Measuring health inequality in the context of cost-effectiveness analysis: Don't concentrate on the concentration index!

Miqdad Asaria, Susan Griffin, Richard Cookson, Karl Claxton, Tony Culyer, Nigel Rice, Mark Sculpher

Centre for Health Economics, University of York

Paper Presented at the Health Economists' Study Group meeting

June 2012, University of Oxford

Abstract: In this paper we develop methods to evaluate public health interventions when there are combined, and potentially conflicting, objectives of increasing population health and reducing population health inequalities. To do this we propose a framework that carefully adapts and combines the techniques of cost-effectiveness analysis and the tools of inequality measurement. We briefly introduce methods for cost-effectiveness analysis and inequality analysis. We then go on to show how these methods can be combined in our proposed framework and explore this framework through a worked example. Finally we compare the analysis produced by our framework with that produced by the more traditional methods used in the health inequalities literature to describe the additional insights made possible by combining these bodies of work. We show that traditional health inequality analysis while useful in exploring specific social determinants of health can be limiting when more sophisticated analysis is required. Our methods drawing from the wider literature on economic inequality allow us to describe health distributions in a manner that is directly informative in the context of public health priority setting, and in so doing clearly outline all underpinning social value judgements.

Correspondence to:

Miqdad Asaria, Centre for Health Economics, University of York, Heslington, York, YO10 5DD, Tel: +44 (0)1904 321973, Email: miqdad.asaria@york.ac.uk

Introduction

Cost-effectiveness analyses of healthcare interventions typically focus on calculating the mean change in health for a subset of the population who have a particular illness, where the intervention is a specific drug or device targeted at this illness. Health inequality measures such as the concentration curve are typically used to describe how unequal the health of the population is in order to compare health inequality across countries and/or over time.

In this paper we aim to develop methods to evaluate public health interventions when there are combined, and potentially conflicting, objectives of increasing population health and reducing population health inequalities. To do this we propose a framework that carefully adapts and combines the techniques of cost-effectiveness analysis and the tools of inequality measurement. We briefly introduce methods for cost-effectiveness analysis and inequality analysis. We then go on to show how these methods can be combined in our proposed framework and explore this framework through a worked example. Finally we compare the analysis produced by our framework with that produced by the more traditional methods used in the health inequalities literature to describe the additional insights made possible by combining these bodies of work.

Some of the key advantages of our framework over existing approaches to cost-effectiveness and health inequality analysis are that it can be used to:

- model the distribution of health benefits in addition to the total health benefits
- explicitly incorporate the health opportunity costs of displaced activities and the distribution thereof
- analyse multiple concepts of health inequality and health poverty with clarity about the underpinning social value judgements
- analyse concepts of distributional dominance in the same framework as concepts of inequality
- analyse tradeoffs between inequality and mean health with clarity about the underpinning social value judgements
- combine multiple unfair determinants of health inequality into a unified analysis, encompassing the concentration index approach as a special case.

Using cost-effectiveness analysis to estimate net health benefits

The role of cost-effectiveness analysis is to provide information to decision makers to help them identify which of a set of investment opportunities would provide the greatest payoff. We take as a starting point a societal decision making framework and describe a typical cost-effectiveness analysis conducted from a health sector perspective. That is, we assume a fixed budget for the provision of health interventions, and that the interventions do not have important non-health benefits or impose costs outside the health sector. When the primary objective of the decision maker is assumed to be maximisation of expected population health levels from available resources, total population health is regularly calculated as the sum total of expected individual discounted health levels. This typifies the use of cost-effectiveness analysis by health technology assessment agencies such as NICE.

Theoretically, cost-effectiveness analysis could take the form of estimating, for all possible interventions for all potential recipients, the expected health outcomes and cost burden and comparing this to the total amount of resources available. The whole allocation problem could then be solved simultaneously by using mathematical programming techniques. However, the informational requirements of this 'global', system-wide approach make it impractical given the need to inform decisions in a timely manner. In general cost-effectiveness analysis informs the choice between the interventions that are relevant to a more narrowly defined decision problem, for example the treatment options available for a specific indication or disease. Should one of the mutually exclusive options be found to offer an improvement in health but at a greater cost than the currently funded alternative, the potential improvement in health outcomes must be compared to the health lost as a result of discontinuing other currently funded activities by reallocating resources to fund the new intervention.

