Handbook of Distributional Cost-Effectiveness Analysis

Chapter 13 Level-dependent equity weights

Fourth draft, October 15, 2019

Ole F. Norheim, Miqdad Asaria, Kjell-Arne Johansson, Trygve Ottersen, Aki Tsuchiya

Word count: 5670

Contents

13.1 Level-dependent SWFs explained 5
13.2 Measures of inequality in health derived from SWFs 14
13.3 Aversion to unfair inequality in health 15
13.4 Equity-efficiency trade-off analysis in practice: sensitivity analysis 17
13.5 Comparison with rank-dependent SWFs 25
13.6 Conclusion 26
13.7 Further Reading 28
References 29

Online only abstract

This chapter explains how you can use level-dependent social welfare functions (SWFs) to evaluate health distributions in a manner that is founded on explicit, challengeable and consistent ethical principles. A level-dependent SWF weights health gains for one person or group relative to another as a function of their absolute health *levels*. By contrast, a rank-dependent SWF weights health gains for one person or group relative to another as a function of their health *ranks* in the wider population. A level-dependent SWF is also called a prioritarian SWF. The basic prioritarian principle is a specific value judgement giving priority to the worse off based on their absolute level of health. The aim of this chapter is to describe level-dependent SWFs and how they can be used to conduct equity-efficiency trade-off analysis. A SWF can be consistent with many different sets of weights, which vary according to the chosen value of an equity parameter. We show how level-dependent equity parameters can be used to perform a systematic sensitivity analysis of how different degrees of concern for the worse off have different policy implications.

Key words

Atkinson, Equally Distributed Equivalent, Kolm-Pollak, Prioritarian, Social welfare function

The previous chapter explained how you can use rank-dependent social welfare functions (SWFs) to evaluate health distributions in a manner that is founded on explicit, challengeable and consistent ethical principles. This chapter explains how you can do this using level-dependent SWFs. A level-dependent SWF uses information about the absolute level of health – for example, a health-adjusted life expectancy (HALE) at birth of 62 – and weights health gains for one person or group relative to another as a function of their respective health levels. By contrast, a rank-dependent SWF weights health gains as a function of the ranking within the full population distribution of health – for example, a HALE of 62 might be the 2nd lowest ranked subgroup in order of health, depending on the health levels of all other subgroups in the population. A level-dependent SWF is also called a prioritarian SWF (Adler, 2012, 2019). A level-dependent SWF gives priority to the worse off based on their absolute level of health compared with others, whereas a rank-dependent SWF gives priority to the worse off based on their relative position in the health distribution compared with others.

The aim of this chapter is to describe level-dependent SWFs and how they can be used to conduct equity-efficiency trade-off analysis. We show how this is done with a numerical inequality aversion parameter within a level-dependent health-related SWF – a level-dependent equity parameter, for short.[1] The basic prioritarian principle is a specific value judgement about giving priority to those who are worse off based on their absolute level of health. This can be consistent with many different sets of weights, which vary according to the chosen equity parameter. We show how level-dependent equity parameters can be used to perform a systematic sensitivity analysis of how different degrees of concern for the worse off have different policy implications.

[1] In the income inequality literature this is usually called an "inequality aversion" parameter. We sometimes use the shorter phrase "equity parameter" because in DCEA the focus is usually on "unfair" inequality associated with equity-relevant characteristics, rather than "pure" inequality between individuals.

The chapter is organised as follows. Section 13.1 introduces the rationale for level-dependent SWFs and explains some differences between classical utilitarian and prioritarian SWFs. Social value judgments over distributions are important characteristics of social welfare and a SWF is therefore inherently a normative evaluation tool (as is the utilitarian SWF). The section also explains the role of equity weights in the analysis. The basic intuition throughout is that there will be some set of weights assignable to people that makes it worthwhile to accept a certain loss of average population health in order to achieve greater equality in health. In section 13.2 we explain how measures of health inequality or inequity can be derived from level-dependent SWFs. Section 13.3 discusses social value judgements about health distributions and how these can be captured in an equity parameter and further used to explore equity-efficiency trade-offs between maximizing health and reducing unfair health inequalities. The basic idea is to use information about a social decision maker's aversion to unfair inequality and make a level-dependent transformation of individual health that displays decreasing marginal social value. Together with information about the existing distribution of health, this then indirectly implies a specific set of equity weights. In section 13.4 we explain equity-efficiency trade-off analysis step by step: how to choose the appropriate SWF, choice of inequality aversion parameter, data inputs needed, how to estimate results and plot them on the equity impact plane, and finally, how to perform sensitivity analysis with different values of the inequality aversion parameter. Section 13.5 compares level-dependent and rank-dependent SWFs and how to choose between them for a particular evaluation. The last section concludes.

13.1 Level-dependent SWFs explained

In this section we introduce one attractive level-dependent SWF that has been applied to health distributions, Atkinson's SWF, which focuses on total health and relative inequality in health (Anand et al., 2001; Asaria et al., 2016; Asaria et al., 2015; Johansson & Norheim, 2011; Norheim, 2013). Box 13.1 describes an analysis from a UK study that used Atkinson's SWF.

> **Box 13.1 Empirical example** (Asaria et al., 2015)
>
> A methodological case study of the UK bowel cancer screening programme compared two redesign policy options: a) the introduction of an enhanced targeted reminder aimed at increasing screening uptake in deprived and ethnically diverse neighbourhoods, and b) the introduction of a basic universal reminder aimed at increasing screening uptake across the whole population. One key result was that the first option would reduce inequities in health, but with loss of total net health gains. The second option would yield higher health gains, but not significantly reduce inequity in outcomes. By using a social welfare function approach, Asaria et al. were able to analyse the trade-off between improving total health gain and reducing inequities in health. The targeted reminder would achieve higher social welfare than the alternative policy with a sufficiently high concern of the worse off (or aversion to inequality).

To understand the structure of such an analysis, consider a simple example, based on table 11.1 in chapter 11.

