# Univariate Assessment of Health Inequalities – Technical Report

**DRAFT, Revised 24 February 2012**

This draft technical report describes how univariate inequality measures from the income inequality and poverty literature may be applied to the measurement of health inequality. The aim of this document is to describe our work in progress on this topic for discussion at a workshop to be held on 12th March 2012 in York.

An accompanying document titled 'Overview' is a more general discussion of a proposed analytical framework for extending cost effectiveness analysis of national health sector programmes to incorporate concerns about health inequality.

**Project title:** Identifying appropriate methods to incorporate concerns about health inequalities into economic evaluations of health care programmes

**Investigators:** Susan Griffin, Richard Cookson, Nigel Rice, Miqdad Asaria, Karl Claxton, Tony Culyer, Mark Sculpher, Centre for Health Economics, University of York

**Funding:** This project is funded through the Public Health Research Consortium under the Department of Health Policy Research Programme

**Contents**

## 1. Introduction

This chapter is a guide through a set of inequality indices and other normative assessment rules that are commonly used in the standard literature on economic inequality and poverty to rank distributions of income or wealth. Here we explore how these rules might be applied to univariate distributions of health using a series of fictional stylised examples. Useful references to the standard economic inequality and poverty measurement literature include (A. K. Sen, 1973), (A. B. Atkinson, 1983), (Lambert, 2001), (Kolm, 1976a), (Cowell, 2011) , (Jenkins & Lambert, 1997), and (Coudouel, Hentschel, & Wodon, 2002) useful references to the standard health inequality literature include (O'Donnell, van Doorslaer, Wagstaff, & Lindelow, 2008), (Mackenbach & Kunst, 1997), (Fleurbaey & Schokkaert, 2009).

In each case we present the relevant rule from the economic inequality literature, the assumptions required for it to apply, and the implications of these assumptions in the health inequalities context. The rules are presented in order of the strength of the value judgements required to accept them. That is, each successive example builds upon the previous one, requiring stronger value judgements than the last. Stronger value judgements allow us to rank a greater number of distributions – but at the cost of making value judgements that may be less generally accepted by decision makers and stakeholders.

*The role of univariate assessment*

In this chapter, we focus on univariate health distributions with the sole variable of interest being health. Restricting ourselves to univariate health distributions only allows us to make a partial assessment of how far one health distribution is "fairer" or "better" than another. A full assessment also requires information about the determinants of health in order to inform social value judgements about how far the resulting distribution of health is "unfair". However, rules for assessing univariate distributions of health can still be a useful tool for clarifying an important sub-set of the social value judgements at stake. As described in our overview paper, we can divide the analysis into two conceptually distinct stages of (1) modelling of health distributions and (2) ranking of health distributions. The modelling stage is about estimating the impact of alternative decision options on the health distribution, and is by far the most difficult stage. Although modelling is partly a "scientific" endeavour, it also involves important social value judgements about the appropriate formulation of the decision options and appropriate ways of characterising the health distribution given the decision context and data availability. Appropriately characterising the health distribution may involve selecting one or more social gradient variables of interest in assessing fairness – for example, income group, ethnic group, gender. It may also involve standardising the distribution of health to allow for determinants of health which are correlated with the social gradient variable but are not considered relevant to assessing fairness in the relevant decision context – for example, age. The ranking stage then involves social value judgements about how to rank (appropriately characterised) distributions of health using inequality indices or other assessment rules for health distributions. This chapter is about ranking of health distributions which are appropriately characterised as a univariate distribution of health, either at individual or group level, and with or without some form of standardisation. Even when there are multiple social gradient variables of interest, and multiple "confounding" determinants of health considered irrelevant to assessing fairness, it will always be possible to characterise unfair health inequality in terms of a single "standardised" univariate health distribution, and then to apply univariate assessment rules to this "standardised" distribution.

*Differences between health and income*

Health differs in important respects from income, and the application and interpretation of income inequality measures to health inequality must be done in light of the differences between the two domains (A. Atkinson, 2010). Health and income both make vitally important contributions to individual wellbeing, and both are "goods" in the sense that more is generally considered to be better. In addition, both can act as stores of future value ("capital") as well as current value ("consumption"). Income can be invested as well as spent on non-durable goods, and health is a form of human capital as well as a consumption good. Furthermore, both concepts can in principle be measured on fixed ratio scale that enables comparisons between all individuals – for example, individual income can be measured using equivalised household consumption in real 2010 US dollars adjusted for currency purchasing power, and individual health can be measured using quality adjusted life years. However, some of the important differences that should be considered include:

1. Levels of income and wealth are defined as a current flow and a current stock respectively, whereas levels of health are often defined as a sum total of past, current and/or expected future flows (e.g. life expectancy)
2. Income is unbounded while health has an upper bound
3. Full equality of income between individuals is in principle achievable but some inequality in health between individuals is irremediable and incompensable
4. Income has only instrumental value, whereas health has intrinsic value (e.g. the value of being alive, mentally alert, free of pain) as well as instrumental value (e.g. as a pre-requisite for undertaking household and employment tasks)
5. Individual income is usually assumed to have diminishing marginal value, whereas any such assumption about individual health is controversial

In addition, there are two further important differences that arise primarily due to additional data limitations on measurement of health compared with measurement of income:

6. Income can usually be measured on a fixed ratio scale, whereas in practice health may often only be measured on an ordinal scale
7. Income distributions can often be measured at individual level, whereas in practice health distributions may often only be measured at group level

*The interpretation of our stylised examples*

We use a series of stylised examples to illustrate the potential application of the different normative assessment rules to support health sector decision making. To keep things simple, we assume the decision maker is faced with a straightforward binary choice of whether or not to implement a particular health sector intervention (or programme or policy). In each case, therefore, we compare a "baseline" health distribution (if the intervention is not implemented) with an "intervention" health distribution (if the intervention is implemented). We initially assume that the intervention does not have any important non-health benefits or opportunity costs falling outside the health sector budget. The "intervention" distribution takes into account not only the health benefits of the intervention but also the health opportunity costs of foregone expenditure within the health sector budget, which would otherwise have been used to deliver alternative health benefits. The difference between the "baseline" and "intervention" distributions thus represents the "net" health effects of the intervention.

For the purposes of developing our analytical framework we have chosen to assume away practical issues of data limitations, in order to focus clearly on the conceptual issues. We will later on consider the implications of data limitations for how our framework can be applied in practice to real cases involving limited data. In effect, this means that for the time being we assume away issues 6 and 7 above, though the first five issues remain in play.

*Assumptions*

1. We assume that it is possible to measure health on a fixed ratio scale. So in what follows, we can think of health gains in terms of quality adjusted life years (QALYs) and health levels in terms of total lifetime QALYs. This assumption about the possibility of measurement is made in order to facilitate conceptual clarity and comparability with the standard framework of cost effectiveness analysis, rather than to impose any particular social value judgement regarding the measure of health or to restrict attention to any particular subset of normative assessment rules. This assumption about the "basic" unit of health does not imply an exclusive focus on equality in achieved levels of health – for example, we can also examine equality in changes in health, in opportunity for health, and in shortfall between achieved health and target health.
2. We assume that it is possible to measure individual level distributions of health. This is intended to make the analytical framework as general as possible, since it is always possible

to move from individual-level analysis to group-level analysis by aggregating individuals to form groups. Starting from individual level measurements keeps our framework flexible, and ensures it is capable of examining within-group inequality as well as between-group inequality, and of exploring the implications of different ways of aggregating individuals into groups.

3. We assume that all measurements of health can be made with no uncertainty. This assumption helps us focus on the conceptual issues involved in assessing fairness in the distribution of health, without being distracted by the further conceptual issues involved in assessing uncertainty. We assume there is no uncertainty about each individual's lifetime experience of health in terms of total lifetime QALYs. This is stronger than merely focusing on a point estimate of ex ante expected lifetime health at the time the decision is made, and setting aside uncertainty around that point estimate. Rather, we are making the counter-factual assumption that the analyst can know for certain each individuals' actual future experience of health and date of death.
4. We assume we have information on the "true" treatment effect on quality adjusted life span for each individual within the relevant population for the decision problem. We want to compare alternative distributions of health within the same population. We want to focus on differences in health that are attributable to the intervention, and not to differences in the population base or changes in health over time that would have happened anyway. To simplify this we assume that the relevant population for the decision maker are "current" individuals alive at the time the decision is made, and that the relevant time horizon for the decision maker is the lifetime of all those individuals. In practice measurements of lifetime health can differ substantially – and can involve different populations of individuals – depending on the time period of the underpinning data. For example, measurements of "life expectancy at birth" are conventionally based on all-age all-cause mortality data for a particular year or handful of years, and so in fact relate to the population of individuals who died during that period rather the population of individuals currently alive in that time period or born in that time period. Such estimates typically substantially under-estimate "true" life span, since they make no allowance for the fact that the people alive in any given year will predictably live longer and healthier lives than people who die in that year, as a result of experiencing better economic and social conditions over their lifecourse, and receiving better health care. It could therefore be potentially misleading to compare a "baseline" distribution of life expectancy (as conventionally defined) at the time of decision – say, in 2015 – with a "post-intervention" distribution at a time after the intervention has had plenty of time to take effect – say, in 2035. This would involve a comparison between two potentially quite different populations of individuals (i.e. those who will die in 2015 and 2035, respectively) who will predictably have quite different life spans irrespective of any intervention. Our

assumption does not preclude analysis of fairness between older and younger groups of individuals within the current population, though it does for the time being set aside important issues involving the balance of interests between current and future generations.

The upshot of all these assumptions is that our stylised examples adopt an individual level perspective with unbiased information on actual lifetime QALYs for the current population of individuals alive at the time of the decision, both with and without the intervention being implemented, and under conditions of complete certainty. We are well aware that in practice the information available to analysts and decision makers is always far more limited than this.

## 2. Actual UK Health Distribution

We start by examining an actual UK health distribution, using the graphical tools we use later on to present our stylised examples. We examine the distribution of age at death from the ONS interim life tables 2007-09 (ONS, 2010). The first diagram shown is the individual health parade inspired by Jan Pen's "parade of dwarfs" diagram often used to describe income distributions (Pen, 1974). This plots the age at death of the population of people who died in the UK between 2007 and 2009, ordered from lowest to highest age at death. The next diagram is the health quintile groups[1] plot. Here we have calculated the mean age at death per quintile group of the population, with the population ordered by health as in the individual health parade. The health quintile groups plot is a group based version of the individual health parade, with the population divided into five equally sized quintile groups. We have chosen to use quintiles, rather than deciles or centiles as is more common in the income distribution literature, because in the health literature it is common to describe bivariate socio-economic distributions of health using socio-economic quintile groups. However, this univariate health quintile group plot differs from a bivariate plot of mean health by socio-economic quintile group, because the groups are based on health rank not socio-economic rank and therefore in general show greater variation with respect to mean health.

[1] The use of quintile groups in the public health literature typically refers to the average across a quintile group rather that the value for the individual at the position in the distribution indicated by the quintile cut point as used in the income literature. This is largely due to data limitations in the public health sphere; in this document we go with the public health convention when dealing with quintiles.

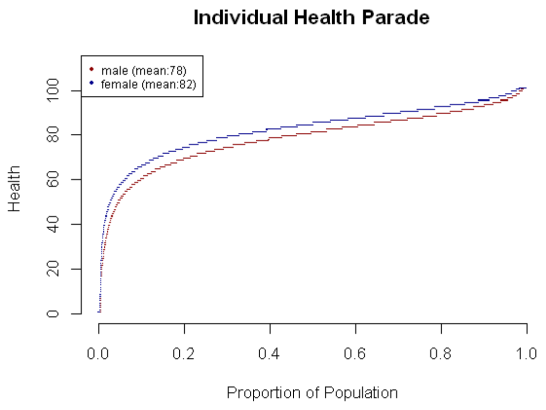


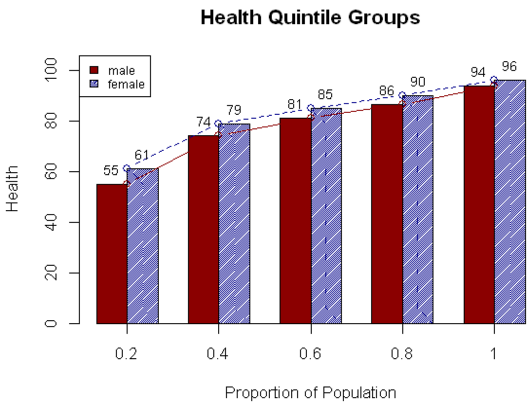


The next diagram is a health histogram. This shows the proportion of the population achieving each of various ranges of life spans. This is another group based version of the individual health parade, except this time with groups based on interval cut points defined in terms of absolute levels of health rather than quantile cut points defined in terms of relative rankings. The interval groups used in this diagram are not of equal size but instead have been selected for their policy significance.

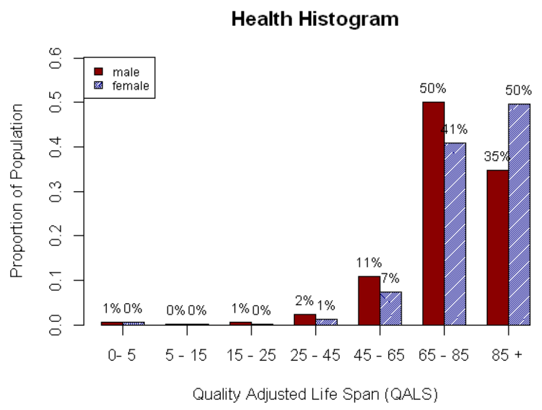


The final two diagrams show the Lorenz curve and the generalised Lorenz curve for health. These diagrams both require that health is measured on a ratio scale, like income, so that it makes sense to add up health across individuals to arrive at the cumulative sum total of population health (the y-axis). The Lorenz curve is used to assess whether one distribution is more equal than another in terms of relative inequality. It is formed by ordering the population by their health from lowest to highest and plotting the cumulative population against the cumulative proportion of total population health that they achieve. The difference between the 45 degree line (representing a perfectly equal distribution of health) and the Lorenz curve represents relative inequality in the distribution of health. As we shall see later, one measure of the degree of relative inequality is twice the area between the 45 degree line and the Lorenz curve – this is the Gini coefficient. The Lorenz curve also has a more general interpretation, however, which does not rely on any specific index of relative inequality. If the Lorenz curve for distribution A lies wholly above Lorenz curve for distribution B, then we can say unequivocally that distribution A is more relatively equal than distribution B. For example, we can see from the diagram that the Lorenz curve for female's lies everywhere above the Lorenz curves for males, suggesting that health is (relatively) more equally distributed among females than males in the UK.

