2018 UNDP Human Development Report
DISCUSSION PAPER
MARCH 2018

On Human Development Indicators

Marc Fleurbaey

Updated version, October 2019

On Human Development Indicators

Marc Fleurbaey

October 26, 2019

Abstract: The Human Development Index (HDI) has evolved over time and Anand (2018) offers a stimulating critique of various aspects. This paper discusses Anand's critique and, in particular, offers an argument in favor of some degree of complementarity between the dimensions of human development (income, health, education), claiming that this does not undermine the priority that should be given to lives in poor countries, and that it reflects the importance of a balanced development at the societal level as well as the individual level.

1 Social welfare framework

This note discusses the old and new indicators used in the Human Development Report around a change introduced in the 2010 report, and devotes special attention to Anand's (2018) critical review. Anand examines the Human Development Index (HDI), the Inequality Adjusted HDI (IHDI), and the Gender Inequality Index (GII). I will focus here on the HDI and the IHDI, but gender equality will also be discussed briefly at the end.

For this analysis, I will take as theoretical background the welfare economic theory of social welfare. This theory is based on the social welfare function (SWF), which combines individual utilities and social welfare into a hierarchy of nested functions:

$$W\left(U_1\left(x_1\right), ..., U_n\left(x_n\right)\right), \tag{1}$$

where x_i describes the situation of individual $i = 1, ..., n$ in terms of the attributes that are relevant (consumption, non-market advantages, functionings, capabilities...), U_i determines not only how good this is for i but also how much it contributes to social welfare via i, and W synthesizes the whole evaluation, taking account of the distribution across individuals.

The utilitarian tradition in economics has largely adopted a subjectivist approach in which individual utility is understood as subjective well-being, but the formal framework of the SWF is much more versatile than that. The utility function can alternatively be an ethical rod that assesses individual situations in terms of whatever notion of how individual situations matter for social evaluation is adopted by the analyst—including capabilities, for instance.

Two features of the SWF framework are particularly useful:

1) It analyzes social situation in terms of individual situations, which reflects the view that nothing can be good for society unless the members of society are themselves benefitting from it. This nowadays seems natural but it may still happen in the current state of the world that governments seek objectives that have little to do with the well-being of the population (such as territorial expansion).

2) The utility functions embody the degree of complementarity or substitutability of the different components of individual situations (how important it is for an individual to be well off in every dimension), while the social welfare function takes care of inequality

aversion, which is nothing else than complementarity between individual situations (how important it is for society that every individual is overall well off). This division of labor between U_i and W is very important to clarify delicate issues such as the comparative value of life years accruing to different individuals, as we will see shortly.

2 Value of increasing life expectancy in rich and poor countries

As well explained in Anand (2018), the main change between the old and new HDI is the move from an arithmetic average formula

$$HDI_{old} = \frac{1}{3}\left(I_{LE} + I_S + I_Y\right) \tag{2}$$

to a geometric average formula

$$HDI_{new} = \left(I_{LE} \times I_S \times I_Y\right)^{1/3}, \tag{3}$$

where I_{LE}, I_S, I_Y are indexes of life expectancy, schooling, and log of GDP, rescaled to fit the 0-1 interval.

Anand forcefully argues against this change, with a series of arguments discussed in this section and the next one. First, an increase in I_{LE} in a particular country always makes the same contribution to HDI_{old}, whereas with HDI_{new}, this contribution is increasing with $I_S \times I_Y$, making an increment in life expectancy more impactful on human development in a rich and educated country than in a poor, less educated country. Anand defends the old approach on the grounds that life should have the same value for rich and poor, for highly educated and for little educated, and that the new measure goes against the "universality of life claims" and against "the tenets of human development."

Since this is a central issue in Anand's critique of the new indicators, it is worth examining it in detail. The universality of life claims is a beautiful principle in the abstract, but the SWF framework is very helpful to analyze what it can and cannot mean. There is one difficulty that must be cleared, however, before proceeding. When looking at the statistics of the HDI with the SWF lens, one may think of two ways of applying the SWF framework to such data:

a) the "country" approach, which treats every country HDI as a sort of SWF by itself, in which the situation of every resident of the country in terms of longevity, education and income contributes to the national assessment;

b) the "world" approach, which seeks to assess how the world as a whole is going, and treats every country HDI as a utility function, as if every country was like a single agent from this standpoint.

Consider the country approach first. It is at the same time natural and awkward to interpret the HDI as a SWF because in standard welfare economics, the SWF is only used as a tool to rank situations, not to grade them with a number. In contrast, the HDI is used for numerical comparisons across countries, not just for ranking them. Anyway, the universal dignity of every individual is embodied in the SWF if the W function is symmetrical (and of course increasing) in its arguments, the individual utilities, and if the computation of utilities is itself unbiased (and suitably increasing in the attributes that depict individual situations). The "universality of life claims" and "the tenets of human development" can be considered to be quite substantially complied with as soon as these conditions are satisfied. But these conditions do not guarantee that adding one life year to any individual's longevity has the same impact on W for all individuals, as will be clear in a moment. This property, however, is actually satisfied by both HDI_{old} and HDI_{new}, because the contribution of every individual to I_{LE} is the same, since I_{LE} is based on life expectancy, which is a computation of average longevity for an artificial cohort that would endure the current mortality rates by age throughout its existence.

Although such a uniformity of the social value of longevity across individuals is indeed popular in debates, it runs against strong principles of fairness that most people also generally endorse. It is commonplace to consider that if a one-year extension of life can be given to one of two dying patients, a middle-aged one and a very old one, the former should have priority. Even if, for practical reasons, medical doctors are not allowed to discriminate among patients on the basis of non-medical characteristics, health systems do give less priority to old patients.