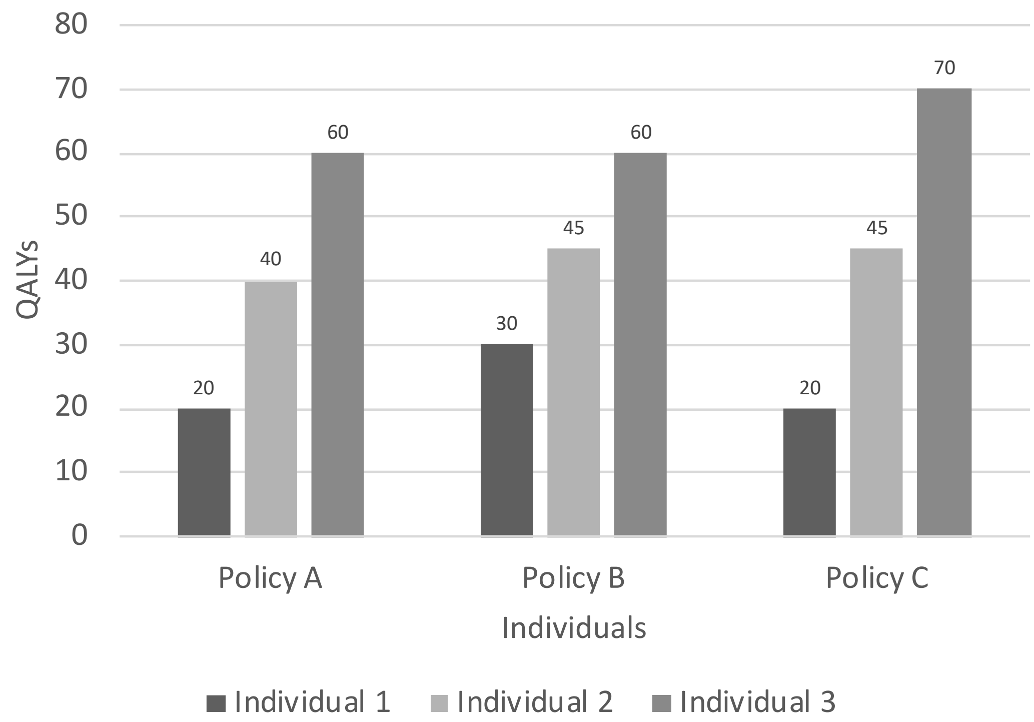

Figure 13.1 Hypothetical distributions of lifetime health generated by two alternative policies (B and C) – compared to A

Figure 13.1 shows the distribution of lifetime health over a population with three individuals under three alternative programmes that cost the same. Programme A is the status quo. With programme B, individual 1, who is the worst off, will gain 10 HALYs, individual 2 will gain 5 HALYs, and individual 3 who is the best off will gain nothing. With programme C, individual 1 will gain nothing, individual 2 will gain 5 HALYs, and individual 3 will gain 10 HALYs. A utilitarian SWF (see below) yields the same net sum of health effects under B and C (15 additional HALYs), and hence they deliver the same gain in social welfare. A utilitarian decision maker will be indifferent between B and C. The baseline inequality in health levels by itself is irrelevant information for a utilitarian decision maker: only the total net health effects compared with the status quo baseline situation under programme A matter for assessing gains in social welfare.[2]

[2] We assume the health gains are measured in a way that reflects a full and final accounting of the individual value of health. In other words, if ***h*** denotes the individual's health profile (or, more precisely, a large vector of numbers representing the probabilities of many different possible multi-dimensional health states occurring in many different periods of time over the course of a whole life) and u(***h***) denotes the value to the individual of

In a level-dependent SWF, more weight is given to health gains for the worse off – in this case individual 1. The total health gain is the same in B and C, so a level-dependent SWF will rank B over C since more weight is given to gains to individual 1. This is straightforward.

Now compare distribution B and E (Figure 13.2). Under programme E, individual 1 will gain 6 HALYs, individual 2 will gain 5 HALYs, and individual 3 will gain 10 HALYs, totalling 21 HALYs. In this case, a utilitarian SWF will rank E over B, since the net sum of health is larger under E compared to B. This is also straightforward.

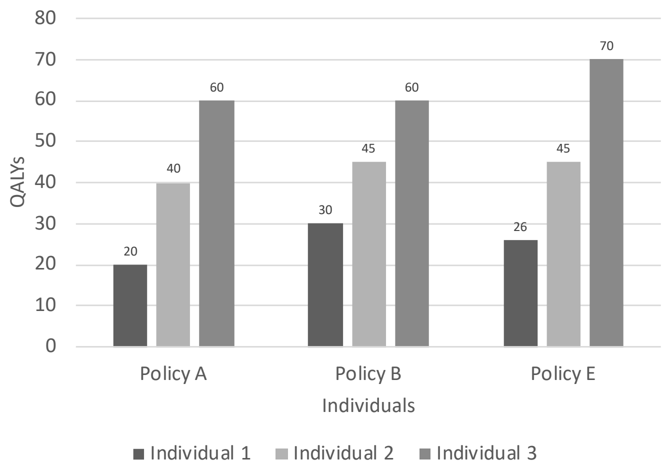

Figure 13.2 Hypothetical distributions of lifetime health generated by alternative policies (B and E) – compared to A

How a level-dependent SWF will rank the two distributions is, however, not so straightforward. It depends on the weights assigned to gains for the worse off. Intuitively, decision makers would rank B over E if their aversion to inequality is sufficiently high: the weight given to gains to individual 1 relative to gains to individual 3 needs to be large enough to cancel out the fact that the health gain to individual 1 is smaller than the health gain to individual 3. If aversion to inequality is low, then the weight given to individual 1 relative to

that health profile, then we assume HALYs represent u(***h***) (= individual value). As explained in chapter 2 on principles of health equity, HALYs can be interpreted in many different ways – e.g. some people interpret them as the individual's own preferences for health whereas others interpret them as a measure of health benefit reflecting general population average value judgements about the individual value of different health states.

individual 3 will not be enough to overcome the difference in health gain, and the decision maker will rank E over B. In what follows, we make explicit this equity-efficiency trade-off and formalise it.