The generalised Lorenz curve can be used to assess whether one distribution is better than another overall, taking into account both concern for relative inequality and concern for efficiency in terms of mean health. It is formed by multiplying the Lorenz curve by the mean level of health. This will prove useful when comparing distributions with unequal mean health levels, where simply looking at relative inequality by comparing Lorenz curves may not be enough to determine which of the distributions is better overall. We examine the assumptions underlying the use of generalised Lorenz curves in the next section. If these assumptions are acceptable, then distribution A can be said to be better than distribution B if the generalised Lorenz curve for distribution A lies wholly above that for distribution B. In this example, this principle of "generalized Lorenz dominance" suggests that the female health distribution is better than the male health distribution.

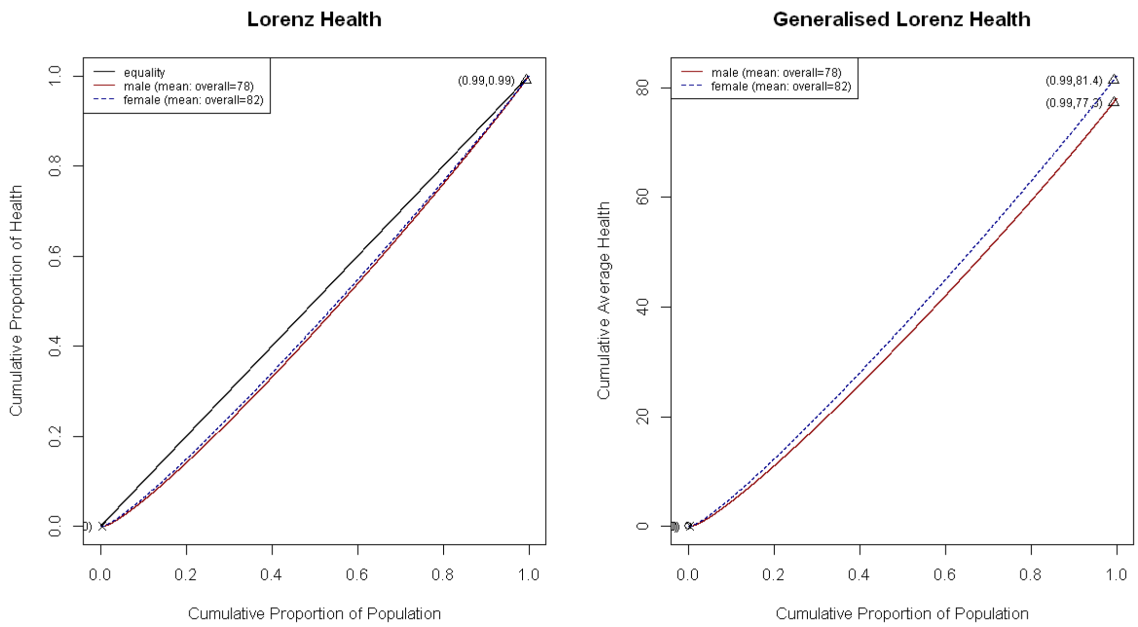


So far we have illustrated our basic graphical techniques using the actual UK distribution of age at death. The stylised examples that follow, however, are completely fictional and have no relationship to the UK health distribution. Rather, they are deliberately contrived to most clearly demonstrate the application of the various inequality indices and normative assessment rules that we present in the following sections.

## 3. Rules for Assessing Health Distributions

Several different approaches to assessing income distributions are discussed in the income inequality literature. We will look at three in this report and apply them to health distributions. First, measures of relative and absolute inequality. Second, measures of poverty, also known as shortfall inequality. Third, measures of social welfare combining concern for inefficiency as well as inequality.

Measures of relative and absolute inequality focus on variation across the entire distribution, whereas measures of poverty or shortfall inequality focus on that part of the distribution that falls below a "poverty line". In the health literature, a natural counterpart to the concept of a "poverty line" might be the concept of a "normal healthy lifespan" or "fair innings"; which may be a fixed target for all individuals or a variable target tailored to individual circumstances such as gender and age. Measures of inequality and poverty are often based on an axiomatic approach, starting from a set of desirable properties and then deriving measures that have these properties. Normative judgements about how concerns for inequality and poverty are to be traded off against concerns for efficiency are not explicitly stated, but instead are made implicitly in the construction of the measures. By contrast, measures of social welfare are based on defining a social welfare function which make explicit normative judgements about such trade offs (A. K. Sen, 1973).

### 3.1 Measures of Relative and Absolute Inequality

The focus with these measures is only to quantify inequality in the health distribution, irrespective of efficiency (Cowell, 2011). We start with a set of inequality "axioms" representing desirable properties for a measure of inequality, and then explore how different inequality measures satisfy these axioms. The first three axioms describe the nature of inequality aversion, while the final two describe desirable properties in terms of transitivity and consistency in decision making.

- **Weak principle of transfers[2]:**
  This axiom expresses inequality aversion (or a preference for equality). It broadly states that the transfer of health from a more healthy to a less healthy person reduces health inequality. More formally it says the following holds: for two individuals having health h and h + x respectively where x is positive. Any positive transfer of health from the more healthy to the less healthy individual will reduce health inequality so long as the amount transferred is less than 2 * x. In the limit repeated transfers that satisfied this criterion would result in a perfectly equal distribution.
  In terms of income inequalities it is possible to remove income earned by one individual in the form of a tax and transfer it directly to another in the form of a benefit. In the context of health care it would not be possible to remove health gained by one individual and transfer it

[2] This is closely relates to the Pigou-Dalton transfer principle often quoted in the income inequalities literature.

to another. Instead it is future expected health gains that would be transferred between individuals.

This first axiom is useful for comparing alternative distributions of a fixed pot of health. The next two axioms discuss how inequality measures react to a change in the size of the pot.

- **Scale independence:**

  This axiom focuses attention on concern for relative inequality between individuals, rather than the size or scale of absolute differences between individuals. The axiom states that any equal proportional change in each individual's level of health should not change the measure of health inequality. While this is relatively uncontroversial when applied to changes in the scale used to measure health – analogous to nominal changes in income due to inflation – it is harder to justify when looking at real differences in health. For example, if everyone's life span doubles then a ten year absolute health gap will grow into a twenty year absolute health gap. It is not self-evident that any reasonable person must deem the new health distribution to be just as "equal" or "fair" as the old one.

- **Translation independence:**

  This axiom focuses on concern for absolute inequality between individuals, rather than relative inequality. It states that any equal absolute change in each individual's level of health should not change the measure of health inequality. A measure cannot fully satisfy both scale independence and translation independence, and it is not self-evident which of the two axioms is more appropriate in the health context. For example, if everyone gains 50 years in life span, a relative gap of twenty percent between 60 and 50 declines into a relative gap of only ten percent between 110 and 100. So although the absolute inequality gap remains the same the relative inequality gap has declined, and there is room for reasonable disagreement about which concept of inequality is the more appropriate.

- **Principle of population:**

  This axiom prescribes that health inequality should be invariant to the size of the population, other things equal. This is desirable to ensure transitivity if comparing populations of different sizes, for example if comparing inequality across different countries or regions or over time for a growing or shrinking population. This axiom states that if a population of size $n$ is combined with another population of size $n$ with the same level of health inequality, then the combined population should have the same level of health inequality.

- **Subgroup consistency and decomposability:**

  This axiom is desirable to ensure consistency between inequality assessments about a population and subsets of the population – for example, between assessments at national and sub-national levels. For subgroup consistency we require that when comparing two different health distributions for a population, inequality rankings within the population as a whole for the two distributions match the inequality rankings within any of the subgroups of which the population is composed. Furthermore, if we can express total inequality in terms of inequality between subgroups and inequality within subgroups then the measure of inequality also satisfies decomposability, allowing us to attribute total inequality to inequality between and within different parts of the population.

Now we turn to the measures of inequality commonly used to characterise health distributions (A Wagstaff, Paci, & van Doorslaer, 1991) and analyse these using the axioms above:

- **Absolute gap or range:**
  An absolute gap or range refers to the absolute difference in health between two individuals or groups, and is a simple measure of absolute inequality. Absolute gaps are often based on extreme groups, and in our stylised examples we use the most healthy and least healthy quintile groups. It is also possible to compute absolute gaps between any two individuals or groups, and a common variant is to compare the least healthy quintile group with all other groups or with the middle quintile group. However, differences in any other parts of the health distribution are not picked up by this measure. This measure does not satisfy the weak principle of transfers, scale independence or decomposability.
- **Relative gap or ratio:**
  A ratio refers to the health of individual or group A divided by the health of individual or group B. A relative gap is a ratio minus one – i.e. the same thing, except that the numerator is the absolute gap between the health of A and B. As with absolute gaps, relative gaps can be computed for any two groups or individuals, but in our stylised examples we focus on the ratio between top and bottom health quintile groups. A relative gap is a relative measure of inequality and hence is scale independent (and therefore not translation independent). It does not however satisfy the weak principle of transfers or decomposability.
- **Slope index of inequality (SII):**
  The slope index of inequality is a simple index of absolute inequality which makes use of information on the whole distribution. The starting point for the slope index of inequality is a plot of individual health ordered from least healthy to most healthy against cumulative population – i.e. an individual health parade. The slope index is then calculated as the least squares line of best fit through these points. The slope index can also be computed for grouped date, using population weighted least squares if the groups are of different sizes. The advantage of the slope index against the previous two measures is that it uses the data from across the entire health distribution. The slope index also satisfies the weak principle of transfers and the principle of population. The slope index however is an absolute index so does not satisfy scale independence; nor does it satisfy decomposability. If health is a linear function of rank order then the slope index should be consistent with the absolute gap

between the most and least healthy individual. However, if an OLS model is inappropriate then these measures can disagree.

- **Relative index of inequality (RII):**
  The relative index of inequality is a relative inequality version of the slope index. It is simply the slope index divided by population mean health. Apart from satisfying scale independence rather than transfer independence, in other respects it shares the same axiomatic properties, advantages and limitations as the slope index.
- **Gini:**
  The Gini coefficient is an index of relative inequality and is calculated as two times the area between the Lorenz curve and the line of perfect equality. The Gini coefficient satisfies the weak principle of transfers, scale independence (but not translation independence) and the principle of population. It uses information from the entire distribution and ranges between 0 and 1. The Gini coefficient however is not decomposable. It has the property that the decrease in health inequality due to a transfer of health between individuals is proportional to the difference in health rank between them rather than the absolute difference in health levels between them.

### 3.2 Measures of Poverty

Poverty measures concentrate on individuals whose income lies below a certain threshold level or "poverty line" deemed necessary for meeting basic needs. Poverty measures typically focus on one or more of the "the three I's of poverty"(Jenkins & Lambert, 1997) :

(1) incidence (i.e. how many people are "poor" in the sense of having an income below the poverty line),
(2) intensity (i.e. how far on average do "poor" people lie below the poverty line), and
(3) inequality (i.e. how much variation there is within the group of "poor" individuals).

In relation to health, one can define analogous concepts of health poverty in relation to a threshold level of health such as a "normal lifespan" or "fair innings". The potential attractiveness of such measures is that they focus on unhealthy individuals who decision makers may be most concerned about. They can also help to avoid the classic "levelling down" objection to inequality measures – i.e. that the best way to reduce inequality is to reduce the health of the better off rather than to improve the health of the worse off. In addition to analogous versions of the inequality axioms discussed above, two further axioms are commonly discussed in the income poverty literature. These are (A. Sen, 1976):

- **Monotonicity:**
  Any reduction in the health of any person below the health poverty line should increase the measure of health poverty.
- **Focus:**
  The health poverty measure should be invariant to changes in health of those above the health poverty line.

Given these and the previously described axioms we will describe some commonly used poverty measures as applied to health:

- **Headcount proportion (H):**
  The headcount proportion is a measure of the incidence of health poverty. It is simply the proportion of the population below the health poverty line. This is the most commonly used and easiest to calculate measure. However, although it focuses on the incidence of health poverty it fails to consider either the intensity of health poverty or inequality in the distribution of health among those with poor health, and therefore violates most of the axioms discussed.
- **Health shortfall proportion (I):**
  The health shortfall proportion is a measure of the intensity of health poverty. It focuses on the population of individuals in health poverty (i.e. below the health poverty line) and computes the mean gap or shortfall between their level of health and the health poverty line. It then presents this as a proportion of the health poverty line. Unlike the headcount ratio, this measure considers the level of health among those with poor health. However, it fails to consider either the incidence of health poverty or inequality in the distribution of health among those with poor health, hence violating many of the axioms discussed.
- **Sen Index (S):**
  The Sen index combines information on all three I's of health poverty – incidence, intensity and inequality. This index can be written as:
  $S = HI + H(1 - I)G_p$
  where H is the headcount proportion, I is the health shortfall proportion, and $G_p$ is the Gini coefficient of inequality among those in health poverty. This index satisfies the weak transfer condition, scale independence, the principle of population, monotonicity and the focus axiom. It does not however, due to its reliance on the Gini coefficient, satisfy the decomposability axiom.
- **The Sen family of poverty indices (Q)**:
  The Sen family of poverty indices are indices formed by replacing the Gini coefficient by alternative inequality measures to calculate the equally distributed equivalent health among

the health poor $h_{ede(p)}$ (see the following section on the social welfare function approach for more details on this concept) in the more general formula:

$$Q = H(1 - \frac{h_{ede(p)}}{z})$$

where z represents the health poverty line. The standard Sen poverty index can then be derived by using the Gini coefficient to derive the equally distributed equivalent health among the health poor $h_{ede(p)}$:

$$h^G_{ede(p)} = \mu_p(1 - G_p)$$

where $\mu_p$ is the mean level of health among the health poor. The selection of the index used in deriving $h_{ede(p)}$ will determine the set of axioms the index satisfies.

### 3.3 Social Welfare Function Approach

The final approach that we will look at is the social welfare function (SWF) approach. A health-related SWF can be used to describe social welfare as a function of the health distribution, other things equal – i.e. setting aside non-health aspects of the social state. A health related SWF may yield either a partial or a complete ordering of health distributions, depending on the strength of the social value judgements it embodies.