If the life extension can be guaranteed to be *in the same conditions and same quality of life* (e.g., in a care center), then one could also extend this line of argument and think that priority should be given to the patient who had the worse life so far, either short (this is the age consideration already mentioned), poor, or with a non-controversial disability (e.g., quadriplegia, which, unlike deafness and blindness, is considered a disability by most people). The reason why the extension from age to poverty and disability of such a priority view is tricky is that there are conflicting considerations in real-life conditions: 1) doctors are generally not allowed to discriminate on such a basis; 2) if one considers two dying patients who had the same life up to now, but one will have an extension in good conditions whereas the other will have a lower quality of life, priority should be given to the one who can benefit more. The second point complicates the definition of the degree of priority for the poor and disabled, because, if they had a worse life beforehand, this increases their priority whereas at the same time, if their life extension will be in worse conditions, this decreases their priority.

These considerations are relevant when one takes the world approach when applying the SWF frame to the problem. In this approach, the HDI is now interpreted as a sort of national "utility" contributing to the world welfare. What worries Anand seems to refer to this context rather than the country context. Indeed, the bone of contention is the fact that an additional year of life expectancy in, say, Zimbabwe has less value for HDI_{new} in Zimbabwe, than an additional year of life expectancy in, say, Italy has for HDI_{new} in Italy. In the SWF framework, this would correspond to a situation in which the contribution of a life extension to a poor person's utility would increase that person's utility less than the same life extension would increase a rich person's utility. It is actually very hard to see what is objectionable about such a lack of uniformity. If two individuals with the same utility function had the same life up to their age of death and their life could be extended by one year at that age, it is actually nonsensical to require that it should have the same value for each of them if the quality of life in the extension is different—for instance, if one is healthy and the other is not, or if one is poor and the other is not poor. Since it would be better for each of them to have the extension with the higher quality of life, it would go against commonsense to adopt a utility function that treats each additional year as having the same value independently of quality of life.

This shows that not only is it not shocking if an additional year of life expectancy is adding less value in Zimbabwe than in Italy, but one should actually require it in order to respect the obvious fact that if one had to live an extra year, one would rather have it with the income and education found in Italy than in Zimbabwe. But, then, why does it appear shocking to many commentators (not just Anand) when a life in Zimbabwe appears to have less value than a life in Italy? The SWF framework may help here. There is nothing shocking about saying that the contribution of a year to utility is increasing with the quality of life enjoyed during this year. But it would be shocking to say that saving lives in Zimbabwe is less of a priority than saving lives in Italy. Now, the latter statement has to do not with utility but with social welfare. And when the function W is inequality averse, it gives greater priority to the worse-off individuals. So, consider the Zimbabwean HDI (=utility) which is lower than the Italian's, and an extension of life expectancy in Zimbabwe or in Italy. One has to recognize that it adds less to the Zimbabwean HDI than to the Italian's, for the quality-of-life reason explained earlier, but one can also rely on W to say that in spite of the Zimbabwean HDI increasing less than the Italian's with an additional year, it has greater priority due to its low level. Arguably, it would make a lot of sense to think that saving lives in Zimbabwe is *more*, not less, important than saving lives in Italy, because the Zimbabwean population is so much worse off and deserves a much greater help in all forms, including life extensions.

Going back to the comparison between the old and new HDI, one sees that the new one appears more reasonable than the old one in considering that the value of an extra year of life expectancy in a country increases with the quality of life in this country, as summarized by income and education. This does *not* imply that saving lives in poor countries is less important than saving lives in rich countries. This would imply it only if the global social welfare objective was to maximize the unweighted sum of the country HDIs. But this would be a very unappealing objective. There is a large consensus in the international community that poor countries deserve special help. The economic system and many policy decisions may not reflect this consensus in practice, due to various institutional failures, but this consensus rules out the idea that the sum of HDI should be maximized. Therefore one should not be worried to acknowledge that certain improvements bring less benefits to poor countries. The benefits may be smaller for the "utilities" of these countries, and still be more important in their impact on the global SWF.

In summary, the idea that longevity should have equal social value for the poor and the rich in a country *is* implemented both by the old and new HDIs via I_{LE}, in spite of this uniform value being disrespectful of a basic principle: for individuals who are identical up to the age of death, a life extension has more individual and social value when it is enjoyed in better conditions. Moving the discussion to the world level, when comparing countries for the assessment of their relative priorities in the global social objective, it is similarly required to acknowledge the lower value, for the national HDI, of extensions of life expectancy in poor (and less educated) countries, but this is compatible with giving greater priority to such extensions in poor countries, on the basis of the greater priority of the worse off.

Now, when the UNDP provides the list of country HDI, it does not explicitly assess how much priority should be given to the less developed countries, and this may confuse readers, suggesting that the list by itself plays the role of a SWF evaluating the world situation. But this is not the case, the world evaluation is actually missing! What the Human Development Report does is equivalent to providing the list $(U_1(x_1), ..., U_n(x_n))$ without providing the full evaluation $W(U_1(x_1), ..., U_n(x_n))$. To remedy this confusing limitation, it would be advised to add some discussion of how the world situation should be assessed in a synthetic way, aggregating the low and high development scores with a degree of priority for the worse-off. In this way, readers of the Report would see that even if life has less value in poor countries in terms of national HDI (as it should, since these nations would themselves prefer extending life in the quality-of-life conditions of the developed countries), it can have more value for the global good, leading to subsequent prioritizing of interventions and help.

3 Complementarity of income, education and life expectancy

Another feature of the new HDI that Anand finds questionable is that the marginal contribution of life expectancy to the HDI is decreasing, unlike the old HDI for which it is constant due to the additive formula. This is closely linked to the fact that the perfect substitutability between I_{LE}, I_S and I_Y in HDI_{old} is replaced with a mild degree of complementarity in HDI_{new}. This complementarity means that for any given fixed level of HDI_{old}, the value of HDI_{new} decreases if the inequality between the three components increases. This reflects a preference for balanced development in which income, health and education make progress together, and it penalizes countries in which income is very high (e.g., due to natural resources exploitation) but social policy is underdeveloped and leaves a good part of the population without adequate health care and education. Such complementarity implies that as life expectancy increases, its marginal contribution decreases, because when it is lower than the other indices, it is a very important priority, whereas when it moves ahead of them, it recedes in the priority order.

