – a signal benefit of a measurement methodology based on data linked by person or household.

4. The AF Method

This section provides a systematic overview of the multidimensional methodology of Alkire and Foster (2007, 2011). We describe how poor people are identified using a ‘dual cutoff’ method. We construct the poverty measures and show how to ‘drill down’ into each measure to unfold distinctive partial indices that can illuminate policy questions. Decompositions are exhibited that explain and clarify the aggregate poverty level. The progression closely mirrors the techniques employed in unidimensional poverty measurement, but with some important differences, most notably in the identification step. In what follows we will assume that the range of dimensional variables has been selected and data are available in the form of an ($n \times d$) data matrix Y for n persons and $d \geq 2$ dimensions.

4.1 Identification

In unidimensional analysis, identification is normally accomplished by the use of a poverty line or threshold, with poor people being identified as those whose resource or achievement variable level falls below the poverty line. In the multidimensional measurement setting, where there are multiple variables, identification is a substantially more challenging exercise. This is the part of our methodology that is most commonly overlooked or misunderstood and so we begin the basic elements of our dual cutoff identification approach.

4.1.1 Deprivation Cutoffs

A vector $z = (z_1, \dots, z_d)$ of deprivation cutoffs (one for each dimension) is used to determine whether a person is deprived. If the person’s achievement level in a given dimension j falls short of the respective

⁶ Let w , x , and y be three random variables. The joint distribution gives the percentage of the population with (w,x,y) or less. The marginal distribution of w gives the percentage of the population with w or less, and similarly for x and y .

deprivation cutoff z_j , the person is said to be deprived in that dimension; if the person's level is at least as great as the deprivation cutoff, the person is not deprived in that dimension.

4.1.2 Weights

A vector $w = (w_1, \dots, w_d)$ of *weights* or *deprivation values* is used to indicate the relative importance of the different deprivations. If each deprivation is viewed as having equal importance, then this leads to a benchmark case where all weights are one and sum to the number of dimensions d . If deprivations are viewed as having differential importance, this is reflected by a vector whose entries sum to d but can vary from one, with higher weights indicating greater importance. Note that the deprivation values affect identification as they determine the minimal combinations of deprivations that will identify a person as being poor; they also affect aggregation by altering the relative contributions of deprivations to overall poverty.

4.1.3 Deprivation Counts

A column vector $c = (c_1 \dots c_n)'$ of *deprivation counts* reflects the breadth of each person's deprivation. The i^{th} person's deprivation count c_i is the number of deprivations experienced by i (in the case of equal weights) or the sum of the values of the deprivations experienced by i (in the general case).

4.1.4 Poverty Cutoff

A *poverty cutoff* k satisfying $0 < k \leq d$ is used to determine whether a person has sufficient deprivations to be considered poor. If the i^{th} person's deprivation count c_i falls below k , the person is not considered to be poor; if the person's deprivation count is k or above, the person is identified as being poor. The title 'dual cutoff' refers to the sequential use of deprivation and poverty cutoffs to identify the poor. Note that when k is less than or equal to the minimum weight across all dimensions we have union identification. When $k=d$, the intersection approach is being used. The deprivation count and poverty cutoff can also be expressed as percentages of d .

4.1.5 Identification Function

The identification function summarizes the outcome of the above process and indicates whether a person is poor in Y given deprivation cutoffs z , weights w , and poverty cutoff k . If the person is poor, the identification function takes on a value of one; if the person is not poor, the identification function has a value of zero.

One of the interesting properties exhibited by our identification approach is that it is applicable even when one or more of the variables are ordinal. All cardinalizations of the ordinal variable (found by applying a monotonic transformation to the variable and its cutoff) yield identical conclusions regarding whether a person is deprived in the dimension and whether the person is poor. This expands the potential reach of the methodology by allowing it to be meaningfully applied to data with lower level measurement properties.

4.2 Censored Matrices

The transition between the identification step and the aggregation step is most easily understood by examining a progression of matrices. The *achievement matrix* Y shows the achievements of n persons in each of d dimensions. The *deprivation matrix* g^0 replaces each entry in Y that is below its respective deprivation cutoff z_j with the deprivation value w_j , and each entry that is not below its deprivation cutoff with zero. It provides a snapshot of who is deprived in which dimension and how much weight the deprivation carries. The *censored deprivation matrix* $g^0(k)$ multiplies each row in the deprivation matrix by the identification function: if the person is poor, then the row containing the deprivational information

of the person is unchanged; but if the person is not poor the information of the person is censored and replaced with zeroes. An example with equal deprivation values is given below for $k = 2$.

$$\begin{array}{c}
 \text{Achievement Matrix} \\
 \text{Dimensions} \\
 Y = \begin{array}{c} \left[\begin{array}{cccc} 13.1 & 14 & 4 & 1 \\ 15.2 & \underline{7} & 5 & \underline{0} \\ \underline{12.5} & \underline{10} & \underline{1} & \underline{0} \\ 20 & \underline{11} & 3 & 1 \end{array} \right] \\ \text{Persons} \end{array} \\
 z = (13 \quad 12 \quad 3 \quad 1) \quad \text{Deprivation Cutoffs}
 \end{array}$$

Deprivation Matrix

$$g^0 = \begin{array}{c} \left[\begin{array}{cccc} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 \end{array} \right] \quad \left[\begin{array}{c} 0 \\ 2 \\ 4 \\ 1 \end{array} \right]
 \end{array}$$

Censored Deprivation Matrix, $k=2$

$$g^0(k) = \begin{array}{c} \left[\begin{array}{cccc} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{array} \right] \quad \left[\begin{array}{c} 0 \\ 2 \\ 4 \\ 0 \end{array} \right]
 \end{array}$$

The original achievement matrix has three persons who are deprived (see the underlined entries) in one or more dimensions. The first person (given in the top row) has no deprivations at all. The next two persons have deprivation counts that are greater than or equal to $k = 2$. They are considered to be poor and hence their entries in the censored deprivation matrix are the same as in the deprivation matrix. The fourth person has a single deprivation and hence is not poor. This single deprivation is disregarded in the censored deprivation matrix, which only displays the deprivations of the poor.⁷

If the entries of Y are all cardinally significant, then the *normalized gap* of a person in a deprived dimension can be defined as the difference between the deprivation cutoff and the person’s achievement, all divided by the deprivation cutoff. The *normalized gap matrix* g^1 replaces each deprived entry in Y with the respective normalized gap multiplied by the deprivation value; it replaces each entry that is not below its deprivation cutoff with zero. The normalized gap matrix provides a snapshot of the depth of deprivation of each person in each dimension, weighted by its relative importance. The *squared gap matrix* g^2 replaces each deprived entry in Y with the square of the normalized gap multiplied by the deprivation value; it replaces each entry that is not below its deprivation cutoff with zero. Squaring the normalized gaps puts relatively more emphasis on larger deprivations. The censored normalized gap matrix $g^1(k)$, the censored squared gap matrix $g^2(k)$, and the censored deprivation counts $c(k)$ are likewise

⁷ This example has identical deprivation values across dimensions; the general case admits a wide variety of identification approaches. For example, if one dimension had overriding importance, and its deprivation value were set above or equal to k , then any person deprived in that dimension would be considered to be poor.

obtained by multiplying through by the identification function, and hence they only display the information of the poor.