Several properties are considered desirable when constructing a SWF. In defining these we will use a notation that mirrors the notation used in the standard income-related SWF literature, but re-interpreting its elements in terms of health rather than income. So we use the terminology $h_{iA}$ to represent the health of individual *i* in health distribution A (rather than the income of individual i in income distribution A), $U_{iA}$ to represent a individual utility function for individual *i* in distribution A, and $W_A$ to represent social welfare in distribution A (A. K. Sen, 1973)(Cowell, 2011). As before, we can think of individual health as being measured in terms of quality adjusted life span. The individual utility function can then be interpreted as representing the value to society of the individual's quality adjusted life span. The suffixes, A and B, are used to distinguish different health distributions.

- **Individualistic:**
  This means the SWF is a function of the individual utilities i.e. the SWF has the form:
  $W_A = W(U_1, U_2 ... , U_n)$
- **Non-decreasing:**
  This expresses a preference for more health. This means given two states A and B, if $h_{iA} >= h_{iB}$ for all i then $W_A >= W_B$ i.e. if every individual has at least as good health in state A as in state B then overall state A is at least as good as state B. A Pareto improvement in health, where every individuals health either increases or stays the same, is always considered as at least as good regardless of any increase in inequality.

- **Symmetric or anonymous:**

  This means that the SWF treats individual utilities anonymously, the value of W does not depend on the particular identity of the individual as a carrier of utility or the order in which individuals appear in the welfare function i.e.

  $W = W(U_1, U_2, ... , U_n) = W(U_2, U_1, ... , U_n) = ... = W(U_n, U_2, ... , U_1)$.

- **Additive:**

  If the social welfare function can be written as a sum of the individual utility functions $U_i$ i.e.:

  $W(h_1, h_2, ... , h_n) = U_1(h_1) + U_2(h_2) + ... + U_n(h_n)$

  This means that any change in individual i's health only affects the utility of individual i, and furthermore the effect of this change on the utility of individual i is independent of the health of the other individuals. For an individualistic, non-decreasing, symmetric, additive health-related SWF we can re-write the above as:

  $W(h_1, h_2, ... , h_n) = U(h_1) + U(h_2) + ... + U(h_n)$

  Where each individual has a common U(.) that increases with health.

- **Concavity:**

  The SWF is strictly concave if the welfare weight always decreases as $h_i$ increases where the welfare weight is defined as:

  $U'(h_i) = dU(h_i)/dh_i$

  This means that when evaluating changes to social welfare we apply lower weight to increases in health to those with higher health than to those with lower health. Note that this has a parallel in the weak principle of transfers.

- **Constant Relative Inequality Aversion:**

  This means that a constant proportionate change in health results in a constant proportionate change in welfare weight i.e. function U(.) takes the form:

  $U(h_i) = (h_i^{1-e} – 1) / (1 - e)$

  Where e is the inequality aversion parameter with a higher level of e implying greater inequality aversion.

We can use these properties to derive rules to help us determine which of two health distributions are socially preferable. The following four rules are listed in order from least restrictive to most restrictive and allow us to partially order health distributions.

- **Rule 1 - Pareto Dominance:** for any individualistic, increasing and additive SWF W, if $h_{iA} >= h_{iB}$ for all i and $h_{iA} > h_{iB}$ for at least one i then $W_A > W_B$ i.e. state A is preferred to state B. Where i represents the same individual in each distribution.
- **Rule 2 – Re-ranked Pareto Dominance:** for any individualistic, increasing, additive and *symmetric* SWF W, if $h_{iA} >= h_{iB}$ for all i and $h_{iA} > h_{iB}$ for at least one i then $W_A > W_B$ i.e. state A is preferred to state B. Where i represents the individual with equivalent health ranking in

each distribution - as the SWF is symmetric, this does not necessarily have to be the same individual under both states. (Cowell, 2011)

- **Rule 3 – Atkinson's Theorem:** for any individualistic, increasing, additive, symmetric, and *strictly concave* SWF W and states A and B with equal mean health, $W_A > W_B$ if and only if the Lorenz curve for A lies wholly inside the Lorenz curve for B. (A. B. Atkinson, 1970)
    - This rule only applies when the mean health in the more relatively equal distribution is more than or equal to the mean health in the less equal distribution.
    - It describes a 'win-win' situation in which the distribution is less relatively unequal and has either the same or greater overall health.
- **Rule 4 – Shorrocks' Theorem:** for any individualistic, increasing, additive, symmetric, and strictly concave SWF W, $W_A > W_B$ if and only if the generalised Lorenz curve for A lies wholly inside the generalised Lorenz curve for B. (Shorrocks, 1983)
    - It is important to note the using this rule we would never prefer a more equal distribution with lower mean health to a less equal distribution with higher mean health.
- **Generalisation of Rules 3 & 4:**
  It has been shown (A. K. Sen, 1973) that rules 3 and 4 can be generalised to not require the SWF to be additive, individualistic and strictly concave. Instead all that is needed for Atkinson's and Shorrocks' theorems to apply is that the SWF is an increasing, symmetric, Schur-concave[3] function of individual levels of health: $W = W(h_1, h_2, \ldots, h_n)$

These four dominance rules will prove to be useful in helping us to rank health distributions and require us to make few restrictions on the nature of our social welfare function. That is, we would not need to specify exactly the nature of the SWF but could describe broad characteristics that encompass whole classes of SWFs, under any of which the welfare rankings of particular interventions would be the same. It is important to recognise that none of these four rules allow us to trade off mean health against equity, we would never prefer a distribution with lower mean health to one with higher mean health using any of these rules.

The dominance rules only provide a partial ranking and we sometimes find ourselves in positions where the rules do not apply, for example where generalised Lorenz curves intersect. In such situations we need to more fully describe the SWF, perhaps by adding a final assumption that the SWF has constant relative inequality aversion (or constant absolute inequality aversion) and more fully define the nature of the SWF by providing a societal inequality aversion level to get an index of social welfare giving us a complete ranking of health states. The key idea used in these SWF based indices is that if health were distributed equally then given the assumptions above less total health

[3] For a definition of Schur-concavity see: (Marshall & Olkin, 1974) for its application to inequality measurement see (Dasgupta & Sen, 1973) and (Shorrocks, 1983)

would be required to produce the same social welfare than if health were distributed unequally. It is important to note that in situations where any of the dominance rules do apply, both relative and absolute families of social welfare functions described below will rank health distributions (in terms of their equally distributed equivalent health) identically to the ranking suggested by the dominance rules.

- **Atkinson Inequality Index:**

  Perhaps the most widely used of these indices is the Atkinson Inequality Index, which is scaled from 0 to 1 and can be calibrated using a single parameter, *e*, representing a social value judgement about the degree of constant relative inequality aversion. This index is based on the family of social welfare functions that are individualistic, symmetric, additive, concave and exhibit constant relative inequality aversion. This family of social welfare functions can be represented using a simple and convenient Atkinson Welfare Index: mean health multiplied by 1 minus the Atkinson Inequality Index. This Atkinson Welfare Index can be interpreted as the "equally distributed equivalent" health: the common level of health in a hypothetical equal distribution of health that has the same level of social welfare as the actual unequal distribution. We shall therefore refer to the Atkinson Welfare Index as the Atkinson EDE Health Index.

  The Atkinson Inequality Index looks at the difference between the mean health in the unequal distribution ($\bar{h}$) and the common level of health in a hypothetical equal distribution ($h_{ede}$) that would provide an equivalent social welfare:

  $$A_e = 1 - \frac{h_{ede}}{\bar{h}}$$

  Given our assumptions that the SWF is additive and exhibits constant relative inequality aversion, this yields:

  $$h_{ede} = \left[\frac{1}{n}\sum_{i=1}^{n}[h_i]^{1-e}\right]^{\frac{1}{1-e}}$$

  $$A_e = 1 - \left[\frac{1}{n}\sum_{i=1}^{n}\left[\frac{h_i}{\bar{h}}\right]^{1-e}\right]^{\frac{1}{1-e}}$$

  where the parameter *e*, which can take any value from 0 to infinity, specifies the level of societal inequality aversion. Higher *e* values make the index more sensitive to changes at the lower end of the health distribution. The higher the *e*, the further the index tilts towards concern for health improvement among less healthy individuals rather than more healthy individuals. A value of 0 represents a classic "utilitarian" view that all that matters is sum total health and not inequality in the distribution of health. Whereas as the value approaches infinity the Atkinson index comes to represent the "maximin" view that all that matters is improving the health of the least healthy individual, irrespective of the health of all other individuals. The resulting value of the Atkinson inequality index indicates the proportion of the mean health that we would be willing to sacrifice to achieve equal health for all, given any particular value of *e*. The proportion of mean health we are willing to sacrifice to achieve equality will increase as our inequality aversion rises.

Some important properties of the Atkinson index are (Kolm, 1976a):

- The Atkinson index is a relative measure of inequality and is scale invariant. The Atkinson index for any particular distribution of health would not change if every individual's level of health was increased by the same proportion.
- Given two distributions A and B where the Lorenz curve for A lies wholly inside the Lorenz curve for B, the Atkinson inequality index will always be lower for distribution A than for distribution B regardless of the relative means of the distributions.
- The Atkinson index is subgroup consistent and decomposable. By this we mean that the ordering of inequalities in the total population is consistent with that in any subgroup of the population, and the level of inequality in the population can be written as a function of a between subgroup and within subgroup level of inequality.
- The Atkinson index is bounded between 0 and 1 with zero representing a perfectly equal distribution.

- **Kolm Index:**
An absolute alternative to the relative Atkinson index also based on in the SWF framework is the "leftist" index proposed by Serge Kolm (Kolm, 1976a) (Kolm, 1976b). Like the Atkinson index the Kolm index compares the mean health in the unequal distribution ($\bar{h}$) to the equally distributed equivalent health in a hypothetical distribution ($h_{ede}$) that would yield equivalent social welfare. The key differences between the two indices is that where the Atkinson index assumes constant *relative* inequality aversion the Kolm index assumes constant *absolute* inequality aversion. The Kolm index can be written as:
$$K_\alpha = \bar{h} - h_{ede}$$
Which given our assumptions on the form of the SWF yields:
$$h_{ede} = -\left(\frac{1}{\alpha}\right) log\left(\frac{1}{n}\sum_{i=1}^{n} e^{-\alpha h_i}\right)$$
$$K_\alpha = \left(\frac{1}{\alpha}\right) log\left(\frac{1}{n}\sum_{i=1}^{n} e^{\alpha[\bar{h}-h_i]}\right)$$

where the parameter α specifies the level of societal inequality aversion, with higher α values making the index more sensitive to changes at the lower end of the health distribution. The value of this index represents the absolute amount by which we would be willing to reduce average health to achieve equal health for all. The amount of mean health that we would be willing to sacrifice to achieve an equal distribution rises with our level of inequality aversion.

Some important properties of the Kolm index are (Kolm, 1976a):

- The Kolm index is an absolute measure of inequality and is translation invariant. The Kolm index for any particular distribution of health would not change if every individual's level of health was increased by the same absolute amount.
- Given two distributions A and B where the Lorenz curve for A lies wholly inside the Lorenz curve for B and the average health in A is not larger than the average health in B, then the Kolm index will always be lower (i.e. more equal) for distribution A than for distribution B[4].
- The Kolm index is both subgroup consistent and decomposable.
- The Kolm index is bounded between 0 and $(\bar{h} - h_{min})$ with 0 representing perfectly equal distribution.

### 3.4 Measures from the Health Inequalities Literature

Most of the health inequality measurement literature focuses on "bivariate" indices of health inequality which focus on the association between health and a social ranking variable such as income. Bivariate health inequality measures are reviewed in (O'Donnell et al., 2008) with recent contributions from Erreygers (Erreygers & Van Ourti, 2011)(Erreygers, Clarke, & Van Ourti, 2011)(Erreygers, 2009). We focus here on "univariate" indices which focus primarily on the distribution of health, though also sometimes allowing for other variables.

This univariate health inequalities literature builds upon the social welfare function approach inspired by Atkinson's work on income inequalities. The most commonly used social welfare function in this literature is the iso-elastic form proposed by (Adam Wagstaff, 1991) which in the simple case of two groups can be written as:

$$W = [\alpha {h_a}^{-r} + (1-\alpha){h_b}^{-r}]^{-\frac{1}{r}}$$

Where $h_a$ and $h_b$ are the respectively the health of group a and group b, α represents the relative weight attached to the health of the two groups (this parameter pivots the social welfare function around the 45 degree line of equal health) and the r parameter represents the degree of absolute inequality aversion (this parameter determines the degree of curvature of the social welfare function).

[4] It is interesting to compare this to Rule 4 Shorrocks theorem on generalised Lorenz dominance.

This social welfare function is empirically calibrated and used to calculate “fair innings” equity weights in (Williams, 1997). The α parameter is assumed to be 0.5 (i.e. groups have equal priority) across a range of inequality aversion values (values of r), these two parameters determine the shape of the social welfare function and an absolute equally distributed equivalent health of 70 QALYs (a fair innings) together with current health endowments determine the position of the social welfare function. The slope of the tangents to this calibrated social welfare function at the points describing the current health endowment give the equity weights for a particular value of r. These weights can then be utilised in a cost-effectiveness analysis to maximise weighted expected lifetime QALYs in a socially optimal manner.

This social welfare function has also been empirically calibrated by (Dolan & Tsuchiya, 2009) to estimate both individual responsibility as represented by α and inequality aversion as represented by r. Here the individual responsibility attempts to deal with equity relevant characteristics to correct for fair differences in health and then any remaining unfair inequality is dealt with by the inequality aversion parameter.

Alternative families of social welfare functions and approaches to use these to derive equity weights based on the rank dependent QALY model are proposed by (Bleichrodt, Diecidue, & Quiggin, 2004), this family of functions has the general form:

$$W = \sum_{i=1}^{n} \pi_i Q_i \ ,$$

$$\pi_i = w\left(\frac{i}{n}\right) - w\left(\frac{i-1}{n}\right), \mathrm{w}(0) = 0 \ , \mathrm{w}(1) = 1 \ , \mathrm{w}(\mathrm{p}) \geq \mathrm{w}(\mathrm{q}) \text{ if } \mathrm{p} \geq \mathrm{q}$$

Where $Q_1 \geq \ldots \geq Q_n$ is the rank ordered QALY profile of the population and w(.) is the weighting function determining the level and nature of inequality aversion implied by the social welfare function. It is important to note that social welfare functions of this form are only concerned with the ranking of groups not with the magnitude of differences between them.