Anand does not raise strong objections against complementarity, which indeed appears quite appealing in light of the "balanced development" preference that it embodies. But one consequence of complementarity, on which Anand spends a lot of attention, is possible difficulties with zero values. To discuss this issue, let us step back a little. Both the new and the old HDIs are part of the family of generalized means, for which the general formula is:

$$HDI_\sigma = \left(\frac{1}{3}\left(I_{LE}\right)^{1-1/\sigma} + \frac{1}{3}\left(I_S\right)^{1-1/\sigma} + \frac{1}{3}\left(I_Y\right)^{1-1/\sigma}\right)^{\frac{1}{1-1/\sigma}}. \tag{4}$$

The coefficient σ measures the "elasticity of substitution" between the components. The old HDI corresponds to the limit of the formula when $\sigma \to +\infty$, whereas the new HDI corresponds to the limit of the formula for $\sigma \to 1$. It is also worth noting that, based on Atkinson (1970), one can interpret $1/\sigma$ as the coefficient of aversion to inequality between the three components. Inequality aversion and complementarity are, in terms of formal concept, the same thing.

When $1 < \sigma < +\infty$, the formula contains some complementarity and is well defined when any component takes the zero value. But when $\sigma < 1$, corresponding to greater complementarity, zero values are a problem and the formula is defined only for strictly positive values of the three indices. When $\sigma = 1$, the formula of HDI_{new} is still well defined for zero values (not for negative values), but takes a degenerate form since $HDI_1 = 0$ as soon as one component is nill, implying that it is then indifferent to the value of the other components. This property is not necessarily shocking, since one could argue that a country with an extremely low longevity is one in which it is not worth living, and the same can be said if schooling is very low or income is below subsistence. This just carries the logic of "balanced development," mentioned earlier, to the case of extreme values. There may be low performances in certain key domains that make everything else irrelevant because they make life too hard for the population.

But Anand notices that it makes the computation very sensitive to the choice of the levels of longevity, schooling and income that correspond to the zero values for the indices I_{LE}, I_S and I_Y, for countries which are close to these levels, because the value of the HDI sunks quickly when one component index approaches zero. This is indeed a challenge, and it implies that one should either have very strong arguments for the selection of a particular zero value, or perform robustness checks to make sure that conclusions are not sensitive to arbitrary choices of parameters.

This challenge should not be exaggerated, however. Indeed, let us compute the partial

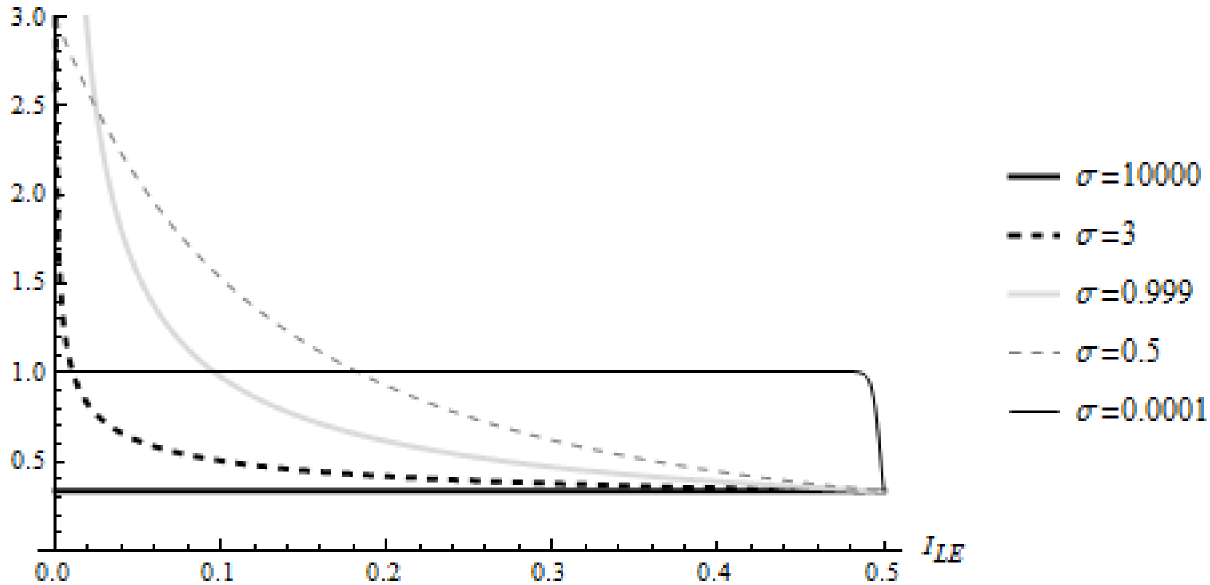

Figure 1: Value of $\frac{\partial HDI_\sigma}{\partial I_{LE}}$ when I_{LE} is small and $I_S = I_Y = 0.5$

derivative of HDI_σ with respect to I_{LE} (similar formulas are obtained for the other components):

$$\frac{\partial HDI_\sigma}{\partial I_{LE}} = \frac{1}{3}\left(\frac{HDI_\sigma}{I_{LE}}\right)^{\frac{1}{\sigma}}. \tag{5}$$

When $\sigma \to +\infty$ (the old HDI), this is always equal to 1/3. When life expectancy is low, so that $I_{LE} < HDI_\sigma$, then the partial derivative is greater than 1/3 but need not become very volatile. For instance, in the case of extreme complementarity ($\sigma \to 0$), then $HDI_\sigma \simeq I_{LE}$ for low values of I_{LE}, so that the partial derivative is close to 1. Figure 1 shows how the partial derivative behaves when I_{LE} approaches zero, for different values of the coefficient σ. The most challenging behavior appears for a value of σ around 1, which corresponds to HDI_{new}. Even in this case, the value of the partial derivative remains below 3 until I_{LE} goes below 0.02.

