

**TECHNICAL CONSULTATION ON
THE MEASUREMENT OF HEALTH INEQUALITIES**

BACKGROUND PAPER

7-8 NOVEMBER 2001

GENEVA, SWITZERLAND



World Health Organization

© World Health Organization 2001

This document is not issued to the general public, and all rights are reserved by the World Health Organization (WHO). The document may not be reviewed, abstracted, quoted, reproduced or translated, in part or in whole, without the prior written permission of WHO. No part of this document may be stored in a retrieval system or transmitted in any form or by any means—electronic, mechanical or other—without the prior written permission of WHO.

Evidence and Information for Policy
World Health Organization
1211 Geneva 27
Switzerland

Table of contents

<i>I. Introduction</i>	5
<i>II. Commentaries on the WHO health inequalities measurement approach</i>	6
A. Social inequalities in health vs. all inequalities in health	6
B. No information on adult inequality	6
C. Data availability	6
D. Index of inequality	6
<i>III. Conceptual Issues.</i>	6
A. WHO framework	6
B. Basic premise: a comprehensive approach	7
C. Voluntary and Genetic Risks	8
D. Best achievable risks?	9
E. Bivariate and multivariate distributions of health with other variables.	10
<i>IV. Summarizing the Distribution of Health Expectancy in a Measure of Health Inequality</i>	10
A. Two Families of health inequality measures	11
1. Individual-Mean Differences	11
2. Inter-Individual Differences	12
B. Choosing a single index of health inequality	12
1. β : Relative versus absolute inequality	13
2. α : Intensity of health gain/loss.	13
3. Inter-Individual versus Individual-Mean Differences	14
C. Scale properties and range	16
D. Alternative summary metrics	16
<i>V. Child survival inequality</i>	18
A. Models used to estimate the distribution of child survival	18
1. The extended beta-binomial model	19
2. The random effects logit model	21
3. Estimating the distribution of expected survival time	22
B. Estimates of child survival inequality	23
C. Decomposition of child survival inequality index	25
D. New datasets available	27
<i>VI. Adult survival inequality</i>	28
A. Models for estimating risk of death distribution in adults	28
B. Individual level data	30
1. Direct record linkage	30
2. Probabilistic record linkage	30
C. An application of the adult survival model to the US data	31
1. Data	31
2. Model used	31
3. Preliminary results	32

D. Household survey data	32
VII. <i>Using Small Area Analysis to Estimate the Distribution of Mortality Risk Across Individuals</i>	33
A. Background	33
B. Principles	33
C. Models to Decompose Observed Variation in Death Rates Across Small Areas into Stochastic and Deterministic Components	34
1. <i>The Beta-Binomial Model</i>	35
2. <i>Normal Mortality Risk-Normal Approximation of the Binomial</i>	35
D. What Fraction of the Total Variance in Mortality Risk Across Individuals is Captured by the Variance in Mortality Risk Across Small Areas?	37
E. Size of Small Areas and Variation in Mortality Risk	39
F. Supplementing Small Area Analysis	39
VIII. <i>Inequality in health states</i>	39
A. Data & estimation methods	39
B. Preliminary results on inequality in health states	41
IX. <i>Inequality in healthy life expectancy (HALE)</i>	42
A. Correlation between risk of mortality and current health state for an individual	43
B. Correlation across ages of the risk of ill-health for an individual	43
X. <i>Health of the poor</i>	45
XI. <i>Discussion</i>	47
XII. <i>References</i>	48
XII. <i>Annexes</i>	52

I. Introduction

This document provides background material for the WHO technical consultation on the measurement of health inequalities. This technical consultation forms part of a wider effort by WHO to engage in a broad-based consultative process on health systems performance assessment concepts, methods and data sources. In this introductory section, we briefly review the approach taken to assessing inequalities in health published in the World Health Report 2000.

The World Health Organization developed a framework for the assessment of health system performance which was utilized as an analytical tool and published in the *World Health Report 2000*¹. The framework identified three intrinsic social goals to which health systems contribute: population health, responsiveness of the health system; and fairness in financial contributions of households to the health system. The goal of population health is defined as improving the average level of health and improving the distribution of health across individuals, in other words reducing inequalities in health.

Levels of health were assessed using Disability Adjusted Life Expectancy, now renamed Healthy Life Expectancy (HALE)². WHO argued that logically, health inequalities should be measured as the distribution of healthy life expectancy across individuals. Because of the limitations of data and methods, for the WHR 2000 health inequality was assessed by measuring the distribution of the probability of survival across children. Survival from birth to two years of age is only one aspect of health, but it is a useful place to start since it is a critical part of health status, particularly in developing countries³.

A parametric model was used to estimate the distribution of the probability of surviving to age 2. This distribution was then summarized with the following measure of inequality:

$$II[3,.5] = \frac{\sum_{i=1}^n \sum_{j=1}^n |s_i - s_j|^3}{2n^2 \bar{s}^{-0.5}}$$

where s_i is the expected survival time of child i from birth to age 2, and \bar{s} is the average survival time in the population. The measure is based on comparing each child to every other child in the population, gives a large weight to the tails of the distribution as all differences are raised to the power of 3 in the numerator, and is a relative measure as the mean is included in the denominator. This measure was selected based on the responses to an internet based survey which included questions on the normative choices involved in the selection of an inequality measure. In the WHR2000 estimates of *equality* in child survival were reported to preserve consistency with the reporting on the other four goals of health system which were all on a positive scale (i.e. a higher number is better). Equality in child survival was simply estimated as one minus the inequality index (II) presented above. Data on child survival came from complete birth histories available through the Demographic and Health Surveys (DHS) program. For developed countries child

survival data from small geographical areas, such as counties or municipalities, were used.

II. Commentaries on the WHO health inequalities measurement approach

Since the publication of the WHR2000, in journal articles and in various regional consultations, a number of important issues have been raised. The main issues can be summarized as follows.

A. Social inequalities in health vs. all inequalities in health

Perhaps the most extensively discussed issue is whether WHO should be measuring inequalities in health by summarizing the entire distribution of health in a population or by reporting on differences in average levels of health across sub-groups of the population defined by a socio-economic attribute such as income, education, race or social class.⁴⁻⁶

B. No information on adult inequality

The measure of health inequality in the WHR2000 was one that reflected only inequality in child survival. This measure is mostly relevant for low- and some middle-income countries where levels of child mortality are still high; however, it is not a good measure of health inequality for high income countries where the level of child mortality is low and there is very little variation in the probabilities of child survival.

C. Data availability

In WHR2000 microdata from 56 countries were used. The remainder of the health inequality figures were estimated using a model relating measured levels of child survival inequality with population characteristics.

D. Index of inequality

The index of health inequality used is not one of the more commonly used ones in the literature. The index was based on responses to a WHO internet-based survey, and thus reflected the opinions of those respondents. The choices involved in selecting a measure of inequality are normative; the more commonly used measures such as the Gini coefficient and the variance represent a particular set of these normative choices.

The remainder of the document is organized around both the main commentaries on health inequality and new analytical developments that are relevant to the assessment of health inequality.

III. Conceptual Issues.

A. WHO framework

WHO proposed a comprehensive approach to the measurement of health inequality that parallels the extensive literature on inequality assessment in other fields such as economics. The conceptual basis for this approach was developed in a detailed argument which will not be repeated here⁷. To summarize the key conclusions, the

most relevant quantity of interest for studying health inequality is the distribution of health expectancy across individuals, constructed for a period, using a clearly defined method for linking the distributions of health risks at different ages. Operationalizing the measurement of the distribution of health expectancy across individuals requires measurement of the distribution of mortality risk at each age and the distribution of health states at each age.


Several aspects of this approach have been challenged. Some of these major points are briefly discussed below.

B. Basic premise: a comprehensive approach

A central debate on health inequality that has emerged since the publication of the World Health Report 2000 has been the extent to which the comprehensive approach proposed by WHO is appropriate⁸. For example, Braveman et al.⁴ have criticized WHO. They argue for a selective approach where only some health inequalities, those correlated with income, social class or race should be quantified. Their central argument takes the form of a moral proposition: only health inequalities correlated with income, social class or race matter and therefore warrant public action. Such a selective approach runs counter to the literature on inequality in other disciplines such as economics and runs the risk of discouraging scientific inquiry into the causes of inequality.

A specific empirical example will help illustrate the differences between the selective and comprehensive approaches to health inequality. Using household survey data, Gakidou and King⁹ have estimated the risk of death for each child in 50 countries (this is discussed in more detail below). Figure 1 illustrates the results of their analysis for Central African Republic. The x-axis shows households divided into five wealth quintiles and the y-axis shows the range of child mortality risk within each wealth quintile. The diagram shows the median risks, the 25 and 75% percentiles, the interquartile range and outliers. Figure 1 shows that the median risk of child death is lower in wealthier households. But it also shows that within each wealth quintile, there is tremendous variation in the risks of death for different children. This variation is systematic or non-random variation in risks of death; it is not simply showing the effects of chance on outcomes.

Figure 1. Distribution of the risk of dying by wealth quintile for the Central African Republic: median, interquartile range and outliers.



In Central African Republic and all other countries that have been studied, there is tremendous inequality across children and adults in the risks of death at the same ages. Some of this variation in the risk of death is correlated with Braveman et al's short-list of socio-economic factors such as income, social class or race. Some of this risk is correlated with community-level factors such as environmental sanitation, water supply, health services and social norms about risk factors. Some of this inequality is correlated with household-level factors such as type of housing, infant feeding practices or birth spacing. Probably a very small proportion of this inequality is correlated with genetic factors; although few genetic factors have been identified so far with a substantial population effect on inequalities. An even smaller component of inequality is correlated with individual's fully informed choices to take health risks such as extreme sports. Given our current state of knowledge, much systematic inequality in mortality risk remains unexplained. It is the challenge for health scientists to try and understand more of this variation in mortality risks.

Other disciplines such as economics tend to use comprehensive approaches to measuring inequality rather than selective approaches. When economists study income inequality, they do not simply report differences in average income for social class or race groups. Rather, they measure the entire distribution of income across individuals or households and summarize that distribution with measures such as the Gini coefficient. With comprehensive measures of income inequality available, it then becomes a scientific question of great interest how much is explained by social class or race. For health, WHO has adopted the same approach. First, measure the full extent of health inequality in a population. This comprehensive measure of health inequality captures all differences both those correlated with known socio-economic factors and those that are uncorrelated. Second, use the tools of science to understand what factors explain this inequality. Third, formulate policies that can act on these causes of inequality. Fourth, monitor and evaluate the impact of these policies on inequality. In this comprehensive approach, an evidence-base can be constructed on the causes of health inequality and the policy options available to address this inequality.

C. Voluntary and Genetic Risks

Some have argued that we should be uninterested in two components of the distribution of health expectancy for a cohort: the component that is not amenable to change and the component that arises from fully informed choices of individuals to decrease their health expectancy through the pursuit of risky activities.

If there were differences in health expectancy that could never be remedied either with current or future technology, one could perhaps persuasively argue that we should be uninterested in this. But which component of the distribution of health expectancy is not amenable to intervention? That due to genes? That due to chance during birth? In both cases, the argument that we cannot intervene to change the effects on the distribution of health expectancy is most likely unsound. With current improvements in technology and future progress, it is likely that these components of the distribution of health expectancy will become amenable to change and thus should not be excluded from a measure of health inequality. Perhaps as important is the argument that there is little evidence of significant cross-population variation in the

contribution of genes etc. The component of health expectancy distribution due to unavoidable factors is likely to be small and completely impossible to assess.

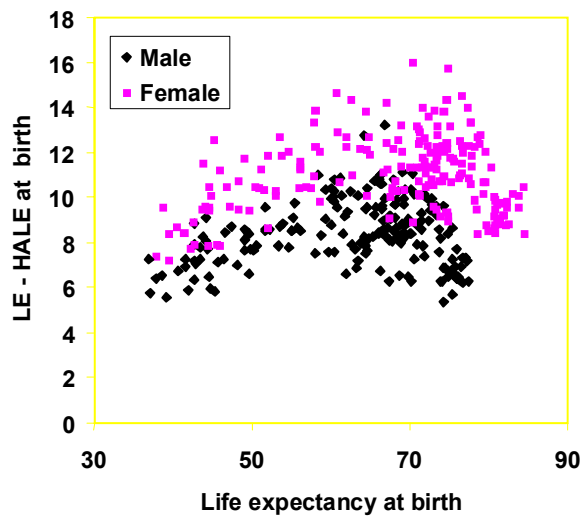
What about volition? How much of the distribution of health expectancy for a population is due to fully informed choices of individuals who have a taste for risky behaviour? This seems like a very slippery slope. Which choices affecting health are fully informed? Would we exclude the effects of tobacco on health expectancy which are likely to be very great because smoking is a choice? Even if we claim that the choice was informed, should it be excluded? We would argue that it should not be excluded. First, in most cases health risks are not adopted because of a love of risky behaviour but rather for other less informed reasons.¹ Second, the true volitional component of the distribution of health expectancy is likely to be very small and can well be ignored. This argument is similar to ones in the field of income inequality, where the variation in the distribution of income due to different leisure income trade-offs in the population is routinely ignored in the measurement of income inequality.

D. Best achievable risks?

In assessing the distribution of healthy life expectancy in a population, it is an interesting and potentially important question to ask what is the counter-factual of the most equal distribution of healthy life expectancy possible in a population? Over the last 100 years of mortality evolution especially in high-income countries, there has been progressive rectangularization of survival curves. Manton¹⁰ results suggest that there has probably also been rectangularization of the curve of living in equivalent full health as well. Mathers et al.¹¹ suggests that there is clear evidence of compression of morbidity for women in populations with life expectancies greater than 70 and for men greater than 65 years. Figure 2 illustrates with the most recent values of HALE for 191 countries a strong effect of compression. What is the implication of this for assessing inequality in healthy life expectancy?

¹ The cost of being fully informed about the health consequences of different choices often is prohibitively high. Most individuals are forced to make choices with incomplete or incorrect information. When the choice to take on risk and the outcome are separated in time, the rate at which individuals discount the future can profoundly influence choices about health. We, for example, would argue that health inequalities resulting from individuals with high discount rates should be included in measuring health inequality. Some philosophers⁵⁷ and economists⁵⁸ have long argued that discounting or myopia is a defect of human judgement and can be self-defeating.

Figure 2. Life expectancy at birth versus years of life lost to time spent in health states less than full health, 191 Member States, 2001.



E. Bivariate and multivariate distributions of health with other variables.


Recognizing that the distribution of health across individuals is of fundamental intrinsic interest, for broader purposes of informing social policy it may also be important to know the bivariate or multivariate distribution of health across individuals with the distribution of levels of attainment on other critical domains of well-being¹². For example, the bivariate distribution of income and health may be particular interest. Not because only the subset of health inequality correlated with income is of intrinsic interest but rather for some purposes the bivariate distribution may be informative. For example, with increasing interest in multi-dimensional constructs of poverty^{13;14}, a poverty line may actually be a two or more dimensional curve defined over the relevant dimensions. If the bivariate distribution of health and income is measured then all relevant measures can be easily calculated to serve the needs of different social policy analysts and decision-makers.