Censored Normalized Gap Matrix, $k = 2$

Censored Squared Gap Matrix, $k = 2$

$$g^1(k) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0.42 & 0 & 1 \\ 0.04 & 0.17 & 0.67 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad g^2(k) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0.42^2 & 0 & 1^2 \\ 0.04^2 & 0.17^2 & 0.67^2 & 1^2 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

The censoring step is key to our methodology, since the censored matrices embody our identification step and are the basic constructs used in the aggregation step. The original deprivation matrices, by comparison, include information on the deprivations of the nonpoor, which should not affect any measure that is focused on the poor. Note that there is one case where this distinction is not relevant: when the poverty cutoff k becomes small enough (no larger than the smallest deprivation value) and our identification method becomes the union approach. Any person who is deprived in any dimension is considered poor and the censored and original matrices are identical.

4.3 Aggregation

The aggregation step of our methodology builds upon the standard FGT technology and likewise generates a parametric class of measures. Each FGT measure can be viewed as the mean of an appropriate vector built from the original data and censored using the poverty line, and the AF measures have an analogous structure. We focus on three main measures corresponding to the key FGT measures.

4.3.1 Adjusted Headcount Ratio

The *adjusted headcount ratio* is defined as $M_0 = \mu(g^0(k))$, or the mean of the censored deprivation matrix. In our example, the sum of the entries of $g^0(k)$ is 6 while the number of entries in the matrix is 16, resulting in an adjusted headcount ratio of $3/8$. Since a completely poor and deprived society has 16 deprivations, $M_0 = 3/8$ can be interpreted as the actual number of deprivations (6) among the poor as a share of the maximum (16).

$$g^0(k) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \begin{array}{l} 0 \\ \underline{2} \quad 2/4 \\ \underline{4} \quad 4/4 \\ 0 \end{array}$$

A second way of viewing M_0 is in terms of partial indices – measures that provide basic information on a single aspect of poverty.⁸ The first partial index is the percentage of the population that is poor, or the

⁸ See for example Foster and Sen (1997) or Foster (2005).

multidimensional headcount ratio H . The second is the *average intensity* A , which calculates the *deprivation share* for each poor person by dividing the deprivation count by d , and then averages across all poor persons. It is easy to see that $M_0 = HA$, which holds in the example since $H = 1/2$ and $A = 3/4$. This breakdown is not unlike the expression $P_1 = HI$ for the unidimensional poverty gap measure, where H is the unidimensional headcount and I is the so-called income gap ratio, which measures the average depth of poverty among the poor. M_0 can be used with ordinal data: any monotonic transformation of a dimensional variable and its cutoff leads to the same censored deprivation matrix and hence the same level of M_0 .

4.3.2 Adjusted Poverty Gap and Adjusted FGT

If all the variables are cardinal significant, then information on the depth of deprivations can be used to construct two additional poverty measures. The *adjusted poverty gap* measure is defined as $M_1 = \mu(g^1(k))$, or the mean of the censored normalized gap matrix, while the *adjusted FGT* measure is $M_2 = \mu(g^2(k))$, or the mean of the censored squared gap matrix. Each can likewise be interpreted using partial indices. Let G denote the *average normalized gap*, which calculates the average value of the normalized gap among all instances where a poor person is deprived (and hence where the censored normalized gap is positive); similarly, let S denote the *average squared gap* among all instances where a poor person is deprived (and hence where the censored squared gap is positive). G provides information on the average depth of deprivations among all poor and deprived states; by taking the square of the normalized gaps, S places relatively greater emphasis on the larger gaps. It can be shown that $M_1 = HAG$ and $M_2 = HAS$, so that each is the product of three intuitive partial indices. In our example, the sum of the nonzero entries in $g^1(k)$ is $(0.04+0.42+0.17+0.67+1+1) = 3.3$ and so $M_1 = \mu(g^1(k)) = 3.3/16$; alternatively, $HA = 3/8$ and $G = 3.3/6$ yield $M_1 = HAG = 3.3/16$.

These expressions for the measures suggest a useful line of analysis when applying the measures to data over time or space. Suppose for example that M_0 has increased over time. It could be useful to know whether the increase was primarily due to an increase in H , the prevalence of poverty in the population, or to an increase in A , the average intensity (or breadth) of poverty among the poor. Analogous breakdowns are available for M_1 and M_2 , with an additional factor being an increase in the average depth or average squared gap across deprived states. In conjunction with our identification approach, the three poverty measures M_0 , M_1 , and M_2 satisfy an increasingly stringent list of axioms, reflecting M_0 's sensitivity to intensity or breadth of deprivation, M_1 's sensitivity to the depth of a deprivation, and M_2 's sensitivity to inequality among deprived states of the poor.⁹ As we will see below, each exhibits two forms of decomposition that are especially useful in empirical applications.

4.4 Decompositions

In developing multidimensional methods, we would not want to lose the useful properties that the unidimensional methods have employed with such success over the years. Prime among them is decomposability, which posits that overall poverty is a population share weighted average of subgroup poverty levels. This requirement has proved to be of great use in analyzing poverty by regions, by ethnic

⁹ The dual cutoff identification combined with measures M_0 , M_1 , and M_2 yield methodologies that satisfy decomposability, replication invariance, symmetry, poverty and deprivation focus, weak and dimensional monotonicity, nontriviality, normalisation, and weak rearrangement; while M_1 and M_2 satisfy monotonicity; and M_2 satisfies a weak transfer property (Alkire and Foster 2011). Our measures are neutral with respect to a positive rearrangement among the poor, and hence just satisfy a weak rearrangement property. Bourguignon and Chakravarty (2003) have emphasized the possibility of alternative directions and strengths of interaction among deprivations, and provide a method that can convert a neutral measure into one that reflects varying degrees of substitutability or complementarity.

groups, and by other subgroups defined in a variety of ways. A related property is subgroup consistency, which requires overall poverty to fall if poverty decreases in one subgroup and is unchanged in the other subgroups, given fixed subgroup populations. If this property did not hold, so that overall poverty could rise at the same time its level falls in subgroups, this could make it difficult indeed to develop appropriate policies to combat poverty. Both properties are satisfied by the traditional FGT measures and also by the AF methodology. In the latter case, this follows immediately from the definition of each measure as the mean of a censored matrix, since the mean is likewise decomposable across subgroups using population share weights.