Moreover, it should be stressed that sensitiveness to low values is unavoidable with complementarity, since complementarity makes it a priority to avoid low performance in any domain, which requires making the indicator very reactive to this domain when low values are observed. Therefore, if $\sigma \leq 1$, which corresponds to a degree of complementarity at least as high as in HDI_{new}, it seems that the zero values should be taken to be exactly those for which the situation is considered so catastrophic that performance in the other domains no longer matters.

If a lower degree of complementarity is chosen, then there is no such thing as a catastrophic level in a domain making other domains irrelevant. Given that longevity and subsistence are preconditions for life, while schooling is a precondition for a dignified life, it seems that a strong complementarity (i.e., at least as strong as with HDI_{new}) is actually recommended for a sensible measure of development. This is also, presumably, what one would obtain if one relied on how the members of the population themselves perceive the value of a good life. For instance, it is not uncommon in the literature (e.g., Murphy and Topel 2006) to adopt utility functions that are multiplicative with respect to longevity:

$$L \times U(c), \tag{6}$$

where L is the number of years lived and c is the level of consumption (assuming it is constant, or taking the constant level that yields the same lifetime utility as the actual life-cycle profile of consumption). This is a natural assumption if one considers that lifetime utility is the sum of instantaneous utility over the years. With such a multiplicative formula, consumption no longer matters when one does not have time to enjoy it, and longevity does not matter if consumption is so low that life is not worth living, making $U(c) = 0$.

Schooling does not appear in this example, because education is usually assumed to have a purely instrumental value in most economic models. But if one wanted to introduce a level of education in the analysis, one could argue that a low education that does not socialize the individual to be a dignified member of society is also something that makes the other domains irrelevant. To make this view compelling, one should include into the measure of education the informal education that individuals receive in their (actual or surrogate) family. The natural zero is then given by the situation of an individual who is not socialized at all and is not able to interact with human beings in a peaceful and articulate way. Being simply illiterate cannot be considered a zero because, at least in certain contexts, one can have a rather rich life without knowing how to read or write. This suggests that the "zero" for the number of formal schooling years should be a negative number of years, in order to take account of the informal education that almost every human being receives. This amounts to adding a positive constant to the years of formal schooling received.

4 Correlation sensitiveness, complementarity and inequality aversion

The example of the utility function in the previous section provides a seamless transition to the question of inequality between individuals, which is adressed in the IHDI (Alkire and Foster 2010) and also discussed by Anand (2018). Indeed, a widespread criticism of the HDI is that it relies on averages and ignores inequalities between individuals inside countries.

There are two aspects of this criticism that should be distinguished. First, GDP is a measure of average income, life expectancy is a measure of average longevity, and the schooling index relies on average years of schooling. By construction, such average magnitudes ignore the distribution among individuals. Second, there is an aspect of inequality that the HDI misses completely because it relies on separate indicators by domains. This missing aspect is the accumulation of disadvantages in the various domains on the worse off in the population. The poor in income also have shorter lives and less education. This is worse, arguably, than a situation in which each of these three domains would have the same level of inequality but without any correlation between domains. In summary, one would like to have an index that is not only inequality averse, but also correlation sensitive.

The SWF framework is useful to understand how these considerations can be incorporated into an index of development. Suppose that both individual utilities U_i and the welfare function W take the same constant-elasticity form, but with possibly different elasticities:

$$\left(\frac{1}{n}\sum_{i=1}^{n}\left(\alpha_i l_i^{1-\gamma}+\beta_i s_i^{1-\gamma}+y_i^{1-\gamma}\right)^{\frac{1-\rho}{1-\gamma}}\right)^{\frac{1}{1-\rho}} \tag{7}$$

where l_i denotes i's longevity, s_i denotes i's schooling, and y_i denotes i's income. The coefficients α_i, β_i weigh longevity and schooling to make the aggregation with income meaningful, and they could be individual-specific. The coefficient γ measures the complementarity between the three attributes, whereas ρ embodies inequality aversion. In principle, one could also make γ individual-specific.

As soon as $\rho > 0$, this formula displays inequality aversion over the individual utilities

$$\left(\alpha_i l_i^{1-\gamma}+\beta_i s_i^{1-\gamma}+y_i^{1-\gamma}\right)^{\frac{1}{1-\gamma}}. \tag{8}$$

What about correlation sensitiveness? Consider the special case in which $\gamma = \rho$. The formula then simplifies into

$$\left(\frac{1}{n}\sum_{i=1}^{n}\left(\alpha_i l_i^{1-\gamma}+\beta_i s_i^{1-\gamma}+y_i^{1-\gamma}\right)\right)^{\frac{1}{1-\rho}} = \left(\frac{1}{n}\sum_{i=1}^{n}\alpha_i l_i^{1-\gamma}+\frac{1}{n}\sum_{i=1}^{n}\beta_i s_i^{1-\gamma}+\frac{1}{n}\sum_{i=1}^{n}y_i^{1-\gamma}\right)^{\frac{1}{1-\rho}}, \tag{9}$$

and one then sees that correlation of inequalities across domains does not matter, since every domain can be aggregated over the population separately. One can compute the formula in two ways, either starting with individual utilities and aggregating over individuals, or starting the aggregation for each domain and then summing over domains. This indifference to the order of aggregation is called path independence in Alkire and Foster (2010), and is not questioned in Alkire and Foster (2010) or Anand (2018). It does have a practical advantage, which is that it makes it possible to rely on separate data sets per domain, a most convenient property given the state of available data in most countries. But path independence kills the possibility of taking account of the correlation of inequalities across domains.

Aversion to a positive correlation of attributes across domains is obtained only if $\rho > \gamma$, i.e., if inequality aversion is greater than complementarity. When $\rho < \gamma$, a positive correlation is actually valued positively, because the strong complementarity across attributes favors putting them in line more powerfully than inequality aversion favors reducing gaps between individuals.