IV. Summarizing the Distribution of Health Expectancy in a Measure of Health Inequality

Figure 3 illustrates the distribution of health expectancies for three populations, A, B and C. Which distribution represents a more unequal distribution of health expectancy? If the x-axis in the graph were income, rather than health expectancy, most people would agree that distribution B is less unequal than C and A. This simple conclusion is based on the concept of decreasing marginal utility of income, namely that an extra dollar generates less utility as income rises. Distribution B has the same variance as A, but a higher mean. In terms of a commonly used measure of inequality, the Gini coefficient, Distributions A and C have the same amount of inequality, while Distribution B has lower inequality than A and C. While for income some people may be in agreement that Distributions A and C have equal amounts of inequality, for health, this finding may be met with less agreement. The notion of declining marginal utility does not apply as persuasively. Some would say C is

clearly worse than A or B and that they cannot distinguish between A and B. The vast literature on measuring income inequality¹⁵⁻¹⁹ is very helpful in the design of a health inequality measure, but this simple example illustrates that health has some fundamental differences from income that require special consideration. To date in the literature on measuring health inequality, there has been little substantive discussion on summary measures of distributions of health.

Figure 3. Distributions of health expectancy.



A. Two Families of health inequality measures

Based on the wide array of measures used to summarize the distribution of income²⁰ and taking into account the fact that absolute, and not just relative to the mean, differences in health expectancies may matter, we propose two families of measures: individual mean differences and inter-individual differences.

1. Individual-Mean Differences

Measures of individual-mean differences compare each individual's health to the mean of the population. The general form is:

$$IMD(\alpha, \beta) = \frac{\sum_{i=1}^n |h_i - \bar{h}|^\alpha}{n\bar{h}^\beta}$$

where h_i is the health of individual i , \bar{h} is the mean health of the population, and n is the number of individuals in the population. The parameter α changes the significance attached to differences in health observed at the ends of the distribution, compared to differences observed near the mean of the distribution. The parameter β controls the extent to which differences from the mean are assessed in absolute units or in units relative to the mean value. A common example of individual-mean differences is the variance when $\alpha=2$ and $\beta=0$. However, many other individual-mean difference measures are possible. When $\beta=1$, the differences are assessed strictly relative to the mean value and when $\beta=0$ it is measuring absolute deviations

from the mean but β could be any value between 0 or 1, reflecting some mix of concern about deviations between a relative to the mean and absolute view².

2. Inter-Individual Differences

Another family of measures is based on comparing each individual's health to every other individual's health rather than comparing each individual to the mean of the population. We propose the general form of these measures to be:

$$IID(\alpha, \beta) = \frac{\sum_{i=1}^n \sum_{j=1}^n |h_i - h_j|^\alpha}{2n^2 \bar{h}^\beta}$$

where h_i is the health of individual i and h_j is the health of individual j , \bar{h} is the mean health of the population and n is the number of individuals in the population. The parameters α and β are the same as for the individual-mean measures described above. A well-known example of this family is the Gini coefficient²¹ often used to measure income distribution, where $\alpha=1$ and $\beta=1$. The Gini is often represented as being derived graphically from the Lorenz curve²² of a population, but in fact is algebraically equal to the equation above. It is worth noting that when $\alpha=2$ the individual-mean difference and the inter-individual difference for any given population distribution are identical. For any other values of α they are different.

B. Choosing a single index of health inequality

For standard comparisons we need to choose a single index of health inequality to summarize the distribution of health expectancy for a population. This choice requires the resolution of three fundamentally normative issues: which family of measures, what should be the value of α and what should be the value of β . Individual's preferences for these normative choices can be elicited through a series of questions that isolate the effect of each on the index of inequality.

To elicit preferences on the normative choices of an inequality measure, eight hypothetical, unrealistic scenarios were designed. Hypothetical populations of 5 to 10 individuals were used in order to depict the implications of the normative choices more clearly. Respondents were asked to pick between two scenarios in which there was only one difference between the two populations. With their choice respondents were indicating their preferences on *one* of the parameters of the inequality measure at a time.

² The term relative measure has been used in the literature on inequality to refer to the property of scalar independence. If every value in a distribution is multiplied by a scalar, s , then scalar independence is when the measure of inequality does not change. From inspection of the equation for the IMD, scalar independence will only occur when α and β are equal and nonzero. In this paper, we do not use the term relative to mean the special property of scalar independence but rather to mean that differences are evaluated relative to the mean value.

1. β : Relative versus absolute inequality

One of the key choices is whether we are more concerned about absolute differences in health, relative differences in health, or a mix of both. For the choice between absolute versus relative measure of inequality four questions were asked. Positive and negative measures of health (life expectancy, probability of survival and probability of death) were used to test the consistency of the preferences. Figure 4 presents an example of scenarios that respondents were asked to choose between. Population A and Population B have the same number of people, distributed in the same way around the mean. The means of the two populations are different, though, with Population A having a life expectancy of 20 years, while Population B has a life expectancy of 65 years. Individuals who are of the opinion that health inequality is the same in the two populations are expressing a preference for measuring absolute inequalities. Those who think that the degree of inequality is different in the two populations reveal a preference of relative measures of inequality, as they think of inequality, relative to the mean of the populations.

Figure 4. An example of a question eliciting respondents' preferences for absolute vs. relative measures of inequality.



2. α : Intensity of health gain/loss.

The second normative choice (illustrated in Figure 5) has to do with whether gains or losses of health that occur at the ends of the distribution should be treated differently from those that occur near the mean. Again there are two Populations, A and B. The initial populations are identical and in each population one individual gains 5 years of life expectancy while another individual loses 5 years of life expectancy. In Population A the individuals affected by the change are at the ends of the distribution, while in Population B they are not. If the respondents thought that the change in A was larger than the change in B, they expressed a preference the ends of the distribution should be weighed more. If respondents thought the two changes were identical, then they were opting for equal weights to all parts of the distribution.

Figure 5. Example of a question eliciting preferences for relative weighting of outliers in a distribution.

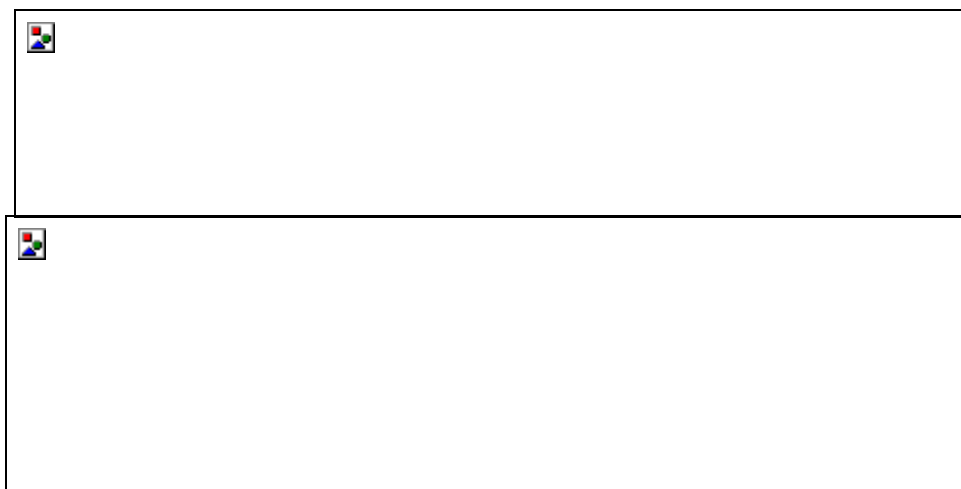


The figure shows two vertically stacked rectangular boxes. Each box is empty and contains a small icon in its top-left corner. The icon consists of four small squares in a 2x2 grid, colored red, green, blue, and yellow.

3. Inter-Individual versus Individual-Mean Differences

The third choice refers to the family of measures: individual-mean or inter-individual comparisons. In Figure 6, the two reductions in health inequality illustrate the choice. In this scenario the two Populations, A and B, have the same mean life expectancy but the distribution of individuals around that mean is different. The same change occurs in Populations A and B: one individual with life expectancy at the mean gets moved to the lower end of the distribution and another individual from the mean gets moved to the upper end of the distribution. Respondents who think of the two changes as equal are primarily concerned with individuals in a population relative to the mean. Those who consider the two scenarios to be different are not only concerned with how an individual fares compared to the average, but also with how an individual fares compared to everyone else in the population.

Figure 6. An example of a question eliciting preferences for IID vs. IMD measures of inequality.



The figure shows two vertically stacked rectangular boxes. Each box is empty and contains a small icon in its top-left corner. The icon consists of four small squares in a 2x2 grid, colored red, green, blue, and yellow.

Each type of question was asked in more than one way to be able to check for internal consistency in the respondents' answers. In addition, responses were checked for internal consistency as they were entered. For each set of inequality questions, a dialog box appeared if the respondents' answer was inconsistent with prior answers they had given. The dialog box had a short explanatory vignette about the questions and urged the respondent to reconsider their answer. Consistency was not forced and respondents had the choice to keep their original answers. The dialog box appeared only once for each set of questions.

The normative choice about the intensity of the transfer and about the family of measures are not completely separable. Inter-individual comparison measures, even when $\alpha=1$, are more sensitive to equivalent transfers of risk farther from the mean.

We do not propose empirical ethics as a blind tool for resolution of normative choices; rather we believe that the results of measuring these values for a broad range of individuals are a useful input to a deliberative process for choosing an index of health inequality.

Table 1 presents the general characteristics of the respondents to the internet survey . Table 2 shows the preferences of the respondents for α , β and the family of measures.

Table 1. General characteristics of respondents

	Developed	Developing	All respondents
Number of respondents	1022	585	1607
Sex -- % Female	55	45	50
Average age (in years)	39	42	40
Average years of education	18	17	18
% respondents who were WHO staff	28	62	32
Average self-reported health status	Very Good	Very Good	Very Good
Time since last visit to physician	Last 6 months	6-12 months	6-12 months
Time since last hospitalization	>3 years	>3 years	>3 years
Language survey was taken in			
% English	72	70	72
% French	11	13	11
% Spanish	17	17	17

Table 2. Responses on the normative choices

	High income countries	Low and middle income countries	All respondents
Inter-individual measures	76	79	77
Individual-mean measures	24	21	23
<i>More weight to outliers</i>	82	80	81
<i>Equal weights for all parts of the dist'n</i>	18	20	19
Relative	56	65	59
Absolute	44	35	41

Based on the results of the survey, the measure used in the WHR was the following:

$$H(3,.5) = \frac{\sum_{i=1}^n \sum_{j=1}^n |h_i - h_j|^3}{2n^2 \bar{h}^{.5}}$$

H is a relative measure that gives a large weight to the tails of the distribution and is based on comparing the health status of each individual to every other individual in the population.

C. Scale properties and range

As the summary measure of health inequality may be used for a number of purposes including as an input to efficiency analysis, its scale properties are also relevant. Both the individual mean and inter-individual types of summary measures have a defined maximum and minimum value which is a function of the values of α and β . Once the range of the summary measure is defined, it can be rescaled for convenience to range from zero to 1.

More problematic, however, is to establish a summary index of a distribution that has interval scale properties. In fact, this can only be established through some form of empirical preference assessment which demonstrates that a change in inequality from a to $a+b$ is valued as much as from c to $c+b$. It is important to note that to date, such interval scaling properties have not been established for any of the well-known summary measures of distributions, such as the Gini coefficient. Needs for an interval-scaled summary measure impose special concerns on health.

D. Alternative summary metrics

A wide variety of alternative summary measures of a distribution could also be proposed that are not part of the two families presented above. This section describes some of the more commonly used measures. These measures have primarily been used by economists in the measurement of income inequality. The descriptions below are based on Sen's summary of income inequality measures.

Standard deviation of logarithms (H), which is calculated in a similar way to the variance, but before calculating the differences, the income levels are transformed to a logarithm. The logarithmic scale attaches greater importance to income transfers at the lower end and also it removes the arbitrariness of units and of absolute levels of income. It is calculated as follows:

$$H = \left[\sum_{i=1}^n (\log \mu - \log y_i)^2 / n \right]^{1/2}$$

One disadvantage with H as a measure of inequality is that it is not compatible with a concave social welfare function.

Dalton's measure is based on a strictly concave utility function, i.e. with diminishing marginal utility and the same function for all. His measure, D is the ratio of actual social welfare to the maximal social welfare based on the mean income of the population.

$$D = \left[\sum_{i=1}^n U(y_i) \right] / nU(\mu)$$

Atkinson's measure, A , is defined as the “equally distributed equivalent income” of a given distribution of total income, which corresponds to the level of income that which if everyone in the population had would make total welfare equal to the total welfare generated by the current income distribution. Atkinson's measure assumes a concave, but not strictly concave utility function, and is bounded by 0 and 1.

$$y_e = y \mid [nU(y) = \sum_{i=1}^n U(y_i)], \text{ and}$$

$$A = 1 - (y_e / \mu)$$

(A more general normative measure can be calculated in a similar way, if instead of using individual utility functions one were to use the social welfare function, and calculate the generally distributed equivalent income as the level of per capita income which would generate the same amount of welfare as the current income distribution. This measure would be equivalent to A if the welfare function used was a utilitarian one.)

Concentration index

A summary measure that has been used in the literature on income-related inequalities in health is the concentration index. The concentration index is different from all the other measures discussed so far in that it is a summary measure of a bivariate distribution. Anand et al.²³ explore the properties of the concentration index and show that in its use in the measurement of income-related inequalities in health, the concentration index measures the covariance between the income ranks of individuals and their health status. This covariance is normalized so that the index always lies between -1 and 1 , just like any statistical correlation coefficient. The concentration index has been shown to equal:

$$CI = 2 \text{ cov}(i, h_i) / n\mu$$

Where i is the rank of each individual with 1 being the poorest and n being the richest person in the population, h_i is the health status of individual i , and μ is the average health status of the population.

As discussed by Anand et al.²³ the concentration index is not a conventional measure of inequality of either the univariate income distribution or the univariate health distribution. Anand et al. also show that the income distribution can worsen significantly with no effect on the concentration index as long as the relative ranks of individuals remain the same. Similarly, the univariate health distribution can worsen

significantly with no effect on the concentration index as long as the covariance between the income ranks and health status of individuals remains the same.

V. Child survival inequality

This section reviews the extended beta-binomial model which was used to estimate child survival inequality for WHR2000 and compares it to a random-effects logit model, which has been used in more recent work. Estimates of child survival inequality are presented for 50 countries, followed by an exploratory decomposition of the inequality index into contributions from major determinants. The last part of this section includes a description of new datasets that have been identified and collected for use in the next round of estimation.

A. Models used to estimate the distribution of child survival

The distribution of child survival was estimated in 50 countries where a Demographic and Health Survey (DHS) had been conducted and the data were available. The DHS is a 20-year project conducting high quality national sample surveys on population and maternal and child health. Funded primarily by the United States Agency for International Development (USAID), DHS is administered by Macro International Inc.²⁴. Low-income country governments and international organizations have long relied on DHS data to monitor a variety of child and maternal health and family planning indicators²⁵. One of the most significant contributions of the DHS is the collection of internationally comparable data on the demographic and health characteristics of populations in developing countries²⁶⁻²⁹

The DHS are conducted through in-person interviews. The samples, which are all above 3,000 households in the countries analyzed in this study, are the result of a multi-stage stratified sampling design³⁰. The DHS sampling weights are used to produce nationally representative estimates³.