The other distinctive idea arising from the health inequalities literature is that of proportional shortfall in health see (Stolk, Pickee, Ament, & Busschbach, 2005) and (Johannesson, 2001). The idea here is that the appropriate lifetime health target or norm for a particular individual or group should depend on their gender, age and other current circumstances that influence how far their lifetime health is amenable to current and future social and medical choices. For example, it may be appropriate to set the target life expectancy for a woman of age 80 higher than the target life expectancy at birth for a man. Hence inequality measures should be applied to changes in health as a proportion of each individual’s shortfall from their current circumstance specific target health rather than to changes in their absolute levels of health. This concept is similar to the intensity index described in the poverty literature with a group specific rather than universal poverty line. This idea has been operationalised in the context of the extended Gini index to create the extended proportional Gini (Norheim, 2010):

$$EPG = \frac{\sum_{i=1}^{n}\left(R_i^v-(R_i-1)^v\right)h_i/h_{imax}}{n^v\mu(h/h_{max})}$$

Here R represents the health rank, v represents the degree of inequality aversion, μ represents the mean health, h represents total health and $h_{max}$ represents the maximum achievable health from which to measure the shortfall.

### 3.5 Inequalities in UK Health Distribution

We now look back to the UK health distribution and see how it fares against the measures that we have defined. We have grouped these measures into five tables. The first table shows the results of applying the dominance rules. From this table we see that generalised Lorenz dominance applies, allowing us to conclude that the female health distribution is preferred to the male health distribution (having considered both inequality and efficiency).

The second table looks at relative inequality indices, here we can see that the indices we have selected to measure relative inequality all judge the female health distribution to be less unequal than the male health distribution.

The third table looks at absolute inequality indices, here again we see that the indices we have selected to measure absolute inequality generally judge the female health distribution to be less unequal than the male health distribution. Only the Kolm index with the highest inequality aversion parameter disagrees.

The fourth table looks at poverty indices, here our focus is no longer on the entire population but instead only on those below a minimum level of health deemed to constitute a reasonable life expectancy. We have chosen to set z=70, representing a health poverty line of 70 years, following the ancient biblical prescription on what constitutes a reasonable life expectancy ("three score years and ten"). These indices agree with the relative and absolute inequality measures in judging the portion of the female health distribution below the health poverty line to be less unequal than the portion of the male health distribution below the health poverty line.

The fifth and final table looks at social welfare indices. Whereas the previous three tables only considered the degree of inequality, in this table we explore different ways in which the two dimensions of the decision problem, level of health and distribution of health, can be combined. These indices become particularly relevant when the two dimensions conflict, in such situations our dominance rules from our first table will not be able to determine the preferred distribution and we will have to make a trade off between more health and a more equal distribution of health. In this table we present results for conceptions of social welfare that incorporate both absolute and relative ideas of inequality at a range of different levels of inequality aversion. For this simple case where our dominance rules apply, we get the expected agreement with the dominance rules with all measures suggesting that the female health distribution is preferred to the male health distribution. It should be noted that in more complex cases where the dominance rules do not apply the choice of suitable inequality aversion parameter becomes critical in deciding between distributions. Determining suitable values for the inequality aversion parameter, particularly in the field of health inequality, has not been dealt with in any detail in the current body of academic literature.

**Dominance Rules:**

| **Rule** | **male** | **female** |
|---|---|---|
| **Pareto Dominant** | FALSE | FALSE |
| **Re-ranked Pareto Dominant** | FALSE | FALSE |
| **Lorenz Dominant** | FALSE | FALSE |
| **Generalised Lorenz Dominant** | FALSE | TRUE |

**Relative Inequality Indices:**

| **Index** | **male** | **female** | **difference ( female - male )** |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 0.70 | 0.57 | -0.13 |
| **Relative Index of Inequality (RII)** | 0.59 | 0.51 | -0.08 |
| **Gini Index** | 0.10 | 0.08 | -0.01 |
| **Atkinson Index (e= 0.5 )** | 0.01 | 0.01 | -0.00 |
| **Atkinson Index (e= 1 )** | 0.04 | 0.03 | -0.01 |
| **Atkinson Index (e= 1.5 )** | 0.14 | 0.12 | -0.02 |
| **Atkinson Index (e= 2 )** | 0.47 | 0.43 | -0.04 |

**Absolute Inequality Indices:**

| Index | male | female | difference ( female - male ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 38.57 | 34.81 | -3.77 |
| **Slope index of inequality (SII)** | 46.00 | 41.52 | -4.48 |
| **Kolm Index (alpha= 0.025 )** | 3.64 | 3.13 | -0.50 |
| **Kolm Index (alpha= 0.05 )** | 10.00 | 9.19 | -0.81 |
| **Kolm Index (alpha= 0.075 )** | 19.64 | 19.56 | -0.08 |
| **Kolm Index (alpha= 0.1 )** | 29.77 | 30.93 | 1.16 |

**Poverty Indices:**

| Index (z= 70 ) | male | female | difference ( female - male ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.22 | 0.14 | -0.07 |
| **Health Gap Ratio (intensity)** | 0.20 | 0.19 | -0.01 |
| **Sen Poverty Index** | 0.07 | 0.04 | -0.02 |

**Social Welfare Indices:**

| Index | male | female | difference ( female - male ) |
|---|---|---|---|
| **Mean Health** | 78.02 | 82.11 | 4.09 |
| **Atkinson EDE (e= 0.5 )** | 76.91 | 81.22 | 4.31 |
| **Atkinson EDE (e= 1 )** | 74.67 | 79.35 | 4.68 |
| **Atkinson EDE (e= 1.5 )** | 67.06 | 72.57 | 5.51 |
| **Atkinson EDE (e= 2 )** | 41.29 | 46.88 | 5.59 |
| **Kolm EDE (alpha= 0.025 )** | 74.39 | 78.98 | 4.59 |
| **Kolm EDE (alpha= 0.05 )** | 68.03 | 72.93 | 4.90 |
| **Kolm EDE (alpha= 0.075 )** | 58.38 | 62.55 | 4.17 |
| **Kolm EDE (alpha= 0.1 )** | 48.25 | 51.18 | 2.93 |

## 4 Univariate Stylised Examples

We now demonstrate the use of these measures and decision rules through our set of stylised examples in the univariate case.

### 4.1 Pareto Dominant Intervention

The first case we will examine is that of a Pareto dominant intervention under which every individual achieves better health.

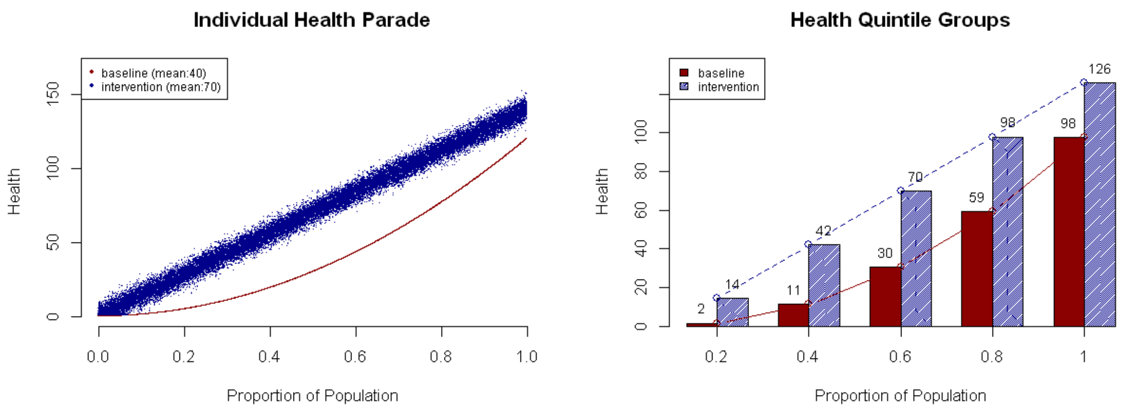


The above diagrams compare the baseline health (red) with that achieved under the intervention (blue). The change in health attributed to the intervention represents its total net effect on health whereby the opportunity cost of the resources used to provide the intervention are characterised as health foregone and subtracted from any health gains. We can see from these diagrams that every individual achieves better health under the intervention and hence mean health by quintile also improves under the intervention.

The diagrams below show how much each individual's health changes from the baseline level under the intervention, the percentage of the population who benefit under the intervention is given in parenthesis under the title. The next diagram shows the proportion of the population achieving the different levels of health. All the diagrams indicate an increase in health across the population. It appears from the individual health benefits diagram that in absolute terms it is the people in the middle of the health distribution that benefit most from the intervention.

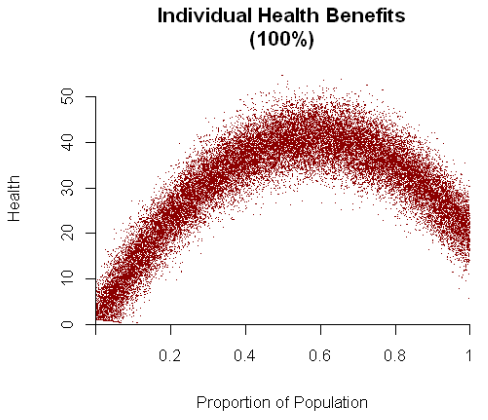


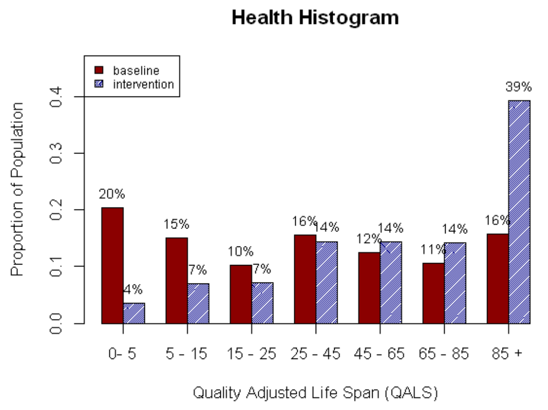


Next we turn to the Lorenz curve to examine the change in relative health inequality. The baseline line is constructed by baseline health ranking, while the intervention line on the diagram is constructed by re-ranking the post intervention population by the health level achieved. Therefore looking at a vertical cross section of the diagram does not compare the same individual under the two policies, instead it compares the two individuals who occupy a specific rank in the health distribution under the two policies. In the diagram below we can see that the intervention in this case results in a more equal relative distribution of health than the baseline.

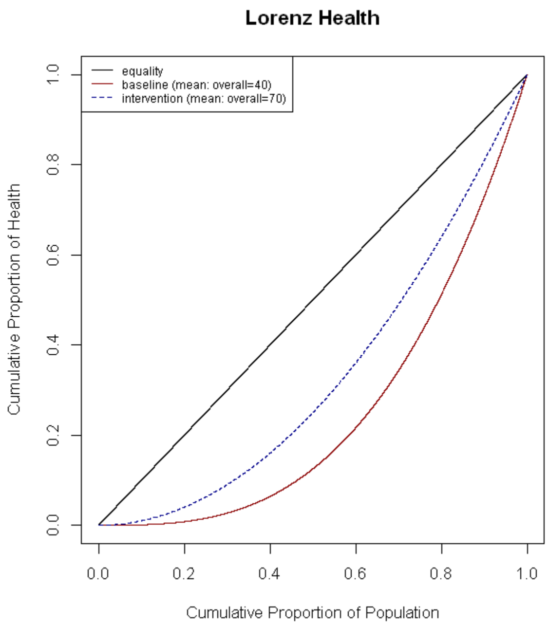


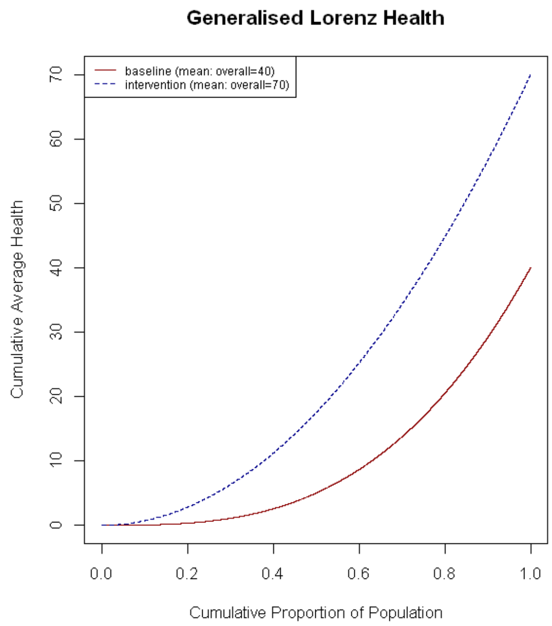


Looking to our battery of measures in the tables below we see that every dominance rule applies indicating that the intervention is preferred the baseline. All relative measures of inequality and poverty measures of inequality judge the intervention to be less unequal than the baseline, whilst all absolute measures of inequality suggest the intervention is more unequal than the baseline. The social welfare measures in both relative and absolute cases agree with the dominance rules in preferring the intervention to the baseline.

**Dominance Rules:**

| Rule | baseline | intervention |
|---|---|---|
| **Pareto Dominant** | FALSE | TRUE |
| **Re-ranked Pareto Dominant** | FALSE | TRUE |
| **Lorenz Dominant** | FALSE | TRUE |
| **Generalised Lorenz Dominant** | FALSE | TRUE |

**Relative Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 59.97 | 8.00 | -51.97 |
| **Relative Index of Inequality (RII)** | 3.00 | 2.00 | -1.00 |
| **Gini Index** | 0.50 | 0.33 | -0.17 |
| **Atkinson Index (e= 0.5 )** | 0.25 | 0.11 | -0.14 |
| **Atkinson Index (e= 1 )** | 0.59 | 0.26 | -0.33 |
| **Atkinson Index (e= 1.5 )** | 0.96 | 0.47 | -0.49 |
| **Atkinson Index (e= 2 )** | 1.00 | 0.69 | -0.31 |

**Absolute Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 96.00 | 112.12 | 16.13 |
| **Slope index of inequality (SII)** | 119.98 | 140.17 | 20.19 |
| **Kolm Index (alpha= 0.025 )** | 12.62 | 18.68 | 6.06 |
| **Kolm Index (alpha= 0.05 )** | 19.65 | 31.08 | 11.43 |
| **Kolm Index (alpha= 0.075 )** | 23.74 | 38.67 | 14.94 |
| **Kolm Index (alpha= 0.1 )** | 26.37 | 43.65 | 17.28 |

**Poverty Indices:**

| Index (z= 70 ) | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.76 | 0.50 | -0.26 |
| **Health Gap Ratio (intensity)** | 0.67 | 0.50 | -0.17 |
| **Sen Poverty Index** | 0.64 | 0.33 | -0.30 |

**Social Welfare Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Mean Health** | 40.00 | 70.00 | 30.00 |
| **Atkinson EDE (e= 0.5 )** | 30.00 | 62.21 | 32.21 |
| **Atkinson EDE (e= 1 )** | 16.26 | 51.60 | 35.34 |
| **Atkinson EDE (e= 1.5 )** | 1.63 | 37.16 | 35.54 |
| **Atkinson EDE (e= 2 )** | 0.03 | 21.74 | 21.71 |
| **Kolm EDE (alpha= 0.025 )** | 27.38 | 51.32 | 23.94 |
| **Kolm EDE (alpha= 0.05 )** | 20.35 | 38.92 | 18.57 |
| **Kolm EDE (alpha= 0.075 )** | 16.26 | 31.33 | 15.06 |
| **Kolm EDE (alpha= 0.1 )** | 13.63 | 26.35 | 12.72 |

### 4.2 Re-Ranked Pareto Dominant Intervention

The next case we will examine is that of a re-ranked Pareto dominant intervention under which while every individual need not achieve better health under the intervention, some permutation of these individuals will dominate the baseline health distribution.