A general decomposition of social welfare into averages, inequality and correlation can be written as follows, if one assumes away individual heterogeneity in the utility functions to simplify the formula (see the appendix for the complete formula):[1]

$$\left(\frac{1}{n}\sum_{i=1}^{n}\left(\alpha l_i^{1-\gamma}+\beta s_i^{1-\gamma}+y_i^{1-\gamma}\right)^{\frac{1-\rho}{1-\gamma}}\right)^{\frac{1}{1-\rho}} = \left(\alpha \bar{l}^{1-\gamma}+\beta \bar{s}^{1-\gamma}+\bar{y}^{1-\gamma}\right)^{\frac{1}{1-\gamma}}\left(1-C_{ineq}\right)\left(1-C_{corr}\right) \tag{10}$$

where $\bar{l}, \bar{s}$ and $\bar{y}$ denote average longevity, schooling and income, while C_{ineq} and C_{corr}

[1] In the literature, Abul-Naga and Geoffard (2006) present a decomposition of a multiplicative index of inequality into an average of inequality per attribute and a correlation term; Bosmans et al. (2015) decompose the same expression as in (10) (up to a multiplicative constant) into a term measuring average individual utility and a term measuring the loss of welfare due to inequality in utilities. They apply it to HDI data (taking countries as individuals), producing a decomposition of the world social welfare into the average HDI over countries and the second term. Here, we are instead interested in making a term similar to the country HDI appear in the computation of social welfare for a country.

measure the cost of inequality and correlation, respectively, and are defined as follows:

$$C_{ineq} = 1 - \left(\frac{\alpha \left(\bar{l} \left(1 - C_{ineq}^{l} \right) \right)^{1-\gamma} + \beta \left(\bar{s} \left(1 - C_{ineq}^{s} \right) \right)^{1-\gamma} + \left(\bar{y} \left(1 - C_{ineq}^{y} \right) \right)^{1-\gamma}}{\alpha \bar{l}^{1-\gamma} + \beta \bar{s}^{1-\gamma} + \bar{y}^{1-\gamma}} \right)^{\frac{1}{1-\gamma}} \tag{11}$$

where

$$C_{ineq}^{l} = 1 - \left(\frac{1}{n} \sum_{i=1}^{n} \left(\frac{l_i}{\bar{l}} \right)^{1-\rho} \right)^{\frac{1}{1-\rho}}, \tag{12}$$

and $C_{ineq}^{s}, C_{ineq}^{y}$ are defined similarly (these are Atkinson inequality indices for the three attributes longevity, schooling and income);

$$C_{corr} = 1 - \frac{\left(\frac{1}{n} \sum_{i=1}^{n} \left(\alpha l_i^{1-\gamma} + \beta s_i^{1-\gamma} + y_i^{1-\gamma} \right)^{\frac{1-\rho}{1-\gamma}} \right)^{\frac{1}{1-\rho}}}{\left(\alpha \left(\frac{1}{n} \sum_{i=1}^{n} l_i^{1-\rho} \right)^{\frac{1-\gamma}{1-\rho}} + \beta \left(\frac{1}{n} \sum_{i=1}^{n} s_i^{1-\rho} \right)^{\frac{1-\gamma}{1-\rho}} + \left(\frac{1}{n} \sum_{i=1}^{n} y_i^{1-\rho} \right)^{\frac{1-\gamma}{1-\rho}} \right)^{\frac{1}{1-\gamma}}}. \tag{13}$$

The formula for C_{ineq} is a weighted average of the three inequality indices $C_{ineq}^{l}, C_{ineq}^{s}, C_{ineq}^{y}$, and therefore if the degree of inequality is similar in the three domains, then C_{ineq} is approximately equal to their levels. The formula for C_{corr} may look cryptic, but it is simply based on the difference between doing the aggregation first over individuals, domain by domain (at the denominator) and doing the aggregation first over domains, individual by individual (at the numerator). It then clearly reflects correlation sensitivity and $C_{corr} = 0$ when path independence is satisfied.

The practical advantage of formula (10) is that even if it is hard to compute the left-hand side in the most exact way, it may be possible to rely on approximate estimations of C_{ineq} and C_{corr} based on more incomplete data than those providing the average levels $\bar{l}, \bar{s}$ and $\bar{y}$. The first term in the formula,

$$\left(\alpha \bar{l}^{1-\gamma} + \beta \bar{s}^{1-\gamma} + \bar{y}^{1-\gamma} \right)^{\frac{1}{1-\gamma}}, \tag{14}$$

is really a HDI-type indicator relying on average achievements by domain, and therefore formula (10) shows that, somehow, only two numbers are needed to go from an HDI-type indicator to an inequality-averse and correlation-sensitive formula. Although it may be hard to find the data to estimate C_{ineq} and C_{corr} reliably in many countries, it may be useful to flag the need for such data in order to make it possible, some day, to move beyond the HDI-type.

Anand (2018) raises two additional issues about the IHDI which should be mentioned here. First, there is ambiguity about the meaning of inequality regarding life expectancy. While life expectancy can be estimated for social subgroups, it is not an individual notion but by construction an average longevity over a group. What seems to be missing both in Anand (2018) and Alkire and Foster (2010) is the simple recognition of the fact that if life expectancy is average longevity, the corresponding individual variable is simply longevity. Now, the literature on inequality has been strangely shy about looking at inequalities in achieved longevity, as if longevity was not a legitimate achievement attribute.[2] The important point here is that with mortality tables it is not difficult at all to compute the distribution of age at death and the corresponding inequality in longevity, although it of course ignores the fact that people who are currently alive will probably benefit from future progress in longevity.