For each country the latest year of available data was used from a nationally representative DHS, ranging from 1987 to 1997. For each mother surveyed the number of children born and the number who survived to age 2 was calculated. The observation period used was ten years, and, to avoid censoring effects, it started two years prior to the interview year. This period is a compromise between providing recent estimates and ensuring enough births to reduce the effects of sampling error. Measuring survival to (or death by) age 5, often referred to as “child mortality,” would involve a longer censoring period, produce older estimates of inequality, and not differ much from the under 2 mortality because on average, 80% of under 5 deaths occurs in the first two years of life^{30,31}

The population of interest includes all children born alive in a country in a given time period. Ideally, we would measure the length of time each child is expected to live

³ To provide an independent validation of the DHS-based results, we also used mortality data by municipality in Mexico⁵⁹ and Brazil⁶⁰ from different data sources.

from birth to two years and then use a measure of inequality to summarize the distribution of these survival expectations. Making the inference from dichotomous data on child survival to health inequality requires several methodological steps.

1. The extended beta-binomial model

The first step is to estimate the distribution of the probability of death across children in each national sample. The chief methodological difficulty here is that for any one child, only the dichotomous variable of survival to two years is measured, and the probability of dying for each child is not observed. These probabilities can be estimated with the extended beta-binomial model³²⁻³⁴. This model has been widely applied in biomedical research, most commonly for modeling animal littermate survival probabilities³⁵. The model compares the observed distribution of child deaths across mothers to that expected on the basis of a binomial distribution with equal risk of dying per child. This comparison allows for the estimation of the underlying distribution of risk of death for children in a population³⁶.

The probability of death, π , for each child follows a binomial distribution. The total number of dead children in a family, Y , is itself distributed binomially with expectation $n\pi$ and variance $n\pi(1-\pi)$, where n is the number of children in the family.

For each family i the binomial distribution can be parameterized as:

$$\Pr(Y_i = y_i | \pi, n) = \binom{n_i}{y_i} \pi^{y_i} (1 - \pi)^{n - y_i}$$

and a simple estimate of π for the whole population is:

$$\hat{\pi} = \sum_i^N y_i / \sum_i^N n_i = \frac{1}{N} \sum_i^N (y_i / n_i)$$

where N represents the number of families. The *estimated* probability of death for each child in the population will be $\hat{\pi}$. And the variance observed would be $n \hat{\pi}(1 - \hat{\pi})$.

Rather than assuming that the probability of a child dying, π , is constant across families, the variability in the unobserved probabilities across families is modeled with a two-parameter beta density, which allows the probabilities to vary even without covariates (or within categories of the covariates).

This is done with a beta-binomial model, allowing the π_i s to follow a beta distribution. The probability π_i for all children belonging to the same mother is modeled as if drawn independently from a common beta distribution. That is each child of a mother, will have the same probability of dying, π_i , but these probabilities will vary across families, following a beta distribution.

The beta distribution is a flexible distribution, bounded by 0 and 1. Conventionally, the beta distribution is represented as:

$$f(\pi | a, b) = \beta^{-1}(a, b)\pi^{a-1}(1-\pi)^{b-1}$$

where a and b are the parameters of the beta distribution and $\beta(a, b) = \Gamma(a)\Gamma(b) / \Gamma(a + b)$

In this analysis, the suggested re-parameterization of Prentice, King and Palmquist³⁷⁻³⁹ was used, where

$$\gamma = (a + b)^{-1}$$

$$\pi = a(a + b)^{-1}$$

The expectation of the beta distribution is π and γ is a dispersion parameter. The variance of the beta distribution is:

$$V(\pi) = \frac{ab}{(a+b)^2} \frac{1}{a+b+1} = \bar{\pi}(1-\bar{\pi}) \frac{1}{1+\gamma^{-1}}$$

A population whose risk distribution followed a beta distribution, would have an average risk of mortality of π and a variance of $V(\pi)$. The actualization of this risk distribution over the course of a time period (two years in the case of child survival) is a binomial distribution where, approximately π percent of children die.

Combining the within-family binomial distribution with the across-family beta distribution produces the extended beta-binomial model³⁷⁻³⁹. It is often a reasonable approximation even when children within the same family have heterogeneous survival probabilities or dependent outcomes.

The simple extended beta-binomial model without covariates estimates *one* beta distribution for each population, i.e. the most likely beta distribution from which the observed distribution of deaths could have come from. The model estimates one π and one γ for each population.

Covariates can also be included in the model. This procedure allows different extended beta-binomial distributions to fit each unique combination of values of the covariates, making for a much more flexible overall model. Covariates were included for the mother's age, number of children, level of education, and average birth interval, all variables that have been shown to affect childhood survival probabilities^{40 41}. The basic model fits well, and adding covariates does not materially affect our estimates of health inequality⁴.

⁴ For Mexico and Brazil, the extended beta-binomial model was also applied to the municipality-level mortality data sets to validate the model. The underlying assumption is that small geographical areas include mostly homogeneous populations for which the risk of death is similar. In both countries, the estimates of inequality from the extended beta-binomial model did not materially differ between the two data sets used.

2. The random effects logit model

The unit of analysis for the EBB model is the mother, and the model estimates the distribution of risk of death for children in the population based on the number of children born and the number of children who have died for each mother. The EBB model estimates the underlying distribution of risk of death for a population. However, it does not estimate child-specific probabilities of death. It uses events – births and deaths of children – per mother to infer the underlying distribution of risk. Because it clusters children by their mother and most of childhood mortality risks are associated with or mediated through the mother, it gives a good estimate of the distribution of risk for each population.

Even though the EBB model fits well there is still a need for individual-level models, i.e. models that will estimate mortality risks for each child and which would also allow for the analysis of the effects of child-level covariates.

Logit and probit models are most commonly used in the case of binary outcome variables, such as death or survival by age 2^{42;43}. Logit and probit models provide estimates of the probability of death for each child. All available variables that are thought to have an effect on childhood mortality risks are included in the model. The parameter estimates for each variable are then used to predict the probability of death for each child. An inequality index can then be calculated based on the distribution for a population of the predicted probabilities.

The accuracy of the logit model predictions of each child's risk of death depends heavily on the quality of the covariates available. In the case of child mortality risk, the information available comes from the Demographic and Health Surveys (DHS)⁴⁴, and therefore is limited to the variables included in the questionnaire. As there are other variables that have an effect on probabilities of death but for which we have no data, any estimate of inequality in child survival based on the predicted distribution from a logit or probit model will be an underestimate of the true inequality in the population.

To deal with the fact that there are unmeasured variables at the household level, a random effect by mother can be added to the logit model. That means that in specifying the logit model, the systematic variation across mothers that is not captured by the available covariates is modeled with a random effect variable. The model does not have the power to detect the magnitude of the effect, i.e. a coefficient, but it gives an estimate of the additional systematic variance across mothers not captured by the existing covariates. In other words, the random effect is a normally distributed variable, with a mean of zero and a standard deviation equal to the systematic variation across mothers not captured by the covariates. The estimated effect has a mean of zero because the model does not have the power to detect the magnitude of the effect on the probabilities and therefore only provides an estimate of the standard deviation. In estimating the predicted probability of death for each child, a random draw from a normal distribution with mean zero and the variance estimated from the model is added to the other covariates. The estimated random effect is at the level of the mother, i.e. children of the same mother are assigned the same random effect; the magnitude of the effect for each mother is a random draw from a normal distribution.

A random effect logit model was run separately for each country. The observed binary variable d is 1 if death or 0 if survival by age 2 for each child. Denote health as d^* , an unobserved variable measuring the health of the child, that generates the observed d 's. For each child i of each mother j , d^* is linearly related to the observed covariates:

$$d_{ij}^* = \mathbf{x}_{ij}\boldsymbol{\beta} + v_j + \varepsilon_{ij}$$

v_j is the random effect for mother j ; it is normally distributed with mean 0 and variance σ^2 ; ε_{ij} is a logistic error term and

$$\mathbf{x}_{ij}\boldsymbol{\beta} = \beta_1 educ_j + \beta_2 asset_j + \beta_3 access_j + \beta_4 rural_j + \beta_5 mage_{ij} + \beta_6 mage_{ij}^2 + \beta_7 birthint_{ij} + \beta_8 birthorder_{ij} + \beta_9$$

where $educ_j$ is the number of years of schooling of mother j , $asset_j$ is the asset index value⁵ for mother j , $access_j$ is mother j 's access to health services (approximated by the percent of her children who had received immunizations), $rural_j$ is whether mother j lives in a rural or urban area, $mage_{ij}$ and $mage_{ij}^2$ is the age and the square of the age of mother j when she had child i , $birthint_{ij}$ is the birth interval preceding the birth of child i of mother j and $birthorder_{ij}$ is the order in which child i was born among the children of mother j .

The probability of death for each child i of each mother j is then computed as follows:

$$\Pr(d_{ij} = 1 | \mathbf{x}_{ij}, \boldsymbol{\beta}, \sigma^2) = \int_{-\infty}^{\infty} \frac{1}{1 + \exp(-\mathbf{x}_{ij}\boldsymbol{\beta} - u_j)} dv, \text{ where the integral averages over the}$$

uncertainty due to the random effect. Note that the probability is a function of the effect parameters $\boldsymbol{\beta}$ and the variance of the random effect σ^2 , as well as the values of the explanatory variables.

3. Estimating the distribution of expected survival time

Once the probability of death for each child has been estimated, the second step is to transform it to the expected survival time in the first two years of life⁴⁵. The probability of death has to be transformed to a positive entity, such as the probability of survival or expected survival time, because inequality measures are designed for positive quantities, i.e. something that one would want more or a higher value of.

Although the results do not change materially, inequality in survival time was measured, instead of probability of survival, as it is analogous to inequality in health expectancy and is more interpretable. Expected survival time can be calculated as

$$S = \frac{1}{{}_2m_0} - \frac{e^{-2{}_2m_0}}{2{}_2m_0}, \text{ where } S \text{ is expected survival time, and } {}_2m_0 \text{ is the mortality rate in}$$

the first two years of life. ${}_2m_0$ can, in turn, be calculated from the probability of


$$\text{dying in the first two years of life, } {}_2q_0 : {}_2m_0 = -\frac{\ln[1 - {}_2q_0]}{2}.$$

As an example, Figure 7 shows the estimated distribution of the probability of dying before age 2 in Benin and the Central African Republic and the corresponding

⁵ The asset index was calculated for each household based on a factor analysis of all available questions on permanent income. A more detailed discussion of estimating permanent income is presented in Section X of this document.

distributions of expected childhood survival time (up to two years) for those countries. These two countries were chosen because they have very similar average probabilities of death (0.13 and 0.12, respectively), and therefore very similar mean survival times (1.86 and 1.87 years, respectively), but markedly different distributions of actual survival time around these means and hence divergent levels of health inequality. For example, in the Central African Republic, about 25% of children born have a probability of death lower than three percent. In contrast, Benin children have risks of death more closely distributed around its mean, with only 4% of its children having a probability of death lower than three percent. Clearly at the lower end of the distributions, Benin does worse, but it does much better at the higher extreme. For example, in Benin less than 1% of children born have a probability of death greater than forty percent, contrasted with the Central African Republic, where more than 4% of children have that probability of death. This is merely one striking example of why summarizing health status with only mean levels is misleading⁶.

Figure 7. Distribution of probability of death between birth and age two ($2q\theta$, in the left graph) and distribution of length of survival between birth and age 2 (in the right graph), for Benin (solid line) and the Central African Republic (dashed line). The curves are density estimates and the vertical lines are the average $2q\theta$ and average length of survival, respectively, for each country.



B. Estimates of child survival inequality

Numerical summaries of the distribution of expected survival time are derived using the WHO health inequality index.

$$II[3,.5] = \frac{\sum_{i=1}^n \sum_{j=1}^n |s_i - s_j|^3}{2n^2 s^{-0.5}}$$

Figure 8 shows a scatterplot of the estimates of inequality in child survival from the extended beta-binomial model and the random effects logit model two models. Table 3 shows for all 50 countries the value of the inequality index and 95% confidence intervals from both models, as well as the difference in the mean value of the index between the two estimates. Countries are ordered by the size of the difference in the estimate of inequality between the two models.

⁶ Because of the covariates used, and parameter estimation, the densities in Figure 7 need not be beta distributions.

Figure 8. Estimates of inequality index from the extended beta-binomial model and the random effects logit model; no covariates used.

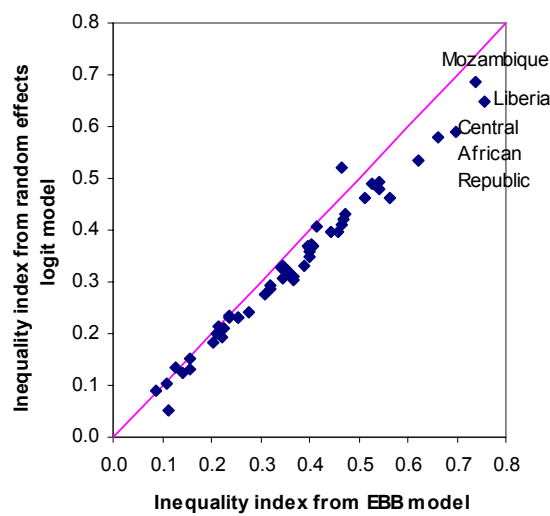


Table 3. Estimates of inequality index and 95% confidence intervals from the extended beta-binomial model and the random effects logit model

Country	EBB	EBB 95% CI	RE logit	RE logit 95% CI	Difference
Mozambique	.74	(.64 - .83)	.64	(.59 - .69)	.10
Rwanda	.56	(.47 - .64)	.47	(.41 - .51)	.10
Central African Republic	.70	(.59 - .80)	.62	(.54 - .68)	.08
Malawi	.62	(.51 - .73)	.54	(.50 - .59)	.08
Namibia	.47	(.37 - .57)	.39	(.32 - .48)	.08
United Republic of Tanzania	.47	(.39 - .55)	.40	(.35 - .45)	.07
Kenya	.34	(.28 - .41)	.28	(.23 - .33)	.06
Ghana	.39	(.29 - .49)	.33	(.27 - .41)	.06
Pakistan	.54	(.47 - .61)	.48	(.42 - .55)	.06
Burkina Faso	.35	(.27 - .42)	.30	(.26 - .33)	.05
Togo	.47	(.40 - .54)	.42	(.37 - .46)	.05
Kazakhstan	.11	(.01 - .22)	.07	(.03 - .11)	.05
Indonesia	.40	(.36 - .44)	.36	(.32 - .39)	.05
Cameroon	.41	(.31 - .50)	.36	(.30 - .43)	.05
Nigeria	.66	(.59 - .74)	.62	(.56 - .67)	.04
Sudan	.41	(.33 - .48)	.36	(.32 - .41)	.04
Cote d'ivoire	.53	(.45 - .61)	.49	(.43 - .53)	.04
Benin	.32	(.24 - .41)	.28	(.25 - .32)	.04
Niger	.54	(.46 - .63)	.51	(.47 - .54)	.03
India	.40	(.38 - .42)	.37	(.35 - .38)	.03
Uzbekistan	.37	(.27 - .47)	.34	(.21 - .51)	.03
Liberia	.76	(.64 - .86)	.73	(.68 - .77)	.03
Burundi	.40	(.30 - .51)	.37	(.31 - .42)	.03
Mali	.51	(.44 - .57)	.49	(.46 - .52)	.02
Bangladesh	.31	(.24 - .37)	.29	(.25 - .32)	.02
Senegal	.23	(.17 - .29)	.21	(.18 - .24)	.02
Trinidad & Tobago	.16	(.09 - .23)	.14	(.06 - .23)	.02
Yemen	.44	(.37 - .51)	.43	(.39 - .47)	.01
Egypt	.36	(.31 - .40)	.34	(.30 - .38)	.01