First, we turn to the theoretical basis for this type of analysis and we start with the normative choice of SWFs (Adler, 2019; Anand, 2002; Anand et al., 2001; Bleichrodt & van Doorslaer, 2006). A health maximising utilitarian SWF is the simple sum of individual health gains:

$$W_u = \sum_{i=1}^{N} h_i ,$$

where W_u is social welfare, h_i is the health of individual *I*, and N the number of people in the population. In principle health could be measured on any cardinal scale that can be added up, and could represent current health, future health or lifetime health. To fix ideas however we can think of health as lifetime health measured in terms of health adjusted life expectancy at birth. Implicitly, standard cost-effectiveness analysis uses such a SWF – it simply adds up individual net health gains (i.e. health gains minus health opportunity costs) to yield net population health gain, implicitly assigning a unitary weight to each.

The level-dependent (prioritarian) SWF is a sum of transformed individual health levels representing each individual's contribution to social welfare. Individual health is transformed into the corresponding social welfare contribution using a concave function, meaning that extra weight is given to health gains to people who are worse off in terms of health. In its more general form, prioritarianism is a consequentialist ethical theory that evaluates outcomes by the transformed sum total of individual wellbeing, with more weight given to wellbeing gains for the worse off relative to the better off (Parfit, 1997, 2002). Prioritarianism has a single objective: to maximise the sum total of well-being values adjusted so as to give "priority" (a larger weight) to the worse off. The ethical justification for assigning extra weight is simply that the worse off a person is, the more important it is, morally speaking, to help that person. In this chapter we are concerned with gains and distribution of lifetime health and define the worse off as those with less lifetime health.[3] In practice, we often speak of two objectives – increasing total health and reducing health inequality – which philosophers call a "pluralist egalitarian" view (Otsuka & Voorhoeve, 2018). In cases where there is a distributive conflict between the two objectives, a trade-off must be made. If a choice requires a trade-off between priority to the worse off and total health, a level-dependent SWF will integrate both concerns by using an 'inequality aversion' parameter that makes this trade-off explicit. The SWF can rank options if the decision maker is willing to specify a value, or by sensitivity analysis if they are willing to specify a *range* of reasonable values.

The level-dependent SWF can be written as:

$$W = \sum_{i=1}^{N} g(h_i),$$

where N = total population size, and $g(h_i)$ is strictly increasing and strictly concave transformation function of health and social welfare.

The transformation of health into social welfare is what distinguishes this SWF from the utilitarian SWF. The Atkinson, isoelastic transformation function can be written as:

$$g(h) = \frac{h^{1-\varepsilon}}{1-\varepsilon}$$

[3] Other ways of defining the worse off are discussed in chapter 6.

with $\varepsilon > 0$. ε is an equity parameter indicating the degree of aversion to health inequality. In the special case of $\varepsilon = 1$, $g(\mathrm{h}) = \log \mathrm{h}$. With $\varepsilon = 0$, the SWF becomes utilitarian. Figure 13.3 illustrates how the transformation function works.

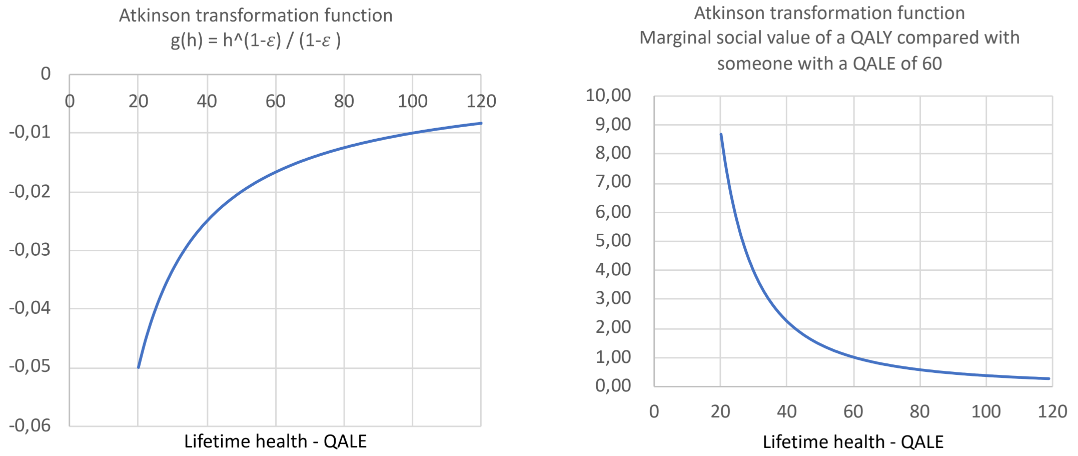

Figure 13.3 Atkinson transformation function and marginal social value of a HALY

Note: The left panel shows the transformation function for a value of ε (=2), the right panel shows the corresponding marginal social value of a HALY compared with a HALY to someone with a lifetime health of 60

The form of the transformation function (left panel in figure 13.3) reflects the relative sensitivity to transfers at different levels of health and is a direct reflection of impartial societal value judgments. Note that the marginal social value of a change in health (HALY) decreases with increasing health (right panel in figure 13.3), and this satisfies the Pigou-Dalton condition ((Bleichrodt & van Doorslaer, 2006), see also chapter 11, this volume).[4]

The concept of 'equally distributed equivalent' (EDE) is key to understanding the Atkinson SWF (Atkinson, 1970). Atkinson introduced EDE as the social welfare of a given income in a population that takes both its total and distribution into account. The EDE is the level of

[4] With Atkinson's transformation function, this works well for lifetime health above 10 to 20 HALYs (which is very often the case). When lifetime health approaches zero, the marginal value approaches infinity. This is not realistic and may restrict the use of this SWF for low levels of lifetime health. One example would be comparisons involving diseases associated with under-five mortality in low-income countries.

income that if equally distributed would generate the same level of social welfare as that generated by the actual distribution of income. For a definition of EDE in terms of health (EDEH), see box 13.2. The value of ε is crucial. As ε approaches infinity, the SWF approaches the 'Rawlsian SWF' that assigns absolute priority to the worst off (as illustrated in figure 13.4).