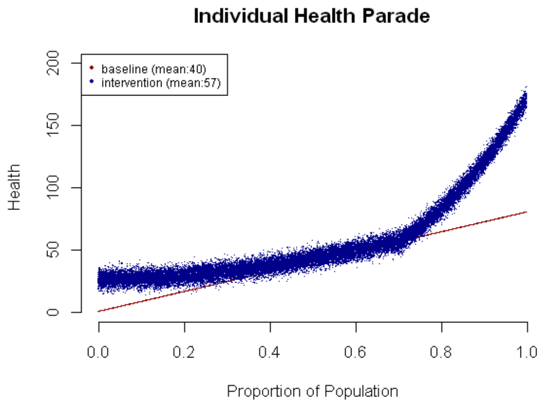


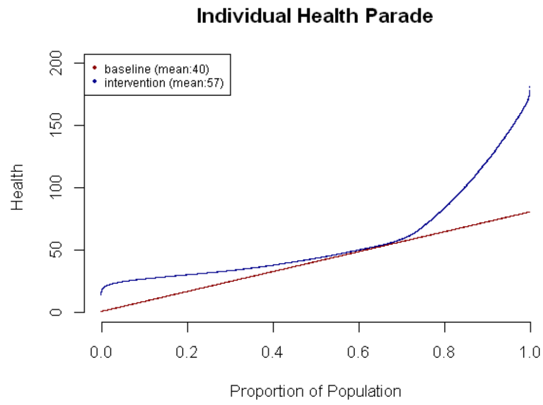


The above diagrams compare the baseline health (red) with that achieved under the intervention (blue). The diagram on the left shows that not every individual achieves better health under the intervention. Re-ranking the population post intervention by health and then re-plotting gives the diagram on the right. Here we see that the use of the symmetry property allows us to use rule 2 (re-ranked Pareto dominance) to show the re-ranked health distribution now Pareto dominates the original health distribution. The diagram below shows the mean health per quintile group after re-ranking.

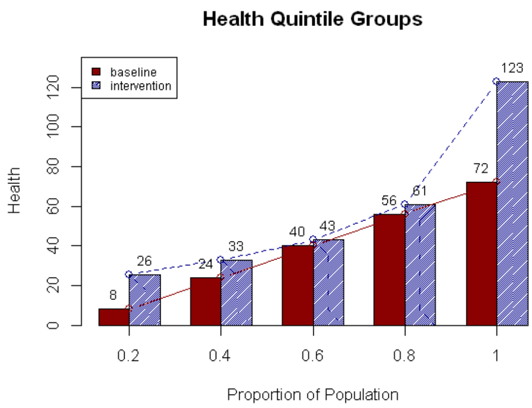


The diagrams below show how much each individual's health changes from the baseline level under the intervention. The next diagram shows the proportion of the population achieving the different levels of health. It appears from the individual health benefits diagram that in absolute terms it is the people in the middle of the health distribution that benefit least from the intervention. It is important to note that even though the re-ranked health distribution is Pareto dominant 11% of people are actually worse off under the intervention. If the magnitude of and position in the distribution of "winners" and "losers" from the intervention may have some impact in terms of policy implementation then it may be important to elucidate how the re-ranking has occurred.  None of the measures discussed in this document rank health gains or changes in inequality as a function of the starting position of individuals.

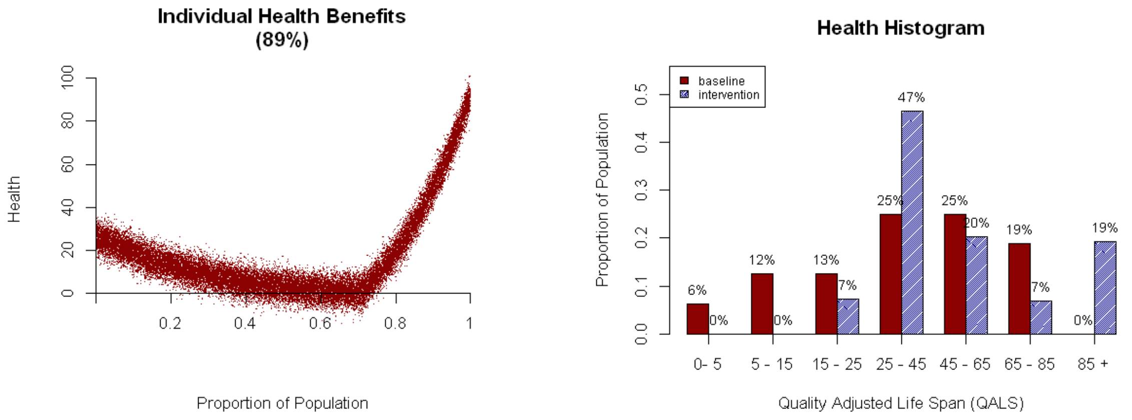


Next we turn to the Lorenz curve to examine the change in relative health inequality.  The baseline line is constructed by baseline health ranking, while the intervention line on the diagram is constructed by re-ranking the post intervention population by the health level achieved.  In the diagram below we can see that the intervention in this case results in a more equal relative distribution of health than the baseline.  Multiplying out by the means we can see that the generalised Lorenz curve for the intervention dominates that for the baseline, the crosses on the generalised Lorenz curves represent the point where the Lorenz curves intersect.

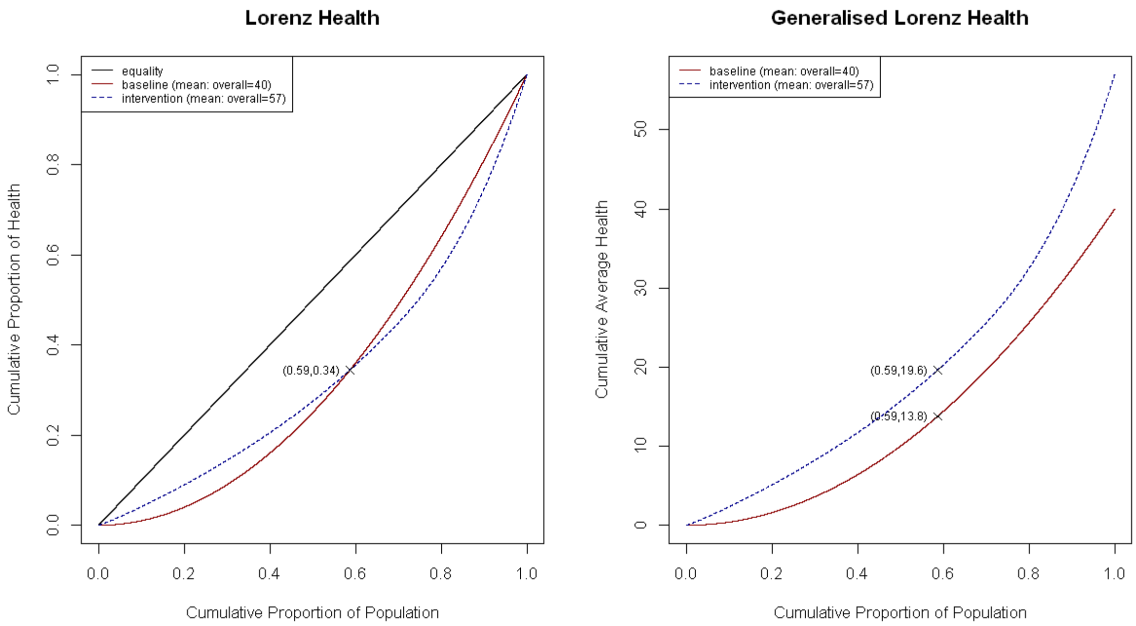

Looking to our battery of measures in the tables below we see that re-ranked Pareto dominance and generalised Lorenz dominance rules apply indicating that the intervention is preferred the baseline. All relative measures of inequality and poverty measures of inequality judge the intervention to be less unequal than the baseline, whilst all absolute measures of inequality suggest the intervention is more unequal than the baseline. The social welfare measures in both relative and absolute cases agree with the dominance rules in preferring the intervention to the baseline.

**Dominance Rules:**

| Rule | baseline | intervention |
|---|---|---|
| **Pareto Dominant** | FALSE | FALSE |
| **Re-ranked Pareto Dominant** | FALSE | TRUE |
| **Lorenz Dominant** | FALSE | FALSE |
| **Generalised Lorenz Dominant** | FALSE | TRUE |

**Relative Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 8.00 | 3.80 | -4.20 |
| **Relative Index of Inequality (RII)** | 2.00 | 1.98 | -0.02 |
| **Gini Index** | 0.33 | 0.33 | 0.00 |
| **Atkinson Index (e= 0.5 )** | 0.11 | 0.08 | -0.03 |
| **Atkinson Index (e= 1 )** | 0.26 | 0.16 | -0.11 |
| **Atkinson Index (e= 1.5 )** | 0.49 | 0.21 | -0.28 |
| **Atkinson Index (e= 2 )** | 0.77 | 0.26 | -0.51 |

**Absolute Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 64.00 | 97.12 | 33.12 |
| **Slope index of inequality (SII)** | 79.99 | 112.63 | 32.64 |
| **Kolm Index (alpha= 0.025 )** | 6.46 | 11.23 | 4.77 |
| **Kolm Index (alpha= 0.05 )** | 11.90 | 16.22 | 4.32 |
| **Kolm Index (alpha= 0.075 )** | 16.07 | 19.07 | 3.00 |
| **Kolm Index (alpha= 0.1 )** | 19.20 | 21.01 | 1.81 |

**Poverty Indices:**

| Index (z= 70 ) | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.88 | 0.76 | -0.12 |
| **Health Gap Ratio (intensity)** | 0.50 | 0.45 | -0.05 |
| **Sen Poverty Index** | 0.58 | 0.41 | -0.17 |

**Social Welfare Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Mean Health** | 40.00 | 57.00 | 17.00 |
| **Atkinson EDE (e= 0.5 )** | 35.56 | 52.20 | 16.65 |
| **Atkinson EDE (e= 1 )** | 29.44 | 48.12 | 18.67 |
| **Atkinson EDE (e= 1.5 )** | 20.39 | 44.75 | 24.36 |
| **Atkinson EDE (e= 2 )** | 9.02 | 42.04 | 33.02 |
| **Kolm EDE (alpha= 0.025 )** | 33.54 | 45.77 | 12.23 |
| **Kolm EDE (alpha= 0.05 )** | 28.10 | 40.78 | 12.68 |
| **Kolm EDE (alpha= 0.075 )** | 23.93 | 37.93 | 14.00 |
| **Kolm EDE (alpha= 0.1 )** | 20.80 | 35.99 | 15.19 |

### 4.3 Lorenz Dominance

In this example we compare two policies with equal mean health. From the diagrams below we can see that the intervention results in higher mean health for the least healthy two thirds of the population and lower mean health for the most healthy third of the population relative to their baseline levels of health.

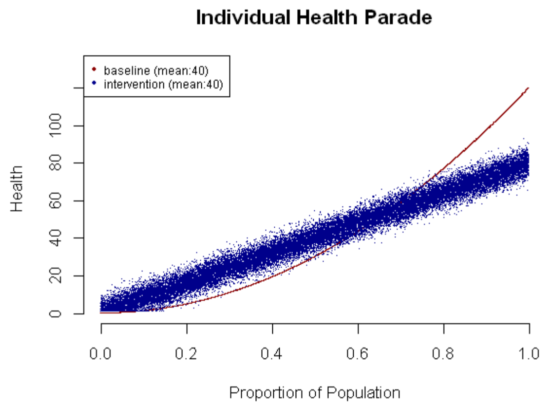


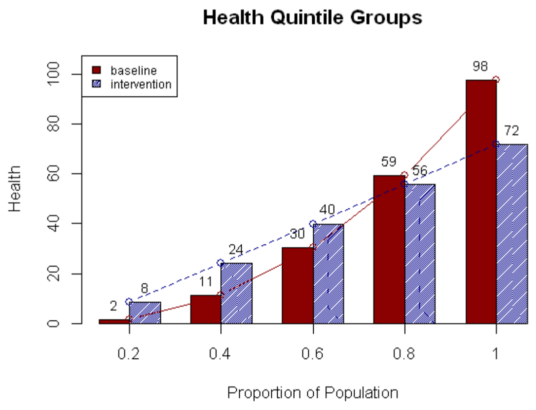


This is further emphasised by the individual health benefits plot showing that health benefits are positive for those with least baseline health and negative for those with greatest baseline health. Overall we can see that 66% of the population benefit from the intervention. The health distribution plot shows how the intervention results in reducing the proportion of the population at the extremes of the health distribution and concentrates the population towards the middle of this distribution.

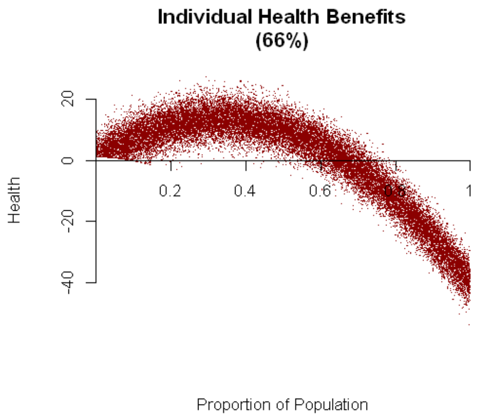


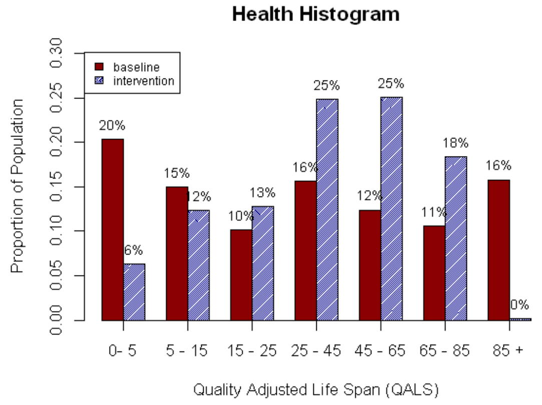


The Lorenz curves show the intervention to be more equal in relative terms than the baseline policy. The triangles drawn on the Lorenz curves show the point where the re-ranked absolute health profiles cross. The proportion of the population beyond this point under the intervention are worse off than their equivalently ranked counterpart in the baseline population.[5] The generalised Lorenz curves in

[5] This does not directly translate to the individuals better off under the intervention due to the re-ranking of the population post intervention

this case where mean health is equal look identical to the Lorenz curves with the appropriate rescaling of the axes.

