Human Development
Research Paper
2010/28

Designing the Inequality-Adjusted Human Development Index (IHDI)

Sabina Alkire
and James Foster

United Nations Development Programme
Human Development Reports
Research Paper

October 2010

Human Development Research Paper 2010/28

Designing the Inequality-Adjusted Human Development Index (IHDI)

Sabina Alkire
and James Foster

United Nations Development Programme
Human Development Reports
Research Paper 2010/28
October 2010

Designing the Inequality-Adjusted Human Development Index (IHDI)

Sabina Alkire
James Foster

Sabina Alkire is Director of the Oxford Poverty and Human Development Initiative at Oxford University. E-mail: Sabina.Alkire@qeh.ox.ac.uk.

James Foster is Professor of Economics and International Affairs at George Washington University and Research Associate at the Oxford Poverty & Human Development Initiative. E-mail: fosterje@gwu.edu.

Comments should be addressed by email to the authors.

Abstract

As a measure of wellbeing, national income misses variations in the things income can and cannot buy. It also misses variations in people's claim on that aggregate income. The Human Development Index attempts to address the first weakness by incorporating two additional dimensions, health and education, into its informational bases. However, the second weakness, inequality, is ignored by the traditional HDI. In practical terms this means that any two countries having the same *mean* achievements will have the same HDI values even if they have very different *distributions* of achievements. This calls into question the accuracy of the HDI as a reflection of people's actual achievements.

This paper proposes a method for adjusting the HDI to reflect the distribution of human development achievements across the population, and across dimensions. We begin with a discussion of the proposed indices in an idealized setting where variables and their scales have been identified and the data are available. We then address the practical issues that must be addressed when applying these methods to real data. The final section presents and evaluates another related approach.

Keywords: Human Development Index, inequality, multidimensional inequality measurement, capability approach, multidimensional welfare.

JEL classification: I0, D63, O15, I3

1. Why inequality?

As a measure of wellbeing, national income has two notable weaknesses. It misses variations in the things that income can and cannot buy. And it misses variations in people's claim on that aggregate income. These two critiques – in terms of the breadth and distribution of human development – are central to the move towards the human development approach and the Human Development Index (HDI).

The HDI responds to the first of these critiques by incorporating two additional dimensions, health and education, into its informational bases.[1] However, the second concern of central importance to the human development approach is notably absent. In practical terms this means that any two countries having the same *mean* achievements will have the same HDI values even if they have very different *distributions* of achievements. Inequality is ignored by the traditional HDI although there are fundamental reasons for associating greater inequality with lower development levels. In response to this criticism, an alternate HDI adjusted for income inequality was reported in the Human Development Report (HDR) from 1991 until 1994, but was discontinued for methodological reasons.[2] Since then, several authors have called for inequality to be incorporated into the HDI in more substantive ways.[3]

This paper presents methods for evaluating the distribution of human development that account for inequality. We begin in Section 2 with a discussion of the proposed indices in an idealized setting where variables and their scales have been identified and the data are available. In section 3, we consider the practical problems that arise when applying these methods to real data. Section 4 presents and evaluates another related approach.

2. Proposed Methods: Theory

The general methodology considered here is drawn from Foster, Lopez-Calva, and Szekely (2005) and is based on the notion of an *equally distributed equivalent* or *ede*[4] as presented by Atkinson (1970).

2.1 Measurement Assumptions. The Foster, Lopez-Calva, Szekely (FLS) approach, like that of the traditional HDI, begins with strong implicit assumptions on the cardinality and commensurability of the three dimensions of human development.[5] The key implication is that after appropriate transformations, all variables are measured using a ratio scale in such a way that levels are comparable across dimensions.[6] In practice, observed values will be bounded by an upper limit, hence without loss of generality we can normalize by this common bound to ensure that normalized observations take values in the interval [0,1]; or, equivalently, they can be normalized to any other convenient common upper limit M, so that all observations for each variable fall within [0,M]. For example, it might prove convenient to measure achievements in each dimension on a percentage scale from 0 to 100.[7]

Given such a normalization, the distribution of the three dimensions of human development in a population can be represented by a matrix X whose rows give the achievements of a person across the dimensions of income, education and health, and whose columns give the distributions of an achievement across the population.

2.2 Human Development Index. In this idealized world,[8] we may view the HDI as a "mean of means" where we first find average income achievement, the average education achievement and the average health achievement, and then take the average across the three to get the average overall achievement level, or HDI level. Alternatively, we could reverse the order of aggregation and construct an average achievement level for each person, then obtain the HDI by averaging across people. A third equivalent definition dispenses with the order of aggregation entirely and views the HDI level as the mean entry in the matrix of achievements; it is the "representative achievement level" or the "equally distributed equivalent (ede) level of achievement" where the use of the arithmetic mean ensures that there is no concern for inequality.[9] If, however, one's evaluation method is "equity preferring," then the typical mean-of-means HDI is best seen as a measure of *potential* human development, rather than actual human development. We discuss this interpretation more extensively below.

2.3 Atkinson's Ede. Foster, Lopez-Calva, and Szekely (2005) suggest the use of a general mean or equally distributed equivalent achievement level to account for inequality in development. Atkinson's (1970) parametric class of edes first raises each achievement level to a given power $\alpha \leq 1$, takes the arithmetic mean of these transformed achievement levels and then takes the aggregate to the reciprocal of this power, namely $1/\alpha$. When the power is less than 1, the transformation is strictly concave and so greater relative weight is being placed on lower achievements. Inequality in achievements lowers the average transformed value, even when the average achievement is unchanged. The final "undoing" of the initial transformation ensures that the value of the ede is located in the zero-one achievement space. The ede value is then interpreted as the level of achievement distributed *without* inequality that would be equivalent to the original distribution of achievements (distributed *with* inequality).[10]

Atkinson (1970) used the parameter $\varepsilon = 1 - \alpha \geq 1$ to index the class of edes; he noted that ε could be interpreted as a level of inequality aversion inherent in the aggregation method (which he took to be welfare), beginning with the zero inequality aversion case $\varepsilon = 0$ and steadily rising. The case $\varepsilon = 1$ (or $\alpha = 0$) is defined separately to be the geometric mean; it is easily shown that as ε tends to 1 (or α tends to 0) the value of the ede tends to the geometric mean of achievements.

2.4 Inequality-adjusted HDI. The inequality-adjusted HDI (or IHDI) is a parametric family of measures H_ε obtained by applying the associated ede to the matrix X of achievements. The limiting case of $\varepsilon = 0$ returns us to the HDI discussed above that is based on the arithmetic mean; it is clearly not sensitive to inequality in achievements.[11] H_1 is an IHDI that employs the geometric mean to evaluate achievements. H_2 is an IHDI that employs the harmonic mean and so forth. For $\varepsilon > 0$, the IHDI discounts for inequality according to the level of inequality aversion indicated by its associated ε.

2.5 Properties. Each member of H_ε for $\varepsilon \geq 0$ can also be obtained by: (i) applying the ede first within dimensions (to obtain an income ede, a health ede, and an education ede) and then across dimensions; or (ii) first across dimensions at the individual level (to obtain an ede achievement level for each person) and then across persons. Using the same Atkinson ede for both steps ensures that the identical final number is obtained along either pathway. The method satisfies a property called *path independence* by Foster, Lopez-Calva, and Szekely (2005). In contrast, using a different Atkinson ede for each step, or employing other forms of ede besides the Atkinson, raises the possibility that the order of aggregation will matter, and hence would have to be decided upon and defended.[12] In addition, the IHDI also satisfies *subgroup consistency* (thanks to its use of an Atkinson ede for aggregation purposes), which ensures that regional changes in human development are consistent with national changes in human development.[13]

2.6 Decomposability and the Evaluation Space. One property satisfied by H_0 that is not satisfied by H_ε for $\varepsilon > 0$ is *additive decomposability*. While the HDI values for the two regions of a country can be weighted by population shares of the regions and summed to obtain the overall HDI value for the country as a whole, the IHDI aggregates across two regions using a general mean or ede that is not additive. Is this a problem for the IHDI?