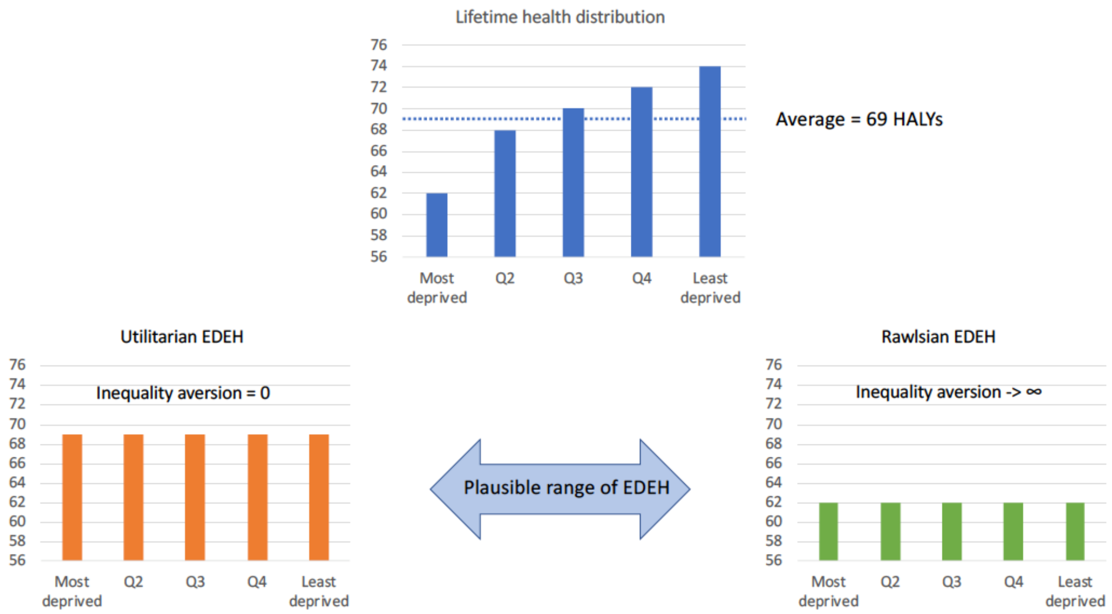

Figure 13.4 EDEH explained

Note: The left panel shows inequality aversion = 0 and the right panel shows inequality aversion → infinity.

In this figure, the upper panel shows one actual distribution of lifetime health for five quintiles. The average is 69 HALYs. Imagine a utilitarian SWF with $\varepsilon = 0$. This would mean that the upper panel and the left-hand panel represent distributions that achieve the same level of social welfare. Alternatively, imagine a Rawlsian SWF, with ε approaching infinity. This would mean that the upper panel and the right-hand panel represent distributions that achieve the same level of social welfare. In other words, a given distribution (in the upper panel) can be equally good as a uniform distribution at 69 or another uniform distribution at 62, depending on the equity parameter.

EDEH is ordinally equivalent to the Atkinson index of social welfare – i.e. it ranks HALY distributions in the same order. It can be written as follows:

$$EDEH = \left(\frac{1}{N}\sum h_i^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}}.$$

EDEH expresses the level of social welfare represented by an Atkinson SWF on the same scale as individual health.

There are several variants of the level-dependent social welfare function. We have so far discussed Atkinson's SWF applied to distributions of health. Let us now briefly compare it to the Kolm-Pollak SWF (Asaria et al., 2016; Kolm, 1976). The Kolm-Pollak SWF can be written as follows:

$$\sum_{i=1}^{N} g(h_i) = -\sum_{i=1}^{N} e^{-\alpha h_i},$$

where α is the aversion to inequality parameter. The ordinally equivalent Kolm-Pollak EDEH form is written as:

$$EDEH = -\left(\frac{1}{\alpha}\right)\log\left(\frac{1}{n}\sum_{i=1}^{N} e^{-\alpha h_i}\right),$$

The most important difference between the two is that the Atkinson EDEH is concerned with relative inequality and the Kolm-Pollak EDHE is concerned with absolute inequality in lifetime health. Compare distributions in Programme X and Y in figure 13.5.

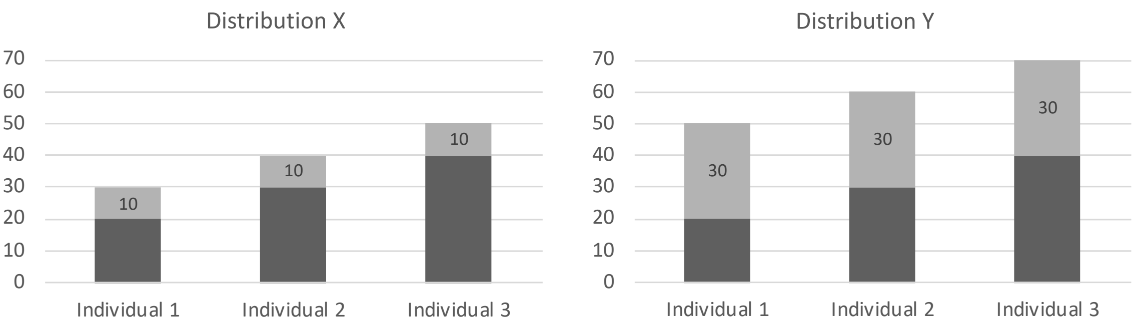

Figure 13.5 Illustration of two distributions with the same absolute inequality, but different relative inequality