The level of poverty in a subgroup may be lower (or higher) than the overall level of poverty, and this has a direct effect on the overall level as specified by the decomposition. Indeed, the subgroup poverty level divided by the overall poverty level, all multiplied by the population share of the subgroup, can be viewed as the subgroup's contribution to overall poverty. The subgroup contributions clearly sum to one.

The measures also exhibit a second form of decomposition – by dimension – which applies after the identification step when censored matrices have been defined. For M_0 , the breakdown is expressed in terms of each dimension's *censored headcount ratio*, which is the percentage of the overall population who are both poor *and* deprived in the given dimension. The dimensional decomposition formula states that M_0 is equal to a weighted average of the censored headcount ratios, where the weight on dimension j is given by w_j/d . The percentage contribution of a given dimension to overall poverty is its weighted censored headcount ratio divided by the overall poverty level.¹⁰ In the above example, the censored headcount ratio for the first dimension is $1/4$ while the weight is $1/4$ so that the contribution of the first dimension to the overall adjusted headcount ratio M_0 is $1/6 = (1/16)/(3/8)$.

We stress that the dimensional decomposition only holds after identification has taken place. Since our identification function cannot be decomposed by dimension (one must check across multiple dimensions to see who is poor), the aggregate poverty level cannot – except in the extreme case of union identification – be reassembled dimension by dimension before identification. This is a fundamental feature of our approach, which arises because our ‘multiple deprivation’ concept of poverty is sensitive to the joint distribution of deprivations. In a very interesting paper Chakravarty, Mukherjee, and Renade (1998) posit a stronger form of factor decomposition where each dimensional term depends only on that dimension's distribution, and show how this property might be useful in practice. However, a cost of this stronger decomposition is that the poverty methodology must use a restricted identification approach; the identification function cannot be sensitive to multiplicities of deprivations or other information derived from the joint distribution. We discuss this further in section 5.1 below.

¹⁰ Alternatively, the weighted censored headcounts are simply the means of the columns of the censored deprivation matrix. So the decomposition says that M_0 is the average of the column means; while the contribution of a given dimension is its column mean divided by the product of d and M_0 .

Figure 1.




Table 1.

		Living Standard					Health		Education		Overall	
		Assets	Cook	Floor	Water	Sanita	Electr	Nutri	Mortal	Attend		Years
	Deprivation Values	0.56	0.56	0.56	0.56	0.56	0.56	1.67	1.67	1.67	1.67	10.00
DRC	Censored Headcount	0.57	0.73	0.65	0.55	0.62	0.69	0.17	0.37	0.40	0.14	0.39
	Contribution	8.1%	10.3%	9.2%	7.8%	8.8%	9.8%	7.2%	15.7%	17.1%	6.1%	100%
MP	Censored Headcount	0.52	0.67	0.57	0.31	0.65	0.25	0.51	0.31	0.32	0.22	0.39
	Contribution	7.4%	9.5%	8.2%	4.4%	9.3%	3.6%	21.6%	13.1%	13.7%	9.2%	100%

Decomposition of MPI for Two Regions¹¹

Dimensional decomposition can be used in conjunction with subgroup decomposability to better understand the patterns of poverty across a population and their differential sources. Alkire and Foster (2011) provide an example which first decomposes a population by ethnic group and then by dimension. They find that one ethnic group's contribution to total poverty is much higher for multidimensional poverty than for income poverty, and then examine how the dimensional contributions vary by ethnic group to see what might lie behind this finding. A second example drawn from Alkire and Seth (2011) provides information on two regions for one cross-country implementation of M_0 called the MPI. The regions have roughly the same population sizes and share M_0 levels of 0.39. Decomposition by dimension reveals how the underlying structure of deprivations differs across the two for the ten indicators (see Figure 1 and Table 1). In Madhya Pradesh, nutritional deprivations contribute the most to multidimensional poverty, whereas in the Congo the relative contribution of nutritional deprivations is much less than, say, deprivations in school attendance. Although the overall poverty levels as measured by M_0 are very similar, decomposition reveals a very different underlying structure of poverty, which in turn could suggest different policy responses.

5. Misunderstandings

Our methodology is not the last word in multidimensional poverty measurement; we hope that there will be improvements in both theory and application. To facilitate this process, we need to be as clear as possible about what our methods are, and are not, saying. Several aspects of the story are easy to miss or

¹¹ Data are drawn from the Demographic and Health Surveys (DHS) for the DR Congo (2007) and Madhya Pradesh (2006), which are nationally-representative household surveys. For details see Alkire and Santos (2010).

misinterpret, and these points could benefit from additional emphasis and clarification. In what follows we present some of these misunderstandings of our approach and of the whole enterprise of multidimensional poverty measurement.

5.1 The Construct: Poverty as Multiple Deprivations

Our methods are based on a concept of poverty as multiple deprivations that are simultaneously experienced. Persons confronted by a broad range of deprivations are poor, while those with limited breadth of deprivation may not be. A censoring process, which limits consideration to the deprivations of the poor, connects our identification and aggregation steps. It follows that our methods are sensitive to the joint distribution of deprivations, a characteristic that is absent from marginal methods and suppressed in unidimensional measures applied to an aggregate variable. This is a key feature of our approach that is most easily seen by contrasting with other methodologies.

A marginal method reflects population deprivations within dimensions, but does not look across dimensions for the same person, and cannot reflect the extent of associations among deprivations. Consider the deprivation matrices given below, and suppose poverty is evaluated using M_0 with $k = 2$. How would conclusions be different if we used a marginal method? Both matrices depict situations in which, for each dimension, one fourth of the population is deprived. However, in Matrix 1 it is the same fourth of the population that is deprived in all dimensions, while in Matrix 2 one fourth is deprived in the first dimension, a second fourth is deprived in the second dimension, and so on, with no one experiencing more than a single deprivation. Our methodology focuses on multiple deprivations and readily distinguishes between the two cases. Indeed, M_0 is $1/4$ in the first case and zero in the second. A marginal method would be unable to detect the difference between the two cases given the limited form of information it admits – the marginal distributions – which are the same in both cases.

$$\begin{array}{ccc}
 \text{Matrix 1} & & \text{Matrix 2} \\
 \\
 g^0 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix} \cdots \begin{bmatrix} 0 \\ 0 \\ 0 \\ 4 \end{bmatrix} & & g^0 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdots \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}
 \end{array}$$