In fact, it would not be difficult to transform H_ε to obtain a class of human development indicators that are both sensitive to inequality and decomposable (and generating the same ranking as H_ε). Instead of using Atkinson's ede, we could use Atkinson's additive welfare function to map the matrix of achievements to its mean welfare level according to his evaluation function $u_\varepsilon(.)$, obtaining a human development index that is both a monotonic transformation of H_ε and decomposable. For example, the mean log attainment or, equivalently, the natural logarithm of the geometric mean achievement, is the index obtained when $\varepsilon = 1$.

However, there are also costs from moving in this direction. Welfare space is less tangible than achievement space and has measurement properties that are less exact. Atkinson's critique of Dalton's approach was of this sort, and it motivated him to move from Dalton's welfare space to Atkinson's ede or achievement space. Also, additive decomposability can actually reduce the intuitive meaning of a measure.[14] Likewise, moving from the geometric means of all achievements (or IHDI for $\varepsilon = 1$) – a reasonably clear measure in attainment space – to the average natural logarithm of achievements, is both a move forward (in terms of decomposability) and a move backward (in terms of conveying a clear picture about the level and distribution of human development). Finally, as we shall see below, the geometric mean has certain advantages from a measurement theory point of view.[15]

2.7 Accounting for Inequality. The IHDI reflects inequalities across all achievements in society, including those within each dimension (health, education, income) across people, and those across dimensions for a given person. Following Atkinson (1970), we can define a natural measure of the aggregate inequalities in a society across all achievements:

$$A_\varepsilon = (H_0 - H_\varepsilon)/H_0$$

H_ε measures human development as the ede achievement level for the society as a whole; it is the level of per capita achievement that, if equally distributed, would produce the same level of human development as in the actual distribution. In contrast, H_0 is the actual level of per capita achievement in the distribution. The inequality measure A_ε represents the share of per capita achievement wasted as a result of inequalities in the distribution of achievements. Alternatively, we may view H_ε as an ede for an underlying multidimensional welfare function and is therefore a transformed welfare function itself. Among all matrices X having the same mean achievement (as given by the per capita achievement level H_0), the matrix whose achievements are completely equal will maximize this welfare function. And this maximum level of H_ε is exactly H_0. Hence H_0 can be seen as a measure of the *potential* human development of the society, or the maximum level of H_ε that is possible if all achievements could be costlessly transferred across dimension and across persons. A_ε is then viewed as the percentage loss in potential human development or welfare (measured by H_ε) arising from inequality.

2.7.1 The Inequality Adjustment. The above expression for A_ε can be stated equivalently as:

$$H_\varepsilon = H_0(1-A_\varepsilon)$$

In other words, the IHDI is the HDI adjusted by the inequality with which achievements are distributed, as measured by the Atkinson inequality measure. For example, in the case of $\varepsilon = 1$, the inequality measure is $A_1 = 1 - g/\mu$, where g is the geometric mean (or actual human development) and μ is the arithmetic mean (or the highest potential human development) associated with the achievement matrix,[16] and the associated IHDI can be expressed as $H_1 = H_0(1\text{-}A_1)$. In what follows, we focus on H_1 as a key example of the IHDI class, and also one that has particularly useful interpretations and properties.

2.7.2 Which Inequalities? H_1 accounts for inequalities both *within* each dimension and *across* dimensions. To see this, let x_1, x_2 and x_3 be the three columns of the achievement matrix X, giving respectively the distributions of income, education and health across people.[17] Then recalling the second definition of the IHDI, we see that

$$H_1(X) = g(g(x_1), g(x_2), g(x_3))$$

where $g(x_i) = \mu(x_i)[1\text{-}A_1(x_i)]$. In other words, H_1 is the geometric mean of the geometric means of income, education and health, and each of the latter can either be calculated directly from the data or constructed from the arithmetic means and inequality levels.[18] If x_i is distributed equally, then inequality $A_1(x_i)$ is 0 and the geometric mean $g(x_i)$ is the arithmetic mean $\mu(x_i)$; if x_i is unequally distributed, then $g(x_i) < \mu(x_i)$ with the difference being due to inequality in the distribution of x_i. This is how the inequality *within* dimensions is incorporated into the H_1.

By the same logic

$$H_1(X) = \mu(g(x_1), g(x_2), g(x_3))[1 - A_1(g(x_1), g(x_2), g(x_3))]$$

and hence H_1 is sensitive to inequality across dimensions. If the three geometric means are equal, then $A_1(g(x_1), g(x_2), g(x_3)) = 0$ and $H_1(X) = \mu(g(x_1), g(x_2), g(x_3))$; if they are unequal, then $H_1(X) < \mu(g(x_1), g(x_2), g(x_3))$ with the difference being due to the inequality across dimensions. This is a second type of inequality – *across* dimensions – that enters into the IHDI. If countries focus on a single dimension of development (say income) to the exclusion of others, this will be reflected in an adjustment of the average $\mu(g(x_1), g(x_2), g(x_3))$ downward in accordance with the inequality across aggregate dimensional achievements.

2.7.3 Suppressing Inequality: An Alternative HDI. In order to assess the impact of within dimensional inequalities on $H_{1,}$ one can define a base human development index H_1* that retains the cross-dimensional inequality but suppresses within-dimension inequalities. Consider the *smoothed* matrix X* in which every entry in x_i is replaced with the dimension i mean $\mu(x_i)$. Define

$$H_1^* = H_1(X^*) = g(\mu(x_1), \mu(x_2), \mu(x_3))$$

which is the level of H_1 human development in the achievement matrix that smoothes out dimensional achievements. Like the HDI, the index H_1* ignores within dimension inequalities; unlike the HDI, it takes into account inequality across dimensions.

Let us suppose that transfers within a given dimension across persons were feasible, but transfers across dimensions were not. Then across all distribution matrices feasible obtained from a given X, the matrix that would lead to the maximum level of H_1 would be X*, the smoothed distribution. Hence $H_1^*(X)$ has the interpretation as a measure of the (maximum) *potential* IHDI level associated with X. This level can then be compared with the actual level $H_1(X)$ of inequality-adjusted human development in X to evaluate the within dimension inequality, namely, $[H_1^*(X) - H_1(X)]/H_1^*(X)$, or the percentage loss in potential IHDI arising from within dimension inequality.

Note that the index H_1^* uses two distinct means – arithmetic and geometric – in its aggregation, and hence the desirable property of path independence is lost. In other words, the potential human development level obtained by first averaging within dimension, and then applying the geometric mean across, is not the same as the level obtained by first applying the geometric mean to assess an individual's development level and then applying the arithmetic mean across the individual levels. This implies that there are two possible ways of evaluating potential human development, and one way must be chosen. As we note below, the absence of individual data with which to compute the latter option leads us to the former option and H_1^*.

2.8 Special Properties of H_1 and H_1^*. As result of the use of the geometric mean in H_1 (the IHDI) and H_1^* (potential IHDI), both measures have key properties and useful interpretations. We now present and discuss some of these attributes.

2.8.1 Individual Scale Invariance. Our basic measurement assumption is that the dimensional variables in X are ratio scales whose levels are comparable across dimension. This allows all the variables to be renormalized by a common factor while preserving the underlying aggregate orderings. On the other hand, a unilateral change in the scale of *single* dimension is *not* included as part of the admissible transformations and, in general, will alter the way that distributions are ordered. For example, if the aggregation function is the arithmetic mean and income is now measured using a scale with units that are doubled, it is as if a weight of ½ were being applied to the original income variable. Equivalently, if the upper bound on the effective range on income is unilaterally made twice as large, the effect is to place a weight of ½ on the income dimension. In general, this single dimensional rescaling will disrupt the ordering of the underlying distribution of human development using H_0.