Country	EBB	EBB 95% CI	RE logit	RE logit 95% CI	Difference
Dominican Republic	.21	(.15 - .28)	.20	(.15 - .25)	.01
Uganda	.35	(.27 - .43)	.34	(.30 - .38)	.01
Comoros	.37	(.25 - .48)	.36	(.27 - .44)	.01
Nepal	.41	(.34 - .49)	.41	(.37 - .44)	.01
Haiti	.40	(.30 - .50)	.39	(.33 - .45)	.01
Morocco	.25	(.19 - .32)	.25	(.20 - .29)	.01
Guatemala	.24	(.19 - .28)	.23	(.20 - .27)	.00
Philippines	.11	(.01 - .22)	.11	(.08 - .13)	.00
Zimbabwe	.22	(.14 - .28)	.21	(.16 - .28)	.00
Madagascar	.46	(.37 - .54)	.46	(.41 - .51)	.00
Peru	.22	(.19 - .26)	.23	(.20 - .26)	-.01
Thailand	.16	(.09 - .21)	.16	(.11 - .24)	-.01
Bolivia	.28	(.21 - .34)	.29	(.25 - .33)	-.01
Colombia	.09	(.05 - .13)	.10	(.07 - .15)	-.01
Nicaragua	.20	(.16 - .25)	.22	(.18 - .26)	-.02
Mexico	.14	(.09 - .19)	.16	(.13 - .20)	-.02
Brazil	.24	(.18 - .30)	.26	(.19 - .31)	-.02
Ecuador	.32	(.23 - .41)	.35	(.27 - .43)	-.03
Tunisia	.26	(.18 - .32)	.29	(.22 - .35)	-.03
Zambia	.47	(.37 - .55)	.50	(.46 - .54)	-.04
Paraguay	.13	(.08 - .18)	.18	(.13 - .24)	-.05

For some countries, such as Mozambique and Rwanda the difference is as high as 0.1; however, the values of the index do not statistically differ for any of the 50 countries. Even for the countries where the mean value of the index differs between the two models, the confidence intervals around those estimates are overlapping, indicating that the difference in the estimates is not statistically significant. Overall, on the DHS data the random effect logit model and the EBB produce very similar estimates of child survival inequality. The correlation coefficient between the mean value of the estimates is .99.

C. Decomposition of child survival inequality index

Once measures of inequality in child survival have been estimated the next step in the analysis is to look at potential determinants. Due to the nature of the inequality index used, a linear decomposition of the contribution of various determinants is not feasible. Therefore, alternative methods have to be developed to quantify the effect of potential determinants on child survival inequality.

In this analysis three main determinants were explored: inequality in assets, inequality in education and access to health services. To get an estimate of the effect of the distribution of income and education on the distribution of child survival time, the following steps are used. Within each country, the asset index for each mother is replaced with the average level of the asset index for that country. This translates into no inequality in assets as all households in the country are set at the mean value. Using the coefficients estimated from the logistic regression model above, estimates of probabilities of death for each child are calculated and then transformed to estimated survival time for the calculation of the inequality index. All other covariates

are left unchanged. Only the asset index value for each mother (and therefore child) is replaced with the average level for the country. This gives us estimates of inequality in survival time in the absence of variation in assets within the country.

Thus the difference between the value of the inequality index estimated this way and the one from the full model provides an estimate of how much inequality in assets contributes to inequality in expected survival time. This approach captures only the direct effect of assets on child survival inequality but does not capture any effect of assets mediated through the other covariates in the model.

The second step is to calculate the further reduction in the inequality index that would result if there were no inequality in educational attainment of mothers in the country. Each mother's educational attainment is replaced with the average number of years of schooling for the country. (The value of the asset index for each mother has already been replaced with the average value for the country.) Using the coefficients estimated from the full logit model, the inequality index is recalculated for each country, in the absence of any variation in assets and in educational attainment across mothers. The difference between the two estimates of the inequality index provides an estimate of the joint effect of education and asset inequality on child survival inequality. Similarly to assets, this approach only captures the direct effect of education on child survival inequality. The independent effect on child survival inequality of inequality in education, after the effect of inequality in assets has been removed can be estimated as the additional reduction in the inequality index resulting from this step of the analysis.

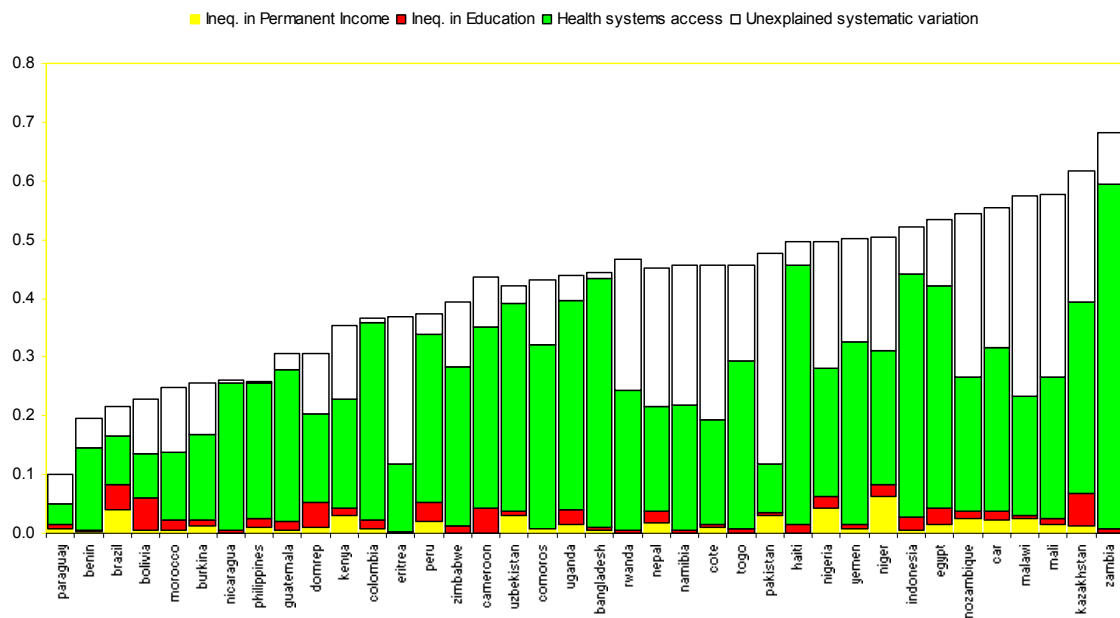
Finally, the effect of not having access to the health system was also quantified. Access to the health system is approximated in this analysis by the percent of children in the household that have received at least one vaccination. This variable may be a better proxy of access to health services in some countries than others; for example, an intensive immunization programme in a country would lead to an overestimation of the percent of households that have access to health services and would therefore underestimate the effect of access to health services on child survival inequality. For child health, however, access to immunization is probably one of the largest determinants of mortality; therefore, the proxy used in this analysis most likely captures the effect that we are most interested in.

The counterfactual scenario is different in the case of access to health services than education or assets. Rather than estimating the effect of inequality in access we quantify the reduction in child survival inequality that would result if the entire population had access to health services. The rationale of quantifying this change both in terms of an increase in the average level of access and in inequality in access is because it is in the mandate of health systems to provide access to health services to the entire population.

In the calculation of the effect of lack of access we replace the percent of children in the household that have received immunizations with 100%, thus simulating a scenario where all children had access to health services.

Figure 9 shows the decomposition of the child survival inequality index according to these 3 determinants.

Figure 9. Contributions of inequality in permanent income, inequality in education, and health systems access to inequality in child survival.



The variation that is labeled “systematic unexplained variation” refers to the part of the inequality index that cannot be attributed to these three covariates; this part includes the contribution of the other covariates in the model, such as mother’s age, birth interval, birth order etc and the contribution of the random effect, i.e. the systematic variation across mothers which is not attributable to the available covariates.

This decomposition analysis is path dependent and the percent attribution of the index might vary if the analysis were performed in a different order. A few different options were tried. If access to health services is explored first, the effects of income inequality and education inequality are minimized. Therefore, access to health services is added to the model last. Interchanging the order in which income inequality and education inequality are controlled for does not materially affect the analysis. Given that a linear decomposition of the index is not feasible, any decomposition analysis will be path-dependent. The order in which the decomposition is done then becomes a normative choice.

D. New datasets available

Since the measurement of child survival inequality for WHR2000 more datasets have been identified and acquired. The Annex includes a list of all available datasets and those expected within a few months that will contribute to the estimation of updated child survival inequality indices.

For the next round of estimation of child survival inequality, in addition to the Demographic and Health Surveys that were used in WHR2000, the birth histories from the Pan-Arab Project on Children (PAPCHILD) and the UNICEF survey program Multiple Indicator Cluster Surveys (MICS) are available. These will help estimate inequality in child survival in most low and some middle-income countries. Similar datasets do not exist for high income countries, as levels of child mortality are low and these countries do not tend to have surveys dedicated to maternal and child health. Survival data from small geographical areas is the only other source of data for these countries. The Annex Table also includes countries with such data available. Where individual-level data from birth histories are not available the extended beta-binomial model will have to be used in place of the logit model.

VI. Adult survival inequality

Unlike with children many fewer datasets are readily available for the estimation of inequality in adult risk of dying. Particularly in low and middle-income countries, information on average levels of mortality is often unreliable, let alone data on the distribution of mortality risk within countries. Intensive efforts have been made to identify datasets that are well suited for the analysis of inequality in mortality risk in adults and models have been modified to fit this analysis.

A. Models for estimating risk of death distribution in adults

Survival analysis models are a good starting point for the estimation of probability of death in age groups other than children. For the most part, data available refer to observations of individuals over the course of several years and dichotomizing this information into survival or death like we do in children would discard valuable data on length of survival.

We opt to use survival analysis models for the estimation of the probability of death in adults. Survival analysis is concerned with analysing the time to the occurrence of an event. The problem with using linear regression to analyse survival data is the assumed normal distribution of residuals. The normal distribution does not apply for events that have a constant risk of occurring every instant, such as survival analysis events. Since results from linear regression are not robust of distributions that are nonsymmetric, distributions likely to be reasonable for time events, linear regression models are not appropriate for analysing survival events.

For adults we want to model survival time in order to estimate baseline hazard (or risk) at time t . This baseline hazard is the component of the hazard that is solely a function of time. The second component of the hazard is a function of covariates. For each adult we have a number of covariates available, such as income, education, age, occupation, race/ethnicity, etc. These covariates can help us estimate the risk of death for individual i ; however there are still many unmeasured covariates and community-level variables that are not captured by individual i 's covariates. Therefore we would like to add a term to capture systematic variation across individuals that is not explained by the available covariates and which operates at the level of the residence/geographic variable.

Therefore for individual i , conditional on their having survived to time t , their hazard (or instantaneous probability) of dying in time t is

$$h(t_i) = e^{\gamma t_i} e^{\beta x_i}$$

The parameter γ is usually positive with mortality data and indicates that the hazard increases with time and as time goes to infinity the hazard goes to 1.

The probability of having survived to time t is given by $S(t)$ and is equal to:

$$S(t) = e^{-\frac{e^{\beta x_i}}{\gamma}(e^{\gamma t} - 1)}$$

The probability density function $f(t)$ is:

$$f(t) = e^{(\beta x_i + \gamma) - \frac{e^{\beta x_i}}{\gamma}(e^{\gamma t} - 1)}$$

The simple Gompertz regression model assumes that the hazard of an individual can be entirely determined by their covariates and the baseline hazard function, based on the γ parameter of the distribution. This parametric specification with the available covariates can only explain part of the variability in observed time to death. Excess systematic variability is known as overdispersion. Overdispersion is caused by misspecification or omitted covariates. “Frailty” models attempt to measure this overdispersion by modelling it as a multiplicative effect on the hazard function. In adult survival data we do not have measures of community-level variables; to capture this unmeasured effect at the community-level we include a group-level effect on the hazard, i.e. the hazard for individual i of group j becomes

$$h(t_{ij} | \alpha_j) = \alpha_j h(t_{ij}) = \alpha_j e^{\gamma t_{ij}} e^{\beta x_{ij}}$$

The hazard of individual i includes his vector of covariates βx_i and the unmeasured variables which are shared within the group j and are estimated by α_j . α_j is a random positive quantity, and for purposes of model identifiability it is assumed to have mean one and variance θ . The Gamma distribution is often used to parameterize this deterministic variability, as it is very flexible and can easily be parameterized to have mean one and variance q . The probability density function of the Gamma($1/\theta, \theta$) is

$$g(x) = \frac{x^{\frac{1}{\theta}-1} e^{-\frac{x}{\theta}}}{\Gamma(\frac{1}{\theta}) \theta^{\frac{1}{\theta}}}$$

Performing the integrations shows that specifying the unmeasured variability as Gamma will result in the survival model where probability of survival to time t , $S(t)$ will be equal to:

$$S_\theta(t) = [1 - \theta \ln\{S(t)\}]^{-\frac{1}{\theta}}$$

This specification of the model is the one that is used in the example of the United States presented below. This model may need to be modified to fit the needs of specific datasets, but it is a good starting point.

B. Individual level data

To estimate the distribution of hazard in each age-sex group in a population individual level data is needed. The ideal dataset would be records of individuals from health surveys or censuses linked to death registration records. This type of dataset would provide multiple years of observation for each individual and a survival outcome at the end of the observation period. The longer the observation period, the better for the model. In some countries it is possible to directly link records from health surveys to death registration; in others techniques have been developed to find probabilistic matches.

1. Direct record linkage

Directly linked death registration and census/ health survey datasets exist for a few countries. A direct record linkage can happen when an individual record has a unique identifier for each person which is the same in the census/ health survey and in the death registration index.

An example of a dataset that is directly linked is the United States data from the National Health Interview Survey from 1987 to 1994 that have been linked to the National Death Index from 1987 to 1995. For individuals in the 1987 National Health Interview Survey data are available for an observation period of 9 years. Survival analysis can be performed on these individuals some of which died during the nine year period, and the majority of whom would be considered “censored” observations, as they were alive at the end of the observation period and we have no further information on them.

Similar datasets exist in a few countries around the world; however, due to confidentiality reasons it not always possible to perform all desired analysis on them. Countries where direct record linkage would be possible include Denmark, Norway, Sweden, the United Kingdom, Finland, Italy, Canada and France.

As the goal of the model is to estimate the distribution of mortality risk by age and sex group in each country, the number of individuals or person years in the dataset has to be rather large.