To illustrate further, recall the dimensional decomposition of M_0 as the weighted average of the censored headcount ratios - where a censored headcount ratio is the share of the population that is both poor and deprived in a particular dimension. For Matrix 1, each censored headcount ratio is $1/4$ and hence M_0 is $1/4$; in Matrix 2, no one is poor, and so the censored headcount ratios and M_0 are 0. What would happen if instead we used uncensored headcount ratios (or the share of the population that is deprived in a dimension) and averaged up to obtain a marginal measure? The average would account for all deprivations in society, not just those that are experienced by poor persons. This is not a problem in the first scenario given above, since all deprivations are experienced by the poor and both forms of dimensional headcount ratios are $1/4$. In the second scenario, though, where no one is poor the uncensored headcount ratios, and hence the marginal measure, would remain at $1/4$. Depending on the joint distribution, the difference between the censored and uncensored values may be small (reflecting few deprivations among nonpoor persons) or large (indicating that many deprivations are scattered

among the nonpoor).¹² A marginal measure, by construction, cannot reflect changes in the joint distribution that leave the marginal distributions unchanged.¹³

The implications are very different for unidimensional measures applied to an aggregate variable. The key question here centrally concerns the aggregation method, which collapses all dimensions and their correlations into a single dimension. If the aggregation method faithfully combines *all* the relevant resources (or achievements) into the right aggregate for *every* person, so that data matrices having the same aggregate value are in fact identical in all important respects, then a multidimensional method is not needed. Dimensional cutoffs and deprivations are irrelevant; only the meaningless variation in the joint distribution is being lost. The difficulty arises when the aggregate is wrong, or when aggregation is used where aggregation is inappropriate (say, the two dimensions are fundamentally incommensurate achievements). If aggregation across achievements is seen as crucial to the entire exercise, this may well drive a careful researcher to focus on dimensions that can be readily aggregated and omit the rest – no matter how fundamental their link with poverty. A better alternative might be to aggregate whenever possible and then to analyze the resulting incommensurate achievements using an appropriate multidimensional methodology.

Our approach seems particularly pertinent when the variables in question relate to capability deprivations. At a technical level, this is a case where aggregation of achievements is an inherently challenging exercise, and perhaps inadvisable, because an achievement above the cutoff in one dimension, like being well educated, may not perfectly compensate for a deprivation in another dimension, like being unemployed. Each capability, and hence each deprivation, has distinct and intrinsic value. Also, when poor people describe their situation, as has been found repeatedly in participatory discussions, part of their description often narrates the multiplicity of disadvantages that batter their lives at once. Malnutrition is coupled with a lack of work, water has to be fetched from an area with regular violence, or there are poor services and low incomes. In such cases, part of the experience and problem of poverty itself is that several deprivations are coupled – experienced together. Our methodology can reflect the joint distribution of disadvantage and the extent of simultaneous disadvantages.

5.2 Data Requirements: Single Survey Sourcing

Our methodology requires data for each variable to be available from the same survey and linked at the individual (or household) level.¹⁴ It cannot be applied to anonymous or unlinked data drawn from different sources. In this it is quite similar to other welfare and empirical analyses at the individual or household level and, in particular, to traditional unidimensional poverty analyses based on expenditure surveys, since clearly all expenditures of the household are needed in order to construct the household's overall expenditure level. We view this as a positive attribute of our methodology – that it makes effective use of the data needed to understand the phenomenon being evaluated.

¹² The degree of 'mismatch' clearly has to do with the degree of positive association among the dimensions. This is also closely related to the discussion of robustness of results to changes in the various parameters. See Foster, McGillivray and Seth (2009, 2011).

¹³ The same logic would apply if uncensored headcount ratios were presented in vector form rather than averaged. Each dimension would typically include instances of non-poor deprivations; if poverty was the focus, then the censored headcount ratios would seem to be a more appropriate vector of dimensional indicators. Put differently, the ministry of health or education would be interested in censored *and uncensored* headcount ratios of health or education; a social development ministry whose main concern is poverty would be interested in the *censored* headcount ratios.

¹⁴ There are exceptions to this – for example if data for the same respondent can be merged from different sources, or if certain variables are uniform across a population subgroup and can be supplied for subgroups from alternative surveys.

The pressing need for surveys that supply the data required to investigate the joint distribution of dimensional achievements was highlighted in the 2009 Commission on the Measurement of Economic Performance and Social Progress set up by President Sarkozy and chaired by Stiglitz, Sen, and Fitoussi. That study recommends that, ‘surveys should be designed to assess the links between various quality of life domains for each person, and this information should be used when designing policies in various fields.’ They elaborate this recommendation as follows:

[T]he consequences for quality of life of having multiple disadvantages far exceed the sum of their individual effects. Developing measures of these cumulative effects requires information on the ‘joint distribution’ of the most salient features of quality of life across everyone in a country through dedicated surveys. Steps in this direction could also be taken by including in all surveys some standard questions that allow classifying respondents based on a limited set of characteristics. When designing policies in specific fields, impacts on indicators pertaining to different quality-of-life dimensions should be considered jointly, to address the interactions between dimensions and the needs of people who are disadvantaged in several domains (pp. 15–16).

Data from the same survey are required for all of the measures that reflect the joint distribution of disadvantage among multidimensionally poor people, regardless of the particular functional form that the measure might take. Hence an investment in the data side could be useful in stimulating the development of alternative measurement methodologies. Furthermore, such data are, as the Sarkozy Commission suggested, required not only for poverty measurement but also for the analysis of quality of life and welfare. Both welfare and poverty analyses need to reflect individuals’ experiences.

While our fully multidimensional poverty measures require data from the same surveys, in a different sense they have far lower data requirements. One of the key attributes of the AF technology for measuring poverty, and one that is highlighted in the MPI example discussed below, is that it is consistent with qualitative data that describe the basic achievements of the poor. The key requirement for ordinal data is that if the cutoff and variables are changed by a monotonic transformation, the level of poverty must remain unchanged, and the same people must be identified as poor. This criterion is satisfied by the dual cutoff identification and the adjusted headcount ratio M_0 . The M_0 methodology can also be used with categorical or dichotomous variables so long as a meaningful deprivation cutoff can be fixed. This is an advantage not only due to the prevalence of these data, but also because of the fact that data for some core functionings are not available in cardinal form.