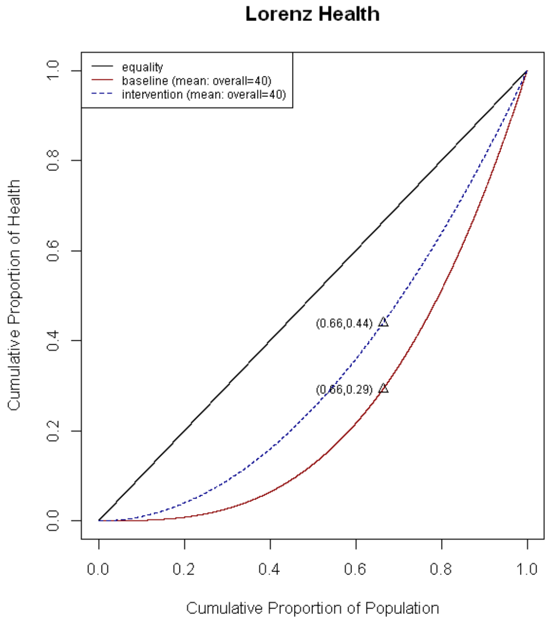


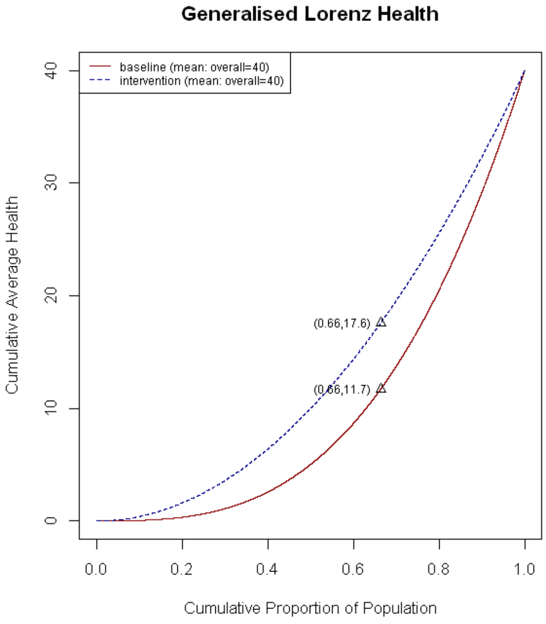


Looking to our battery of measures in the tables below we see that the Lorenz dominance and the generalised Lorenz dominance rules apply indicating that the intervention is preferred the baseline. All measures of inequality (with the exception of the headcount ratio) judge the intervention to be less unequal than the baseline. The social welfare measures in both relative and absolute cases agree with the dominance rules in preferring the intervention to the baseline.

**Dominance Rules:**

| Rule | baseline | intervention |
|---|---|---|
| **Pareto Dominant** | FALSE | FALSE |
| **Re-ranked Pareto Dominant** | FALSE | FALSE |
| **Lorenz Dominant** | FALSE | TRUE |
| **Generalised Lorenz Dominant** | FALSE | TRUE |

**Relative Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 59.97 | 8.11 | -51.86 |
| **Relative Index of Inequality (RII)** | 3.00 | 2.01 | -0.99 |
| **Gini Index** | 0.50 | 0.34 | -0.16 |
| **Atkinson Index (e= 0.5 )** | 0.25 | 0.11 | -0.14 |
| **Atkinson Index (e= 1 )** | 0.59 | 0.26 | -0.33 |
| **Atkinson Index (e= 1.5 )** | 0.96 | 0.46 | -0.50 |
| **Atkinson Index (e= 2 )** | 1.00 | 0.66 | -0.34 |

**Absolute Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 96.00 | 64.37 | -31.63 |
| **Slope index of inequality (SII)** | 119.98 | 80.44 | -39.54 |
| **Kolm Index (alpha= 0.025 )** | 12.62 | 6.51 | -6.11 |
| **Kolm Index (alpha= 0.05 )** | 19.65 | 11.98 | -7.67 |
| **Kolm Index (alpha= 0.075 )** | 23.74 | 16.17 | -7.57 |
| **Kolm Index (alpha= 0.1 )** | 26.37 | 19.30 | -7.06 |

**Poverty Indices:**

| Index (z= 70 ) | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.76 | 0.88 | 0.11 |
| **Health Gap Ratio (intensity)** | 0.67 | 0.50 | -0.17 |
| **Sen Poverty Index** | 0.64 | 0.59 | -0.05 |

**Social Welfare Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Mean Health** | 40.00 | 40.00 | 0.00 |
| **Atkinson EDE (e= 0.5 )** | 30.00 | 35.51 | 5.51 |
| **Atkinson EDE (e= 1 )** | 16.26 | 29.46 | 13.20 |
| **Atkinson EDE (e= 1.5 )** | 1.63 | 21.55 | 19.92 |
| **Atkinson EDE (e= 2 )** | 0.03 | 13.51 | 13.48 |
| **Kolm EDE (alpha= 0.025 )** | 27.38 | 33.49 | 6.11 |
| **Kolm EDE (alpha= 0.05 )** | 20.35 | 28.02 | 7.67 |
| **Kolm EDE (alpha= 0.075 )** | 16.26 | 23.83 | 7.57 |
| **Kolm EDE (alpha= 0.1 )** | 13.63 | 20.70 | 7.06 |

The dominance rules indicate that while there is no form of Pareto dominance. Atkinson's and therefore Shorrocks' theorem rank the intervention as better than the baseline.

This shows that assessing inequality independent of efficiency may be straightforward when there is no change in population health. With no change in population health we would be indifferent between the intervention and baseline if applying a decision rule to maximise population health. In these cases the level of inequality could represent an additional criterion that differentiates between the interventions in terms of social preference.

## 4.4 Generalised Lorenz Dominance

In this example we will see that we no longer have Lorenz dominance (the Lorenz curves cross) however we do have a greater mean health under the intervention than under the baseline policy. We will therefore need to turn to the generalised Lorenz curves, formed by multiplying the Lorenz curves by their respective mean levels of health.

The diagrams below show that the most healthy and least healthy people benefit from the intervention with those in between losing out.

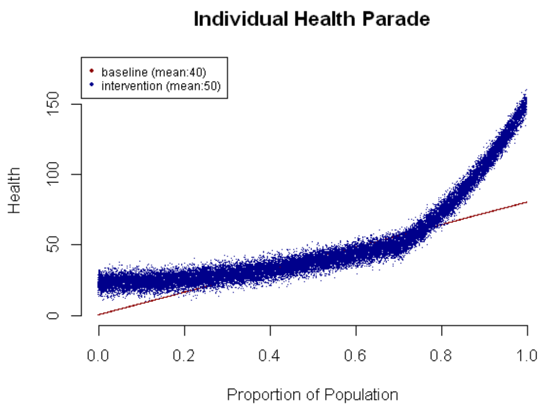


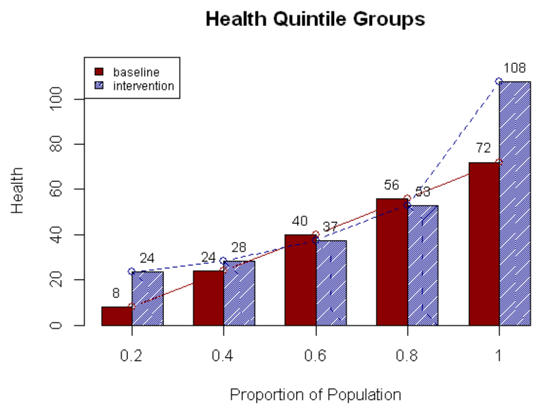


We see this even more clearly in the individual health benefits diagram in which we see that 68% of the people are better off with the intervention than they would be in the baseline case. The health histogram paints an interesting picture suggesting a large increase in population health for the least healthy and also the emergence of a super healthy group under the intervention.

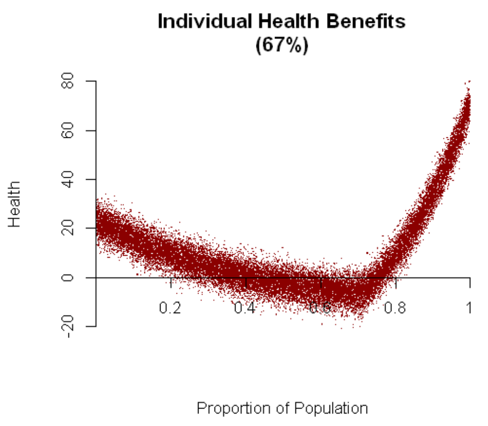


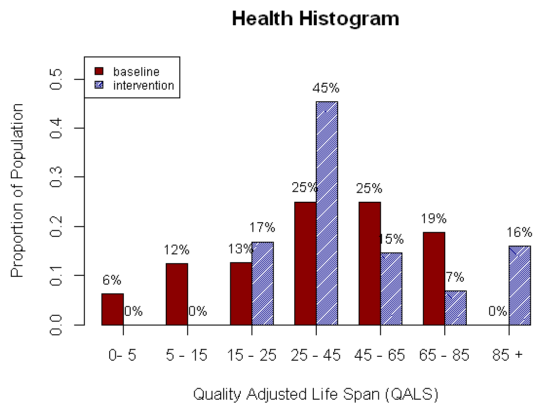


We see from the Lorenz curves that it is not straight forward to decide which of the policies is more equal in relative terms as the Lorenz curves cross. We also see from the two pairs of triangles on these curves that the re-ranked absolute health of the populations also crosses twice. The health of equivalently ranked individuals is better below the 42nd percentile and above the 76th percentile and worse in between. The Lorenz curve plot is annotated with the mean health under each policy as well as the mean health below and above the points that the curves cross. We can see that the intervention has greater mean health overall as well as both below and above the point of intersection[6].

Given that the Lorenz curves cross and the mean health for the intervention is greater than the mean health for the baseline we turn to the generalised Lorenz curves to see if Shorrock's theorem can be used to rank the policies. The Generalised Lorenz curve for the intervention does indeed dominate that of the baseline hence using Shorrocks's theorem we conclude that the intervention is preferred to the baseline. The crosses on the generalised Lorenz curves show the point that the Lorenz curves cross, below this point the intervention is more equal than the baseline policy. The triangles on the generalised Lorenz curves show the points where the re-ranked absolute health achieved under the policies cross. The points below the first pair of triangles and above the second pair of triangles are where equivalently ranked individuals under the intervention have more health than under the baseline policy. The points below the first triangles show where the intervention is both more equal and equivalently ranked people are better off under the intervention.

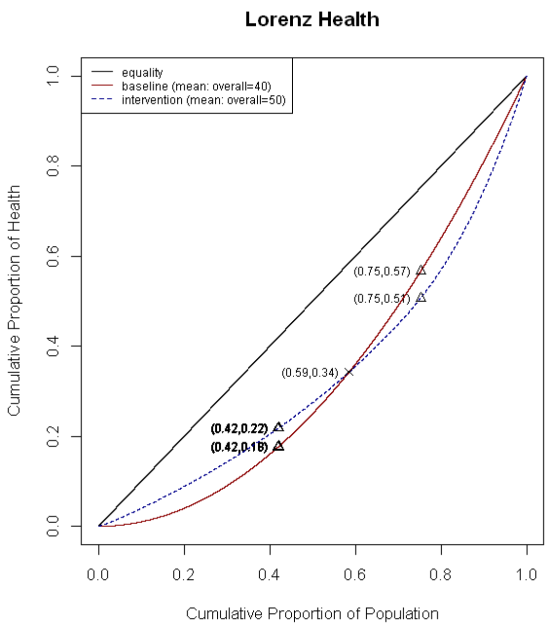


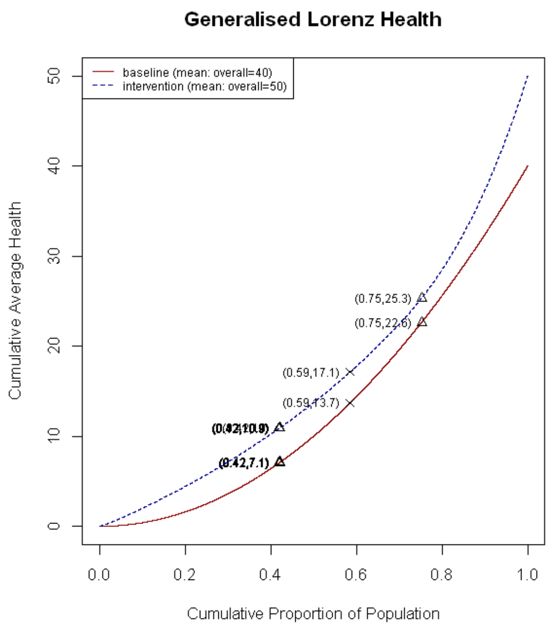


Looking to our battery of measures in the tables below we see that the generalised Lorenz dominance rule applies indicating that the intervention is preferred the baseline. All relative measures of inequality and poverty measures of inequality judge the intervention to be less unequal than the baseline, whilst all absolute measures of inequality particularly at lower levels of inequality aversion suggest the intervention is more unequal than the baseline. The social welfare measures in both

[6] If the overall mean for a policy is higher, then the mean above and below the cross point must also be higher for that policy and similarly if the overall mean is lower for a policy then the mean above and below the cross point must also be lower for that policy (see note on means at the end of this document for an explanation of why this must be the case).

relative and absolute cases agree with the dominance rules in preferring the intervention to the baseline.