The baseline levels (shaded dark grey) are the same in X and Y. In distribution X, each individual gets another 10 HALYs. In distribution Y, each individual gets another 30 HALYs. Absolute differences between the resulting levels of health (10 HALYs) are the same between individuals in X and Y, while the relative differences are not (in distribution X, individuals 2 and 3 have 33% and 67% more health than individual 1; in distribution Y, the gaps are 20% and 40% respectively). For a given α, the Kolm-Pollak index of inequality will be the same in the two distributions, while Atkinson's' index of inequality would be different. Another way to put it is that the Atkinson index satisfies the axiom of scale invariance: any equal proportional change in each individual's level of health should not change the measure of inequality. In the measurement of income inequality, scale invariance is often seen as a desirable property. Others, including Atkinson, have argued that scale invariance may not be the right property of a health inequality index (Atkinson, 2013). Since age at death has an upper limit of, say, about 120 years, decision makers may be more concerned about absolute inequalities than relative. If the analyst is concerned with absolute inequality, the Kolm-Pollak index may be appropriate. However, the Kolm-Pollak index is rarely used and most analysts seem to be most concerned with relative inequalities (for discussions about relative and absolute inequality in health see (Gakidou et al., 2000; Harper et al., 2010)). Whether decision makers in general think absolute or relative inequality is most important is unfortunately not known. In the absence of any reason for presuming one thing rather than another, some attempt to elicit actual decision makers' views in any specific application seems desirable.

The Kolm-Pollak SWF transformation function can be written as follows:

$$g(h_i) = -e^{-\propto h_i}.$$

The form of the function is illustrated in figure 13.6.

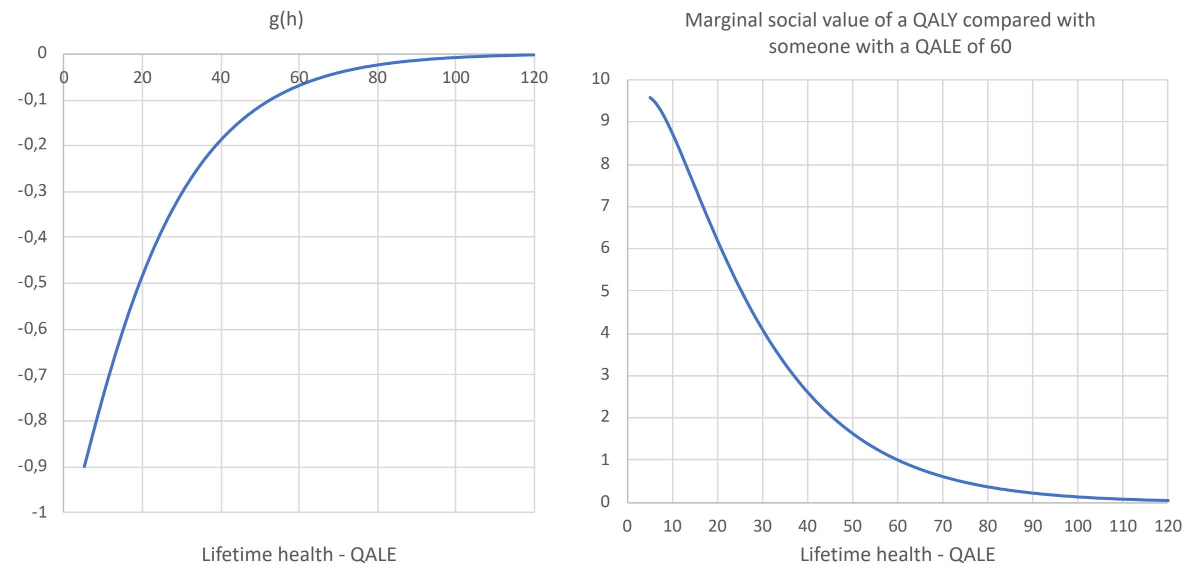

Figure 13.6 Kolm-Pollak transformation function and marginal value of a HALY

Note: The left panel shows the function for a given value of α (=0.025); the right panel shows the marginal social value of a HALY difference compared with someone with a HALE of 60.

We see from the right-hand panel in the figure that the marginal social value of a HALY difference is decreasing as health increases, and – unlike the Atkinson function – the marginal value does not approach infinity as health approaches zero. Hence, the Kolm-Pollack function may be useful in analyses involving diseases associated with child mortality.

13.2 Measures of inequality in health derived from SWFs

A level-dependent SWF can be represented as the combination of aggregate health and an inequality metric applied to the distribution of health (Sen, 1997). An abbreviated form of EDEH can then be written as:

$$\text{EDEH} = \mu * (1 - \text{Inequality}) ,$$

where μ is mean health in the population and 'Inequality' is a measure of inequality or inequity in health. For absolute measures, such as Kolm-Pollak, the abbreviated form is: EDEH = μ - $K(\alpha)$. It is easy to see that if health is distributed equally, EDEH equals mean health in the population. For an equal distribution, EDEH = μ holds by definition, regardless of the aversion for inequality. If there is inequality (measured on a scale from 0-1, where zero is no inequality), the level of welfare is discounted according to the level of inequality and the disvalue assigned to this level of inequality. A distribution with high mean health and high health inequality thus can have lower social welfare than a distribution with lower mean health and lower health inequality if inequality aversion is sufficiently strong.

Atkinson index of inequality A(ε), where 0 represents no inequality and 1 represents full inequality, can be written as follows:[5]

$$A(\varepsilon) = 1 - \frac{EDEH}{\mu}$$

The abbreviated form allows a convenient split between equity and health impact and these two objectives can be represented on the equity impact plane. The split also helps us perform trade-off analysis between health maximization and equity.