5.3 Methodologies and Measures: AF compared with the MPI (and the HDI)

Our methodology is a general framework for multidimensional poverty measurement, which can be filled in different ways. The dimensions and cutoffs could vary, as could the weights and poverty cutoff. The measure could be applied at different levels. For example, a poverty measure could be implemented at the village, state, or national levels. And the specific choice of measures might vary: one institution might implement a measure with cardinal data to reflect the depth of poverty or inequality among the poor underlying M_1 or M_2 , whereas another could only have ordinal data available, and so would report M_0 and the breadth of poverty. In sum, our methodology is a very flexible framework and can give rise to a number of concrete applications whose shapes depend upon the purpose for which they are designed.

One notable example of the AF methodology can be found in recent work by Alkire and Santos (2010) on acute multidimensional poverty in developing countries. The Multidimensional Poverty Index (MPI)

implements M_0 for 104 developing countries using a set of ten indicators for which internationally comparable data are available.¹⁵ It identifies people as poor depending upon achievements of household members. It applies cutoffs that are broadly related to international standards such as the MDGs. Nested weights are applied to the indicators, with the health, education and living standard categories receiving equal shares of the aggregate weights, and each indicator being allocated an equal share of the category's weight. Consequently, each of the six living standard indicators has a deprivation value of 10/18, while each of the two health (and the two education) indicators has a deprivation value of 10/6. A poverty cutoff k is chosen that identifies a person as poor if she is deprived in any two health or education indicators, in all six standard of living indicators, or in three standard of living and one health or education indicator. Note that the MPI is one particular application of the AF methodology that was designed for cross-country purposes. Naturally it reflects the myriad of constraints that are often seen when constructing internationally comparable measures, yet does create meaningful comparisons across countries.

The MPI demonstrates the operability of the AF methodology even in the presence of rather limited datasets. The aim in implementing the MPI was not to suggest that this particular set of indicators, dimensions, cutoffs, and weights is appropriate in *every* application. Rather, it was to apply the AF methodology fully with respect to one particular problem (cross-country evaluations of acute poverty based on DHS, MICS, and similar datasets). The existence of a practical cross country implementation can aid in the development of local, national, and institutional methods that better reflect the particularities of poverty in different contexts and inform poverty reduction policies.

Because the AF methodology is multidimensional, there is a chance it might be confused with the Human Development Index (HDI), which aggregates across achievements in health, education, and standard of living. In fact, the two measure very different things. The AF methodology (and its particular example of the MPI) measures poverty: it identifies who is poor and ignores the data of the nonpoor. In contrast the HDI is a welfare index based on three marginal distributions that combines the aggregate dimensional achievements of *all* people (not just the poor) into one overall score. While the HDI may be limited in terms of data, dimensions, and methodology, it has helped bring into view people's achievements in non-monetary spaces, and made it possible for other categories of multidimensional measures (such as poverty measures) to be envisioned.

5.4 Underpinnings: Poverty and Welfare

Our methodology has its roots in the literature on axiomatic poverty measurement (Sen, 1976), which – like related work in inequality and welfare indices – employs axioms to discern between measures. Within the axiomatic approach, a method is judged to be acceptable or unacceptable according to the properties it satisfies. To critique a method of measurement, one identifies desirable properties it fails to exhibit. Non-axiomatic arguments, by contrast, are considered less relevant. The axiomatic approach provides the foundational theoretical structure for the AF methodologies, and the paper devotes substantial effort to formulating and discussing the axioms that they satisfy.

There is a clear difference between evaluating the axiomatic characteristics of a measurement method and questioning the calibrations underlying a *particular* implementation of the approach. For example, there may be broad agreement on the use of the squared gap FGT measure, but disagreement over using

¹⁵ The MPI was produced collaboratively by OPHI researchers with the UNDP Human Development Report Office and was included in the 2010 *Human Development Report* as an experimental index. See UNDP (2010). Note that in the MPI literature, the term 'indicator' corresponds to 'dimensions' or columns in the data matrix, while the term 'dimension' is reserved for the three broader categories – health, education, and standard of living – containing these indicators.

income or consumption as the variable, or about the specific poverty line, or which PPP rates to use. In the multidimensional case, there may be disagreement over deprivation values and cutoffs, even if the AF methodology is selected. Any critique that relies on a specific set of implementation choices is contingent on those choices and may not reveal much about the general methodology – which is more properly evaluated in axiomatic terms.¹⁶

Another way that poverty measures have been interpreted is through their link with some concept of welfare. For example, the Watts unidimensional poverty measure is related to the geometric mean – one of Atkinson's social welfare functions – in the following way. All incomes above the poverty line are censored to the poverty line level to create the censored distribution. The geometric mean is applied to the censored distribution to obtain the censored welfare level. The Watts measure is the difference between the log of the poverty line and the log of the censored welfare level. Perhaps an analogous link could be invoked to evaluate the AF methodology or other multidimensional measures.¹⁷

One difficulty in grounding poverty measurement in welfare is that welfare itself is so challenging to measure in practice. A welfare function must be able to make meaningful evaluations at all levels of achievements across all persons. This requires strong assumptions about the measurement properties of data¹⁸ and on the functional form to get anywhere; and although data analysis might help constrain the possibilities somewhat over certain dimensions (e.g., those traded in markets), much of the exercise of selecting a welfare function is by definition normative and has many degrees of freedom. There will likely be a multiplicity of acceptable functions, and, even if a unique welfare function could be agreed upon, there is no unique transformation from welfare function to poverty measure. Moreover, if one *could* settle on a specific function to measure welfare, with a meaningful welfare indicator for each individual, this would more naturally suggest the unidimensional approach to poverty measurement with welfare as the aggregator and a welfare cutoff as the target. It is the absence of such an aggregator that encourages consideration of the multidimensional approach.

An alternative exercise might be to see whether the tradeoffs implied by a poverty measure are broadly consistent with some underlying notion of social welfare. This is indeed a reasonable approach, but one whose conclusions are often ignored in practice in deference to other favorable aspects of a measure. For example, the headcount ratio, so commonly used in traditional poverty measurement exercises, has the interesting property that any decrease in the income of a poor person (no matter how large a decrease) paired with any increase in income for a nonpoor person (no matter how small the increase) will leave poverty unchanged. This, of course, is rather untenable from a welfare perspective. Likewise, any decrease in the income of a poor person (no matter how large the decrease) paired with an increase in the income of another poor person sufficient to lift the person to the poverty line income (no matter how small the increase) will decrease poverty. Again, this would appear to go against the judgement of any reasonable welfare function censored at the poverty line. Note, though, that the fact that these tradeoffs are not justified in welfare terms has not forced the removal of this measure from consideration. The headcount ratio is a remarkably intuitive, if somewhat crude, measure that takes the identification process very seriously and then reports a meaningful number: the prevalence of poverty.

¹⁶ It would be another matter to assert that a methodology is *impossible* to implement. This would indeed be a type of general critique, but would require an evaluation of *all* calibrations, not just one.