Health outcomes are quantified in terms of a single cardinal measure of health that reflects mortality and morbidity, for example a quality adjusted life year. The purpose is to have a common measure that can be applied to all evaluations, allowing comparisons to be made across a range of health conditions. The resources of interest, which determine the cost burden of the interventions, are simply those that fall on the budget constraint, i.e. resources within the health sector. The budget is fully allocated, so introducing a new health intervention displaces another. The health gains forgone by displacing existing health interventions represent the opportunity cost of introducing the new intervention. If the health gains from the new intervention exceed this opportunity cost then it is regarded as cost-effective. Box 1 describes the basic principles of evaluation.

Box 1. Evaluating the cost-effectiveness of new health interventions

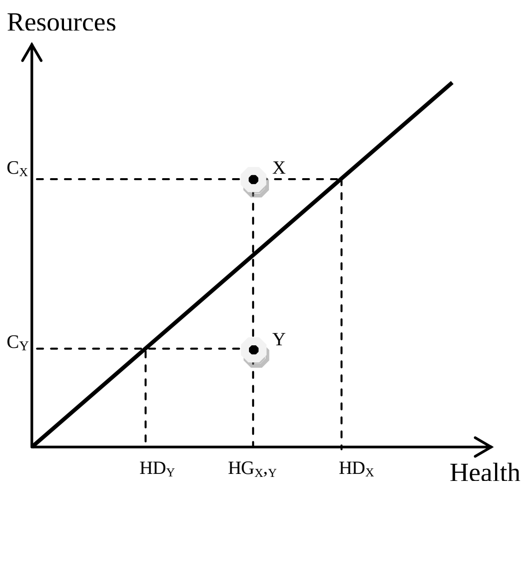

The slope of the diagonal line represents the amount of health that can be generated per unit of resources used to provide existing health interventions. Two new interventions (X and Y) are evaluated to determine whether they should be provided within the health sector.

Intervention X offers a health gain HG_X at a cost of C_X. Reducing provision of existing services to fund X would displace amount of health HD_X, resulting in a loss of health overall as $HG_X - HD_X < 0$

Intervention Y offers a health gain HG_Y at an additional cost of C_Y. Reducing provision of existing services to fund Y would displace amount of health HD_Y, resulting in a gain of health overall as $HG_Y - HD_Y > 0$

Identifying and evaluating displaced health service activity is not often undertaken. Instead an estimation of the shadow price of the budget constraint provides a threshold for assessing cost-effectiveness. This threshold represents an estimate of the impact of the displacement that would occur by diverting resources to the new intervention in terms of health forgone, i.e. the opportunity cost of resources in health terms. When the amount of additional resources demanded by each intervention has only a marginal impact on existing services the cost-effectiveness threshold will not alter with each reallocation of resources, and a common estimate can be used within each budgetary period. Thus a common cost-effectiveness threshold can be used to describe the cost burden of interventions in terms of health, and these can be subtracted from the expected health gains attributed to an intervention in order to describe its net health benefit. This approach means that cost-effectiveness analyses identify and characterise the eligible patient population for the intervention being evaluated, but the characteristics and health levels of non-recipients are left undetermined.

Changes in health levels between the interventions being compared are typically calculated from the point at which the treatment decision is made, leaving health achieved up to that point unspecified. From the point of the treatment decision outcomes that do not differ according to the intervention received may also be left unspecified. So, interventions are evaluated by focusing on changes in health, without explicitly characterising the resulting levels of health.

Extending cost-effectiveness analysis to estimate "fairness" adjusted health distributions

In order to extend this analysis to allow us to evaluate health inequalities we require a description of the distribution of health. The distribution of changes in health attributed to an intervention is informed by the distribution of the health gains among recipients of the intervention, and the distribution of the opportunity costs among those who would have received the displaced activities. If we wish to go further and assess changes in the distribution of health levels, we require further information on total health levels without the intervention. Finally, a distinction can be made between any variation in health and only variation that is considered to be unfair by identifying equity relevant subgroups of both recipient and non-recipient populations. To do all of this we require our cost-effectiveness analysis to characterise any variation explained by subgroup characteristics in:

1. treatment effect in terms of health gains and resource use;
2. incidence or eligibility for the interventions being evaluated;
3. health opportunity cost;
4. total health levels without the intervention

This would allow the extent of each of these to be estimated for each sub-group within a population. A value judgement is required as to which inequalities in the distribution of health are regarded as unfair. If the inequality concern does not apply to all sources of variation in health – for example, if some determinants of individual ill health are deemed to be a matter of unavoidable bad luck or individual responsibility – then further analysis is required in order describe the distributional impact on unfair variation in health outcomes. For example, differences in health outcomes could be presented for relevant sub-groups between whom equality in health is desired. Alternatively, this could take the form of a multivariate analysis in which estimates of total health are adjusted to control for fair differences in order to describe the unfair distribution of health. This requires that additional data be collected on equity relevant characteristics alongside data on treatment efficacy, health-related quality of life and resource use.