13.3 Aversion to unfair inequality in health

Before moving to the trade-off analysis, more needs to be said about the inequality aversion parameter. The level-dependent SWF forces the analyst to be explicit about social value

[5] For absolute inequality measures, such as Kolm-Pollak, we get $K(\alpha) = 1 - \frac{EDEH}{\mu}$.

judgments: the weights to be assigned to transfers from the better off to the worse off or vice versa. As we have seen, when $\varepsilon = 0$, the SWF becomes 'utilitarian'. As ε approaches infinity, the SWF approaches the 'Rawlsian SWF' that would assign all weight to the worst-off. When Atkinson's SWF is used for evaluation of income distributions, analysts typically use a range of ε from 0 to 4 (Anand et al., 2001). Social preferences over health distributions are far less explored though there is a growing literature where aversion to inequality in health has been empirically explored (Ali et al., 2017; Baker et al., 2010; Bansback et al., 2012; Cookson et al., 2018; Dolan, 2003; Dolan et al., 2008; Dolan & Tsuchiya, 2009; Edlin et al., 2012; Lancsar et al., 2011; Lindholm & Rosen, 1998; Norman et al., 2013; Petrou et al., 2013; Robson et al., 2015; Rowen et al., 2016a; Rowen et al., 2016b; Shah et al., 2014, 2015; Tsuchiya & Dolan, 2005, 2007). Most of these studies ask members of the public about aversion to unfair health inequalities associated with socioeconomic status. Few studies ask decision makers or compare their preferences to the general public (Lindholm & Rosen, 1998; Tsuchiya & Dolan, 2007). Only a few studies ask about aversion to "pure" inequality in health (Dolan et al., 2008; Dolan & Tsuchiya, 2012; Ottersen et al., 2014). As a general pattern, the degree of inequality aversion identified in these studies tend to be weaker (albeit still inequality averse) compared to those found in studies that focus on inequalities in health across socioeconomic groups (Costa-Font & Cowell, 2019).[6]

Results vary and this may be due to different survey methodologies, countries and settings. Some are also based on convenience samples and cannot be seen as representative. Yet, these preliminary findings may indicate that aversion to health inequalities are stronger than for income inequalities. Aversion to health inequalities associated with socioeconomic status may also be stronger than aversion to pure health inequalities. The range for Atkinson's ε varies

[6] More work is also needed to distinguish clearly between studies measuring aversion to inequality from a perspective of a) individual self-interest, b) veil of ignorance, and c) impartiality.

from 1.4 to more than 10 in this sample of surveys, indicating that the range of values for the inequality aversion parameter is wide. More empirical studies with standardized methods are clearly warranted (Costa-Font & Cowell, 2019).

If the analyst has credible evidence on a social decision maker's aversion to inequality in the form of a specific value or range of ε, the level-dependent SWF can be directly applied to the analysis of health distributions. If not, an alternative is to do sensitivity analysis over a wider range of values. This is the topic we explore in the next section.

13.4 Equity-efficiency trade-off analysis in practice: sensitivity analysis

We now have enough background and tools to explore how distributions of health can be ordered according to a level-dependent social welfare function. First, we show the result from such an analysis for the simple 3-person example illustrated in figure 13.2. Next, we outline a stepwise procedure for a trade-off analysis with data from the nicotine replacement therapy (NRT) exercise.

Returning to figure 13.2, recall that we compared two policies (B and E) where one had higher impact on health than the other, while the latter reduced health inequality more than the former. The question is: for what values of ε would a level-dependent SWF rank E over B and for what values would it rank B over E? By using EDEH as a measure of social welfare, we can plot social welfare from different distributions as a function of the inequality aversion parameter ε (figure 13.8).

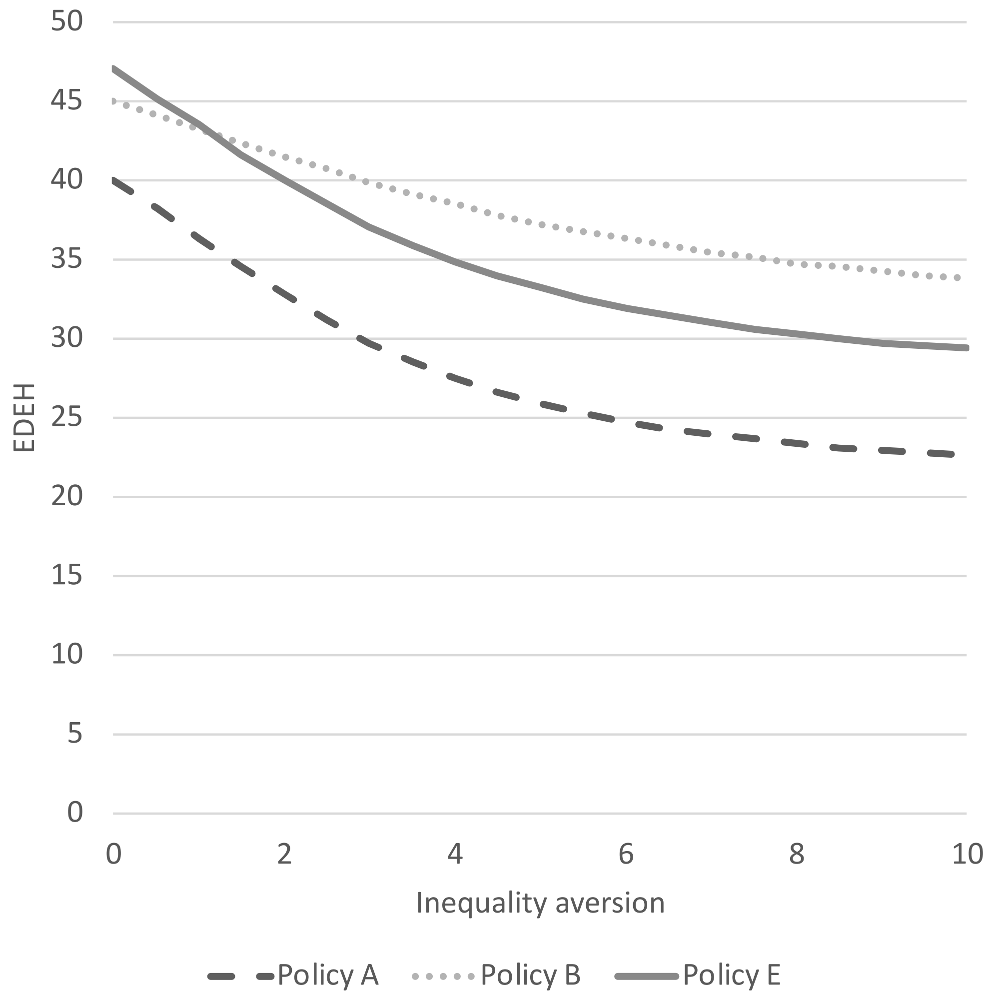

Figure 13.8 Social welfare from different health distributions (A, B, and E) as a function of aversion to inequality (range 0 to 10)

We see from figure 13.8 that both B and E is ranked above the baseline (Programme A), while E is ranked over B when ε is low, below approximately 1.3, and B is ranked over E when ε is higher, for all values above 1.3.