2. Probabilistic record linkage

In some countries direct record linkage is not feasible because unique identifiers do not exist that could link records from the census or surveys to death registration. Since the 1960s, techniques on probabilistic record linkage have been developed that allow for matches to be made between records from different datasets within certain constraints and probability structures.

Probabilistic record linkage involves assigning agreement and disagreement weights (or odds) for each value of each matching variable. The matching variables commonly used as they are often present in both mortality and census data include geocodes (codes for area of residence – county, municipality, etc), sex, age, ethnic/racial group, country of birth, date of birth, education, and occupation. A larger number of variables common to both records leads to a better overall matching of records between the two datasources.

When records are probabilistically matched a balance has to be found between maximizing the number of links obtained and minimizing the estimated percentage of false-links. Commercial software packages are available that conduct probabilistic record linkage. Statistics Canada have developed a software named *Generalized Record Linkage Software* (GRLS) and another commonly used commercially available package is *Automatch*⁷.

C. An application of the adult survival model to the US data

1. Data

The data used in this example come from the National Health Interview Survey from 1987 to 1994. The records of the interviewed individuals have been linked using a unique person identifier to the National Death Index for the years 1987 to 1995. There are 415158 records in the dataset. Individuals interviewed by the NHIS are a nationally representative sample of the non-institutionalized population of the US aged over 18. The questionnaire includes questions on their socio-demographic characteristics, including age, sex, occupation, education, income, race, as well as questions on how individuals rate their own health and whether they are limited in their usual activities. The questionnaire also has information on the geocode of the place of residence of the respondents; for confidentiality reasons the public files include randomized geocodes, i.e. the users of the data know which individuals lived in the same area, but do not have information on where this area is.

2. Model used

The model used in this analysis was a survival analysis model with a Gompertz distribution on the baseline hazard and a gamma distribution on the random effect. The Gompertz distribution on the baseline hazard means that the risk of death rises monotonically with age; the parameter of the Gompertz distribution γ has been set equal to 0.09 for males and 0.08 for females. The value of this parameter was taken from the life table for the US; it reflects the rate of increase of the mortality risk with age.

Based on the parameterization described above, the hazard of individual i living in area j at time t is being estimated as:

$$h(t_{ij} | \alpha_j) = \alpha_j h(t_{ij}) = \alpha_j e^{\gamma t} e^{\beta x_{ij}}$$

where α_j is the “overdispersion” or excess systematic variability (called random effect in our child survival model) for area j in which individual i lives, γ is set to the value

⁷ Generalised Record Linkage System (GRLS) Software was developed by Statistics Canada. Automatch Generalised Record Linkage System was developed by MatchWare Technologies, Inc.


of 0.09 for males and 0.08 for females, and βx_{ij} is the vector of covariates for individual i which include: age, education race, income, whether they are above or below the poverty line, whether they live alone or not, whether they were employed or not, if they have any difficulties in performing major daily activities, and their rating of their own health status.

3. Preliminary results

A preliminary analysis of variation in predicted risk of death was done using the model described above. The model was run separately for each age-sex group. There are not enough observations in the data to group the population in five year age groups, as would be ideal; the groupings used in this analysis were: 25-34 years, 35-44, 45-54, 55-64, and over 65 years.

Figure 10 shows one example of the output of the survival analysis model. Based on the model a predicted probability of death for each individual can be estimated, based on the individual's set of covariates and the geographic area in which they reside. Figure 10 shows the predicted distribution of risk of death for males and females aged 45-54 years. The distribution of risk of death seems to be less unequal for females than for males; this finding is consistent with the findings of the small area analyses⁴⁶ which reports much greater inequality in adult males than females. Patterns of inequality in risk of death across ages need to be explored further and the causes of these inequalities also need to be researched.

Figure 10. Predicted distributions of risk of death for males and females, aged 45-54 years, US.



This type of analysis should also be performed on individual-level datasets from other countries. A cross-national comparison of inequalities in risk of death at different ages is likely to yield very interesting results.

D. Household survey data

WHO has engaged in a World Health Survey programme conducting household, postal, and telephone surveys in nationally representative samples of the population on health. The results from the first round of the World Health Survey on health state distribution in different populations are described in Section VIII below. Primarily the

World Health Survey has been used to collect information on current health of the interviewees along with their socio-demographic characteristics.

WHO is considering adding questions to collect information on household deaths in the 12 months preceding the interview. Such questions would increase the data available on adult mortality particularly in low income countries. Information on the socio-demographic characteristics of the deceased household members could also be collected. If the sample size were large enough, this would result in an individual-level dataset that could be used to estimate inequality in adult mortality.

In addition to enquiring about deaths in household members, a more ambitious project would entail linking individuals from the World Health Survey probabilistically to census records and to future death registration records. This would result in a dataset of individuals with baseline information on health state and follow-up for mortality.

VII. Using Small Area Analysis to Estimate the Distribution of Mortality Risk Across Individuals

A. Background

Where sufficiently large linked individual record datasets are available, it is possible to estimate the distribution of mortality risk across individuals in an age-group. Unfortunately, for many countries such linked datasets are not currently available. The computerization of the latest round of decennial censuses in many countries holds out the prospect that such linked datasets may in the near future become more common. Nevertheless, for the vast majority of countries alternative methods are needed to quantify the extent of inequalities in mortality risks. For nearly 80 countries in the world with complete vital registration systems and another 40 countries with partial vital registration data, the analysis of mortality rates for small areas is feasible. In this section, we explore the extent to which the analysis of the distribution of mortality rates across small areas can be used as an estimator of the distribution of mortality risks across individuals.

The well-known large variation in mortality risks across small areas⁴⁶⁻⁴⁹ suggests that the distribution of mortality risks across small areas does reveal a considerable fraction of the variance in mortality risks across individuals. For example, within the United States the variation in life expectancy across counties is 61 to 77.5 for males and 70.5 to 83.5 for females. To use small area variation in mortality risk as a valid and reliable indicator of the variation of mortality risk across individuals would require a number of assumptions. In this section, we explore these more carefully and use the United States as an empirical illustration of the strengths and weaknesses of this proxy approach.

B. Principles

If the probability of living in a given small area is independent of mortality risk for any individual then we would observe the same expected value of average mortality risk across all small areas¹². The only observed variation across small areas would be due to chance or measurement error. As the unit for small area analysis was

decreased in size, the observed stochastic variation in death rates would increase. Given random association with appropriate statistical methods (see below), however, no systematic variation in expected average mortality rate would be detected. It is clear, however, in all countries studied that people with similar socio-economic status and other health related covariates are more likely to live together than at random. In addition, community level factors such as environmental quality, quality of health services and community health interventions can also determine mortality risks. For both of these reasons, we expect that there is considerable deterministic variation in expected average death rates across small areas. The survival analysis models that include a random effect by location developed in the previous section show the potential importance of location and also the relatively high covariance between important predictors of individual mortality risk and location.

For deterministic variation in average mortality risk across small areas to be used as a proxy for the variation in mortality risk across individuals, we must address three key questions. First, what fraction of the variance in individual mortality risks is captured by the deterministic variance in average death rates across small areas? Second, in order to make meaningful comparisons across countries, we would also need to know how the measured deterministic variance across small areas is affected by the size of the small areas included in the analysis? Third, is the relationship between deterministic variance in death rates across small areas and the variance in mortality risks across individuals consistent across populations and overtime? If it is consistent then the observed variation across small areas can be used to make comparable estimates of the variation in mortality risks across individuals.

In the following sub-sections, we first explore models that can be used to decompose the observed distribution of mortality rates for an age-sex group across small areas into the stochastic and deterministic components. The next section compares small area assessments for the United States with the linked individual record data analysis presented earlier for the US. The influence of the size of small areas on the results is explored next. Finally, further extensions of small area studies requiring supplementary data are discussed.

C. Models to Decompose Observed Variation in Death Rates Across Small Areas into Stochastic and Deterministic Components

The observed variation in death rates across small areas is due to a combination of stochastic variance which is in proportion to the number of individuals in an age-sex group in each small area and deterministic variation in the average death rate across small areas. Two strategies have been traditionally used to deal with this problem in small area studies. First, small area data have been combined overtime or across adjacent locations to achieve a threshold population size that provides ‘acceptable’ confidence intervals for the mortality parameter being estimated. The main advantage of this approach is that it allows for estimates with a known uncertainty interval for each particular small area or grouping of small areas. The distribution across small areas is, nevertheless, still a function of both stochastic and deterministic variation in the average death rate across small areas. If we are not concerned about obtaining estimates of the expected mortality risk for each small area but only in the

distribution of mortality risks across small areas, then a variety of statistical models can be used to decompose the observed variation in death rates across small areas into stochastic and deterministic components.

1. *The Beta-Binomial Model*

The beta-binomial model presented earlier to estimate the distribution of mortality risk across mothers given data on the survival of their children can also be applied to small area data for an age-sex group. Each small area has a certain number of individuals observed in an age-sex group and a certain number of deaths. Assuming that each individual in the small area experiences the same mortality risk, then the number of deaths observed will be distributed binomial. If we assume that mortality risks across small areas are distributed beta, then the beta-binomial model presented earlier can be used to estimate the deterministic variation in mortality risk across small areas. The main limitation of this model is that the computational time for each age-sex group is very long and depends on the number of individuals in each small area. Given available computing power, it may not be feasible to run this model on datasets that include multiple small areas with large populations.

2. *Normal Mortality Risk-Normal Approximation of the Binomial*

A much more efficient model can be used to estimate the distribution of mortality risk across small areas if some simplifying assumptions are made. First, we assume that for each small area, the observed death rate y_i is distribution normally with mean μ_i and standard deviation σ_i :

$$y_i \sim N(\mu_i, \sigma_i^2)$$

σ_i is assumed to be the normal approximation of the binomial with where σ_i is equal to:

$$\sigma_i = \sqrt{\frac{pq}{N}}$$

where p is the observed death rate and q is one minus the death rate and N is population size in the age-sex group.

Second, we assume that μ_i is distributed normally across small areas with mean μ and standard deviation θ :

$$\mu_i \sim N(\mu, \theta^2)$$

It follows that the distribution of observed death rates is:

$$y_i \sim N(\mu, \sigma_i^2 + \theta^2)$$

Although the assumption that expected mortality risk across small areas is distributed normally cannot be true as mortality risks cannot be less than zero or greater than one, it appears that the approximation gives reasonable answers. Estimates of μ and θ can be obtained using maximum likelihood methods. In other terms, for every age and sex group, the observed death rate has a small area-specific deterministic component – which is the same as saying there fundamental variability or a small area-specific random effect with a mean μ of standard deviation θ – and a stochastic component with mean 0 and (known) standard deviation σ_i .

Using data from the US Burden of Disease study for 2077 counties or country clusters, this model has been estimated for each age-sex group. Table 4 shows the

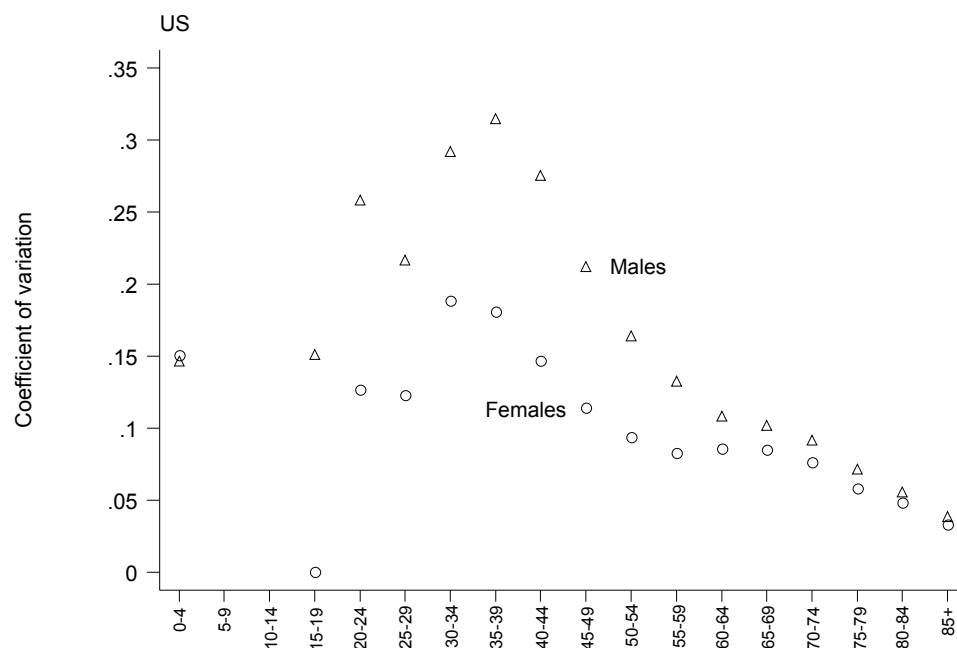
mean death rate across small areas, the observed standard deviation of the death rates and the estimated θ or the standard deviation of the expected death rate.

Table 4. Expected death rates and estimated standard deviation of the expected death rate across small areas in the US.

Age Group	Expected Death Rate		Std Deviation of Expected Death Rate	
	Females	Males	Females	Males
0-4	0.00189	0.00242	0.00028	0.00035
15-19	0.00041	0.00112	1.99E-07	0.00017
20-24	0.00046	0.00155	5.87E-05	0.00040
25-29	0.00055	0.00158	6.84E-05	0.00034
30-34	0.00071	0.00181	0.00013	0.00052
35-39	0.00101	0.00222	0.00018	0.00069
40-44	0.00148	0.00289	0.00021	0.00079
45-49	0.00252	0.00454	0.00028	0.00096
50-54	0.00424	0.00742	0.00039	0.00121
55-59	0.00684	0.01228	0.00056	0.00162
60-64	0.01049	0.01896	0.00090	0.00205
65-69	0.01553	0.02842	0.00132	0.00289
70-74	0.02473	0.04468	0.00188	0.00409
75-79	0.03803	0.06616	0.00220	0.00473
80-84	0.06319	0.10154	0.00305	0.00565
85+	0.10682	0.15262	0.00353	0.0059

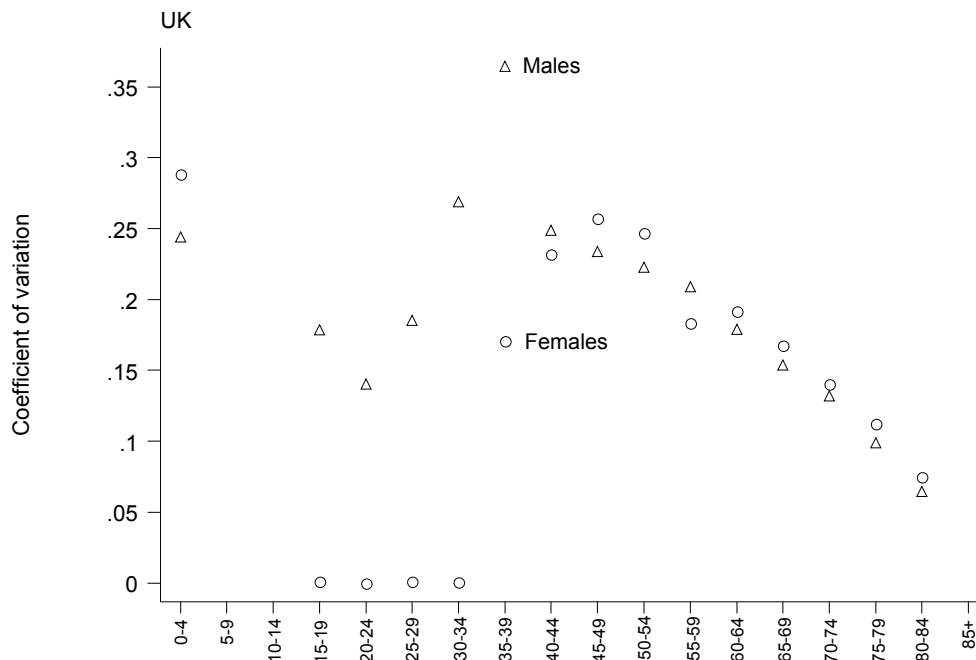
Figure 10 shows for males and for females for the US the coefficient of variation in the expected death rate as a function of age. As can be seen in the graph, the largest coefficient of variation is age-group 35-39 for males and 30-34 for females.