The situation is very different if we change the scale of a given variable (or, equivalently, the range over which achievements in that dimension are measured) while measuring actual and potential inequality-adjusted human development using indices H_1 and H_1*. Suppose that country A has the distribution (x_1, x_2, x_3) and country B has (y_1, y_2, y_3). Now suppose that instead of the distribution x_1 and y_1 of the first variable these distributions become $x'_1 = \alpha x_1$ and $y' = \alpha y_1$ for some positive α. What happens to H_1 and H_1*? Clearly, the geometric mean (or arithmetic mean) for the first dimension becomes α times the previous level. The overall geometric mean associated with each measure (H_1 and H_1*) then falls to a level of $\alpha^{1/3}$ times its original value; hence, the initial ranking across countries is preserved. In practical terms, this means that the researcher does not need to be overly concerned with the selection of the appropriate scale for each variable in applying the measure. The properties of the measure ensure that the ranking will not be affected by the choice of scale.

The indices H_1 and H_1* have an even stronger invariance property, which preserves *percentage* increases or decreases, and not just the ranking or direction of change. Suppose that country A with the distribution (x_1, x_2, x_3) has a level of H_1 (or H_1*) that is p% above the respective level of country B with (y_1, y_2, y_3). Now if x_1 becomes $x_1' = \alpha x_1$ and y_1 becomes $y_1' = \alpha y_1$, as before, then since the respective aggregate levels are $\alpha^{1/3}$ times their original values, country A’s level is still p% above country B’s. Changing the scale of a single variable has no impact on percentage differences. The practical impact of this observation is that even if the objective of the research is to evaluate *percentage* changes in H_1 (or H_1*), the researcher does not have to be overly concerned with the selection of the appropriate scale for any particular variable.

2.8.2 Independence of Standardized Values. The traditional HDI presentation provides a table of four values for each country: the overall aggregate value and the three individual component values. Each component variable is scaled so that all component values lie between 0 and 1, and thus the aggregate HDI is also in this range.[19] This presentation table can be easily replicated for both actual H_1 and potential H_1* inequality-adjusted human development. One could use the observed maximal values of the arithmetic mean in each dimension for scaling purposes: simply divide all individual achievements in the dimension by this level. The resulting component and aggregate values of H_1 and H_1* would also lie between 0 and 1. In addition, the ranking and percentage comparisons across countries (overall and within each dimension) would be invariant to the choice of scale, so this selection would entail no loss of generality.

There is another way of depicting the data that may be very useful for certain purposes, and this method can best be presented in example form. The tables below give results on H_1 for three countries. The top table presents the data for a traditional scaling of the variables that yields components and overall H_1 level between 0 and 1. The middle table rescales all of the entries of each country's distribution matrix X in such a way that: all values for dimension 1 are divided by the component value for country A; all values for dimension 2 are divided by the component value for country A; and all values for dimension 3 are divided by the component value for country A. In other words, all the data are standardized to country A levels. Then it is clear that H_1 for Country A will also be 1, while the other Countries' values are rescaled as depicted in the table. Each column in the middle table is proportional to the respective column in the top table, indicating that the component and overall H_1 rankings (and percentage comparisons) are unchanged by standardizing by Country A values. And the numbers listed in the middle table have a concrete meaning as the percentage of the respective values for Country A. The final table has the same information using Country B as the basis for standardization.

Original scaling for 0-1 variables

	H_1	g(x)	g(y)	g(z)
Country A	0.69	0.74	0.55	0.80
Country B	0.77	0.66	0.90	0.77
Country C	0.34	0.33	0.54	0.23

Standardized to Country A

	H_1	g(x)	g(y)	g(z)
Country A	1.00	1.00	1.00	1.00
Country B	1.12	0.89	1.64	0.96
Country C	0.50	0.45	0.98	0.29

Standardized to Country B

	H_1	g(x)	g(y)	g(z)
Country A	0.89	1.12	0.61	1.04
Country B	1.00	1.00	1.00	1.00

Country C	0.45	0.50	0.60	0.30

The first table gives the most natural way of reporting overall results for all countries; the other two tables would likely be helpful for individual countries in calibrating their component and overall achievements.

2.8.3 Consistency over Time. As achievements rise over time, the scale may need to be revised if one wants to preserve a particular range (say 0-1) for the component and overall values. The properties of the indices H_1 and H_1* ensure that their rankings and even percentage differences will be unchanged after the updating occurs, whether all variables are rescaled by a common factor or each is rescaled by its own factor. In other words, the conclusions of each subsequent HDR will be consistent with previously obtained results (unless, of course, there are substantive revisions in the data). This consistency has been lacking in previous results and Reports.

2.9 A Digression on the Differential Treatment of Income. The link between the above approach and the traditional HDI methodologies is more easily seen if we take the log transform versions of H_1 and H_1*:

$$\ln H_1(X) = \mu(\ln g(x_1), \ln g(x_2), \ln g(x_3)) = \mu(\ln X)$$

and

$$\ln H_1*(X) = \mu(\ln \mu(x_1), \ln \mu(x_2), \ln \mu(x_3)).$$

The log version of the H_1 can be viewed as a mean of the logarithms of the distribution-specific geometric means, or as a mean of the logarithms of the entries in the distribution matrix X. As noted above, this index is additively decomposable across populations (and dimensions);

however, the interpretation of the values obtained is not so straightforward as for H_1 itself.[20] The log version of H_1* is the mean of the logarithms of the component arithmetic means. It is similar to H_0 in that it is based on component means; however, it applies a log transform of these values before taking their average.[21]

Since 1999 the traditional HDI has varied from the simple "mean of means" form H_0 in a way that is analogous[22] to the following hybrid of H_0 and ln H_1*:

$$H = \mu(\ln \mu(x_1), \mu(x_2), \mu(x_3))$$

where x_1 is the income distribution. The mean income is subjected to a log transform before it is averaged with the other two indicators. What is the effect of a log transformation? What is its purpose? Is this justification unique to this dimension or is it also applicable to the other two dimensions? We now consider these questions.

2.9.1 Impact of the Income Transformation. The log transformation in H ensures that the marginal impact of an additional unit of (per capita) income is higher at low levels and lower at high levels. In other words, the "marginal product" of income as a driver of measured human development diminishes as income rises. More importantly, since the marginal impact of each of the other two achievements is fixed, this implies that the marginal rate of substitution (MRS) between income and a second achievement level decreases as income rises (the amount of education needed to compensate for the loss of a unit of income falls as income rises). While the MRS changes as income changes, it is entirely independent of the levels of the other two achievements.

2.9.2 Justification of the Income Transformation. Anand and Sen justify the asymmetry between income and the other two indicators in this way: “Of the three, both life expectancy and literacy can be seen to be valuable in themselves (even though they may also be useful for pursuing other ends too). Income, however, is quintessentially a means to other ends” (1993:3). The distinction between means and ends is central to human development, and this argument has been consistently carried in the literature on the HDI. The 1990 *Human Development Report* also gave a second argument for the treatment of income: it should be transformed in order to “reflect the diminishing returns to transforming income into human capabilities. In other words, people do not need excessive financial resources to ensure a decent living” (HDR 1990:12).

These general justifications do not point to a particular transformation for income such as the log transformation. Indeed the transformation of the income variable used in the construction of the HDI has not been stable over time, with substantial changes occurring in 1990, 1991, and 1999 (Anand and Sen 2000). The log transformation is a familiar one to economists and is often invoked labor economic discussions of earnings, or as a simple way of converting income into utility. However, no specific justification has been found as to why it should be used, either in its natural logarithm or common logarithm form.

2.9.3 Transforming Health and Education. Do the arguments given above apply only to the income component, or might similar arguments be put forth for the other two dimensions? The fact that income is “quintessentially” a means while health and knowledge are undoubtedly ends as well does not imply that their marginal product must be constant. Indeed there are at least two reasons why the health and education variables should be transformed to reflect diminishing returns.