**Dominance Rules:**

| Rule | baseline | intervention |
|---|---|---|
| **Pareto Dominant** | FALSE | FALSE |
| **Re-ranked Pareto Dominant** | FALSE | FALSE |
| **Lorenz Dominant** | FALSE | FALSE |
| **Generalised Lorenz Dominant** | FALSE | TRUE |

**Relative Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 8.00 | 3.87 | -4.13 |
| **Relative Index of Inequality (RII)** | 2.00 | 1.98 | -0.02 |
| **Gini Index** | 0.33 | 0.33 | 0.00 |
| **Atkinson Index (e= 0.5 )** | 0.11 | 0.08 | -0.03 |
| **Atkinson Index (e= 1 )** | 0.26 | 0.16 | -0.11 |
| **Atkinson Index (e= 1.5 )** | 0.49 | 0.22 | -0.27 |
| **Atkinson Index (e= 2 )** | 0.77 | 0.27 | -0.51 |

**Absolute Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 64.00 | 85.52 | 21.53 |
| **Slope index of inequality (SII)** | 79.99 | 99.15 | 19.17 |
| **Kolm Index (alpha= 0.025 )** | 6.46 | 9.10 | 2.65 |
| **Kolm Index (alpha= 0.05 )** | 11.90 | 13.49 | 1.59 |
| **Kolm Index (alpha= 0.075 )** | 16.07 | 16.07 | -0.01 |
| **Kolm Index (alpha= 0.1 )** | 19.20 | 17.84 | -1.36 |

**Poverty Indices:**

| Index (z= 70 ) | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.88 | 0.79 | -0.08 |
| **Health Gap Ratio (intensity)** | 0.50 | 0.50 | 0.00 |
| **Sen Poverty Index** | 0.58 | 0.47 | -0.11 |

**Social Welfare Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Mean Health** | 40.00 | 50.00 | 10.00 |
| **Atkinson EDE (e= 0.5 )** | 35.56 | 45.76 | 10.21 |
| **Atkinson EDE (e= 1 )** | 29.44 | 42.14 | 12.70 |
| **Atkinson EDE (e= 1.5 )** | 20.39 | 39.15 | 18.75 |
| **Atkinson EDE (e= 2 )** | 9.02 | 36.71 | 27.69 |
| **Kolm EDE (alpha= 0.025 )** | 33.54 | 40.90 | 7.35 |
| **Kolm EDE (alpha= 0.05 )** | 28.10 | 36.51 | 8.41 |
| **Kolm EDE (alpha= 0.075 )** | 23.93 | 33.93 | 10.01 |
| **Kolm EDE (alpha= 0.1 )** | 20.80 | 32.16 | 11.36 |

## 4.5 More Cost Effective Less Equal

In this example we will see that the intervention policy increases mean health but also increases health inequality in relative and absolute terms. This example represents an extreme case of the intervention generated inequalities evident in many real public health interventions. In such cases none of our dominance rules apply and we need instead to turn to the inequality measures we have defined to help us decide which policy we prefer.

The diagrams below show that while everybody benefits from the intervention the most healthy seem to benefit more than the least healthy in both relative and absolute terms.

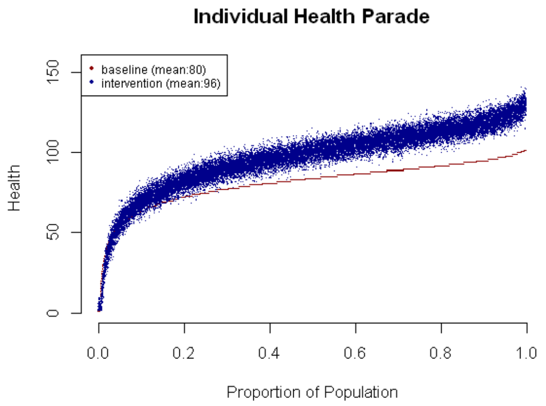


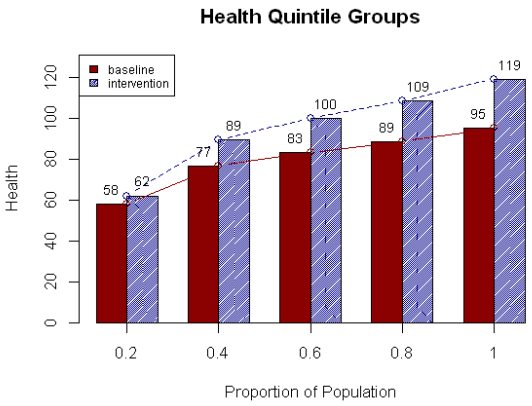


We see this even more clearly in the individual health benefits diagram in which we see that 95% of the people are better off with the intervention than they would be in the baseline case with health benefit increasing with baseline health.  The health histogram further confirms this conclusion.

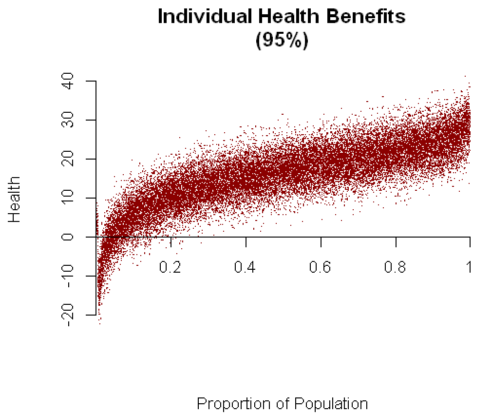


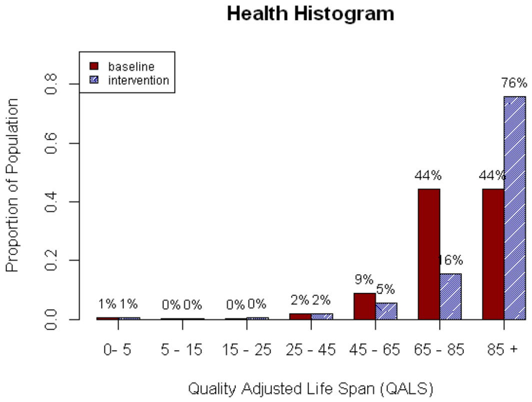


We see from the Lorenz curves that the baseline health looks relatively more equal than the intervention health, though due to the difference in mean health between the two Atkinson's rule does not apply.

Given mean health for the intervention is greater than the mean health for the baseline it is unsurprising to see the dominance relationship between the curves partially reversed, however the generalised Lorenz curves cross (the intersection point is marked by the circle in the figure), hence Shorrocks' theorem does not apply either.

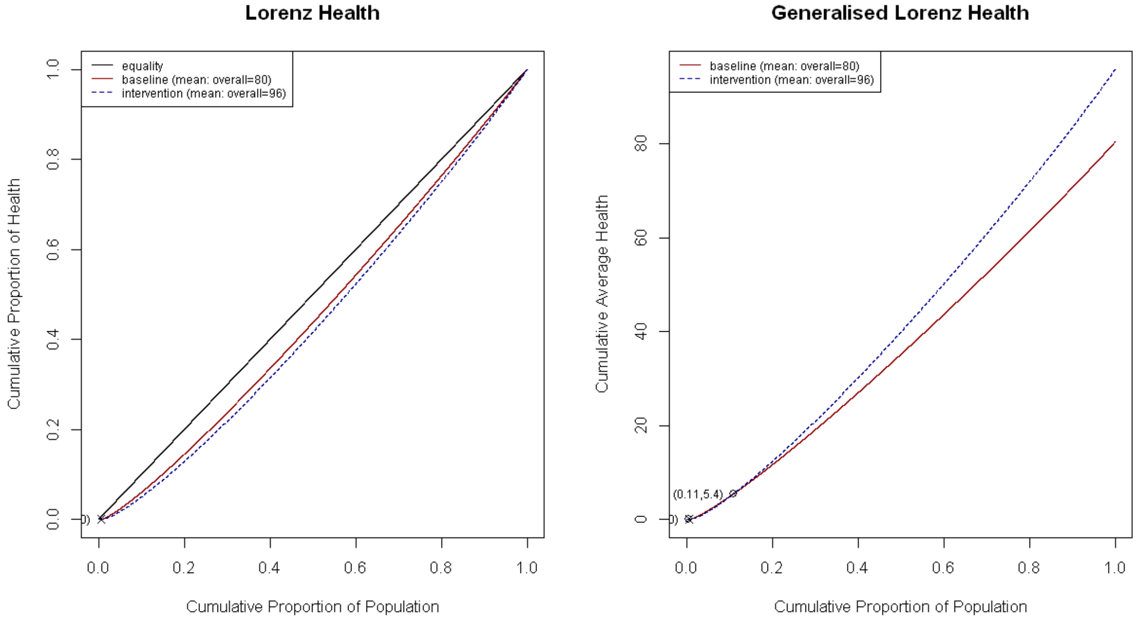


Looking to our battery of measures in the tables below we see that none of our dominance rules apply, hence we are unable to unambiguously prefer one distribution to the other, we will instead need to turn to our social welfare indices to decide between the distributions. Relative, absolute and poverty measures of inequality all strongly suggest that the intervention is more unequal that the baseline. The social welfare indices however recognise the greater level of mean health generated by the intervention and except for at the higher levels of inequality aversion prefer the intervention to the baseline. It should be noted that the choice of suitable inequality aversion parameters is critical in making this decision.

**Dominance Rules:**

| Rule | baseline | intervention |
|---|---|---|
| **Pareto Dominant** | FALSE | FALSE |
| **Re-ranked Pareto Dominant** | FALSE | FALSE |
| **Lorenz Dominant** | FALSE | FALSE |
| **Generalised Lorenz Dominant** | FALSE | FALSE |

**Relative Inequality Indices:**

| Index | baseline | intervention | difference<br>( intervention - baseline ) |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 0.64 | 0.94 | 0.31 |
| **Relative Index of Inequality (RII)** | 0.55 | 0.73 | 0.18 |
| **Gini Index** | 0.09 | 0.12 | 0.03 |
| **Atkinson Index (e= 0.5 )** | 0.01 | 0.02 | 0.01 |
| **Atkinson Index (e= 1 )** | 0.04 | 0.05 | 0.01 |
| **Atkinson Index (e= 1.5 )** | 0.13 | 0.13 | 0.01 |
| **Atkinson Index (e= 2 )** | 0.45 | 0.36 | -0.09 |

**Absolute Inequality Indices:**

| Index | baseline | intervention | difference<br>( intervention - baseline ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 36.94 | 58.22 | 21.28 |
| **Slope index of inequality (SII)** | 44.13 | 69.72 | 25.59 |
| **Kolm Index (alpha= 0.025 )** | 3.44 | 8.04 | 4.60 |
| **Kolm Index (alpha= 0.05 )** | 9.75 | 21.32 | 11.57 |
| **Kolm Index (alpha= 0.075 )** | 19.82 | 36.40 | 16.58 |
| **Kolm Index (alpha= 0.1 )** | 30.57 | 48.15 | 17.58 |

**Poverty Indices:**

| Index (z= 70 ) | baseline | intervention | difference<br>( intervention - baseline ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.17 | 0.11 | -0.06 |
| **Health Gap Ratio (intensity)** | 0.20 | 0.29 | 0.09 |
| **Sen Poverty Index** | 0.05 | 0.05 | 0.00 |

**Social Welfare Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Mean Health** | 80.37 | 95.84 | 15.47 |
| **Atkinson EDE (e= 0.5 )** | 79.36 | 93.98 | 14.62 |
| **Atkinson EDE (e= 1 )** | 77.31 | 90.79 | 13.48 |
| **Atkinson EDE (e= 1.5 )** | 70.14 | 83.01 | 12.87 |
| **Atkinson EDE (e= 2 )** | 44.36 | 61.26 | 16.90 |
| **Kolm EDE (alpha= 0.025 )** | 76.93 | 87.80 | 10.88 |
| **Kolm EDE (alpha= 0.05 )** | 70.63 | 74.52 | 3.90 |
| **Kolm EDE (alpha= 0.075 )** | 60.55 | 59.45 | -1.10 |
| **Kolm EDE (alpha= 0.1 )** | 49.80 | 47.69 | -2.11 |

### 4.6 Less Effective More Equal

In this example mean health under the intervention is lower than under the baseline, however the distribution of health under the intervention is more equal than that under the baseline. We can see this example as one where a policy decision now impacts the distribution and magnitude of future improvements in health.

The diagrams below show that the while health is lower for all under the intervention the distribution of health in both absolute and relative terms is more equal.

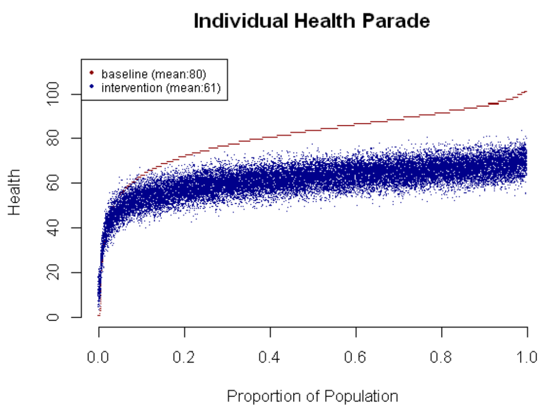


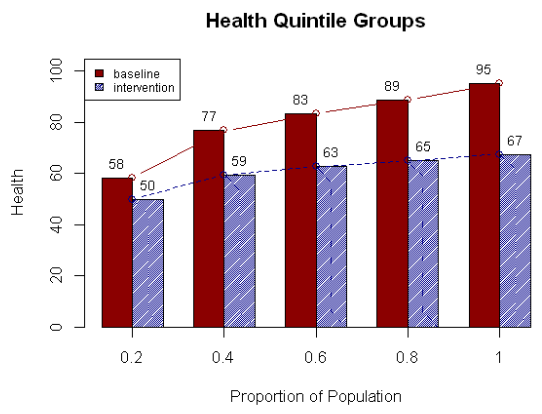


The individual health benefits diagram shows that only 2% of the people are better off with the intervention than they would be in the baseline case.  The health histogram supports the idea of inequality reduction showing a much greater concentration of health around the mean under the intervention as compared to the baseline.