Figure 10. Coefficient of variation in expected death rate, males and females, US.



For comparison, the same model has been estimated for district-level data for the UK, 1997. Data from the UK (illustrated in Figure 11) were available for 403 districts. One notable difference is that, in the UK data, there appears not to be a big difference in the coefficient of variation of the expected death rate for males versus females. In the data from the US, on the other hand, males tended to have a higher expected death rate for almost all age groups.

Figure 11. Coefficient of variation in expected death rate, males and females, United Kingdom.



Comparing age-group 65-69, it appears that the normal-normal model gives a similar standard deviation of mortality risk across small areas as did the more computationally intensive beta-binomial model. One advantage of the beta-binomial model is that it can handle when there are zero deaths in an age-sex group in a small area. Zero deaths in an age-sex group presents some problems for the normal-normal model. In this analysis for districts with zero deaths, the average death rate for that age group across all districts was used in place of a zero death rate.

D. What Fraction of the Total Variance in Mortality Risk Across Individuals is Captured by the Variance in Mortality Risk Across Small Areas?

There is clearly no theoretical answer to the question what is the fraction of total variance in mortality risk across individuals captured by the variance in mortality risk across small areas. If small area analysis captures a predictable fraction of the variance this must be established empirically. If generalizable statements on this relationship are to be made, they will need to be based on careful empirical assessments in a wide range of settings. Here, we present only one comparison based on the individual analysis for adult mortality in the US and small area analysis for the US.

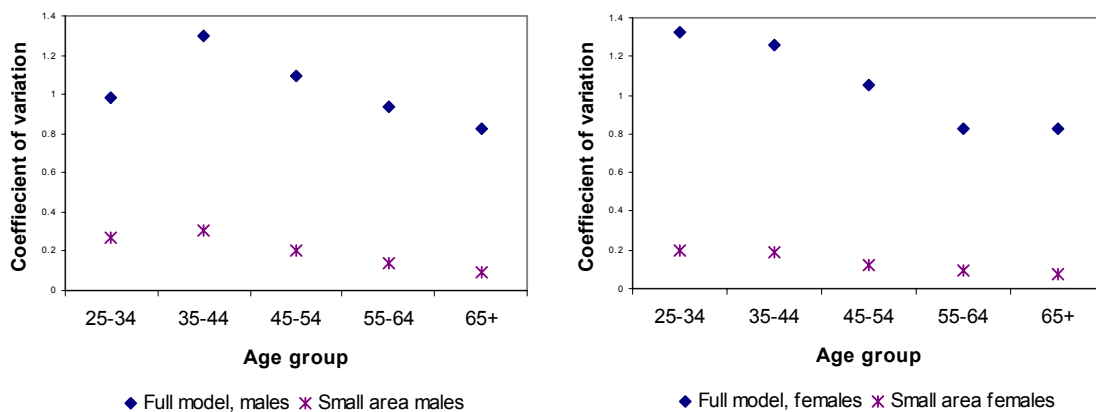
Table 5 shows the standard deviation of the estimated mortality risk across individuals in the US by age and sex using the full survival model and the standard deviation estimated across small areas⁸.

Table 5. Standard deviation of estimated risk of death across individuals from two analysis: full survival analysis model with random effect and small area analysis.

	Full model	Small area analysis
Males		
25-34	0.0010	0.0005
35-44	0.0020	0.0008
45-54	0.0041	0.0012
55-64	0.0092	0.0021
65+	0.0255	0.0050
Females		
25-34	0.0005	0.0001
35-44	0.0013	0.0002
45-54	0.0028	0.0004
55-64	0.0051	0.0009
65+	0.0181	0.0027

The standard deviations increase with age in both analysis, and are larger in the individual-level analysis than in the small area analysis. The individual-level survival analysis is picking up more inequality in risk than the small area analysis. Figure 12 shows for males and for females the coefficient of variation of risk of death with age. The same age pattern is observed from the survival analysis model and the small area analysis in both males and females; as shown in Table 5 above the survival analysis model is showing more inequality in risk of death than the small area analysis.

Figure 12. Coefficient of variation of risk of death from survival analysis model and small area analysis, US, males and females



Further empirical studies in settings where linked individual datasets and small-area data are available will help establish if there is a consistent relationship between the

⁸ Note that the age groups are different from those in the previous section because there are not enough observations in the individual level data to group individuals in smaller age groups.

variation in mortality risk across small areas and the variation across individuals. If such a relationship does exist, it implies that there is some common process that underlies patterns of human aggregation. If true, patterns of human aggregation where individuals with similar mortality risk tend to live close to each other is an interesting and important phenomenon in its own right. Perhaps, the stability overtime of small area deviations in mortality rates observed in the US and the UK over decades (citations) is related to this process. Alternatively, community level factors may be more important determinants of mortality risk than has generally been appreciated.

E. Size of Small Areas and Variation in Mortality Risk

The magnitude of the variance in expected mortality risk across small areas must depend on the size of the small areas being analyzed. In other words, the systematic variation in mortality risks across states in the US is much smaller than across counties. Further research needs to be conducted into the ideal size of small areas that is best suited for this type of analysis so that there are enough person-years in the data to estimate the distribution of risk of death, but the small areas are still homogeneous enough to capture the effect of unmeasured covariates on risk.

F. Supplementing Small Area Analysis

Small area analysis will always underestimate the total variance in mortality rates across individuals. The key problem lies in the fact that individuals within a small area experience different mortality risks. This inter-individual variation is not captured in the small area data. In some countries, such as the Russian Federation, linked individual data is available for a few small areas. This type of data can provide a direct assessment of within small area variance in risk. Individual-level data from a few small areas can be used to improve estimates of the variation in risk achieved with small area data only.

VIII. Inequality in health states

In addition to accurately estimating the distribution of risk of death, the distribution of non-fatal health outcomes also needs to be estimated. Data on health states that is comparable across populations, socio-economic and demographic groups has not traditionally been available. The section below briefly outlines the WHO approach to obtaining cross-population comparable data on health states from the World Health Survey and presents preliminary findings on the distribution of health states for 52 countries.

A. Data & estimation methods

The distribution of health state by age and sex group in any given population is estimated using data from the WHO Multi-Country Household Survey Study. Estimates of health states are obtained using self-report responses to questions on several core domains of health such as mobility, cognition, affect, pain, usual activities, and self-care. One example is the core self-report question on the domain of mobility: individuals are asked to respond to the following question: “Overall in the

past 30 days, how much difficulty did you have with moving around?”. Respondents are asked to pick one of five response categories: “1 = Extreme/Cannot do; 2 = Severe; 3 = Moderate; 4 = Mild; 5 = No problems”.

These self-report categorical response data are not comparable across countries, or even across socio-demographic groups within countries. This is because of response category cut-point shifts⁵⁰. Conceptualizing the observed data as a mapping from an underlying latent variable (e.g., ability on the domain of mobility) to observed categorical responses: response category cut-points mark ranges on the latent variable corresponding to each categorical response. Due to differing norms, expectations, and experiences, these cut-points are likely to be different across countries. Hence, the observed categorical responses are not cross-population comparable. In other words, someone saying they are in “good” health in Denmark may not mean the same thing as someone saying they are in “good” health in Ethiopia in terms of the underlying latent variable that we are trying to measure.

In order to make these data cross-population comparable, we make use of vignettes and the hierarchical ordered probit (HOPIT) model. A vignette is a description of a concrete level of ability on a given domain that respondents are asked to evaluate with relation to the same main question and on the same categorical response scale as the main self-report question. The vignette fixes the level of ability such that variations in categorical responses are attributable to variations in response category cut-points. This introduction of exogenous information in the form of responses to vignettes allows us to identify the effects of a set of socio-demographic covariates (such as age, sex, education, country of residence, etc.) on both the level of the underlying latent variable that is being estimated as well as on the cut-points. In other words, vignettes allow us to estimate cut-point differences by socio-demographic groups and these are then used to calibrate the self-report responses so as to make the cross-population comparable.

After applying the HOPIT model to the survey data, estimates are derived of individuals’ levels of attainment for each of the six domains of health measured in the survey. Attainment on each of the six domains is mapped into a health state valuation using a multivariate valuation function. This valuation function has been developed from data collected in ten nationally representative household surveys where respondents were asked to value a series of hypothetical health states. The primary question used in these surveys was the visual analogue. To be able to convert responses on the visual analogue scale to a strength of preference scale, a multi-methods approach was implemented in a smaller sample in each country. The multi-methods approach includes the standard gamble, time trade-off, person-trade-off and visual analogue. The collection, analysis and interpretation of the valuation data is more fully described elsewhere⁵¹.

A more detailed account on the data, methods used to analyse them and the models used to estimate health states for the populations in the 56 surveys is presented elsewhere.^{52 53 50 51}

B. Preliminary results on inequality in health states

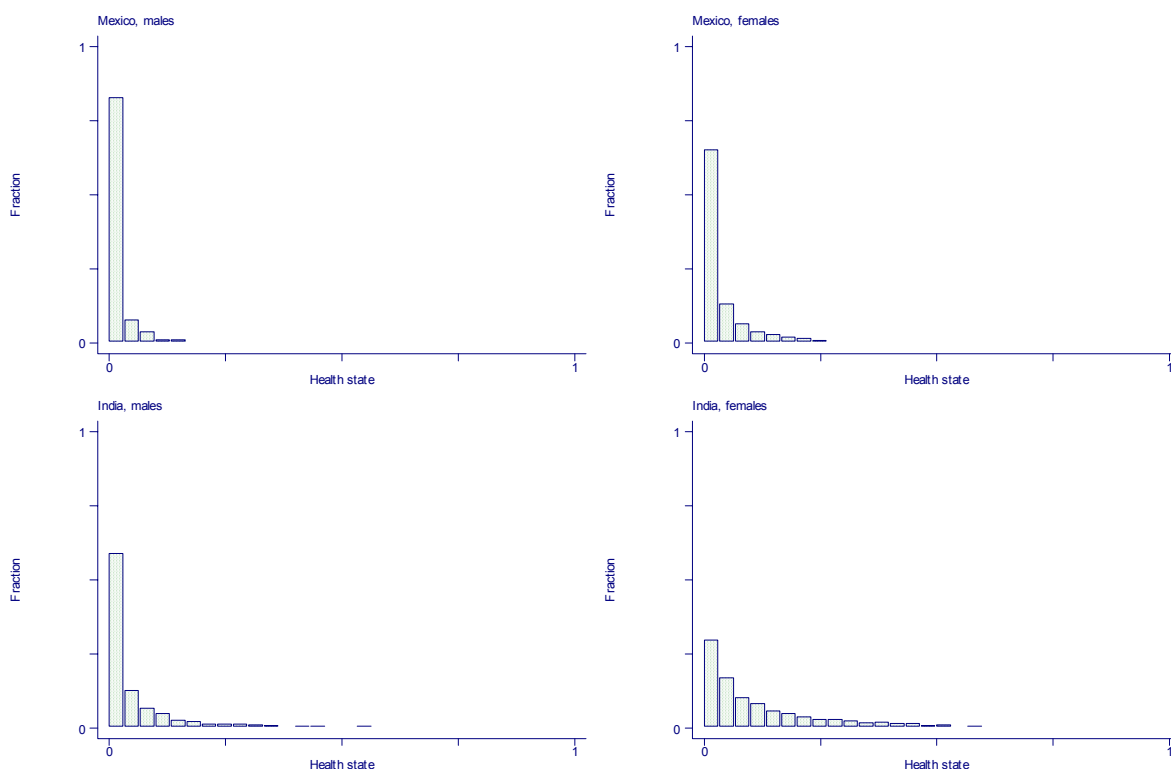
The graphs and tables below demonstrate preliminary findings from the analysis of the World Health Survey. Table 6 presents the standard deviation of health states for the populations in the 56 countries in which surveys have been conducted. The countries are listed in order of decreasing standard deviation in health states.

Table 6. Standard deviation in health states for 56 countries

Country	Standard deviation	Country	Standard deviation
1 Portugal	0.247	29 Sweden	0.149
2 Romania	0.237	30 Austria	0.148
3 Russia	0.210	31 Ukraine	0.147
4 Netherlands	0.205	32 Cyprus	0.146
5 Georgia	0.205	33 United Kingdom	0.144
6 Morocco	0.193	34 Thailand	0.143
7 Bulgaria	0.192	35 USA	0.143
8 Estonia	0.190	36 New Zealand	0.141
9 Slovakia	0.188	37 Costa Rica	0.139
10 India	0.181	38 Korea	0.136
11 Malta	0.179	39 Bahrain	0.134
12 Latvia	0.179	40 China	0.134
13 Croatia	0.178	41 Greece	0.133
14 Egypt	0.176	42 Colombia	0.133
15 Belgium	0.176	43 Argentina	0.131
16 Hungary	0.173	44 Canada	0.130
17 Spain	0.171	45 Ireland	0.130
18 Germany	0.171	46 Island	0.128
19 Poland	0.169	47 Venezuela	0.123
20 Czech Republic	0.167	48 Luxemburg	0.122
21 Kyrgyzstan	0.164	49 Turkey	0.121
22 Lithuania	0.159	50 Australia	0.120
23 Italy	0.159	51 Indonesia	0.119
24 Chile	0.155	52 Denmark	0.119
25 Finland	0.153	53 United Arab Emirates	0.119
26 Trinidad & Tobago	0.152	54 Oman	0.116
27 Jordan	0.151	55 Mexico	0.115
28 France	0.150	56 Switzerland	0.096

Figure 13 below illustrates for adults (aged 15-59) the distribution of health states in two countries, Mexico and India. Health states shown are on a disutility scale with 0 being equivalent to perfect health and 1 being equivalent to death.

Figure 13. Health state distribution for adults aged 15-59, Mexico and India, 2001.



These figures indicate that there is more inequality in health states in adults in India than in Mexico. Mexico has a larger fraction of the population at health states close to 0, i.e. high level of health and a small tail of individuals in worse health states. India has a smaller proportion of adults in high levels of health and a much longer tail in the distribution with a larger proportion of adults in health states worse than .25, while Mexico has a very small fraction in health states less than .25. These graphs also indicate that there is more inequality in health states in adult females than males. In both Mexico and India the distribution of health states in females looks more unequal than in males. This is different from the distributions of mortality risk (at least evidence from small area studies) which indicate that there is more inequality in adult male mortality than in adult female mortality.