It is important to note that because we are operating within the framework of cost-effectiveness analysis we assume that health is measured on a fixed cardinal scale such as in QALYs or quality adjusted life expectancy. We will see that the extent of any reduction in inequality may need to be compared to the extent of improvement in total health. While it is possible to compare ordinal health distributions using a slightly different set of inequality measurement tools than those described in the following section, without a fixed cardinal scale health variable it may not be possible to estimate net health changes in distributions and in particular to allow for health opportunity costs appropriately. This inability to incorporate distributions of opportunity costs into the analysis would be a serious drawback, and due to this limitation we suggest sticking with fixed cardinal scale health variables when using our framework.

Comparing health distributions

In this section we draw heavily from the literature on income inequalities. 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 scales that enable comparisons between all individuals. However, there are also some important differences that should be considered, and the application and interpretation of income inequality measures to health inequality must be done in light of these differences between the two domains (A. Atkinson, 2010).

- 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)
- Income is unbounded while health has an upper bound
- Full equality of income between individuals is in principle achievable but some inequality in health between individuals is irremediable and incompensable
- 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)

- Individual income is usually assumed to have diminishing marginal value, whereas any such assumption about individual health is controversial

Once we have identified the health distribution of interest (i.e. decided which sources of inequality should be consider unfair and which health concept should be the focus of distributional concern such as expected lifetime health measured in QALYs), we finally need to decide how to measure and quantify the level of inequality. This last decision requires that we understand some key properties of inequality measures as outlined below.

Weak principle of transfers[1]*:*

The most universally recognised concept of what we mean by inequality is the weak principle of transfers. 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 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 to another. Instead it is future expected health gains that would be transferred between individuals.

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

Scale independence – relative inequality measures:

Scale independence focuses attention on concern for relative inequality between individuals, rather than the size or scale of absolute differences between individuals. It 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.

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

Translation independence – absolute inequality measures:

Translation independence 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. 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.

Entire distribution versus poverty measures:

Next we need to decide whether we are interested in the entire distribution of health or just a specific part of this distribution. Poverty measures concentrate only on individuals whose health lies below a certain threshold level or "poverty line" deemed necessary for meeting basic needs[2]. The implication for such measures being that any reduction in the health of any person below the health poverty line should increase the measure of health poverty, and that the health poverty measure should be invariant to changes in health of those above the health poverty line (A. Sen, 1976). In the context of health, one can define 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 easiest way to reduce inequality is to reduce the health of the better off rather than to improve the health of the worse off. 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 being 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).

While these concepts help us to understand the variety of inequality measures and what we mean when we say one distribution is more equal than another according to one particular conception of inequality, they may tell us little about which distribution we prefer. This is an issue even with a single objective of reducing health inequality, but even more so where we have a conflict between dual goals of increasing overall health and making the distribution of health more equal. To address these trade-offs and inform decisions about which intervention should be reimbursed we need to turn to measures of health related social welfare.

[2] Note that the use of a health poverty line is one way of defining an equity relevant sub-group, as was discussed in the previous section.

Social welfare functions combining concern for mean health and the distribution of health

A health-related social welfare function (SWF) can in principle 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 useful when constructing a SWF. In describing these properties we use the terminology h_{iA} to represent the health of individual *i* in health 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. We can think of individual health as being measured in terms of quality adjusted life expectancy. The individual utility function can then be interpreted as representing the value to society of the individual's quality adjusted life expectancy. The suffixes, A and B, are used to distinguish different health distributions. (A. K. Sen, 1973)(Cowell, 2011)

Individualistic:
This means the SWF is a function of the individual utilities i.e. the SWF has the form:
$W_A = W(U_1, U_2 \ldots , U_n)$

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. .

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

Additive:
If the social welfare function can be written as a sum of the individual utility functions U_i i.e.:
$W(h_1, h_2, \ldots , h_n) = U_1(h_1) + U_2(h_2) + \ldots + U_n(h_n)$
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:
$W(h_1, h_2, \ldots , h_n) = U(h_1) + U(h_2) + \ldots + U(h_n)$
Where each individual has a common U(.) that increases with health.