The first is that the indicators for health and education represent the intrinsic value of these dimensions imperfectly. The 2010 HDI uses years of schooling – a resource – rather than literacy, which seems to reflect a basic functioning. It is not clear how to relate years of schooling to the marginal gain in knowledge that is of intrinsic value, although one could advance arguments that the marginal product at lower levels of education is higher.[23] Also, the knowledge and educational capabilities of people depend fundamentally on the quality of education, the safety of school, home support, and other factors not captured in years of schooling, echoing the way that income only imperfectly connects to final capabilities. Similar arguments could be made for life expectancy, which omits the quality of health, and also only captures one aspect of health. Thus given the obvious data limitations, it is arguable that the two achievements correspond imperfectly and in a concave fashion to capabilities. A second reason is that health and education are not only of intrinsic value; they, like income, are instrumental to other dimensions of human development not included in the HDI (Sen 1999) and their ability to be converted into other ends may likewise incur diminishing returns. Given these considerations, it would seem advisable to apply a transformation to the health and education components of the HDI to reflect diminishing returns. [24]

2.9.4 Which Transformations? The empirical basis for deciding on the relationship between each dimension and human development is not well developed. The above arguments suggest a concave relationship between each dimension and human development, but they do not pinpoint a specific transformation for each. The log transformation advanced for the income dimension would certainly fulfill the basic requirement of diminishing returns. If there were no convincing information differentiating across dimensions, it would be advisable to use a transformation for each component that was substantively the same. If this were done, the measure obtained would be a log transformation of the potential human development index H_1*. Replacing the arithmetic mean of each dimension with the geometric mean would then lead us back to the inequality adjusted human development index H_1.

2.10 Summary In this section, we presented the class of inequality-adjusted human development indices H_ε that are obtained by applying a standard Atkinson ede formula to the matrix of human development achievements. We discussed the properties of the IHDI class and showed how they account for inequality by adjusting the usual HDI downward using an Atkinson inequality measure. We focused on the index H_1 based on the geometric mean and defined an associated "potential human development" index H_1* by ignoring within distribution inequality. H_1* is higher than H_1 and indicates the maximum level of H_1 that would be possible if each dimension were distributed evenly. The special properties of H_1 and H_1* were discussed, including the invariance of their implied rankings *and* the associated percentage comparisons to changing the scale of a single dimension. Likewise, it was shown how achievements can be standardized to component levels in any particular country of interest while preserving all rankings and percentage comparisons. The methods also allow updating of variable ranges over time without disrupting results obtained in previous years. A simple log transform of H_1 and H_1* yields forms that are comparable to the way that the traditional HDI has been implemented using a logarithm of the income component. We discussed the impact and justification of this differential treatment of income and provided reasons for transforming the other two variables. This leads back to the two indices H_1 and H_1* implemented in the next section.

3. Proposed Methods: Implementation

The methodology sketched above directly shapes the treatment of the data, both normatively and technically. This section describes how the methodology can be implemented given existing data possibilities. It links normative choices – such as that of the upper and lower cutoffs – with the human development approach and justifies these. It describes how to interpret the choice of cutoffs, the dimensional indices, marginal rates of substitution, and changes across time. It signals the choices to which the measure is particularly sensitive. It draws attention to technical considerations – such as the need to replace zero and negative values for H_1 and H_1*, identifies assumptions about the data and raises concerns regarding data quality.

3.1 Constructing Variables. Transforming variables from their raw state to normalized variables fitting the above measurement assumptions requires the selection of lower cutoffs (corresponding to 0% achievement levels) and upper cutoffs (corresponding to 100% achievement levels). The resulting achievement index is taken to be linear between the two extremes, so that the selection of cutoffs within this framework effectively calibrates the variables for use in H_1 and H_1*.[25]

The assumption that the normalized variables are measured using a ratio scale such that their levels are completely comparable across dimensions is, indeed, a strong one and difficult to justify within the capability framework, where empirical variables and their underlying capabilities can be linked together in complicated ways. It is adopted for simplicity. The selection of cutoffs can be informed by a capability approach, with the zero point of each indicator interpretable as an equivalently low or absent level of capability and the maximal cutoff occurring at a level of capability achievement that is comparably high for each of the variables. If a dimensional variable has a natural zero level that is readily associated with zero capabilities, then half of the exercise would be complete. Moreover, due to the properties of the H_1 and H_1* indicators, the selection of the upper cutoff is not as important as in the case of H_0 where it effectively influences component weights. To be sure, the choice of upper cutoff will still determine the equivalence scale across dimensions and the nominal values of the component and overall indicator. However, with the geometric mean structure, the rankings and percentage changes will not be affected. One can conclude from this that the choice of upper cutoffs is somewhat less important in this case and that the use of data-generated cutoffs, such as the highest observed geometric or arithmetic mean, might be reasonable as a first approximation.

3.1.1 Lower cutoffs. The lower cutoff is selected to correspond to a zero capability level, and there are at least two different viewpoints on what zero capability entails. First, there may be a natural zero point in the raw data; second, there may be a nonzero value that corresponds to a state in which life becomes unsustainable.[26]

For education, the combined mean years of schooling and school life expectancy would appear to have a natural zero corresponding to non-attainment of any education at all. Of course, this variable is not a perfect reflection of underlying knowledge: schooling does not reflect learning outside school, nor the quality of school, nor different learner achievements in the same school. Furthermore the constructions of school life expectancy and of mean years of schooling will surely be measured with some error. However, given the variable being used, a lower cutoff at zero would appear to be a reasonable choice.[27]

In the dimension of health, one natural lower cutoff for life expectancy is zero. However this is not the only possibility. Recall that life expectancy data are fundamentally group-based data, drawn from life tables rather than an individual's own objective state. If a society or a subgroup from society has a life expectancy below the normal age of reproduction, that society or group would tend to diminish; an inability to transition into adulthood dampens the very basic capability to survive and also has a cost in terms of the survival of the species. One could imagine, then, selecting a lower cutoff at that reflects the expectation that a person will not survive into adulthood. This points to an alternative cutoff in the range of 14 to 20 years although this may vary across societies and time. The 2010 *HDR* uses a lower bound of 20 years for aggregate HDI data.

For the income dimension, one lower cutoff is the natural zero point in income (or consumption) space. This is indeed the lower cutoff used in many evaluations of wellbeing and inequality; salience and consistency with this literature argue in favor of this choice. A second possible way of constructing a lower cutoff in resource space is to consider some percentage of a minimum level of income. The 'zero' point in capability space of command over resources might be calibrated via an extreme poverty line equivalent to a food basket providing minimum caloric intake for short term sustainability or via some higher line that includes an increment for basic shelter and clothing for the longer term.[28] To the extent that the data do not include in-kind income, self-production, and public provision, there is greater justification for locating the cutoff at a lower value and perhaps even at the actual zero in the data. The 2010 HDR uses the lower bound of $163 for aggregate income.[29]

The precise choice of lower cutoffs may well have a large influence on the subsequent measured levels of human development, so the choice should not be taken lightly. And regardless of the particular selections, tests should be run to evaluate the robustness of conclusions to lower cutoffs.[30] On balance, we would suggest that the lower cutoffs be set to coincide with the zeroes of the respective variables. To be sure, this may not be an entirely accurate way of calibrating a zero level of capability, especially in the health dimension. However, there is considerable value in selecting lower cutoffs that are focal points in the analysis (as zero income and zero education level are) and are commonly adopted in other analyses of dimension-specific achievements. For example, being able to directly link income inequality as usually measured in income space to the loss in human development can help clarify the impact of various the policy options. In the other direction, the already salient notion of income inequality can be directly invoked in subsequent discussions of human development.

3.1.2 Upper cutoffs. The final step in constructing the ratio scale variables for measuring human development is to select upper cutoffs for normalizing the observed data. Actually, the key part of this step is to select values (whether sufficiently high or not) that reflect comparably fulfilling levels of achievement, in capability space. These can then be converted to upper cutoffs by scaling all three values up (or down) until the observed values fall the below the cutoffs by a desired amount. The appropriate and comparable ratio scale variables are constructed by dividing each variable by its upper cutoff.

Note that there is some ambiguity here between selecting a cutoff value above the highest country (or relevant subgroup) aggregate and selecting a value above the highest individual value. We would argue that the former is the correct approach: given the inevitable inequality with which each variable is distributed, adopting the more restrictive individual perspective would compress the aggregate human development levels into a much smaller range in [0,1]. There is benefit from having a larger effective range from [0,1] for comparing component values and for evaluating overall values. Moreover, the comparability across dimensions is preserved by a simultaneous rescaling; as are the within-component rankings and aggregate rankings for all H_ε (as well as for H_1*). Thus, once the initial comparable levels of the three variables have been found, the selection of upper cutoffs can be made as desired without loss of generality.