These are just two illustrative examples of how the distribution of health states varies across countries. There appear to be different patterns in inequality in health states than in inequality in risk of mortality. Further research needs to be conducted into inequality in health states and potentially measures of inequality in health states should be calculated.

IX. Inequality in healthy life expectancy (HALE)

Ultimately we would want to quantify inequality in HALE. To estimate the distribution of HALE we first need to estimate HALE for individuals in a population. This requires knowing the multivariate distribution of risk of death and health state prevalence at each age. From this distribution we can draw a vector of age-sex

specific risks of death and health states for individual i , estimate their healthy life expectancy and repeat the process to get the distribution of HALE in the population.

The sections above have discussed methods that can be used to estimate the age-sex specific distributions of risk of death and of health states. To be able to estimate the multivariate distribution, two other pieces of information are necessary: the correlation of risk of death and health state for an individual and the correlation of risk of ill-health across ages for an individual. The sections below describe proposed methods that could be used to estimate these correlations.

A. Correlation between risk of mortality and current health state for an individual

Empirically it would be very difficult to estimate the mortality risk and the current health state for a large number of individuals as this would entail following up large cohorts of individuals who have participated in comparable health surveys and looking at their survival status at the end of the follow-up period. Currently datasets do not exist that would allow us to estimate both the current health state and the mortality risk for the same group of individuals. The distribution of mortality risk is most often estimated from a longitudinal study of census records linked to death registration records while the distribution of health states at the same age-sex group is derived from nationally representative health surveys.

One possible way to obtain data on risk of death and risk of non-fatal health outcomes for the same individuals would be to link the World Health Survey programme to death registration over the course of a few years and observe the survival of individuals who participated in the World Health Survey. This would be an ideal dataset.

Until we can have access to such data methods to estimate the correlation between risk of death and current health state need to be developed. Proposals include using the currently observed correlation across countries between age-sex specific mortality rates and health state preferences.

During the course of this consultation we hope to be able to discuss this issue in more detail and make progress on methods that could be used to estimate this correlation.

B. Correlation across ages of the risk of ill-health for an individual

The second methodological challenge in estimating the distribution of HALE for a population entails approximating the correlation of risk of ill-health across time (or age groups) for an individual. For example, if an individual at age 40 is in the lowest fifth percentile of the distribution of risk of ill-health, which part of the distribution would they be in when they are at the age of 45 or 50?

Empirically it would be very difficult to design and estimate a study that would help estimate the correlation across ages of risk of ill-health for an individual. Methods

need to be developed to use all currently available information from different groups of individuals and come up with an estimate of the quantity of interest.

One possible way to do this would be use the correlation of death rates across small areas. As described above small area data can be very powerful in approximating the quantity of interest. Taking the correlation of death rates across age-sex groups for each small area in the country would not result in the correct correlation of risk of death across ages, as there is a large stochastic component in the death rates available for small areas that is due to sampling.

We would want to estimate the correlation of the true death rate across age groups for a small area; thus we need to develop a model that estimates the deterministic and the stochastic component of a death rate and then estimates the correlation across age groups of the deterministic component of the death rate across small areas.

An illustration of the proposed method for two age groups is as follows. Let y_{1i} and y_{2i} be the observed death rates for two age groups 1 and 2 in any given small area i . These can be assumed to be distributed bivariate normal with means μ_{1i} and μ_{2i} , standard deviations σ_{1i} and σ_{2i} , and correlation coefficient $\rho = 0$ for the stochastic component:

$$y_{1i}, y_{2i} \sim N(\mu_{1i}, \mu_{2i}, \sigma_{1i}^2, \sigma_{2i}^2)$$

The means of the expected death rate μ_{1i} and μ_{2i} are themselves assumed to be distributed bivariate normal with means μ_1 and μ_2 and deterministic standard deviations θ_1 and θ_2 . In addition, there is also some correlation between the deterministic variations, ρ which can be estimated as well. The parameters can be estimated using the condition that, given the above structure, the data are distributed bivariate normal:

$$y_{1i}, y_{2i} \sim N(\mu_1, \mu_2, \sigma_{1i}^2 + \theta_1^2, \sigma_{2i}^2 + \theta_2^2, \rho^*)$$

where the correlation between the deterministic components ρ can be recovered from the estimated correlation ρ^* using the following formula:

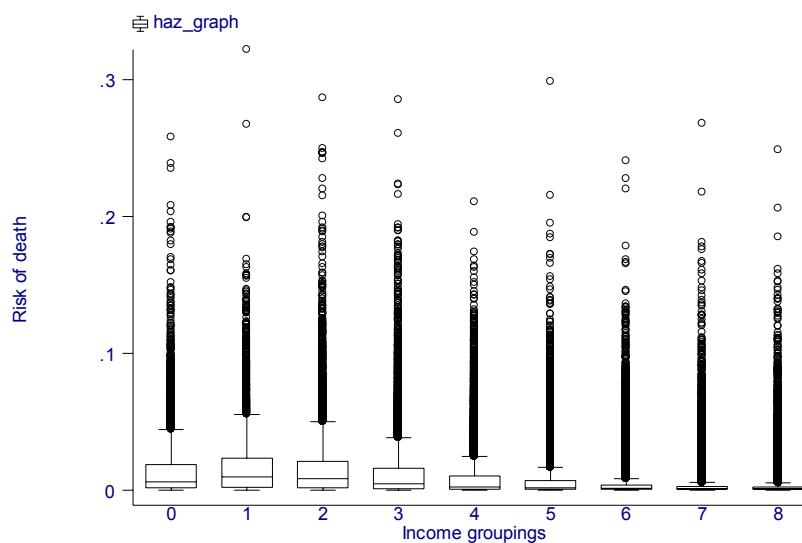
$$\rho = \frac{\sqrt{\sigma_{1i}^2 + \theta_1^2} \cdot \sqrt{\sigma_{2i}^2 + \theta_2^2}}{\theta_1 \theta_2} \cdot \rho^*$$

Obviously the model would be much more powerful if we estimated the correlations between all age groups simultaneously, rather than two at a time. Therefore, we would want to have a multivariate distribution of the death rate for a small area across different age groups. This may not be computationally feasible though and the less desirable method of estimating the correlations two at a time may need to be used.

X. Health of the poor

WHO and many governments have a particular interest in health outcomes in the poor. This interest is also manifested in the international development targets which make explicit reference to health outcomes in the poor^{54;55}. In countries with linked individual data on mortality and various individual covariates including household income, it is a simple matter to calculate and report HALE in the poor. For example, the record linkage data for the US allows the comparison shown in Figure 14 of the probability of death versus income per capita of the household. Average probabilities of death can be reported for the choice of any poverty line. By using the probability distributions for each age-group, the distribution of non-fatal health outcomes and the variance-covariance matrix of each of these age-sex specific distributions, the HALE for the poor can be calculated.

Figure 14. Risk of death by income group, US. Income group 0 is the poorest and 8 the richest.



Such a bivariate relationship also points the direction for the operationalization of multivariate constructs of poverty. For example, the poverty line might be defined in the space of income and health – a two-dimensional health and income poverty line. Although there has been considerable discussion of multi-dimensional constructs of poverty, there has been no cross-population comparable attempts at assessing poverty taking into account health.

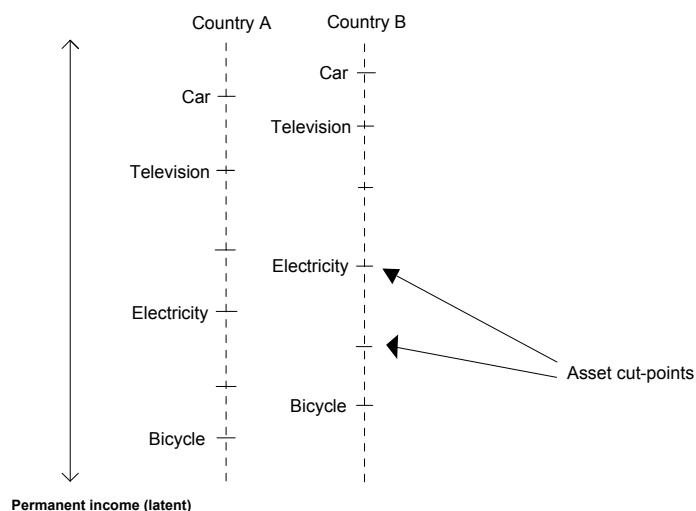
For many countries, however, where the primary source of data for the assessment of levels and inequalities in health will be household survey data, a main problem is valid, reliable and comparable assessment of income or expenditure. Full income and expenditure modules such as used in the World Bank's Living Standards Measurement Studies (LSMS) are long and cannot easily be adapted to surveys where the primary focus is the collection of health data. Despite huge improvements in the information base to assess levels of child mortality through the Demographic and Health Surveys, there is little or no direct information on child mortality in the absolute poor.

Filmer and Pritchett ⁵⁶ have proposed that a short-list of questions on assets and indicators of permanent income such as electricity can be used to provide a reasonable assessment of income. Gwatkin et al. have used this asset index to examine differentials in child mortality in various DHS surveys. The decomposition of inequality in child mortality presented in the section above has also made use of this type of asset index.

Incorporating simple questions on indicators of permanent income in surveys is a promising direction for increasing the information available to assess levels of health in the poor. The asset index developed so far has made use of asset or indicator questions that were included in household surveys primarily for other purposes. The primary method used to develop the index has been factor analysis which has one of two key limitations. First, if analyzed on a country by country basis, the factor score for a household does not provide an estimate of permanent income that is comparable across countries. Second, if the factor analysis is undertaken on data for multiple countries at once, it assumes that the contribution of each asset or indicator variable towards the factor score is the same across countries.

As part of the development of the World Health Survey, WHO is using the Hierarchical Ordered Probit model (HOPIT) to map the level of permanent income at which a household is likely to have an asset or a particular indicator, such as running water or electricity. Figure 15 shows the principle that on the latent variable of permanent income, there are cutpoints at which for a given household the probability of having an asset or indicator is 50%. This asset or indicator ladder then allows for a survey instrument to have the most appropriate set of asset or indicator variable questions. Figure 15 illustrates that ownership of a certain asset does not necessarily imply the same level on the latent variable of permanent income in different countries. This is due to cut-point shifts that may be due to different prices for assets or indicators, or they may be due to differences in norms and needs across populations⁵⁰. This is the same construct as was presented earlier in Section VIII.A where the use of HOPIT was demonstrated in estimating the latent variable of health using self-reported categorical survey data.

Figure 15. Differences in cut-points across countries for a set of asset or indicator variables.



As many indicator variables as are available can be mapped to the latent variable of permanent income. Given a set of responses on these questions, the household level of permanent income can be estimated. Units of permanent income can be converted into local currency or \$I units taking advantage of the linear relationship between the latent variable and income in International dollars. The indicator variable approach to estimating household permanent income can be validated by comparing estimates to recorded household income and expenditure.

By analyzing existing data from LSMS, DHS, MICS and other household surveys, it should be possible to map the asset/indicator ladder in a large number of countries. The permanent income latent variable estimated for each country can be made comparable in terms of international dollars if the true variance of permanent income in each country is known or can be estimated. With country specific asset/permanent income indicator ladders available, then surveys such as the World Health Survey can use an appropriate set of asset questions to develop valid, reliable and comparable measures of permanent income. This information will make it much more feasible to report on levels of health in the poor in a wide range of countries.

XI. Discussion

This background document has presented the work undertaken by WHO in the measurement of health inequalities for the WHR2000 and since. Ultimately WHO aims to measure inequality in HALE. As demonstrated in Annex III there are many inputs to the estimation of the distribution of HALE. This paper has presented progress in implementing these measurement approaches and some of the remaining challenges.

In the area of child survival inequality new models have been developed which allow for the decomposition of the inequality index according to potential determinants, such as income, education and access to health services. Understanding the determinants of child survival inequality is a critical step in the formulation of policy recommendations that will lead to the reduction in inequalities.

In extending the measurement of inequalities from child survival to inequalities in HALE, one promising avenue of research is the use of survival analysis models. These models have been modified to measure the excess systematic variation, not captured by the available covariates, which improve the estimation of the probability of dying by age and sex for adults. One of the main methodological challenges remaining is to account for inter-age group correlations in the systematic variation in risk of death, as well as intra-age group correlation between risk of death and current health state.

Finally, new methods are introduced for the measurement of permanent income using asset and other indicator variables. These methods can lead to a more accurate and cross-population comparable measurement of permanent income which will feed into the better estimation of the bivariate relationship between health and income, and help report on health of the poor as a first step in designing targeted health policies.

XII. References

1. World Health Organization. World Health Report 2000. Health Systems: improving performance. Geneva, World Health Organization, 2000.
2. World Health Organization. The World Health Report 2001. Mental Health 2001: New Understanding, New Hope. Geneva: World Health Organization, 2001.
3. Wagstaff A. Socioeconomic inequalities in child mortality: comparisons across nine developing countries. *Bulletin of the World Health Organization* 2000;78:19-29.
4. Braveman P, Starfield B, Geiger HJ. World Health Report 2000: how it removes equity from the agenda for public health monitoring and policy. *BMJ* 2001;323:678-81.
5. Houweling TA, Kunst AE, Mackenbach JP. World Health Report 2000: inequality index and socioeconomic inequalities in mortality (letter). *Lancet* 2001;357:1671-2.
6. Almeida C, Braveman P, Gold MR, Szwarcwald CL, Ribeiro JM, Miglionico A *et al.* Methodological concerns and recommendations on policy consequences of the World Health Report 2000. *Lancet* 2000;357:1692-7.
7. Gakidou E, Murray CJL, Frenk J. Defining and measuring health inequality: an approach based on the distribution of health expectancy. *Bulletin of WHO* 2000;78:42-54.
8. Murray CJ. Commentary: comprehensive approaches are needed for full understanding. *BMJ* 2001;323:678-81.
9. Gakidou E, King G. An individual-level approach to health inequality: Child survival in 50 low and middle-income countries. GPE discussion paper series, No. 18. 2001. Geneva, World Health Organization.
10. Manton KG. Changing Concepts of Morbidity and Mortality in the Elderly Population. *Milbank Memorial Foundation Quarterly/Health and Society* 1982;60:183-244.
11. Mathers, C., Murray CJL, Salomon, J. A., and Lopez, A. D. Estimates of healthy life expectancy for 191 countries in the year 2000: methods and results. GPE discussion paper series, No.38. 2001. Geneva, World Health Organization.
12. Wolfson M, Rowe J. On measuring inequalities in health. *Bulletin of the World Health Organization* 2001;79:553-60.
13. World Bank. World Development Report 2000/2001: Attacking Poverty. Oxford: Oxford University Press for the World Bank, 2001.
14. Sen A. Development as Freedom. New York: Random House, Inc., 1999.
15. Sen A. Inequality reexamined. Cambridge: Harvard University Press, 1992.
16. Kakwani NC. On the estimation of income inequality measures from grouped observations. *The review of econometrics studies* 1976;43:483-92.