The second point raised by Anand is that zero values are even more a problem for the IHDI than for the HDI, because some individuals in the data do have zero values (e.g., zero schooling, or zero income). In formula (10), zero individual values are a problem for C_{ineq} if $\rho \geq 1$, and for C_{corr} if either $\rho \geq 1$ or $\gamma \geq 1$. Anand rightly objects to the manipulation of zero values by replacing them with a small positive constant. This is not, however, as far-fetched as one might think in certain cases. For instance, people with zero income may receive informal help in consumption that does not appear

[2] Fleurbaey et al. (2014) explore this issue and argue that it is an important component of inequality.

in the data, so that raising the level of their income without touching that of better off people may provide a better approximation of standards of living. Presumably, a better solution to this problem consists in giving zero its true meaning when there is strong inequality aversion or strong complementarity, which is that it really makes the situation unbearable. This implies choosing as the zero of the variables a value that is strictly below every observation, if we think that the situation deserves to be assessed and should not just be branded as "unbearable." Another option is to identify the subpopulation for whom the situation is unbearable and compute the indicator on the remaining population. This requires publishing the indicator in the following form: for x percent of the population, the situation is below the acceptability threshold, and for the remaining population, the indicator value is...

5 Selecting the weights and shape of the indices

Anand (2018) does not examine the selection of the weights of longevity, schooling and income in the HDI, and this is not an issue that has been agitated around the 2010 HDR. Yet, it is important in the eyes of many critics of synthetic indicators, who question the arbitrariness of the weights of the various domain indices in the aggregating formula. I have written extensively on this issue elsewhere, detailing various methods for choosing the weights (Fleurbaey and Blanchet 2013), but here it may be useful to relate this issue to those discussed in the previous sections.

When one looks at formula (10), the coefficients ρ (inequality aversion) on the one hand and γ (complementarity), α and β on the other hand, depend, arguably, on very different considerations. Inequality aversion is about adjudicating the interests of different individuals, which is eminently an ethical issue. Such ethical issues should be determined by rational deliberation and cannot be left to the whims and vagaries of individuals. In contrast, the relative weight of longevity and schooling relative to income and the degree of complementarity between the attributes depends on what is good for any given individual. There would be nothing shocking in allowing these three parameters (α, β, γ) to be individual specific (although it would complicate the decomposition formula (10)). For some individuals, living long may be more important than for others; for some individuals succeeding in the three domains may be important whereas for others succeeding in one domain may be enough; and so on.

When one looks at a synthetic indicator like (14), it is therefore helpful to refer back to the SWF framework and ask where the parameters would belong in a full SWF formula like (1). In (14), the three parameters (α, β, γ) belong to the individual utility function, not to the W function which aggregates over individuals. Therefore it would be natural to elicit the population's view about the relative weights and about the degree of complementarity. And it would be ideal not just to take the average view in the population, but even to allow for heterogeneity across individuals, in order to respect their values and preferences over the components of their own lives.

While the debate about synthetic indicators has generally focused on the weights, this discussion shows that the shape of the indicators should also be part of this debate. Therefore, in the assessment of HDI_{old} and HDI_{new}, which differ in the complementarity parameter (γ in the previous section), one should not simply do armchair philosophy about whether development should be balanced or can succeed in one dimension while neglecting the others, but ask how the population members themselves would trade-off longevity, education and income. This is why the typical multiplicative form of individual utilities over income and longevity was relevant to bring to bear on the choice between HDI_{old} and HDI_{new} in section 3.

The perspective of relying on the population views for the selection of the indicator's weights and shape may seem to be even more challenging in terms of data requirements than anything that has been considered so far, but in fact this is not necessarily the case. Subjective well-being surveys accompanied with questions about objective attributes may provide ways to elicit the population's priorities in life in a cost-effective way. If registering the correlation between attributes requires a survey on the joint distribution of the attributes, it is not diffcult to add a few subjective questions. One could at least imagine country-specific indicators. Decancq and Neumann (2016) compare different measures of "utility" and show that taking account of preferences, and especially of heterogeneous preferences, does matter the ranking of individuals and in particular the identification of the poor.

One common objection against allowing for heterogeneous indicators for different populations is that it makes comparisons between them impossible. This is not true, fortunately. If that were true, it would also be impossible to compute social welfare in a country with heterogeneous preferences among its residents (Fleurbaey and Tadenuma 2014). Consider for instance the old HDI and imagine that different countries had different weights, but always adding up to one:

$$HDI_k = a_k \times I_{LE} + b_k \times I_S + c_k \times I_Y \tag{15}$$

where k denotes the country's name, and $a_k + b_k + c_k = 1$. One way to understand how such weighting operates is to ask in what circumstances the weights no longer play a role in comparing country indicators. Clearly, when $I_{LE} = I_S = I_Y$, this common level is equal to HDI_k. Therefore, when countries have a fully balanced development, they are compared in a way that is independent of their specific weights. Now, consider a country for which $I_{LE} = I_S = I_Y$ does not hold. One can see that its value for HDI_k is actually the level of balanced development that would give the same indicator according to its own weights. In other words, this type of index compares countries in terms of their “equivalent” balanced development, where the equivalence is respectful of their specific weights.

Exactly the same can be said for any generalized-mean indicator of the form:

$$HDI_k = \left(a_k \times I_{LE}^{1-\gamma} + b_k \times I_S^{1-\gamma} + c_k \times I_Y^{1-\gamma}\right)^{\frac{1}{1-\gamma}}, \tag{16}$$

including the case of a geometric indicator:

$$HDI_k = I_{LE}^{a_k} \times I_S^{b_k} \times I_Y^{c_k}, \tag{17}$$

provided that $a_k + b_k + c_k = 1$. Therefore, heterogeneous weights do not make inter-country comparisons meaningless, they can still be quite transparent. Likewise, one can allow heterogeneity among individual utilities in a SWF as in (7) and make sense of the interpersonal comparisons.