One implication to draw from this example is that the conclusion is quite clear if we believe the social decision maker has a moderate to high aversion to inequality. Programme B is ranked above E in terms of EDEH when there is high aversion to inequality. A more definitive conclusion is not possible without better information on a reasonable range of ε for the social decision maker.

Next, let us turn to the NRT exercise. Recall that we compare two programmes, universal NRT and proportional NRT, to no publicly financed NRT (status quo). Net health impact from the two policies on 10 groups is illustrated in figure 13.9. Again, the task is how to rank the two policies. Universal NRT has higher health impact than Proportional Universal NRT, while the latter has higher equity impact.

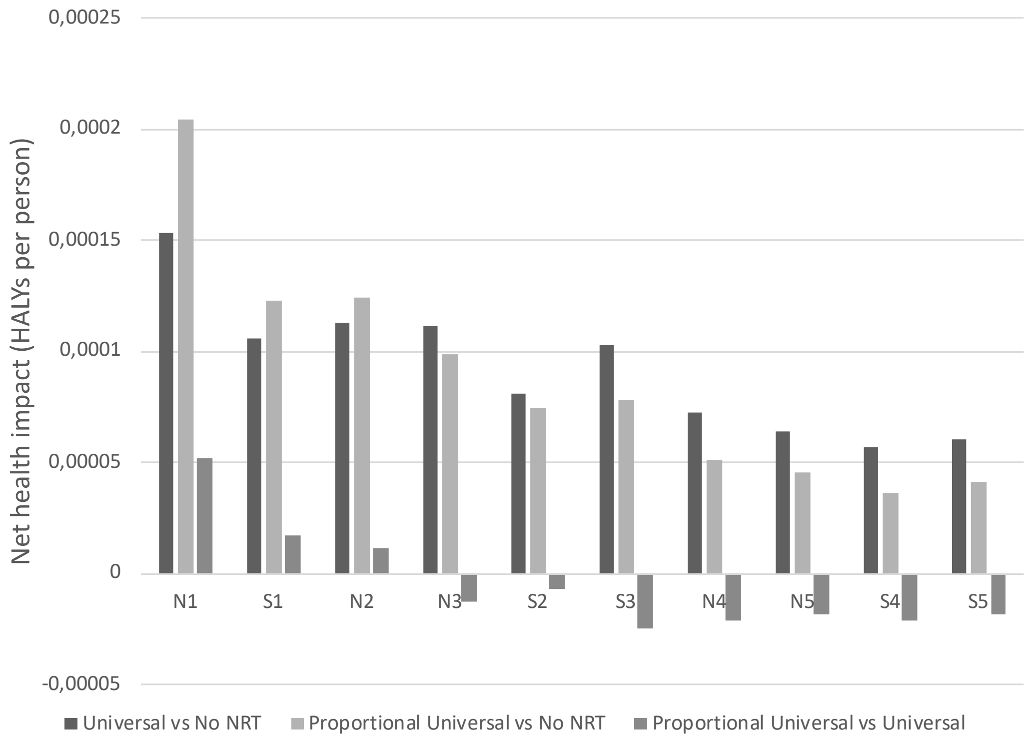

Figure 13.9 Net health impact of Universal NRT and Proportional Universal NRT versus no NRT for 10 social groups, ranked by health

This example uses reported data and results for groups defined by geography and socioeconomic status. To do our analysis with the data at hand, we must make three simplifying assumptions:

1. Our analytical focus in this textbook is on health-related wellbeing.
2. If we have the data, it is always preferable to estimate health-related wellbeing from individual data. In this case we do not have individual data, only data grouped by socioeconomic status and area of living. We also know the number of people in each

group. We must assume, and this is a strong assumption, that all individuals in each group have the same baseline health and the same health increment. We therefore disregard within-group inequalities.

3. Finally, level-dependent SWFs are by definition concerned with pure health distributions (univariate analysis), regardless of cause or association (bivariate or multivariate analysis). With the data at hand, we have only grouped data based on non-health characteristics. No other inequality is observed or modelled. To overcome this limitation, we identify the worse off by ranking groups from lowest to highest lifetime health, disregarding within-group inequalities. In this simplified univariate analysis we necessarily underestimate overall actual inequalities in lifetime health.

We now explain a stepwise procedure for trade-off analysis using a level-dependent SWF. The task is to order policies according to their impact on total health and health distribution.

Step 1: Choose data input

One advantage of level-dependent SWFs is that they do not require more information than is typically available in a standard economic evaluation: baseline levels of health and incremental health. Information on levels of health are often available since it is needed to calculate incremental health gains in most models. If the lifetime health perspective is chosen, though, some additional modelling of the baseline may be needed. Table 13.2 illustrates typical data needs for evaluating distributions, based on the NRT exercise.

Table 13.2 Example of data needs: population size and HALE under the different programme options (based on the NRT exercise)

Subgroups	Population size	No NRT (HALE)	Universal (HALE)	Proportional Universal (HALE)
1	7279927	61.925214	61.925367	61.925418
2	3959316	65.953455	65.953561	65.953577
3	5050982	66.816952	66.817065	66.817076
4	4616011	68.975036	68.975148	68.975135
5	6331048	70.344928	70.345009	70.345002
6	6474305	72.213544	72.213647	72.213622
7	4701015	72.648215	72.648287	72.648266
8	4159448	74.561156	74.561220	74.561201
9	6194904	75.597859	75.597916	75.597895
10	6501111	77.323144	77.323204	77.323185
Average HALE		70.626197	70.626290	70.626287

In this example, we are interested in lifetime health and health expectations. One single intervention will typically not substantially change life expectation in large groups. The average health impact in each group is therefore small, while actual benefits for some individuals may be large. The direction of change is therefore key.