One approach employed in constructing the HDI uses maximum observed countrywide averages across all countries and across a time span of 30 years as both the comparable and the upper cutoff levels. Using an empirical maximum effectively regards a country or person who has reached that level of achievement as fulfilled and those having a lesser amount as proportionally lacking. The difficulty with this approach is that it uses an observed value to set a variable that should be justified normatively. In particular, comparability across dimensions is not assured. An 80% achievement level for one dimension need no longer be equivalent to an 80% level for another; this simply means that the achievement is 80% of the highest observed value in the given dimension. Nonetheless, in the case of H_1 and H_1*, rankings and percentage changes will be robust to this choice, so there is some cost but much convenience in taking this route.[31]

Another way forward, which has some appeal with respect to other simultaneous conversations related to the environment and to wider approaches to measuring wellbeing, is to work to identify an upper cutoff that normatively reflects *sufficient* achievement in each dimension.

Bhutan's Gross National Happiness Index applied such sufficiency cutoffs across all variables. Like a poverty line in income space and multidimensional poverty lines across several achievements, a collection of comparable sufficiency standards would have the advantage of focusing on salient levels of achievement – ensuring that normative values underlie comparisons across dimensions. An upper bound could then be set as some common increment of the sufficiency levels. Of course, value judgments about the sufficient (and feasible) levels of income for a society will differ, so implementation may be conceptually and practically difficult. While there would likely be a positive gain in the understanding of the level of each component (both absolute and relative to other levels) and overall levels (reflecting values derived using reasoned processes) given the invariance properties of H_1 and H_1*, there would be no effect on the ranking and percentage changes across countries and across time. We therefore suggest using the existing method and recommend further study of the normative setting of the upper cutoffs. In the 2010 *HDR*, the upper cutoff is 83.2 years for life expectancy, 12.6 years of schooling, and $51,200 per capita.

3.2 Estimating Inequalities. It would be natural to construct the IHDI using similar data to the traditional HDI, with the geometric mean in each dimension being calculated at the country level and then aggregated across dimensions using the geometric mean once again. The resulting IHDI would be a true equally distributed equivalent achievement level – indicating the overall level of achievement in each country, but adjusted for inequality.

However, data underlying the three components of the traditionally estimated HDI are not typically available at the population level. The income per capita figures are drawn from national accounts which clearly do not contain information on distribution. Life expectancy data are not available at the individual level, nor are they available for all relevant population subgroups. Data on years of schooling *are* typically available but not in combined form with the school life expectancy data that can be used to account for the achievements of children in school. The practical problem then is how to estimate the IHDI across a broad range of countries – say more than 140 countries – when the standardly used data do not allow this.

This section addresses the problem of constructing the H_1 and, in particular, the component geometric means, when the base data are not sufficiently disaggregated to permit direct computation. There are two general approaches, both of which involve the estimation of inequality as measured by $A_1 = 1 - g/\mu$. The first approach estimates A_1 using household surveys. The second approach uses aggregate data over population subgroups to obtain a lower bound estimate for A_1. Given the estimates for A_1 within each dimension, the component average can be adjusted to the respective ede or geometric mean using the formula given above and then the overall geometric means of geometric means or H_1 can be obtained. The basic and important assumption made in each dimension is that the inequality in the chosen variable is an acceptable proxy of inequality in that dimension, or, at the very least, moves in a similar direction.

3.2.1 Income. Data on income or on consumption drawn from nationally representative household surveys are used to estimate inequality A_1 for the income dimension. Income would seem to correspond more naturally to the underlying notion of "command over resources" (Foster and Szekely). Consumption data are arguably more accurate in developing countries, are less skewed by high values, and reflect the conversion of resources directly (Atkinson, Grosh and Glewwe). Income data also pose technical challenges because of the greater presence of zero and negative values. In an ideal world, one would be consistent in the choice of income or consumption data to estimate inequality. However in order to obtain sufficient country coverage, it will be necessary to use both types of data. Either income can be converted into an approximation of consumption, or the actual values can be used, and the statistical tables can indicate the type of data used. It should be noted though, that the final inequality estimates will likely be influenced by: a) whether the data are income or consumption; b) the quality and accuracy of the data; c) the extent to which the data actually represent command over resources in a given society.

3.2.2 Education. The education variable for the 2010 report is a combination of two variables. First, for the adult population, the years of schooling completed; second, for children, school life expectancy. For simplicity, the estimate of inequality in education is based only on the first: the distribution of years of schooling across the population, drawn from nationally representative household surveys. To be sure, years of schooling of adults is not an adequate or full measure of knowledge; it does not include quality of education, access to information and technology, or the year requirements for comparable degrees (countries vary widely in the time it takes to obtain the same qualification). Furthermore, in some countries we do not have years of schooling but only broader levels of schooling; years of schooling must be assigned to the levels data, which clearly masks within-level inequality. Also, measuring inequality in knowledge using years of schooling is subject to other criticisms, particularly related to very high achievements. A person who takes 6 years to do a PhD (or a country in which the average year requirement for a PhD is 6 years) has, by the "years of schooling" measure, a greater achievement than a person (or country) who completes an equivalent PhD in 4 years. It might be desirable to use actual years of schooling through some level, such as up to fifteen years, then to standardize years of schooling according to the degree(s) obtained.

3.2.3 Health. The most difficult domain in which to measure inequality of achievement is health. Child mortality data, often used to represent health inequality in developing countries, are not available for all countries. Nor is there an alternative international indicator of general health functionings which is regularly updated and of which we could take a general mean. Life expectancy data are commonly used in aggregate indicators but are not available at the individual level, nor by population subgroups in all countries. It is possible to estimate a lower bound of inequality by constructing the arithmetic and general means of the distribution of life expectancy for different age cohorts of the population, relying on data from life tables. Of course, this measure of "between-group" inequality is only as accurate as the tables from which it is drawn. It smoothes inequality within each age cohort used, and it does not reflect disability or morbidity – only the existence of physical life. But given the absence of other data sources with sufficient coverage across countries, this seems the approach that will generate the most realistic approximation of health inequality.

3.3 Examples

Let us construct the IHDI by first constructing the HDI with aggregate data and then applying the inequality adjustment from unit level data to its components. Suppose that two countries, Norway and Haiti, have the achievements noted below:

	Indicators			
	Life Expectancy at Birth	**Mean years of schooling**	**Expected years of Schooling**	**GNI per capita**
	(years)	(years)	(years)	(PPP US$)
Norway	81.0	12.6	17.3	58,810
Haiti	61.7	4.9	6.8	949

Suppose the cutoffs are defined as follows:

	Upper	Lower
Income	51,200 $	0[32]
Life Expectancy	83.2 yrs	0[33]
Mean Yrs School	12.6 yrs	0
School Life Expectancy	20.5 yrs	0

We might construct three dimensional indices as follows.[34]

$$\mu_x = \frac{I}{51200}$$
$$\mu_z = \frac{H}{83.2}$$
$$\mu_y = \left(\frac{\left(\frac{E_m}{12.6}\right) + \left(\frac{E_{sl}}{20.5}\right)}{2} \right)$$

Example: *Haiti*

$$\mu_x = \frac{949}{51200} = 0.019$$
$$\mu_z = \frac{60.7}{83.2} = 0.741$$
$$\mu_y = \left(\frac{\left(\frac{4.9}{12.6}\right) + \left(\frac{6.8}{20.5}\right)}{2} \right) = 0.361$$

Next, we adjust these normalized arithmetic means by estimated inequality across the distribution to obtain an estimated geometric mean as follows:

$$\hat{g}_x = \mu_x(1 - \hat{A}_x)$$
$$\hat{g}_y = \mu_y(1 - \hat{A}_y)$$
$$\hat{g}_z = \mu_z(1 - \hat{A}_z)$$

Example: *Haiti*

$$\hat{g}_x = 0.019 \times 0.748 = 0.014$$
$$\hat{g}_z = 0.741 \times .672 = 0.498$$
$$\hat{g}_y = 0.361 \times .587 = 0.212$$