¹⁷ Note that certain commonly used unidimensional poverty measures, such as the FGT squared gap index, do not have a direct welfare interpretation - hence this is not viewed as a requirement within the literature. On the other hand, the FGT class is closely linked to welfare dominance; see Foster and Shorrocks (1988), Atkinson (1987), or Ravallion (1994).

¹⁸ See for example the discussion in Alkire and Foster (2010a) on the assumptions underlying the Inequality Adjusted Human Development Index. In particular, ordinal or categorical variables can be especially troubling.

The fact that it is at variance with notions of welfare would appear to be of second order importance to practitioners who continue to highlight it as the ‘headline’ statistic.¹⁹

Similarly, the poverty gap implies a constant and unitary marginal rate of substitution between any two poor incomes (i.e., a marginal equal-sized transfer between any two poor persons leaves the poverty gap unchanged). One might rather expect that a progressive marginal transfer should result in a lower poverty level; or equivalently, that the amount of income that needs to be given to a poorer poor person to compensate for a \$1 decrease in a richer poor person’s income would be less than \$1. This welfare-based argument would direct one instead to the distribution sensitive poverty indices like the FGT, the Watts or the Sen index. And yet, the poverty gap, which ignores the distribution or severity of poverty, is reported with much greater frequency. The revealed preference evidence about practitioners’ use of unidimensional measures suggests that simplicity or informational content of certain poverty measures often trumps the favorable ‘marginal rate of substitution’ properties of others.

The main margin of analysis for unidimensional poverty measures is across people. For multidimensional poverty measures there is another margin to be considered: across the different dimensions of poverty. But how can we credibly account for the *deprivations* of the poor if there is no aggregation system across dimensional *attainments* (as was posited in the multidimensional case)? The AF methodology takes the simplest way forward: it requires a positive value to be set for each deprivation and then adds up (or averages) across these values to determine the breadth of deprivations for the identification step and also for the adjusted headcount ratio M_0 . The other two measures figure in the depth or severity of deprivations, in cases where the variables allow it. The values or weights afforded each deprivation are straightforward in interpretation and lead to measures whose numerical values convey meaning. As noted in the next section, they can be set in a variety of ways suitable to the purpose of the poverty measurement methodology. Information on how individuals themselves value deprivations can be part of the mix, as might a broader perspective of society’s priorities, in arriving at this key subset of parameters used to calibrate the AF methodology.

5.5 Calibration: Who Will Choose Parameters?

Our methodology is a general framework for measuring multidimensional poverty – an open source technology that can be freely altered by the user to best match the measure’s context and evaluative purpose. As with most measurement exercises, it will be the designers who will have to make and defend the specific decisions underlying the implementation, limited and guided by the purpose of the exercise and by commonly held understandings of what that purpose entails. Traditional unidimensional measures require decisions that are qualitatively similar. For example, should the variable be expenditure or income? What should the poverty cutoff be? Other implementation choices are less apparent but can likewise be important for final results, as we have seen with the recent updating of PPP values and the subsequent impact on world poverty counts. Robustness tests are crucial both for ensuring that the results obtained are not unduly dependent upon the calibration choices and for allowing these choices to be made in the first place.²⁰

For example, the multidimensional measure could seek to reflect capability poverty. In this case then, following Sen (1987, 1992), the selection of relevant functionings is a value judgement, as is the selection of weights and cutoffs. There are various procedures for making such value judgements. Ideally,

¹⁹ When data have lower level measurement properties, such as when the variable is ordinal or categorical, we may be forced to use the headcount ratio of some other counting measure that is still meaningful in this context.

²⁰ For example, the variable poverty line robustness methods such as those found in Foster and Shorrocks (1988) have helped make the fundamentally arbitrary \$1.25 a day poverty line much more palatable.

recurrent participatory and deliberative processes could be used to choose them, because done (and if inequalities in power and voice are managed well) these procedures have the potential to facilitate the exchange of reasons, build consensus, and create legitimacy. In other cases the choices might be made by a technical committee, and might reflect participatory processes in other ways – for example by drawing on recent participatory studies, a national plan, or the constitution. In any case, Sen stresses the need to communicate the chosen parameters explicitly in the media and other outlets, so they could be improved by public discussion and debate in the future.

The calibration choices will depend upon the purpose of the measure, such as the space in which poverty is evaluated, the relevant comparisons across time or populations that the measure will inform, or the particular programmes or institutions which will be evaluated. Calibration choices will also reflect data and resource constraints.

The flexibility in choice of parameters makes the AF methodology particularly useful at the country level where measurement decisions can be made locally to embody prevailing norms of what it means to be poor. For example, if dimensions, weights, and cutoffs are specified in a legal document such as the Constitution, the identification function might be developed using an axiomatic approach, as was done in Mexico.²¹ The weights can be developed by a range of processes: expert opinion, coherence with a consensus document such as a national plan, or the MDGs. And the poverty cutoff, which is analogous to poverty lines in unidimensional space, could be chosen so as to reflect policy needs and resources.²²

It is helpful to subject measures to a series of robustness tests that check whether the results obtained for an initial set of parameters are upheld for other plausible parameter values. The tools can include formal dominance orderings, which indicate when agreement exists for all values of one or more parameters, or simple empirical robustness checks, which re-evaluate results for a finite number of alternative values.²³ Clearly, the initial choice of parameters would be more difficult if important comparisons were sensitive to small adjustments in them. By applying robustness tests this sensitivity can be explored transparently.

In our view, enabling people to choose parameters according to a range of processes provides an essential flexibility and adaptability to allow the measures to be tailored to institutional, cultural, and data-specific circumstances. Also, the AF methodology is relatively transparent, and this feature can be helpful when parameters are set by (or at least opened to) public debate. It uses explicit indicators, weights, and cutoffs, so that serious shortcomings in the choice of parameters could be debated and changed. To counterbalance and inform this flexibility, we suggest the use of dominance results and of robustness and sensitivity tests, which will show whether the key points of comparison are robust to a range of plausible parameter choices.

²¹ See Alkire and Foster (2010b).

²² On the choice of dimensions see Alkire (2002, 2008); on the choice of weights see Decancq and Lugo (forthcoming); on the choice of poverty cutoffs see Alkire and Foster (2011).

²³ For example see Alkire and Foster (2011) and Lasso de la Vega (2010); see also Alkire and Santos (2010) and Alkire, Santos, Seth, and Yalonetzky (2010). A stringent and full set of dominance conditions that ensure the robustness of comparisons to widely varying weights, deprivation cutoffs, and poverty cutoffs, has been derived in Yalonetzky (2010). Statistical tests for these conditions are available in Yalonetzky (forthcoming) for discrete variables; and in Anderson (2008), for continuous variables. Bennett and Singh (2010) propose a test for multiple hypotheses that allows the researcher to check the robustness of the cutoffs and compare two groups. Yalonetzky (2011) derives asymptotic standard errors for the basic statistics of the AF family (H, A and M₀) and their percentage changes for various sampling procedures.