Note that with these first four properties social welfare is not a function of inequality or dispersion in the distribution of health. A Pareto improvement in health, where every individual's health either increases or stays the same, is always considered as at least as good regardless of any increase in inequality. The next axiom introduces inequality aversion into the social welfare function.

Concave:

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:

$$U(h_i) = \frac{{h_i}^{1-\varepsilon}}{1-\varepsilon}, \qquad \varepsilon \neq 1$$

$$U(h_i) = \ln h_i, \qquad \varepsilon = 1$$

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

Constant Absolute Inequality Aversion

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

$$U(h_i) = -\frac{1}{\alpha} e^{-\alpha h_i}$$

Where α is the inequality aversion parameter with a higher level of α 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 determine a partial ordering of 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 that using this rule we would never prefer a more equal distribution with lower mean health to a less equal distribution with higher mean health.

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[3]. 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 equality in the distribution of health, we would never rank a distribution with lower mean health higher than one with greater mean health. This does not mean that we always prefer the distribution with higher mean health, as there will be cases where higher mean health combined with a higher level of inequality leaves us unable to rank distributions.

The dominance rules only provide a partial ranking and we will be left with a set of distributions that are unranked in terms of social welfare even though they may differ meaningfully in overall health and in health inequality. In order to obtain a complete ranking we could 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 is distributed unequally then, given that we have an aversion to inequality, more total health would be required to produce the

[3] It has been shown (A. K. Sen, 1973) that rules 3 and 4 can be generalised to relax the additive, individualistic and strictly concave restrictions on the SWF. Instead all that is needed for Atkinson's and Shorrocks' theorems to apply is that the SWF is an increasing, symmetric, Schur-concave function of individual levels of health: $W = W(h_1, h_2, ... , h_n)$. For a definition of Schur-concavity see: (Marshall & Olkin, 1974) for its application to inequality measurement see (Dasgupta & Sen, 1973) and (Shorrocks, 1983)

same social welfare than if health were distributed equally. 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 (EDE) health) identically to the ranking suggested by the dominance rules.

Atkinson Inequality Index – constant relative inequality aversion:

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, ε, 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 "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.

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:

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

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

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

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

where the parameter ε, which can take any value from 0 to infinity, specifies the level of societal inequality aversion. The higher the ε, 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 ε. 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 zero and one, with zero representing a perfectly equal distribution.

Kolm Index- constant absolute inequality aversion:

Like the Atkinson index the Kolm "leftist" index is also based in the SWF framework and 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 (Kolm, 1976a) (Kolm, 1976b). 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.

Social Welfare Functions used in the health inequalities literature:

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:

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

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).

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. Individual responsibility attempts to deal with equity relevant characteristics in an unadjusted health distribution to correct for fair differences in health and then any remaining unfair inequality is dealt with by the inequality aversion parameter.

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

The Sen family of poverty indices:

The Sen family of poverty indices focus measures of social welfare on that part of the distribution below a poverty line. An appropriate inequality measure is selected to calculate the equally distributed equivalent health among the health poor $h_{ede(p)}$ in the general formula:

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

where H represents the proportion of the population below the health poverty line and z represents the health poverty line. The standard Sen poverty index can 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 μ_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 properties of the index.

Having understood methods for cost-effectiveness analysis and measures of inequalities, and explored a variety of health related social welfare functions that would allow us to combine these methods to make decisions about which health distribution we prefer, we next turn to a stylised example to look at how we would conduct this analysis in practise.

Stylised Example

In this example we use purely stylised data to explore a baseline health distribution and a corresponding post intervention health distribution. The mean quality adjusted life expectancy (QALE) in the baseline health distribution (which can be thought of as a counterfactual to the intervention) is 40 years whereas if the intervention is carried out this mean QALE increases to 50 years. However, the distribution of health between the baseline and intervention distributions differs. We will use the tools described in this paper to explore these two distributions in more detail to see what conclusions can be derived for a decision maker considering whether or not to carry out the intervention. The netting out of any opportunity cost has already been carried out, the distributions discussed here are distributions showing only variation in net QALE. For the purposes of this example we deem all variation in health to be unfair, we discuss how to adjust the health distribution appropriately for variation in health due to variation in fair and unfair determinants in the next section.