Using the estimated geometric means we can now create H_1 by taking the geometric mean across them:

$$H_1 = \left(\hat{g}_x \times \hat{g}_y \times \hat{g}_z\right)^{1/3}$$

Haiti

$$H_1 = (0.014 \times 0.498 \times 0.212)^{1/3} = 0.114$$

Similarly, *Norway*:

$$H_1 = (0.865 \times 0.935 \times 0.965)^{1/3} = 0.901$$

H_1 is the Inequality Adjusted HDI, IHDI.[35] To interpret IHDI, we remove inequality within each dimension to create H_1*. H_1* reflects the *potential* HDI that a country could enjoy if it distributed its current mean achievements equally across the population.

$$H_1^* = \left(\mu_x \times \mu_y \times \mu_z\right)^{1/3}$$

$$H^*_{Norway} = (1.0 \times 0.974 \times 0.924)^{1/3} = 0.965$$

$$H^*_{Haiti} = (0.019 \times 0.741 \times 0.361)^{1/3} = 0.171$$

The value of the IHDI becomes apparent when we compare H_1 and H_1*

Norway: $H_1 = 0.901$ $H_1^* = 0.965$

Haiti: $H_1=0.114$ $H_1^* = 0.171$

The percentage change between H_1 and H_1* represents the inequality adjustment:

Haiti's actual IHDI is 66% of its potential HDI; Norway's is 93%. Thus Norway, having far less inequality across its distribution, is far closer to achieving its potential HDI given its current achievement levels than Haiti is.

3.4 Robustness Analysis

As was mentioned above, it would be useful to subject the final IHDI to a series of stress tests in order to ascertain its robustness or sensitivity to certain aspects of its composition. These might include:

- Sensitivity to a change in the lower bound (e.g., of 15 vs 20 years for LE)
- Sensitivity to a change in the upper bound (e.g., of GNI)
- Sensitivity to transformations of income (e.g., using log GNI)
- Sensitivity to alternative forms of generating the educational index (using arithmetic vs geometric mean of educational achievements)
- Sensitivity to choice of replacement for zero and negative values (e.g. by adding 0.1 rather than 1 year to years of schooling. ($E_M=E_o+1$; $E_{M'}=E_o+0.1$)

4. Conclusions and Future Research

Many of the issues for further research relate to challenges in implementing the IHDI for a large set of countries. Given data limitations, choices must be made among imperfect alternatives, using a combination of empirical investigation and normative rationale. The aim is to apply the measurement technique in a way that is as accurate as possible and is distorted by data constraints as little as is feasible. This section briefly identifies some of the crucial issues that require further investigation.

4.1 Income versus Consumption. A key question for IHDI is whether to use income or consumption or some related measure such as an asset index as the variable whose inequality best represents relevant inequality in the dimension of living standards or ‘command over resources’. A related question is how to include countries that do not report that variable. For example, if consumption is the preferred variable but some country only collects income data, should income data a) be converted into consumption or b) used directly or c) should that country be excluded? In practice, income data are commonly collected by some countries and regions and consumption data in others. As is well known, the mean and distribution of these data differ, and the definitions employed in distinct consumption and income surveys differ. This is a challenge for comparative analyses.

4.2 Data Quality for Health and Education. Data constraints twenty years after the HDI was launched remain surprisingly pervasive. Ideally, the health indicator used would reflect health functionings more generally and not just longevity. Similarly, the education variable would reflect quality of schooling and the knowledge attained, not just years of schooling. Data constraints prevent improvements in the global HDI. However it would be extremely useful to explore enriched indicators in a subset of countries for which better data are available. For example, it would be useful to compare the Atkinson measure of inequality in alternative variables, such as the 2000 *World Health Report’s* measure of health inequality for 191 countries based on the risk of child mortality, the under-5 child mortality rates that are often used to reflect health inequality, and multidimensional health indices constructed at the individual or household level. Such comparisons would clarify if the chosen HDI variables departed in predictable ways from alternative justifiable measures. This knowledge would both clarify the strengths and weaknesses of the IHDI indicators, and also contribute to the following research question. An additional concern is the accuracy of the data. For example, the mean years of schooling and the related Atkinson measure of inequality for different datasets for the same country and similar years can vary, and scrutiny of these divergences is warranted.

4.3 Inequality of Proxy Versus Inequality in Dimension. The IHDI intends to measure inequality in some domain such as health. The assumption is made that inequality in the focal variable (life expectancy for different cohorts, years of schooling) can be used as a proxy of inequality in capabilities related to health and education more generally. It is necessary to subject this assumption to further conceptual scrutiny and to empirical tests. It could be interesting to compare the Atkinson inequality measures obtained from those data with Atkinson inequality measures for alternative indices of health functioning and relatedly for knowledge-related functionings and for wealth and permanent income, to see the extent of similarity between inequality using alternative variables.

4.4 Zero Value Replacements and Robustness. The geometric mean is highly sensitive to the lowest values in the distribution, and particularly to the lower bound. When data are well defined, this is a signal strength of the measure: it emphasizes the situation of the poorest poor. However in situations where data do not have a natural zero, or where the lowest values are not well defined, the sensitivity of the final measure to these values is problematic. The IHDI in its current form is not immune from problems. Income data have zero and negative values, which must be replaced by some low value, and the final inequality measure will be sensitive to those replacement values. A similar situation exists for years of schooling, in which many zero values are present. The current IHDI has made some specific replacements; however, sensitivity analysis reveals that the rankings are indeed changed by different replacements. Hence while the zero value replacement should be informed by the sensitivity tests, it may best be chosen by normative logic in order to create a ratio scale variable.

For example, in income space, the present procedure replaces the zero and negative values by some fixed amount such as the income of the person at the 2.5 percentile of the population, whatever that may be. Alternatively, the replacement could be by a fixed value that is translated into the comparable PPP value for each country, or by a fixed nominal value. Careful consideration of this issue and of parallel issues in education is warranted.

4.5 Natural Zero, Normative Cutoffs, and the Ratio Scale Assumption.

It is crucially important to choose the lower cutoff well – given that the geometric mean relies on a ratio scale assumption. The lower cutoff is selected to correspond to a zero capability level. The upper bound is designed to obtain comparability across components. The assumption of the ratio scale means that the comparability across dimensions travels down such that 50% of the highest in one variable is the same as 50% of the other two. In combination with the geometric mean, the ranking of countries is not dependent upon the particular upper cutoff chosen. Doubling the upper cutoff is like cutting the variable in half, but since the geometric mean is multiplicative, this means that the original geometric mean is just a constant multiplied by the new geometric mean and the ranking is preserved. We have discussed choosing natural zeros in life expectancy and in income to, perhaps, depart from the zero value observed in the data. Because of the importance of such choices, they must be subject to strict consideration.

A further assumption implicit in the discussion above is that each variable measures capabilities in a linear fashion. This assumption may require further analysis or scrutiny. In particular, there may be grounds for considering a more concave transformation for the income variable to reflect its diminishing marginal rate of substitution. This awaits further research.

---

[1] The three dimensions are alternatively called *command over resources*, *longevity*, and *knowledge*.

[2] Anand and Sen (2000).

[3] See Hicks (1997), Sagar and Najam (1998), Anand and Sen (2000), Foster, Lopez-Calva, and Szekely (2005), Grimm, Harttgen, Klasen, and Misselhorn (2008), Seth (2009) among others.

[4] In Atkinson (1970), the income distribution is evaluated with the help of an "equity preferring" welfare function. The welfare level of an initial unequal distribution is less than the welfare level of the completely equal distribution having the same mean. By scaling down the completely equal distribution one lowers welfare and eventually finds a level of per capita income, which when distributed equally, is equivalent in welfare to the original distribution. This is the "equally distributed equivalent" income level.

[5] The assumptions are never made clear either by Foster, Lopez-Calva, and Szekely (2005) or in the various Human Development Reports. Below we discuss the transformations used in practice.

[6] A ratio scale requires the variable to have a natural 0 and requires ratios of values to be meaningful.