6 Gender inequality

Anand (2018) examines various approaches to gender inequality. There is the gender-related development index (GDI) proposed by Anand and Sen (1995):

$$GDI = \frac{1}{3}\left[\left(p^F\left(I_{LE}^F\right)^{1-\rho} + p^M\left(I_{LE}^M\right)^{1-\rho}\right)^{\frac{1}{1-\rho}} + \left(p^F\left(I_S^F\right)^{1-\rho} + p^M\left(I_S^M\right)^{1-\rho}\right)^{\frac{1}{1-\rho}} + \left(p^F\left(I_S^F\right)^{1-\rho} + p^M\left(I_S^M\right.\right.\right. \tag{18}$$

in which female and male domain indicators I_X^F, I_X^M (for $X = LE, S, Y$) are combined (weighted by population proportions p^F, p^M) in a way that applies a discount when they are unequal. There is Dijkstra and Hanmer’s (2000) Relative Status of Women

(RSW):

$$RSW = \frac{1}{3}\left(\frac{I^F_{LE}}{I^M_{LE}} + \frac{I^F_S}{I^M_S} + \frac{I^F_Y}{I^M_Y}\right) \tag{19}$$

and Klasen and Schüler's (2011) Gender Gap Measure (GGM):

$$GGM = \left(\frac{I^F_{LE}}{I^M_{LE}} \times \frac{I^F_S}{I^M_S} \times \frac{I^F_Y}{I^M_Y}\right)^{1/3}, \tag{20}$$

as well as Permanyer's (2013) generalization and variant of GGM. Anand is by and large positive about these approaches, distinguishing between measures of (gender neutral) gender inequality as in (18) and measures of (female focused) gender disparity as in (19)-(20). He also notes that (18) actually corresponds to a symmetrical measure of disparity, since one can write:

$$\left(p^F\left(I^F_{LE}\right)^{1-\rho} + p^M\left(I^M_{LE}\right)^{1-\rho}\right)^{\frac{1}{1-\rho}} = I_{LE}\left(\frac{p^F}{\left(p^F + p^M\frac{I^M_{LE}}{I^F_{LE}}\right)^{1-\rho}} + \frac{p^M}{\left(p^F\frac{I^F_{LE}}{I^M_{LE}} + p^M\right)^{1-\rho}}\right)^{\frac{1}{1-\rho}}. \tag{21}$$

The Gender Inequality Index (GII) introduced in UNDP (2010) receives a much more critical treatment in Anand (2018) as in Permanyer (2013). The complexity of the index and the fact that it violates basic monotonicity properties with respect to maternal mortality rate and adolescent fertility seem in particular quite problematic.

Here I will limit myself to two issues. The first is that all the above formulae miss the question of the correlation over domains of female disadvantage. The fact that women are disadvantaged in the three domains in many countries is worse than the sum of three disadvantages for three independent halves of the population, because being a woman means that you are more likely to *accumulate* these gaps. To fully capture gender inequality, one could apply formula (10) to a data set in which C_{ineq} and C_{corr} are computed by replacing every individual's attributes (for C_{ineq}) and or every individual's utility (8) (for C_{corr}) by the average value of these magnitudes in the gender group of the individual. When this is done in this way, Anand and Sen's formula (18) can be shown to be equivalent to computing the first two terms of (10) and neglecting the third term involving C_{corr}. Data limitations may prevent this third term from being easily estimated, but the issue should not disappear from the radar of statisticians and policy-makers.

The second issue is that maternal mortality and teenage pregnancy raise the very interesting challenge of specifically female health and social issues for which comparison with men is very hard. The willingness of the conceptors of GII to use these data is laudable, but the challenge is there. In the GII, men are represented with an arbitrary constant for these domains (as noted by Anand, the constant takes the maximum possible value of these rates, which is strange for dimensions that do not exist for men). A closer examination of the issue, however, suggests that some of the difficulty is only apparent. Consider maternal mortality. It is specific to women to die in delivery, but dying is a similar experience for all humans, and therefore maternal mortality seems already counted as an additional mortality risk for women, just like men are submitted to a greater risk of death by gang violence. In other words, longevity statistics should capture the phenomenon to a large extent.

Teenage pregnancy is more complex, because there is nothing directly comparable for men, unlike mortality. However, the relevant consequences of teenage pregnancy include increased drop-out rates at school, health issues for mother and child, and subsequent loss of income, all of which can be compared to male problems. One aspect of teenage pregnancy that is harder to compare to men's issues is the responsibility and loss of freedom that befall the young mother at a premature age. In an equal society, the father should share this burden, but the fact that most of this burden falls on the lone mother is itself a symptom of gender discrimination. But for such costs, one can reason in terms of equivalent cost. It is possible to ask what other event, which can happen to men, would similarly harm the mother, where "similarly" refers not to the multidimensional and qualitative aspect of the situation but to a perception of loss of similar magnitude in the eyes of the individual. We are back here to relying on individual values and preferences to assess complex situations.

7 Conclusion

This paper has relied on the social welfare framework as a compass to navigate the debate around the new HDI, the IHDI, and the GII. It has defended the introduction of some degree of complementarity between the dimensions of achievements, arguing that: 1) this does not undermine the priority that should be given to lives in poor countries; 2) it reflects the importance of balanced development at the societal level as well as the individual level; 3) it potentially paves the way to taking account of two issues that have been neglected in the debate around the HDR: a) one should record not only inequalities within each domain but also of the correlation across domains that is particularly harmful for the most deprived populations; b) one should consider relying on the populations' values and preferences in order to choose the degree of complementarity and the weights of the domains. It has argued that the latter point can also be helpful in enriching the way in which gender inequality measures can capture the specificity of female disadvantages such as teenage pregnancy.

This paper has proposed a decomposition formula (10) which separates a HDI-like indicator (based on averages by domain) from the correction due to inequalities by domain, and the correction due to correlation among domains leading to cumulative disadvantage for the worse off populations. This formula shows that even an ambitious approach based on a social welfare function can make room for a HDI-like indicator, which is then seen as capturing a part of the picture of social welfare, and needing to be supplemented by a record of inequalities by domain (producing an IHDI-like indicator) and correlations across domains. Relying on the social welfare framework, which is versatile enough to accommodate a variety of approaches (including an approach in terms of functionings and capabilities), makes it natural to wonder why the populations' own views about how to assess their personal situations are conspicuously ignored in these indicators.