6 Concluding Remarks

The literature on unidimensional poverty measurement provides the bedrock for wider approaches. This paper opened by observing how multidimensional measures build upon and depart from these roots, for example through the application of deprivation cutoffs in each dimension prior to the identification step. We then introduced one particular methodology of multidimensional poverty measurement. The AF methodology introduces a dual-cutoff identification method, while its aggregation methodology builds on the traditional FGT approach. The overall measures as well as their sub-indices are intuitive and easy to interpret, and satisfy a set of desirable properties such as decomposability. Although the AF methodology has a specific structure for identification and aggregation, its implementation is flexible: parameters such as dimensions, cutoffs, and weights can be chosen to reflect the purpose of the measure and its context.

The value of having a poverty measure like M_0 is that it can show change over time in a unified and internally consistent framework. This overall picture of change can be coherently and consistently deepened and sharpened by the more specific insights contained in decompositions and partial indices. Understanding this combination – an overview indicator that can give rise to a range of more specific analyses – is key to our measurement approach.

We also clarified five misunderstandings that may arise when the AF methodology is initially encountered. The first pertains to our fundamental use of the joint distribution in constructing our measures. We require information on all of a person's achievements in order to determine if the person is poor. The joint distribution is likewise incorporated into our measures of poverty, each of which may be disaggregated by population subgroup or (after identification) by dimension to better understand the structure of poverty. Note, though, that the dimensional decomposition creates dimensional sub-indices (like the censored headcount ratios) that also rely fundamentally on achievements and deprivations across all dimensions – or the joint distribution. In contrast, marginal methods aggregate up from the marginal distributions, and decompose by dimension back to dimensional indices (like the uncensored headcount ratios) where they began.

A second misunderstanding is related to implementation and data sources: because each variable must be linked to a particular household, all variables normally must come from the same survey source. We echo, therefore, the call of the Stiglitz, Sen and Fitoussi Committee for better multi-topic household survey data. We also note that our methods can incorporate qualitative or ordinal data, which are commonly available.

In 2010 our M_0 methodology was used by Alkire and Santos (2010) to develop a Multidimensional Poverty Index (MPI) reflecting acute poverty in 104 developing countries. That exercise revealed the applicability of this methodology, and stimulated interest in similar analyses. But it is easy to fall into a third misunderstanding: that the *particular* MPI dimensions, indicators, cutoffs, and weights, are part of the general *methodology*, instead of recognising the range of calibration choices that are possible. The parameters chosen for the international MPI reflect acute poverty across developing countries within strongly binding data constraints; other specifications could be considered – for example to measure national poverty, target beneficiaries, or evaluate a particular intervention.

Fourth, it might be assumed that multidimensional poverty methods have been constructed without reference to a theoretical framework. In fact there is a mature literature on axiomatic poverty measurement that has been brought over into the multidimensional environment for evaluating prospective methodologies. While links with welfare can provide a helpful way of interpreting poverty measures, requiring such links can be detrimental: it drives poverty measurement towards the economic variables that are the arguments of traditional welfare functions and unidimensional poverty measures

and away from new approaches to well-being and capability. Conceptions and empirical measures of multidimensional welfare must be improved before their potential implications for multidimensional poverty measurement become clear. In the meantime, axioms can imbue poverty methodologies with welfare-related considerations, and likewise such methodologies can be criticised if they fail to exhibit key properties.

Fifth and finally, multidimensional poverty measures can and often do take form in active relationship with participatory and deliberative processes about what poverty is and what current priorities might be. In particular, because the dimensions, indicators, weights and cutoffs of our methodology are flexible, they can be filled in in many ways. Robustness checks should normally be implemented to ensure that key points of analysis are robust to a range of plausible parameters.

‘Measurement...’ as Sen put it, ‘is not an easy task’. In this paper we elucidate the strengths, limitations, and misunderstandings of multidimensional poverty methods in order to clarify the debate and to catalyse further research. Of particular importance is ongoing clarification of what a multidimensional poverty methodology is, what value-added it provides, and how axiomatic multidimensional poverty methodologies are best evaluated. Naturally, there are many other topics of interest including those related to preferences, weights, functional forms, data sources, and statistical methods. One can anticipate a sustained and vigorous exchange of ideas.

7. References

- Alkire, S. 2002. *Valuing Freedoms. Sen's Capability Approach and Poverty Reduction*. Oxford: Oxford University Press.
- Alkire, S. 2008. ‘Choosing Dimensions: The Capability Approach and Multidimensional Poverty’. In N. Kakwani and J. Silber (eds.), *The Many Dimensions of Poverty*. New York: Palgrave Macmillan. 89–119.
- Alkire, S. and Foster, J. 2007. ‘Counting and Multidimensional Poverty Measures’. *OPHI Working Paper 7*.
- Alkire, S. and Foster, J. 2010a. ‘Designing the Inequality-Adjusted Human Development Index’. *OPHI Working Paper 37*.
- Alkire, S. and Foster, J. 2010b. ‘An Axiomatic Approach to Identification and Measurement of Multidimensional Poverty’, *OPHI Research in Progress*, 21a.
- Alkire, S. and Foster, J. 2011. ‘Counting and Multidimensional Poverty Measurement’, *Journal of Public Economics*, 95(7-8), 476-487.
- Alkire, S. and Santos, M. E. 2010. ‘Acute Multidimensional Poverty: A New Index for Developing Countries’, *OPHI Working Paper Series*, 38.
- Alkire, S., Santos, M.E., Seth, S. and Yalonetzky, G. 2010. ‘Is the Multidimensional Poverty Index Robust to Different Weights?’ *OPHI Research in Progress*.
- Alkire, S. and Seth, S. 2009. ‘Determining BPL Status: Some Methodological Improvements’, *Indian Journal of Human Development*, 2(2): 407–24.
- Alkire, S. and Seth, S. 2011. ‘Decomposing India’s MPI by State and Caste: Examples and Comparisons’, *OPHI Research in Progress*.
- Anand, S. and Sen, A. 1997. *Concepts of Human Development and Poverty: A Multidimensional Perspective*. New York: UNDP.
- Anderson, G. 2008. ‘The Empirical Assessment of Multidimensional Welfare, Inequality and Poverty: Sample Weighted Multivariate Generalizations of the Kolmogorov–Smirnov Two Sample Tests for Stochastic Dominance’, *Journal of Economic Inequality*, 6(1): 73–87.
- Apablaza, M., Ocampo, J. and Yalonetzky, G. 2010. ‘Decomposing Changes in Multidimensional Poverty for Ten Countries’. *mimeo*.
- Atkinson, A. 1987. ‘On the Measurement of Poverty’, *Econometrica*, 55(4): 749–764.
- Atkinson, A. 1989. *Poverty and Social Security*. New York: Harvester Wheatsheaf.
- Atkinson, A. 2003. ‘Multidimensional Deprivation: Contrasting Social Welfare and Counting Approaches’, *Journal of Economic Inequality*, 1(1): 51–65.
- Bennett, C. and Singh, S. (Forthcoming.) ‘Multidimensional Poverty: Measurement, Estimation and Inference’, *Econometric Reviews*.