The individual health parade, shown below, plots the health of each individual in the population in order of baseline QALE starting with the least healthy on the extreme left of the plot. Baseline health for each individual is shown in red and the corresponding health that each individual would have should the intervention be carried out is shown in blue. We can see in the diagram that mean health rises due to the intervention raising the health of the most and least healthy groups in the population. The individual health benefits diagram, plotting the difference between health under the intervention and baseline health, shows that while the health under the intervention improves for 67% of the population it decreases for the middle third of the population.

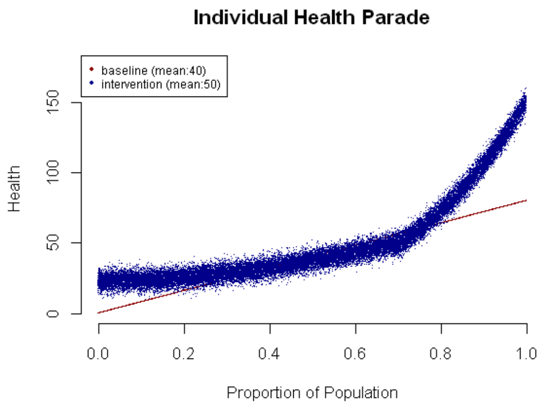

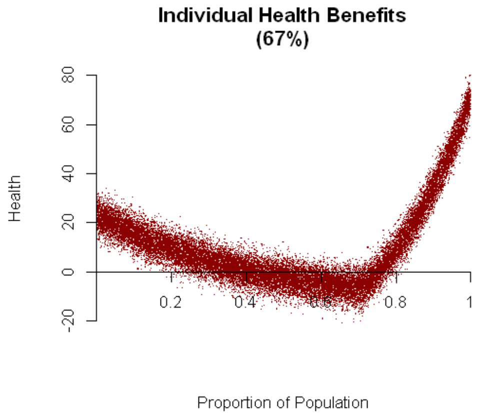

This relationship also holds at the aggregate level when we compare the average health per quintile groups for the baseline and intervention distributions as shown in the plot below.

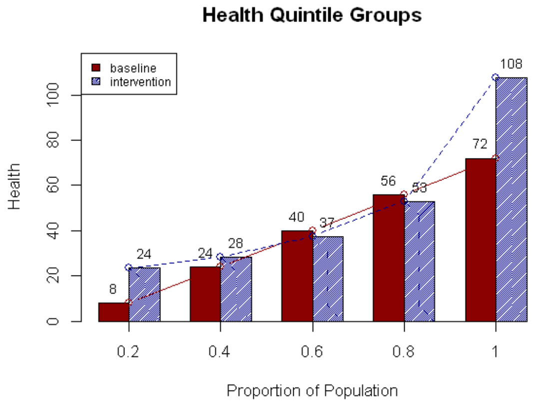

While it is clear through this analysis that mean health increases due to the intervention it is not clear whether or not we would deem the distribution of health to be more or less equal. To start to examine this we turn to the Lorenz curves for the distributions. Lorenz curves plot the cumulative proportion of health against the cumulative proportion of the population, with the population ordered by health from least healthy to most healthy. The closer the Lorenz curve for a distribution lies to the 45 degree line the closer the distribution is to being perfectly equal. If the Lorenz curve for one distribution lies everywhere inside (closer to the line of perfect equality) the Lorenz curve for a second distribution we can conclude the first distribution is more equal than the second in relative terms. If a more relatively equal distribution in this sense has at least as high a mean as the distribution it is being compared to we can conclude, using Atkinson's theorem, that this more equal distribution is to be preferred in terms of social welfare.

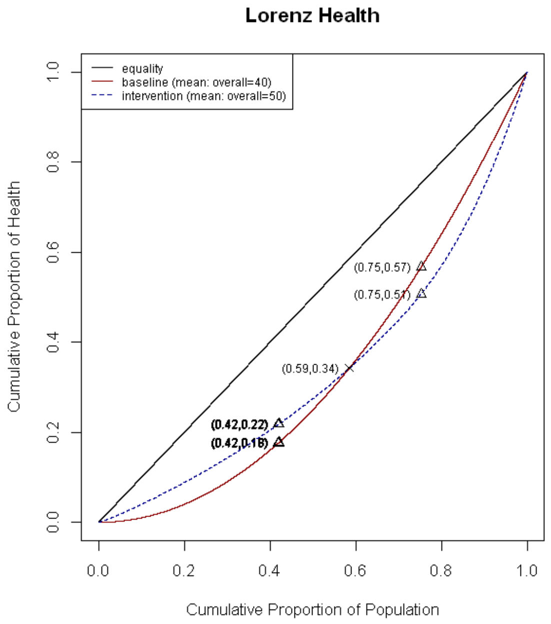

We see from the Lorenz curves for this example, shown above, that due to the fact that the Lorenz curves cross it is not straight forward to decide which of the distributions is more equal in relative terms. Atkinson's theorem does not apply in cases where Lorenz curves cross. We also see, as indicated by the two pairs of triangles on these curves, that the re-ranked absolute health of the populations also cross twice. The health of equivalently ranked individuals is better under the intervention below the 42nd percentile and above the 76th percentile and worse in between.