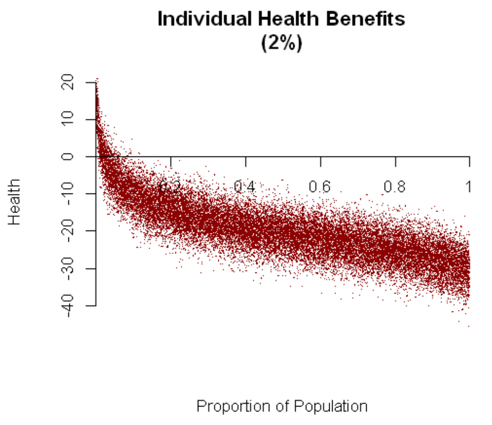


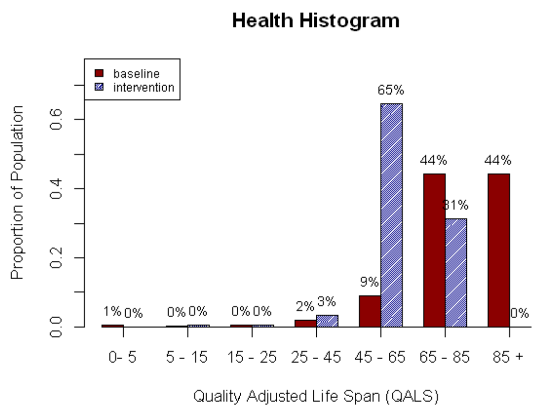


The Lorenz curve for the intervention almost entirely dominates that for the baseline, however the difference in means does not allow us to use Atkinson's theorem.  Given the higher baseline health, the generalised Lorenz curves reverse this relationship, through the fact that they cross prevents us from using Shorrocks' theorem.

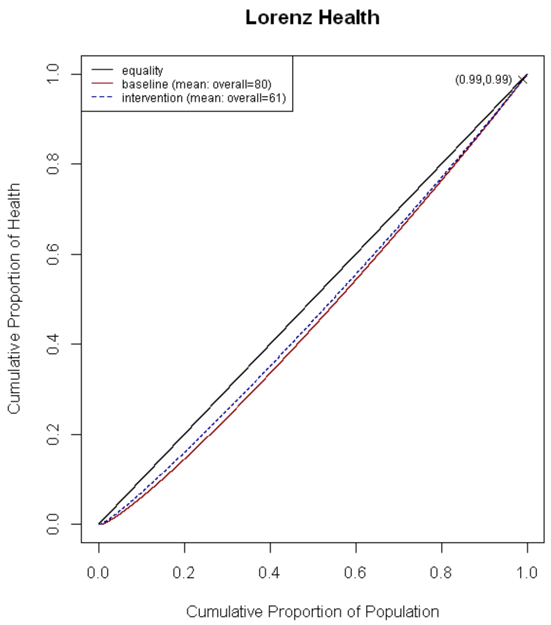


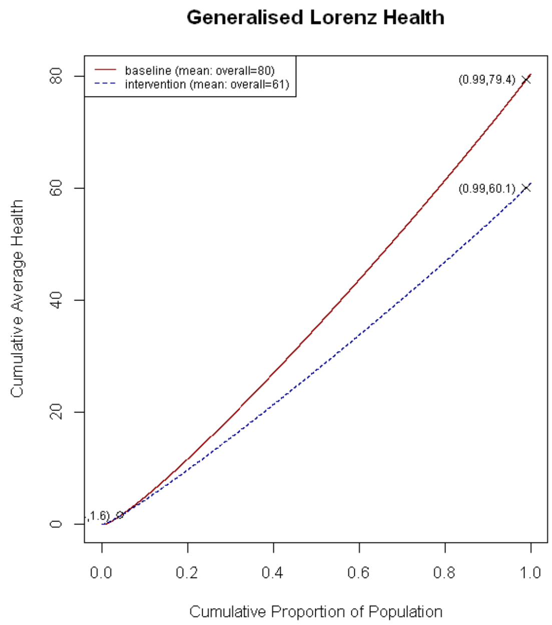


Looking to our battery of measures in the tables below we see that as in the previous example none of our dominance rules apply, hence we are unable to unambiguously prefer one distribution to the other, we will instead need to turn to our social welfare indices to decide between the distributions. Relative, absolute and poverty measures of inequality all strongly suggest that the intervention is less unequal that the baseline. The social welfare indices however recognise the greater mean level of health in the baseline and except for at the higher levels of inequality aversion prefer the baseline to the intervention. Again it should be noted that the choice of suitable inequality aversion parameters is critical in making this decision.

**Dominance Rules:**

| Rule | baseline | intervention |
|---|---|---|
| **Pareto Dominant** | FALSE | FALSE |
| **Re-ranked Pareto Dominant** | FALSE | FALSE |
| **Lorenz Dominant** | FALSE | FALSE |
| **Generalised Lorenz Dominant** | FALSE | FALSE |

**Relative Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Relative Gap Index (ratio)** | 0.64 | 0.44 | -0.19 |
| **Relative Index of Inequality (RII)** | 0.55 | 0.42 | -0.13 |
| **Gini Index** | 0.09 | 0.07 | -0.02 |
| **Atkinson Index (e= 0.5 )** | 0.01 | 0.01 | -0.01 |
| **Atkinson Index (e= 1 )** | 0.04 | 0.01 | -0.02 |
| **Atkinson Index (e= 1.5 )** | 0.13 | 0.02 | -0.10 |
| **Atkinson Index (e= 2 )** | 0.45 | 0.04 | -0.41 |

**Absolute Inequality Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Absolute Gap Index (range)** | 36.94 | 21.50 | -15.44 |
| **Slope index of inequality (SII)** | 44.13 | 25.74 | -18.39 |
| **Kolm Index (alpha= 0.025 )** | 3.44 | 1.02 | -2.43 |
| **Kolm Index (alpha= 0.05 )** | 9.75 | 2.46 | -7.29 |
| **Kolm Index (alpha= 0.075 )** | 19.82 | 4.64 | -15.18 |
| **Kolm Index (alpha= 0.1 )** | 30.57 | 7.87 | -22.70 |

**Poverty Indices:**

| Index (z= 70 ) | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Head Count Ratio (incidence)** | 0.17 | 0.92 | 0.74 |
| **Health Gap Ratio (intensity)** | 0.20 | 0.15 | -0.05 |
| **Sen Poverty Index** | 0.05 | 0.19 | 0.13 |

**Social Welfare Indices:**

| Index | baseline | intervention | difference ( intervention - baseline ) |
|---|---|---|---|
| **Mean Health** | 80.37 | 60.80 | -19.58 |
| **Atkinson EDE (e= 0.5 )** | 79.36 | 60.44 | -18.93 |
| **Atkinson EDE (e= 1 )** | 77.31 | 59.98 | -17.33 |
| **Atkinson EDE (e= 1.5 )** | 70.14 | 59.38 | -10.76 |
| **Atkinson EDE (e= 2 )** | 44.36 | 58.50 | 14.15 |
| **Kolm EDE (alpha= 0.025 )** | 76.93 | 59.78 | -17.15 |
| **Kolm EDE (alpha= 0.05 )** | 70.63 | 58.34 | -12.29 |
| **Kolm EDE (alpha= 0.075 )** | 60.55 | 56.15 | -4.40 |
| **Kolm EDE (alpha= 0.1 )** | 49.80 | 52.93 | 3.13 |

## 5 Conclusion

The income inequality literature provides a rich set of tools for incorporating distributional concerns into economic analysis. In this document we have explored ways in which these methods can be applied to health distributions to help address the growing concern for health inequalities. We have seen through the use of stylised examples how we can use these graphical and numerical tools to better understand health distributions, and how this understanding can help us address our concerns for health inequality.

Our examples have shown that the various tools for exploring health distributions do not always agree in their conclusions about which health distribution should be considered more equal and we have outlined the key assumptions underlying the different measures used and emphasised the importance of assessing a range of measures rather than just picking a “best” measure.

Our examples have further demonstrated that where we can apply dominance rules, we can make very general decisions about which health distribution we prefer, requiring very few assumptions on the form of our social welfare function. However where our dominance rules do not apply we need to turn to our various social welfare functions parameterised with our assessment of the societal level on inequality aversion to make a judgement. This is a much more contentious process requiring judgements to be made on the appropriateness of different forms of social welfare functions as well as on different levels of inequality aversion. In such situations it is advisable to look at a range of functions across a range of inequality aversion parameters to help inform the deliberation process.

## 6 References


Atkinson, A. (2010). Atkinson Brocher Lecture revised October 2010. *Measurement and Ethical Evaluation of Health Inequalities*.

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory*, *2*(3), 244–263.

Atkinson, A. B. (1983). *The Economics of Inequality* (Second., p. 330). Oxford University Press.

Bleichrodt, H., Diecidue, E., & Quiggin, J. (2004). Equity weights in the allocation of health care: the rank-dependent QALY model. *Journal of health economics*, *23*(1), 157-71. doi:10.1016/j.jhealeco.2003.08.002

Coudouel, A., Hentschel, J. S., & Wodon, Q. T. (2002). Poverty Measurement and Analysis. *World Bank*, 27-74.

Cowell, F. (2011). *Measuring Inequality* (p. 304). Oxford University Press.

Dasgupta, P., & Sen, A. (1973). Notes on the Measurement of Inequality. *Journal of economic theory*, *187*, 180-187.

Dolan, P., & Tsuchiya, A. (2009). The social welfare function and individual responsibility: some theoretical issues and empirical evidence. *Journal of health economics*, *28*(1), 210-20. doi:10.1016/j.jhealeco.2008.10.003

Erreygers, G. (2009). Can a single indicator measure both attainment and shortfall inequality? *Journal of health economics*, *28*(4), 885-93. doi:10.1016/j.jhealeco.2009.03.005

Erreygers, G., & Van Ourti, T. (2011). Measuring socioeconomic inequality in health, health care and health financing by means of rank-dependent indices: a recipe for good practice. *Journal of health economics*, *30*(4), 685-94. Elsevier B.V. doi:10.1016/j.jhealeco.2011.04.004

Erreygers, G., Clarke, P., & Van Ourti, T. (2011). “Mirror, mirror, on the wall, who in this land is fairest of all?”-Distributional sensitivity in the measurement of socioeconomic inequality of health. *Journal of health economics*. Elsevier B.V. doi:10.1016/j.jhealeco.2011.10.009

Fleurbaey, M., & Schokkaert, E. (2009). Unfair inequalities in health and health care. *Journal of health economics*, *28*(1), 73-90. doi:10.1016/j.jhealeco.2008.07.016

Jenkins, S., & Lambert, P. J. (1997). Three “I”s of poverty curves, with an analysis of UK poverty trends. *Oxford Economic Papers*, *49*, 317-327.

Johannesson, M. (2001). Should we aggregate relative or absolute changes in QALYs? *Health Economics*, *10*(7), 573–577. Wiley Online Library. doi:10.1002/hec.646

Kolm, S.-C. (1976a). Unequal inequalities. I. *Journal of Economic Theory*, *12*(3), 416-442. doi:10.1016/0022-0531(76)90037-5

Kolm, S.-C. (1976b). Unequal inequalities. II. *Journal of Economic Theory*, *13*(1), 82–111. Elsevier.

Lambert, P. J. (2001). *The Distribution and Redistribution of Income* (Third., p. 313). Manchester University Press.

Mackenbach, J. P., & Kunst, A. E. (1997). Measuring the magnitude of socio-economic inequalities in health: an overview of available measures illustrated with two examples from Europe. *Social Science & Medicine*, *44*(6), 757–771. Elsevier.

Marshall, A. W., & Olkin, I. (1974). Majorization in multivariate distributions. *The Annals of Statistics*, 1189–1200. JSTOR.

Norheim, O. F. (2010). Gini Impact Analysis: Measuring Pure Health Inequity before and after Interventions. *Public Health Ethics*, *3*(3), 282-292. doi:10.1093/phe/phq017

ONS. (2010). England and Wales, Interim Life Tables, 1980-82 to 2007-09. *Life Expectancies, Reference* .

O'Donnell, O., van Doorslaer, E., Wagstaff, A., & Lindelow, M. (2008). *Analysing health equity using household survey data: a guide to techniques and their implementation*. *Bulletin of the World Health Organization* (Vol. 86). World Bank Institute.

Pen, J. (1974). *Income distribution* (p. 437). Penguin Books.

Sen, Amartya. (1976). Poverty: An ordinal approach to measurement. *Econometrica*, *44*(2), 219-231.

Sen, A. K. (1973). *On economic inequality* (p. 118). Norton.

Shorrocks, A. F. (1983). Ranking Income Distributions. *Economica*, *50*(197), 3. doi:10.2307/2554117

Stolk, E. a, Pickee, S. J., Ament, A. H. J. a, & Busschbach, J. J. V. (2005). Equity in health care prioritisation: an empirical inquiry into social value. *Health policy (Amsterdam, Netherlands)*, *74*(3), 343-55. doi:10.1016/j.healthpol.2005.01.018

Wagstaff, A, Paci, P., & van Doorslaer, E. (1991). On the measurement of inequalities in health. *Social science & medicine (1982)*, *33*(5), 545-57.

Wagstaff, Adam. (1991). QALYs and the equity-efficiency trade-off. *Journal of Health Economics*, *10*(1), 21–41. Elsevier.

Williams, A. (1997). Intergenerational equity: an exploration of the "fair innings" argument. *Health economics*, *6*(2), 117-32.

**Appendix A: Univariate Assessment – Additional Dominance Principles**

The income inequality literature includes a number of further dominance principles, in addition to the Atkinson and Shorrocks principles described in the main text. We do not think these are likely to be helpful in the case of health inequality, as they require value judgements which are hard to interpret and potentially controversial in a health context. So we prefer to stick to the relatively uncontroversial dominance principles, and then move directly on to inequality, poverty and social welfare indices to analyse the implications of stronger social value judgements.

**Appendix B: Note on means**

Common total population: $P$
Mean health baseline: $\overline{H}_B$
Mean health intervention: $\overline{H}_I$

Total health baseline: $H_B = P * \overline{H}_B$
Total health intervention: $H_I = P * \overline{H}_I$

Common proportion of population at cross point: $P_{cross}$
Common proportion of health at cross point: $H_{cross}$

Total health below cross point baseline: $H_{(Low)B} = H_B * H_{cross}$
Total health below cross point intervention: $H_{(Low)I} = H_I * H_{cross}$

Total health above cross point baseline: $H_{(High)B} = H_B * [1-H_{cross}]$
Total health above cross point intervention: $H_{(High)I} = H_I * [1-H_{cross}]$

Ordering on mean health implies ordering of total health as population is common across policies
$H_I > H_B$ if and only if $\overline{H}_I > \overline{H}_B$

Ordering for total health implies ordering for total health above and below the cross point
If $H_I > H_B$ then $H_{(Low)I} > H_{(Low)B}$ and $H_{(High)I} > H_{(High)B}$

Ordering for total health above and below cross points implies ordering for mean health above and below the cross points as we have a common population and population cross point.

This implies that if mean health is higher overall then it must be higher both above and below the cross point. For the same reason if mean health is lower overall it must be lower both above and below the cross points.