[7] Given the assumption on comparability across dimensions, one cannot simply renormalize individual dimensions independently. In general, using a different upper bound for different dimensions can lead to reversals in rankings with respect to overall human development. As we will note below, there are special cases where the ranking is independent of the choice of upper bound, which could be a favorable property in the event that the selection of bounds (or equivalently, of the comparability relationship across dimensions) was subject to error.

[8] Idealized in that we have data on individual achievements that have been appropriately scaled. We are also ignoring the fact that income is treated somewhat differently in the classic implementation of the HDI. See section 2.9.

[9] The ede was defined in Atkinson (1970), who explicitly allowed the no inequality aversion case.

[10] Other names for this aggregation method include "general mean" and "r-order mean". See Hardy, Littlewood, and Polya (1967), Marshall and Olkin (1979), Blackorby, Donaldson, and Auersperg (1981), Foster and Sen (1997) and Foster, Lopez-Calva, and Szekely (2005).

[11] The classic implementation of the HDI treats income somewhat differently. See section 2.9.

[12] A second consideration for using a path independent measure is that information on income, health and education are typically not available in the same dataset; path independence is allows the measure to be calculated from different sources and yet be interpretable as an average level of individual human development in society.

[13] An alternative approach of Anand and Sen (1994) and Hicks (1999), uses the Sen welfare function within each dimension and the arithmetic mean across dimensions, but is not path independent nor subgroup consistent. See also Foster, Lopez-Calva and Szekely (2005) for the other properties satisfied by the IHDI class.

[14] The Atkinson (1970) parametric class of inequality measures is arguably more intuitive than the decomposable generalized entropy measures. The decomposable FGT poverty indices, may well be less intuitive than a simple transformation; see Foster, Greer, Thorbecke (2010).

[15] The key reference in measurement theory is Roberts (1979).

[16] A_1 is also a transformation of Theil's second measure of inequality, also called the mean log deviation.

[17] Note that the variables may have to be changed to move from population-based indicators to individually distributed indicators of wellbeing. See the discussion below.

[18] This expression in terms of the dimensional arithmetic means and dimensional inequality measures facilitates computation of the IHDI in practice. See section 3 below.

[19] The resulting arithmetic sum is clearly and crucially dependent on the particular scale chosen, which is clearly an important normative choice for the HDI. However, the choice of scale is apparently being set with regard to the highest observed country levels in a given year without consideration being given to whether the implied variables are normatively meaningful.

[20] If the original variables were normalized to 0-1, then ln H_1 would be negative. If they were normalized to 0-100, the value of ln H_1 would likely be positive, but would not be constrained to fall below 1.

[21] An alternative implementation based on this formula is discussed in section 4 below.

[22] The analogy will not be exact since the method of constructing “goalposts” importantly alters the picture. This is discussed at greater length in section 4.

[23] See for example Murthi, Guio and Dreze (1995).

[24] Each of these empirical relationships merit careful scrutiny, as the relative impact of different achievements in income, health and education on achievements in other dimensions is likely to vary depending upon the other achievements present, the time period considered, and so on.

[25] Calibrating pairs of points does not by itself lead to complete comparability of dimensions. We are assuming here that only affine transformations from the raw data to the normalized variable are admissible. Alternative transformations might be admitted at

this stage, but it is important to remember that H_1 will effectively be transforming the normalized variables by the log before aggregating them.

[26] This latter term suggests a dynamic justification, which in turn requires a horizon to be specified. Longer-term sustainability could require a higher achievement level.

[27] Note that data concerns for education and the other variables are detailed in section 3.2, and hence are not mentioned further in this section.

[28] There may be some conceptual overlap or duplication between a “survival” poverty line in resource space (seen as a means) and a lower cutoff in health space (seen as an end).

[29] There are significant conceptual problems in adjusting individual data to non-zero ‘natural zeroes’. There is the question of consistency: if numerous individual observations fall below the ‘natural zero’ and these observations are taken to be valid, this would call into question the validity of that designation as a ‘natural zero’. If the observations are considered to be errors, then there is the question of how to transform the observed data to deal with this. One possibility would be to replace all values equal to and lower than the natural zero with a positive value. For example the 2010 HDI income index uses a ‘natural zero’ of $163, so a consistent approach would be to subtract $163 from each observed value, then replace all zero or negative values with a small increment above that value. The same procedure could be followed for life expectancy using a ‘natural zero’ of 20 years in cohort data. Note that these adjustments will alter both the arithmetic and the geometric means in unpredictable ways, which places into doubt the use of untransformed data to estimate the inequality level.

[30] We include appropriate tests below. Setting the lower cutoff is especially important in the case where H_1 (or H_1*) is used since it serves to establish the base from which all variables are measured. If the cutoff is lower, the initial percentage values are lower and it takes a larger change in the untransformed variable to achieve a 1% change in the transformed variable.

[31] Note that the upper and lower cutoffs of the usual implementation of the HDI have both been set using the data-driven method. Hence the same critique – that the same value across different dimensions may not be normatively comparable – applies. Moreover, the arithmetic mean does not share the invariance properties of the geometric mean, and hence the robustness of rankings to specific cutoffs is also problematic.

[32] Note that the 2010 HDI uses a lower bound of $163 on income data.

[33] Note that the 2010 HDI uses a lower bound of 20 years on aggregate data.

[34] Note that the 2010 HDI aggregates the two educational variables by taking the geometric rather than the arithmetic mean across them.

[35] Note that the 2010 HDR has one further adjustment: it applies the ratio H_1/H_1^* to the 2010 HDI to generate an inequality-adjusted index that relates to the HDI. The 2010 HDI is constructed as the geometric mean of normalized indices for health, education and the *log* of income.

References

Alkire, S. (2002). ‘Dimensions of human development’, *World Development*, 30(2), 181-205.

Anand, S., & Sen, A. (2000). ‘The income component of the human development index,’ *Journal of Human Development*, 1, 83–106.

Anand, S., & Sen, A. (1997). ‘Concepts of human development and poverty: A multidimensional perspective’. In *Poverty and Human Development: Human Development Papers 1997*, New York: UNDP, 1-20.

Anand, S., & Sen, A. (1995). ‘Gender inequality in human development: theories and measurement’, Occasional Paper 19. New York: UNDP.

Anand, S., & Sen, A. (1994). ‘Human Development Index: Methodology and Measurement’, *Human* Development Report Office Occasional Paper. New York: UNDP, 138-151.

Atkinson, A.B. (2009). ‘Income, Health and Multi-dimensionality’. In K. De Boyser, C. Dewilde, D. Dierckx, and J. Friedrichs. *Between the Social and the Spatial: Exploring the Multiple Dimensions of Poverty and Social Exclusion*. Farnham: Ashgate, 21-34.

Atkinson, A. B. (1970). ‘On the measurement of inequality’, *Journal of Economic Theory*, 2, 244-263.

Atkinson, A. B., Atkinson, T., Marlier, E., Nolan, B., & Vandenbroucke, F. (2002). *Social Indicators: The EU and Social Inclusion*. Oxford: Oxford University Press.

Atkinson, A. B., & Micklewright, J. (1983). ‘On the reliability of income data in the family expenditure survey’, *Journal of the Royal Statistical Society*, 146(1), 33-61.

Atkinson, A. B. & Zaidi, S. (2002). ‘Guidelines for constructing consumption aggregates for Welfare Analysis’, LSMS Working Paper 135. Washington DC: The World Bank.

Bardhan, K., & Klasen, S. (1999). ‘UNDP's gender-related indices: A critical review’, *World Development*, 27(6), 985-1010.

Booysen, F. (2002). ‘An overview and evaluation of composite indices of development’, *Social Indicators Research, 59*, 115–151.

Cahill M. B. (2005). ‘Is the human development index redundant?’ *Eastern Economic Journal*, 31(1), 1-5.

Chakravarty, S.R. (2003). ‘A generalized human development index’, *Review of Development Economics*, 7, 99–114.

Chatterjee, S.K. (2005). ‘Measurement of human development: an alternative approach’, *Journal of Human Development*, 6, 31–53.

Cherchye, L., Ooghe, & E. Puyenbroeck, T.V. (2008). ‘Robust human development rankings’, *Journal of Economic Inequality*, 6, 287–321.