17. Nygard F, Sandstrom A. Income Inequality Measures Based on Sample Surveys. *Journal of econometrics* 1989;42:81-95.
18. Dagum C, Zenga Me. Income and wealth distribution, inequality and poverty: Proceedings of the Second International Conference on Income Distribution by Size: Generation, Distribution, Measurement and Applications, held at the University of Pavia, Italy, September 23-30, 1989. New York; Berlin; London and Tokyo: Springer, 1990.
19. Sen A. On economic inequality. Oxford: Oxford University Press, 1997.
20. Sen A. On economic inequality. Oxford: Oxford University Press, 1997.
21. Gini C. Variabilita e mutabilita. Bologna: 1912.
22. Lorenz MO. Methods for measuring the concentration of wealth. *Journal of the American Statistical Association* 1905;9.
23. Anand S, Diderichsen F, Evans T, Shkolnikov VM, Wirth M. Measuring Disparities in Health: Methods and Indicators. In Evans T, Whitehead M, Diderichsen F, Bhuiya A, Wirth M, eds. *Challenging inequities in health: From ethics to action*, pp 49-67. New York: Oxford University Press for the Rockefeller Foundation, 2001.
24. USAID & Macro International, Inc. Demographic and Health Surveys. www.measuredhs.com . 1984.
25. Stanton C, Abderrahim N, Hill K. An assessment of DHS maternal mortality indicators. *Studies in Family Planning* 2000;31:111-23.
26. Sullivan, JM, Rutstein, SO, and Bicego, GT. Infant and Child Mortality. 15. 1994. Calverton, Macro International, Inc. 1994.
27. Curtis, S. L. Assessment of the quality of data used for direct estimation of infant and child mortality in DHS-II surveys. 1995. Calverton, Maryland, Macro International. Demographic and Health Surveys.
28. Boerma JT, Sommerfelt AE. Demographic and Health Surveys (DHS): contributions and limitations. *World Health Statistics Quarterly* 1993;46:222-6.
29. Hill K, Pande R, Mahy M, Jones G. Trends in child mortality in the developing world: 1960-1996. New York: UNICEF, 1999.
30. Rutstein S. Factors associated with trends in infant and child mortality in developing countries during the 1990s. *Bulletin of the World Health Organization* 2000;78:1256-70.
31. Macro International Inc. An assessment of the quality of health data in DHS surveys. 2. 1993. Calverton, Maryland, Macro International Inc.
32. Prentice RL. Binary regression using an extended beta-binomial distribution, with discussion of correlation induced by covariate measurement errors. *Journal of the American Statistical Association* 1986;81:321-7.

33. King G. Unifying political methodology: the likelihood theory of statistical inference. Cambridge: Cambridge University Press, 1989.
34. Palmquist, B. Analysis of proportions data. 1999. College Station, Texas.
35. Griffiths D. Maximum likelihood estimation for the beta-binomial distribution application to the household distribution of the total number of cases of a disease. *Biometrics* 1973;29:637-48.
36. Brooks SP, Morgan BJ, Ridout MS, Pack SE. Finite mixture models for proportions. *Biometrics* 1997;53:1097-115.
37. Prentice RL. Binary regression using an extended beta-binomial distribution, with discussion of correlation induced by covariate measurement errors. *Journal of the American Statistical Association* 1986;81:321-7.
38. King G. Unifying political methodology: the likelihood theory of statistical inference. Cambridge: Cambridge University Press, 1989.
39. Palmquist, B. Analysis of proportions data. 1999. College Station, Texas.
Ref Type: Conference Proceeding
40. Bicego, G. and Ahmad, OB. Infant and Child Mortality. 20. 1996. Calverton, Macro International, Inc. 1996.
41. Sullivan, JM, Rutstein, SO, and Bicego, GT. Infant and Child Mortality. 15. 1994. Calverton, Macro International, Inc. 1994.
42. King G. Unifying political methodology: the likelihood theory of statistical inference. Cambridge: Cambridge University Press, 1989.
43. Greene W.H. Econometric analysis. New York: Macmillan Publishing Company, 1990.
44. Demographic and Health Survey. <http://www.macrint.com/dhs/>.
45. Preston SH, Keyfitz N, Schoen R. Causes of death: Life tables for national populations. New York and London: Seminar Press, 1972.
46. Murray, C. J. L., Michaud, C., McKenna, M., and Marks, J. M. U.S. patterns of mortality by county and race: 1965-1994. 1998. Cambridge, Harvard School of Public Health and National Center for Disease Prevention and Health Promotion.
47. Benach J, Yasui Y, Borrell C, Saez M, Pasarin MI. Material deprivation and leading causes of death by gender: evidence from a nationwide small area study. *Journal of Epidemiology & Community Health* 2001;55:239-45.
48. Reading R, Colver A, Openshaw S, Jarvis S. Do interventions that improve immunisation uptake also reduce social inequalities in uptake? *BMJ* 1994;308:1142-4.
49. Galal OM, Qureshi AK. Dispersion Index: measuring trend assessment of geographical inequality in health--the example of under-five mortality in the Middle East/north African region, 1980-1994. *Soc Sci Med* 1997;44:1893-902.

50. Tandon A, Murray C., and Salomon, J. A. Statistical Methods for Enhancing Cross-Population Comparability. GPE discussion paper series [in progress] . 2001. Geneva, World Health Organization.
51. Salomon, J. A., Tandon A, and Murray CJL. Using Vignettes to Improve Cross-Population Comparability of Health Surveys: Concepts, Design, and Evaluation Techniques [in progress]. GPE discussion paper series, 2001. Geneva, World Health Organization.
52. Sadana R, Salomon, J. A., Tandon A, Chatterji S, and Murray C. Health state vignettes: design, empirical analysis and critical assessment. GPE discussion paper series [in progress]. 2001. Geneva, World Health Organization.
53. Murray CJL, Tandon A, Salomon, J. A., Mathers, C., and Sadana R. Cross Population Comparability of Evidence for Health Policy. GPE discussion paper series [in progress]. 2001. Geneva, World Health Organization.
54. World Bank. World Development Indicators. Washington, DC: The International Bank for Reconstruction and Development/The World Bank, 2000.
55. World Bank. World Development Report 2000/2001: Attacking Poverty. Oxford: Oxford University Press for the World Bank, 2001.
56. Filmer D, Pritchett L. Estimating wealth effects without expenditure data - or tears: An application to educational enrollments in states of India [forthcoming]. *Demography* 2001.
57. Parfit D. Reasons and persons. Oxford: Oxford University Press, 1984.
58. Pigou AC. The economics of welfare. London: Macmillan, 1932.
59. Fundacion Mexicana para la Salud and Instituto Nacional de Estadística Geografía e Informática. Deaths and population by municipality in Mexico. 1998.
60. Instituto Brasileiro de Geografia e Estadística (IBGE). Sistema de Informacao sobre Mortalidade. <http://www.datasus.gov.br>, <http://www.ibge.gov.br> . 1994. <http://www.datasus.gov.br>, <http://www.ibge.gov.br>.

Annex II. Data available for the measurement of health inequalities.

Country	Dataset type*	Year	Sample Size
Afghanistan	MICS	2000-01	21923
Albania	MICS	2000-01	3113
Algeria	PAPCHILD	1992	
	MICS	2000-01	30774
Angola	MICS	2000-01	12479
Argentina	Small area		
Australia	Small area	1990-93	
Austria	Small area	1988-94	
Azerbaijan	MICS	2000-01	7697
Bangladesh	DHS	1994	
	DHS	1997	9127
	MICS	2000-01	126947
Benin	DHS	1996	5491
Bhutan	MICS	2000-01	2064
Bolivia	DHS	1989	
	DHS	1994	8603
	DHS	1998	11187
	MICS	2000-01	8142
Bosnia and Herzegovina	MICS	2000-01	3839
Botswana	DHS	1988	4368
	MICS	2000-01	1597
Brazil	DHS	1986	
	DHS	1991	
	DHS	1996	12612
	Small area	1979-97	
Bulgaria	Small area	1993-98	
Burkina Faso	DHS	1993	6354
	DHS	1998	6445
	MICS	2000-01	11616
Burundi	DHS	1987	3970
	MICS	2000-01	6565
Cameroon	DHS	1991	
	DHS	1997	5501
	MICS	2000-01	14693
Canada	Small area	1995-97	
	Adult, direct		
Central African Republic	DHS	1995	5884
	MICS	2000-01	3550
Chad	DHS	1997	
	MICS	2000-01	7458
Chile	Small area	1997	
China	Small area	1990	
Colombia	DHS	1986	
	DHS	1990	
	DHS	1995	11140
	DHS	2000	11585
	Small area	1979-96	
Comoros	DHS	1996	3050
	MICS	2000-01	676
Congo	MICS	2000-01	50335
Costa Rica	Small area		

Country	Dataset type	Year	Sample Size
Cote d'Ivoire	DHS	1994	8099
	MICS	2000-01	14526
Cuba	MICS	2000-01	11160
Denmark	Small area		
	Adult, direct	1970-86	
Dominican Republic	DHS	1986	
	DHS	1991	
	DHS	1996	8422
	MICS	2000-01	8364
Ecuador	DHS	1987	4713
	MICS	2000-01	12411
Egypt	DHS	1993	
	DHS	1995	14779
	PAPCHILD	1991	
El Salvador	DHS	1985	5207
	MICS	2000-01	6154
England and Wales	Small area	1991-97	
Equatorial Guinea	MICS	2000-01	442
Eritrea	DHS	1995	5054
Ethiopia	DHS	2000	
Finland	Adult, direct	1975-2000 (by 5 year)	
France	Small area	1990-97	
	Adult, direct	1975-89	
Gambia	MICS	2000-01	1268
Georgia	MICS	2000-01	5006
Germany	Small area		
Ghana	DHS	1988	
	DHS	1994	4562
	DHS	1998	4843
Guatemala	DHS	1987	
	DHS	1995	12403
	DHS	1998	6021
Guinea	DHS	1999	
	MICS	2000-01	7360
Guinea Bissau	MICS	2000-01	1187
Guyana	MICS	2000-01	855
Haiti	DHS	1995	5356
Hungary	Small area	1990-97	
India	DHS	1993	89777
	MICS	2000-01	998056
Indonesia	DHS	1981	
	DHS	1991	
	DHS	1994	28168
	DHS	1997	28810
	MICS	2000-01	209255
Iraq	MICS	2000-01	22450
Italy	Adult, direct	1981	
Japan	Small area	1995-97	
Kazakhstan	DHS	1995	3771
	DHS	1999	4800

Country	Dataset type	Year	Sample Size
Kenya	DHS	1989	
	DHS	1993	7540
	DHS	1998	7881
	MICS	2000-01	29549
Korea	MICS	2000-01	23702
Kyrgyzstan	DHS	1997	
	MICS	2000-01	4669
Lao PDR	MICS	2000-01	5297
Lebanon	PAPCHILD	1996	
	MICS	2000-01	3236
Lesotho	MICS	2000-01	2108
Liberia	DHS	1986	5239
Libyan Arab Jamahiriya	PAPCHILD	1995	
	MICS	2000-01	5471
Madagascar	DHS	1997	7060
	MICS	2000-01	15497
Malawi	DHS	1992	4849
Malaysia	Small area	1991-99	
Maldives	MICS	2000-01	278
Mali	DHS	1987	
	DHS	1996	9704
Mauritania	PAPCHILD	1990	
Mexico	DHS	1987	9310
	Small area	1990-96	
	Adult, probabilistic		
Moldova, Republic of	MICS	2000-01	4380
Mongolia	MICS	2000-01	2621
Morocco	DHS	1987	
	DHS	1992	9256
	DHS	1995	4753
	PAPCHILD	1997	
	MICS	2000-01	27867
Mozambique	DHS	1997	8779
Myanmar	MICS	2000-01	45059
Namibia	DHS	1992	5421
Nepal	DHS	1996	8429
	MICS	2000-01	23385
Netherlands	Small area	1991-98	
New Zealand	Small area	1991-96	
	Adult, probabilistic		
Nicaragua	DHS	1998	13634
Niger	DHS	1992	
	DHS	1995	7577
	DHS	1998	
	MICS	2000-01	10400
Nigeria	DHS	1990	8781
	DHS	1999	7647
	MICS	2000-01	108945
Norway	Small area	1996	
	Adult, direct	1970,1980, 1990	
Pakistan	DHS	1991	6611

Country	Dataset type	Year	Sample Size
Panama	Small area		
Papua New Guinea	DHS	1998	13016
Paraguay	DHS	1990	5827
Peru	DHS	1986	
	DHS	1992	
	DHS	1996	28951
Philippines	DHS	1993	
	DHS	1998	13983
	MICS	2000-01	74454
	Small area	1997	
Poland	Small area	1994-96	
Portugal	Small area		
Romania	MICS	2000-01	22402
Russia	Small area	1990-98	
Rwanda	DHS	1992	6551
	MICS	2000-01	7235
Sao Tome & Principe	MICS	2000-01	144
Scotland	Small area		
Senegal	DHS	1986	
	DHS	1993	
	DHS	1997	8593
	DHS	1999	17189
Senegal	MICS	2000-01	9240
Sierra Leone	MICS	2000-01	4717
Somalia	MICS	2000-01	9672
South Africa	Small area	1968-90, 1995	
Sri Lanka	DHS	1987	
	Small area		
Sudan	DHS	1990	5860
	PAPCHILD	1993	
	MICS	2000-01	28883
Suriname	MICS	2000-01	415
Swaziland	MICS	2000-01	980
Sweden	Small area	1979-97	
	Record linkage, direct	1960-85	
Syria	PAPCHILD	1993	
	MICS	2000-01	15725
Tajikistan	MICS	2000-01	6104
Thailand	DHS	1987	6775
	MICS	2000-01	60856
	Small area		
Togo	DHS	1988	
	DHS	1995	
	DHS	1998	8569
	MICS	2000-01	4512
Trinidad and Tobago	DHS	1987	3806
	MICS	2000-01	1289
Tunisia	DHS	1988	4184
	PAPCHILD	1995	
	MICS	2000-01	9460
Turkey	DHS	1998	8576
	MICS	2000-01	65546

Country	Dataset type	Year	Sample Size
Uganda	DHS	1989	
	DHS	1995	7070
Ukraine	MICS	2000-01	50658
United Kingdom	Record linkage, direct	1970,80,90, >1993	
United Republic of Tanzania	DHS	1992	
	DHS	1996	8120
	DHS	1999	4029
United States	Record linkage, direct	1987-1994	
Uruguay	Small area		
Uzbekistan	DHS	1996	4415
	MICS	2000-01	23942
Venezuela	MICS	2000-01	23706
Viet Nam	MICS	2000-01	78705
Yemen	DHS	1992	6010
	PAPCHILD	1992	
Yugoslavia	MICS	2000-01	10637
Zambia	DHS	1992	
	DHS	1996	8021
	MICS	2000-01	8976
Zimbabwe	DHS	1988	
	DHS	1994	6128
	DHS	1999	5907

* Dataset type key

DHS = Demographic and health survey

PAP CHILD = Pan-Arab Project on Children

MICS = Multiple indicator cluster surveys

Small area

Record linkage = Individual level data for adults, linked either directly or probabilistically

Annex III. Inputs to the measurement of inequality in HALE