References

Abul-Naga R.H., P.Y. Geoffard 2006, "Decomposition of bivariate inequality indices by attributes," *Economics Letters* 90: 362–367.

Alkire S., J.E. Foster 2010, "Designing the Inequality-Adjusted Human Development Index (IHDI)," *Human Development Reports Research Paper 2010/28*, United Nations Development Programme.

Anand S. 2018, "Recasting human development measures," *UNDP Discussion Paper*, March.

Anand S., A.K. Sen 1995, "Gender Inequality in Human Development: Theories and Measurement", in UNDP, *Background Papers: Human Development Report 1995*, United Nations Development Programme.

Bosmans K., K. Decancq, E. Ooghe 2015, "What do normative indices of multidimensional inequality really measure?" *Journal of Public Economics* 130: 94–104.

Decancq K., D. Neumann 2016, "Does the Choice of Well-Being Measure Matter Empirically?" in M. Adler and M. Fleurbaey (eds.), *Oxford Handbook of Well-Being and Public Policy*, Oxford U. Press.

Dijkstra A.G., L.C. Hanmer (2000): “Measuring Socio-Economic Gender Inequality: Toward an Alternative to the UNDP Gender-Related Development Index,” *Feminist Economics* 6: 41–75.

Fleurbaey M., D. Blanchet 2013, *Beyond GDP*, Oxford U. Press.

Fleurbaey M., M.L. Leroux, G. Ponthière 2014, “Compensating the dead,” *Journal of Mathematical Economics* 51: 28-41.

Fleurbaey M., K. Tadenuma 2014, “Universal social orderings,” *Review of Economic Studies* 81: 1071-1101.

Klasen S., D. Schüler 2011, “Reforming the Gender-Related Index and the Gender Empowerment Measure: Implementing Some Specific Proposals,” *Feminist Economics* 17: 1–30.

Murphy K.M., R.H. Topel 2006, “The Value of Health and Longevity,” *Journal of Political Economy* 114: 871-904.

Permanyer I. 2013, “A Critical Assessment of the UNDP’s Gender Inequality Index,” *Feminist Economics* 19: 1–32.

UNDP 2010, *Human Development Report 2010 – 20th Anniversary Edition. The Real Wealth of Nations: Pathways to Human Development*, United Nations Development Programme.

Appendix

The decomposition (10) can be extended to the case of heterogeneous preferences, as follows:

$$\left(\frac{1}{n}\sum_{i=1}^{n}\left(\alpha_i l_i^{1-\gamma_i}+\beta_i s_i^{1-\gamma_i}+y_i^{1-\gamma_i}\right)^{\frac{1-\rho}{1-\gamma_i}}\right)^{\frac{1}{1-\rho}}=\left(\bar{\alpha}\bar{l}^{1-\bar{\gamma}}+\bar{\beta}\bar{s}^{1-\bar{\gamma}}+\bar{y}^{1-\bar{\gamma}}\right)^{\frac{1}{1-\bar{\gamma}}}\left(1-C_{ineq}\right)\left(1-C_{corr}\right)\left(1-C_{pref}\right) \tag{22}$$

with

$$C_{ineq}=1-\left(\frac{\bar{\alpha}\left(\bar{l}\left(1-C_{ineq}^{l}\right)\right)^{1-\bar{\gamma}}+\bar{\beta}\left(\bar{s}\left(1-C_{ineq}^{s}\right)\right)^{1-\bar{\gamma}}+\left(\bar{y}\left(1-C_{ineq}^{y}\right)\right)^{1-\bar{\gamma}}}{\bar{\alpha}\bar{l}^{1-\bar{\gamma}}+\bar{\beta}\bar{s}^{1-\bar{\gamma}}+\bar{y}^{1-\bar{\gamma}}}\right)^{\frac{1}{1-\bar{\gamma}}} \tag{23}$$

$$C_{corr}=1-\frac{\left(\frac{1}{n}\sum_{i=1}^{n}\left(\bar{\alpha}l_i^{1-\bar{\gamma}}+\bar{\beta}s_i^{1-\bar{\gamma}}+y_i^{1-\bar{\gamma}}\right)^{\frac{1-\rho}{1-\bar{\gamma}}}\right)^{\frac{1}{1-\rho}}}{\left(\bar{\alpha}\left(\frac{1}{n}\sum_{i=1}^{n}l_i^{1-\rho}\right)^{\frac{1-\bar{\gamma}}{1-\rho}}+\bar{\beta}\left(\frac{1}{n}\sum_{i=1}^{n}s_i^{1-\rho}\right)^{\frac{1-\bar{\gamma}}{1-\rho}}+\left(\frac{1}{n}\sum_{i=1}^{n}y_i^{1-\rho}\right)^{\frac{1-\bar{\gamma}}{1-\rho}}\right)^{\frac{1}{1-\bar{\gamma}}}}. \tag{24}$$

and

$$C_{pref}=1-\frac{\left(\frac{1}{n}\sum_{i=1}^{n}\left(\alpha_i l_i^{1-\gamma_i}+\beta_i s_i^{1-\gamma_i}+y_i^{1-\gamma_i}\right)^{\frac{1-\rho}{1-\gamma_i}}\right)^{\frac{1}{1-\rho}}}{\left(\frac{1}{n}\sum_{i=1}^{n}\left(\bar{\alpha}l_i^{1-\bar{\gamma}}+\bar{\beta}s_i^{1-\bar{\gamma}}+y_i^{1-\bar{\gamma}}\right)^{\frac{1-\rho}{1-\bar{\gamma}}}\right)^{\frac{1}{1-\rho}}}.$$

One observes that C_{ineq} and C_{corr} are the same as before, but relying on the average values of the parameters α,β,γ. The new term C_{pref} captures the impact of heterogeneous preferences, and could be further decomposed into terms distinguishing the variation in preference parameters (bracketing out the variations in attributes) from the correlation between preferences and attributes.