To explore the level of inequality in the distribution further, we can turn to commonly used indices of relative and absolute inequality. From these we see that where our concern is with relative inequality (see Table 1) the indices suggest that the intervention should be judged to be at least as equal if not more equal than the baseline health distribution. While if our concern is with absolute inequality (see Table 2) then our indices in general suggest that the intervention should be judged to be more unequal than the baseline health distribution, with exceptions at the higher levels of inequality aversion on the Kolm index. The disagreement between these sets of indices, driven by our different views of inequality, indicates that it is far from obvious which of the distributions is unambiguously more equal.

Table 1: Common 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

Index	baseline	intervention	difference (intervention - baseline)
Atkinson Index (ε = 0.5)	0.11	0.08	-0.03
Atkinson Index (ε = 1.0)	0.26	0.16	-0.11
Atkinson Index (ε = 1.5)	0.49	0.22	-0.27
Atkinson Index (ε = 2.0)	0.77	0.27	-0.51

Table 2: Common 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 (α = 0.025)	6.46	9.10	2.65
Kolm Index (α = 0.050)	11.90	13.49	1.59
Kolm Index (α = 0.075)	16.07	16.07	-0.01
Kolm Index (α = 0.100)	19.20	17.84	-1.36

The analysis so far has been looking at inequality across the entire distributions, if instead we were only interested in looking at the difference in the distributions for individuals below some minimum level of health, we can turn to our poverty indices. For this example we set our health poverty line to a QALE of 40 years, the baseline mean level of health, to give the results in Table 3 below. We can see that our incidence measure suggests that more of the population fall under the health poverty line in the intervention distribution than in the baseline health distribution while our intensity and inequality measures suggest that these people who fall under the health poverty line fare better in terms of intensity of poverty and inequality between the health poor under the intervention than under the baseline. Again, the disagreement among the various measures leaves us unable to prefer unambiguously one distribution of health over the other.

Table 3: Poverty Indices

Index (z= 40)	baseline	intervention	difference (intervention - baseline)
Head Count Ratio (incidence)	0.50	0.54	0.04
Health Gap Ratio (intensity)	0.50	0.29	- 0.21

Index (z= 40)	baseline	intervention	difference (intervention - baseline)
Censored Gini (inequality)	0.33	0.12	- 0.21
Sen Poverty Index	0.33	0.21	-0.13

While we do have a ranking of the distributions in terms of mean health we do not have an unambiguous ranking of the distributions in terms of inequality. We next proceed to combine these two concepts explicitly to see if we can conclude anything about social welfare. The most general way to start this analysis is by examining the Generalised Lorenz curves. Generalised Lorenz curves are constructed by multiplying out Lorenz curves for distributions by their respective means. If the Generalised Lorenz curve for one distribution lies everywhere above that for another distribution, then under very general conditions we can say that we prefer the first distribution to the second in terms of social welfare. This is a consequence of Shorrocks' Theorem described as rule 4 above. The Generalised Lorenz curves for this example are shown below:

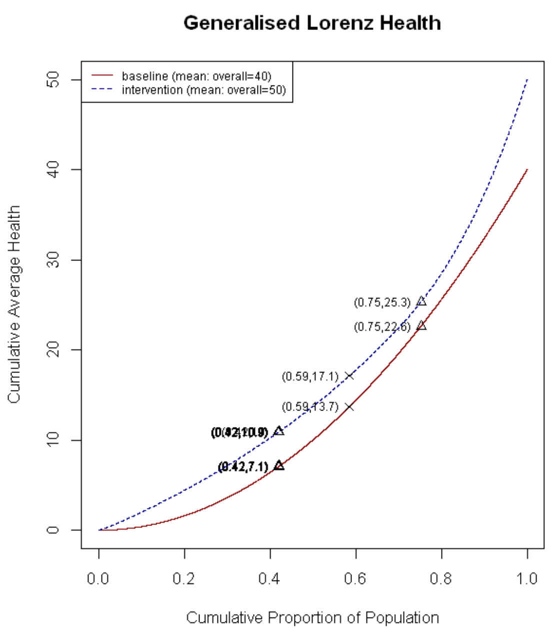

We can see in the plot above that the Generalised Lorenz curve for the intervention does indeed dominate that of the baseline, hence using Shorrocks' theorem we conclude that the intervention is preferred to the baseline. The substantial increase in mean health under the intervention allows us to prefer the intervention distribution in terms of social welfare despite the inconclusive ranking of the distributions in terms of inequality measures. The crosses on the generalised Lorenz curves show the point that the Lorenz curves cross, the triangles on the generalised Lorenz curves show the points where the re-ranked absolute health achieved under the policies cross. The points between the first pair of triangles and above the second pair of triangles are where equivalently ranked individuals under the intervention have less health than in the baseline health distribution. The points below the first triangles show where the intervention is both more equal and equivalently ranked people are better off under the intervention. Given that our dominance rule applies we need make no further assumptions and might conclude that the intervention should be reimbursed. However, were this not the case, had the Generalised Lorenz curves crossed for example, we would need to proceed to look at our social welfare measures (see Table 4) of equally distributed equivalent levels of health for the distributions compared at suitable levels of inequality aversion.

Table 4: Social Welfare Measures – Equally Distributed Equivalent Levels of Health

Index	baseline	intervention	difference (intervention - baseline)
Mean Health (ε = 0 or α = 0)	40.00	50.00	10.00
Atkinson EDE (ε = 0.5)	35.56	45.76	10.21
Atkinson EDE (ε = 1.0)	29.44	42.14	12.70
Atkinson EDE (ε = 1.5)	20.39	39.15	18.75
Atkinson EDE (ε = 2.0)	9.02	36.71	27.69
Kolm EDE (α = 0.025)	33.54	40.90	7.35
Kolm EDE (α = 0.050)	28.10	36.51	8.41
Kolm EDE (α = 0.075)	23.93	33.93	10.01
Kolm EDE (α = 0.100)	20.80	32.16	11.36
Sen Poverty Index EDE (z = 40)	13.33	24.94	11.61

In this case, given that our dominance rules do apply, we note that all our social welfare measures prefer the intervention to the baseline health distribution at all levels of inequality aversion[5]. In cases where dominance rules do not apply, the choice of inequality aversion parameters can become a critical value judgement in the decision making process.

Comparison with the concentration curve for health inequalities

The standard way to analyse health inequalities in the literature is by way of bivariate methods such as concentration curves. Concentration curves are analogous to Lorenz curves with individuals ordered according to some other socio-economic variable apart from health, most commonly income. The concentration curve allows one to see how the distribution of health varies with a variable that is considered to be an unfair determinant of health. Where ranking by income and ranking by health coincide the Lorenz curve and the concentration curve for a distribution will be identical. The

[5] Social welfare measures must by definition rank distributions consistently with the dominance rules where these rules apply.

diagram below shows the Lorenz curves and concentration curves for our stylised example. We can clearly see that these differ substantially, allowing us to conclude that inequalities in the health distribution are not completely explained by inequalities in the income distribution.

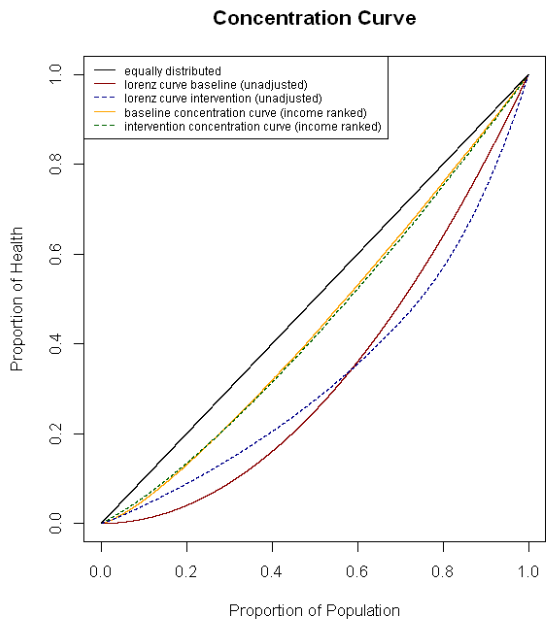

Our adjustment process, that we outlined at the start of the paper, carried out through the estimation of a reduced form equation for health, allows us to understand more fully just how much of the inequality in health is due to independent contributions of multiple different determinants of health, and allows us to standardise for those that we consider fair and in so doing isolates only difference in health that we consider unfair. In the stylised example we estimated the following reduced form equation for health:

$$h_i = \alpha + \beta_1 * \log(income_i) + \beta_2 * \log(income_i)^2 + \beta_3 * ethnicity + \beta_4 * \log(income_i) * ethnicity + \beta_5 * \log(income_i)^2 * ethnicity + \varepsilon_i$$

and used this to generate an adjusted health distribution where differences in health due to either differences in income or differences in ethnicity are considered unfair. The Lorenz curves for this adjusted distribution are plotted against the unadjusted concentration curve ranked by income in the diagram below. We see in this diagram that we consider more of the inequality unfair in our adjusted distribution, as depicted in the Lorenz curves for the adjusted distribution, than is suggested by the concentration curves. This is unsurprising as in addition to income determined differences in health we also consider ethnicity determined differences in health to be unfair. Our framework allows us to explore multiple unfair determinant of health simultaneously. Additionally our reduced form model allows us to identify more accurately the effects that any particular determinant has on health, isolating the independent impact of each determinant from those of other correlated determinants. Furthermore once we have generated our adjusted health distributions our framework allows us to analyse these distributions using a wide array of measures encompassing relative, absolute and poverty concepts of inequality and social welfare.

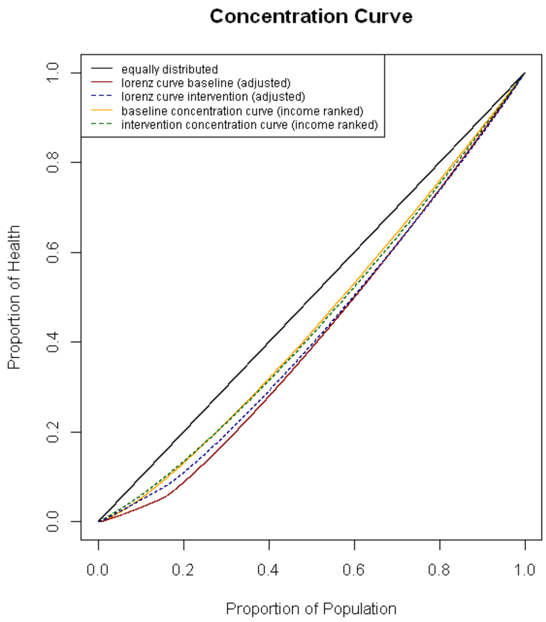

It should be obvious that should we wish to we could, by estimating a suitable reduced form equation, exactly replicate the concentration curve based analysis using our framework.

Conclusion

The income inequality literature provides a rich set of tools for incorporating distributional concerns into economic analysis. In this paper we have explored ways in which these methods can be applied to assess the cost-effectiveness and health inequality impact of public health interventions. We have demonstrated through the use of a stylised example how we can use these graphical and numerical tools to gain a better understanding of health distributions, and how this understanding can help us decide how to rank interventions being evaluated.

Our worked example has 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. Our example has 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 measures, parameterised with the societal level of 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. The social welfare function may be unidentifiable, or the nature of the inequality aversion may be complex taking into account both relative and absolute inequality. In such cases the analysis may best inform decision makers by outlining the key assumptions underlying the different measures used and assessing a range of measures. Little work has been done to establish empirical values for inequality aversion parameters or the trade off between improvements in overall health and reduction in health inequality, and so it may be advisable to look at a range of functions across a range of inequality aversion parameters to help inform the deliberation process.

We have seen that traditional health inequality analysis while useful in exploring the social determinants of health can be limiting when more sophisticated analysis is required. Our methods drawing from the wider literature on economic inequality allow us to get past some of these limitations and more fully describe health distributions in a way that is directly informative in the context of public health priority setting.

In order to use these methods we have seen that we need to extend cost-effectiveness analysis to calculate both the distribution of health benefits due to particular interventions and the distribution of health opportunity costs due to activities displaced by the interventions. We have also seen how the resulting distribution of health, as produced by this extended cost-effectiveness analysis, can be adjusted to incorporate different views on what constitute fair and unfair determinants of health. Further empirical work is required to identify the distribution of opportunity costs of displaced activities more fully in specific contexts.

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.