Cherchye, L., Moesen W., & Van Puyenbroeck T. (2004). ‘Legitimately diverse, yet comparable: on synthesizing social inclusion performance in the EU’, *Journal of Common Market Studies*, 42, 919–955.

Decanq, K. (2009). ‘Copula-based Measurement of Dependence Between Dimensions of Well-being’, HEDG Working Paper 09/32. University of York.

Decanq, K., Decoster, A. & Schokkaert, E. (2009). ‘The Evolution of World Inequality in Well-being’, *World Development*, 37(1), 11-25.

Dar, H.A. (2004). ‘On making human development more humane’, *International Journal of Social Economics,* 31, 1071–1088.

Desai, M. (1991). Human development: concept and measurement, *European Economic Review,* 35, 350-357.

Despotis, D.K. (2005). ‘A reassessment of the human development index via data envelopment analysis’, *International Journal of Social Economics,* 56, 969–980.

Dowrick, S., Dunlop, Y., Quiggin, J. (2003). ‘Social indicators and comparisons of living standards’, *Journal of Development Economics,* 70, 501–529.

Erikson, R. (1993). ‘Descriptions of inequality: the Swedish approach to welfare research’. In M. Nussbaum & A. Sen (Eds.), *The Quality of Life*. Oxford: Clarendon Press, 67-83.

Fields, G. S. (1994). ‘Data for measuring poverty and inequality changes in the developing countries’, *Journal of Development Economics*, 44(1), 87-102.

Fleurbaey, M., Hagneré, C., Trannoy, A. (2003). ‘Welfare comparisons with bounded equivalence scales’, *Journal of Economic Theory*, 110, 309–336.

Foster, J. McGillivray, M. & Seth, S. (2009). ‘Rank robustness of composite indices’, Working Paper 26. Oxford Poverty and Human Development Initiative, University of Oxford.

Foster, J., Lopez-Calva, L., & Szekely, M. (2005). ‘Measuring the distribution of human development: methodology and an application to Mexico’, *Journal of Human Development*, 6(1), 5-29.

Foster, J. and A. Sen (1997). ‘After a quarter century’, In A. K. Sen (Ed.), *On Economic Inequality*. Oxford: Clarendon Press, Appendix.

Fukada-Parr, S., & Shiva Kumar, A.K. (2003). *Readings in Human Development.* Oxford University Press, Oxford.

Fukada-Parr, S. (2001). ‘Indicators of human development and human rights – overlaps, differences and what about the human development index?’ *Statistical Journal of the United Nations*, 18, 239–248.

Grimm M., Harttgen, K., Klasen, S. & Misselhorn, M. (2008). ‘A Human Development Index by Income Groups’, *World Development,* 36.

Haq, M. ul (1995). *Reflections on Human Development*. New York: Oxford University Press.

Hardy, G. H., Littlewood, J.E. & Polya, G. (1967). *Inequalities.* Cambridge: Cambridge University Press.

Hicks, D. A. (1997). ‘The inequality adjusted human development index: A constructive proposal’, *World Development*, 25, 1283-1298.

Hopkins, M. (1991). ‘Human development revisited: A new UNDP report’, *World Development*, 19, 1469–1473.

Jolly, R. (2004). ‘Global development goals: the United Nations experience’, *Journal of Human Development and Capabilities*, 7(1), 69-95.

Kelly, A. (1991). ‘The Human Development Index: Handle with Care’, *Population and Development Review*, 17, 315-324.

Lind, N. (2004). ‘Values reflected in the human development index’, *Social Indicators Research,* 66, 283–293.

Lind, N. C. (1992). ‘Some thoughts on the human development index’, *Social Indicators Research*, 27(1), 89-101.

Lüchters, G., & Menkhoff, L. (2000). ‘Chaotic signals from HDI measurement’, *Applied Economic Letters*, 7, 267–270.

Mahlberg, B., & Obersteiner, M. (2001). ‘Remeasuring the HDI by data envelopment analysis’, *International Institute for Applied Systems Analysis Interim Report,* 01-069.

Marshall, A. W., & Olkin, I. (1979), *Inequalities: Theory of Majorization and Its Applications*. New York: Academic Press.

Mazumdar, K. (2003). ‘A new approach to human development index’, *Review of Social Economy*, 61, 535–549.

McGillivray, M. (1991). ‘The Human Development Index: yet another redundant composite development indicator?’ *World Development*, 19(10), 1461-1468.

McGillivray, M., & Shorrocks, A. (2005). ‘Inequality and multidimensional well-being’, *Review of Income and Wealth*, 51(2), 193-199.

McGillivray, M., & White, H. (1993). ‘Measuring development? The UNDP's human development index’, *Journal of International Development*, 5(2), 183-192.

McNeill, D. (2007). ‘Human Development: the power of ideas,’ *Journal of Human Development and Capabilities*, 8(1), 5-22.

Murthi, M., Guio, A.-C., & Dreze, J. (1995). ‘Mortality, Fertility and Gender Bias in India’, *Population and Development Review*, 21: 745–82.

Nardo M., Saisana M., Saltelli A., Tarantola S., Hoffman A., & Giovannini E. (2005). ‘Handbook on Constructing Composite Indicators: Methodology and User’s Guide’, OECD Statistics Working Paper (JT00188147).

Noorbakhsh, F. (1998). ‘A modified human development index’, *World Development*, 26, 517–528.

Ogwang, T. (1994). ‘The choice of principal variables for computing the Human Development Index’, *World Development*, 22, 2011–2014.

Panigrahi R., & Sivramkrishna, S. (2002). ‘An adjusted human development index: robust country rankings with respect to the choice of fixed maximum and minimum indicator values’, *Journal of Human Development*, 3(2), 301–311.

Ranis, G., Stewart, F., & Samman, E. (2006). ‘Human development: beyond the human development index’, *Journal of Human Development and Capabilities*, 7(3), 323-358.

Robeyns, I. (2003). ‘Sen's capability approach and gender inequality: selecting relevant capabilities’, *Feminist Economics*, 9(2), 61-92.

Robeyns, I. (2006). ‘Measuring gender inequality in functionings and capabilities: findings from the British Household Panel Survey’. In P. Bharati & M. Pal (Eds.), *Gender Disparity: Manifestations, Causes and Implications.* New Delhi: Anmol Publishers, 236-277.

Sagara, A. D., & Najam, A. (1998). ‘The human development index: a critical review’, *Ecological Economics*, 25(3), 249-264.

Saisana, M., Saltelli, A., & Tarantola, S. (2005). ‘Uncertainty and sensitivity analysis as tools for the quality assessment of composite indicators’, *Journal of the Royal Statistical Society, Series A*, 168, 1–17.

Sen, A. (2000). ‘A Decade of Human Development’, *Journal of Human Development*, 1(1), 17-23.

Seth, S. (2009). ‘Inequality, interactions, and human development’, *Journal of Human Development and Capabilities*, 10, 375–396.

Smits, J. and Monden, C. (2009). ‘Length of life inequality around the globe’, *Social Science & Medicine*, 68: 1114–1123.

Srinivasan, T.N. (1994). ‘Human development: A new paradigm or reinvention of the wheel?’, *American Economic Review* (Papers and Proceedings), 84, 238–243.

Stewart, F. (1985). *Planning to Meet Basic Needs*. London: Macmillan.

Streeten, P. (1994). ‘Human development: Means and ends’, *American Economic Review* (Papers and Proceedings), 84, 232–237.

Streeten, P. (1995). ‘Human development: The debate about the index’, *International Social Science Journal*, 143, pp. 25–37.

Tisdell, C., Roy, K., & Ghose, A. (2001). ‘A critical note on UNDP's gender inequality indices’, *Journal of Contemporary Asia*, 31(3), 385-399.

Trabold-Nubler, H. (1991). ‘The human development index - a new development indicator?’ *Intereconomics*, September/October.

UNDP. (1990). *Human Development Report 1990: Concept and Measurement of Human Development*. New York: Oxford University Press.

UNDP. (1995). *Human Development Report 1995: Gender and Human Development.* New York: Oxford University Press.

UNDP. (1999). *Human Development Report 1999: Globalization with a Human Face*. New York: Oxford University Press.