- Bossert, W., Chakravarty, S. and D'Ambrosio, D. 2009. 'Multidimensional Poverty and Material Deprivation', *Working Papers* 129. ECINEQ, Society for the Study of Economic Inequality.
- Bourguignon, F. and Chakravarty, S. 2003. 'The Measurement of Multidimensional Poverty', *Journal of Economic Inequality*, 1(1): 25–49.
- Brandolini, A. and D'Alessio, G. 1998. *Measuring Well-being in the Functioning Space*. Rome: Banco d'Italia Research Department.
- Chakravarty, S., Mukherjee, D. and Renade, R. 1998. 'On the Family of Subgroup and Factor Decomposable Measures of Multidimensional Poverty', *Research on Economic Inequality*, 8: 175–194.
- Chakravarty, S. and D'Ambrosio, C. 2006. 'The Measurement of Social Exclusion', *Review of Income and Wealth*, 53(3): 377–398.
- Chakravarty, S. and Silber, J. 2008. 'Measuring Multidimensional Poverty: The Axiomatic Approach', In N. Kakwani and J. Silber (eds.) *Quantitative Approaches to Multidimensional Poverty Measurement*. Basingstoke: Macmillan.
- Decancq, K. and Lugo, M. (Forthcoming.) 'Weights of Multidimensional Indices of Well-Being: An Overview', *Econometric Reviews*.
- Deutsch, J. and Silber, J. 2005. 'Measuring Multidimensional Poverty: An Empirical Comparison of Various Approaches', *The Review of Income and Wealth*, 51: 145–174.
- Duclos, J.-Y., Sahn, D. and Younger, S. 2006. 'Robust Multidimensional Poverty Comparisons', *The Economic Journal*, 116(514): 943–968.
- Foster, J. 2005. 'Poverty Indices in Poverty, Inequality and Development', In A. de Janvry and R. Kanbur (eds.) *Essays in Honor of Erik Thorbecke*. Amsterdam: Kluwer Academic Publishers.
- Foster, J., Greer, J. and Thorbecke, E. 1984. 'A Class of Decomposable Poverty Measures', *Econometrica*, 52(3): 761–766.
- Foster, J., McGillivray, M. and Seth, S. 2009. 'Rank Robustness of Composite Indices', *OPHI Working Paper Series*, 26.
- Foster, J., McGillivray, M. and Seth, S. 2011. 'Composite Indices: Rank Robustness, Statistical Association, and Redundancy'. *mimeo*.
- Foster, J. and Shorrocks, A. 1988. 'Poverty Orderings and Welfare Dominance', *Social Choice Welfare*, 5(2–3): 179–198.
- Foster, J. 2007. 'A Class of Chronic Poverty Measures'. *Working Paper* 07-W01. Department of Economics Vanderbilt University.
- Foster, J. and Sen, A. 1997. 'After a Quarter Century', In A. Sen (ed.), *On Economic Inequality*. Oxford: Clarendon Press.
- Kakwani, N. and Silber, J. 2008a. *The Many Dimensions of Poverty*. Basingstoke: Palgrave MacMillan.
- Kakwani, N. and Silber, J. 2008b. *Quantitative Approaches to Multidimensional Poverty Measurement*. Basingstoke: Palgrave Macmillan.
- Lasso de la Vega, C. 2010. 'Counting Poverty Orderings and Deprivation Curves'. *Research on Economic Inequality*, 18: 153–172.
- Ravallion, M. 1994. *Poverty Comparisons*. Switzerland: Harwood Academic Publishers.
- Ravallion, M. 2010. 'Mashup Indices of Development', *World Bank Policy Research Working Paper*, 5432.
- Ravallion, M. 2011. 'On Multidimensional Indices of Poverty', *World Bank Policy Research Working Papers*, 5580.
- Rippin, Nicole 2010. 'Poverty Severity in a Multidimensional Framework: The Issue of Inequality between Dimensions', *Courant Research Centre (CRC) Discussion Paper* No. 47.
- Sen, A. 1976. 'Poverty: An Ordinal Approach to Measurement', *Econometrica*, 44(2): 219–231.
- Sen, A. 1987. *The Standard of Living: The Tanner Lectures*. Cambridge: Cambridge University Press.
- Sen, A. 1989. 'Social Progress Index: Some Methodological Issues', First Report to UNDP/Latin America Bureau. *mimeo*.
- Sen, A. 1992. *Inequality Re-examined*. New York: Russell Sage Foundation.
- Stiglitz, J., Sen, A. and Fitoussi, J.-P. 2009. *Report by the Commission on the Measurement of Economic Performance and Social Progress*. The Commission on the Measurement of Economic Performance and Social Progress. Paris.
- Thorbecke, E. 2008. 'Multidimensional Poverty: Conceptual and Measurement Issues', In N. Kakwani and J. Silber (eds.), *The Many Dimensions of Poverty*. New York: Palgrave MacMillan.
- Tsui, K. 2002. 'Multidimensional Poverty Indices', *Social Choice and Welfare*, 19(1): 69–93.
- UNDP 1997. *Human Development Report 1997: Human Development to Eradicate Poverty*. New York: Oxford University Press.
- UNDP 2010. *Human Development Report 2010: The Real Wealth of Nations: Pathways to Human Development*. New York: Palgrave Macmillan.
- World Bank. 2000. *World Development Report 2000/2001: Attacking Poverty*. New York: Oxford University Press.
- Yalonetzky, G. 2010. 'Conditions for the Most Robust Poverty Comparisons Using the Alkire-Foster Family of Measures'. *mimeo*.
- Yalonetzky, G. 2011. 'A Note on the Standard Errors of the Members of the Alkire-Foster Family and its Components', *OPHI Research in Progress* 25.

Yalonetzky, G. (Forthcoming). 'Stochastic Dominance with Ordinal Variables: Conditions and a Test', *Econometric Reviews*