Step 2: Choose Social Welfare Function

The main choices are Atkinson's SWF as an attractive example of level-dependent SWFs the Kolm-Pollak SWF if absolute inequality is of key concern. If the analyst is *only* interested socioeconomic health inequality, bivariate or multivariate equity metrics should be chosen (as discussed in chapters 11 and 12).

Step 3: Choose inequality aversion parameter (ε) or range of ε

The sensible procedure here, in the absence of information about the kind of inequality aversion actually felt by decision makers will be to choose a reasonable range of values for ε.

We suggest as a rule of thumb to start with a range between 0 and 20 and reduce the range if high values do not change the ranking. Yet another, possibly more politically demanding procedure would be to conduct in-project experiments designed to elicit the relevant features of decision makers' values.

Step 4: Estimate net health benefit and equity impact

Calculating net health benefits is straightforward and follows standard methods, where incremental health gain equals total health gains for the choice option in question minus total health from the alternative programme or the baseline (or status quo).

To calculate equity impact, choose the appropriate measure of equality. If the Atkinson framework is chosen, first calculate EDEH for the policies being compared and then calculate inequality by using $A(\varepsilon) = 1 - \frac{EDEH}{\mu}$. Equity impact is calculated as the difference in A(ε) between the policies being compared Use a reasonable range of values for ε and choose a plausible value for presentation of your main result. Table 13.3 illustrates results for health impact per person, A(ε) and EDEH using data from table 13.2. Here we have used ε = 5.

Table 13.3 Health distributions summarised as average HALE, A(ε), and EDEH for the NRT example (ε = 5)

	No NRT	Universal	Proportional Universal
Average HALE	70.626197	70.626290	70.626287
A(ε) [bigger is worse]	0.0122070	0.0122068	0.0122067
EDEH [bigger is better]	69.7640611	69.7641664	69.7641719

In our case study, the results in table 13.3 (with high inequality aversion) would imply that Proportional Universal would be ranked over Universal (EDEH is bigger) though differences are really small. In practical applications, early evidence that even high assumptions about aversion generate no substantive difference in conclusions than would be reached from a regular CEA might enable a major simplification of the study design and a reputable defence for making it.

Step 5: Plot results on equity impact plane

Results on net health impact and net equity benefit can also be plotted on the equity impact plane and is useful for understanding whether there is a trade-off between net health impact and net equity benefit, and the magnitude of potential impacts. It is also useful (in later steps) for communication with decision makers. To avoid information overload, it is convenient to choose a reasonable value for ε.[7] Testing out different values of ε is also helpful.

For the NRT example, net health and equity impact is calculated directly from table 13.3 (below), see figure 13.10.

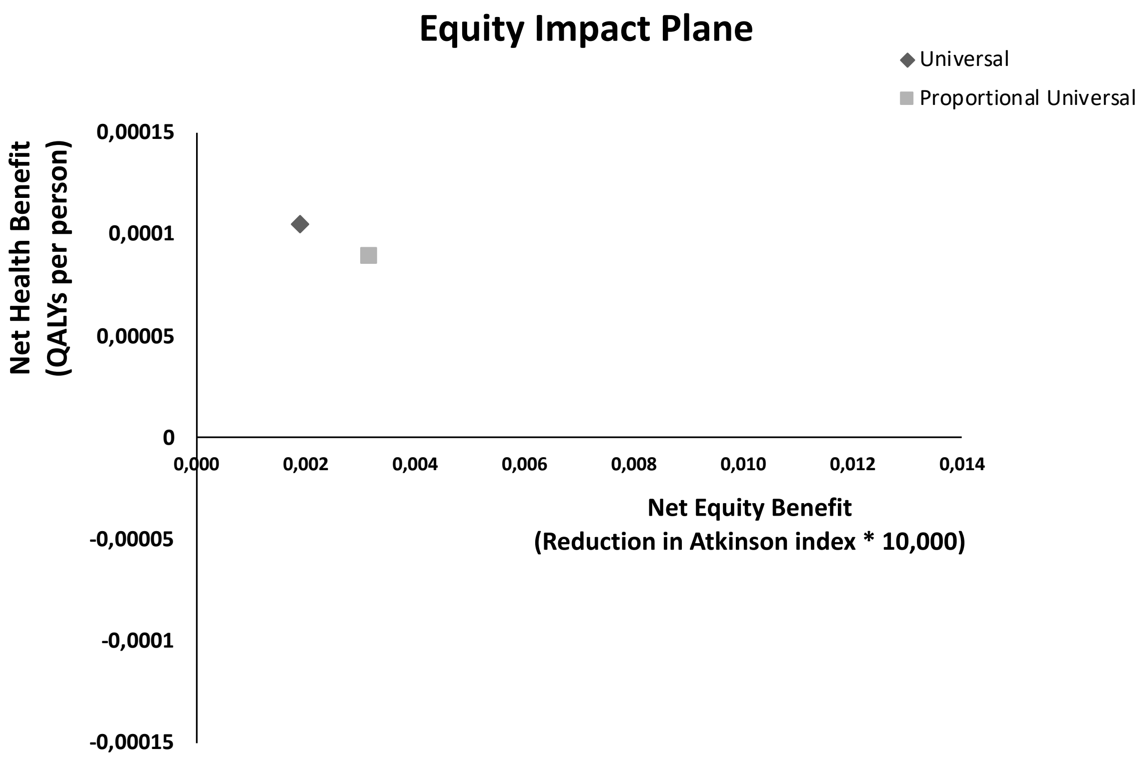

Figure 13.10 Equity impact plane, NRT example (ε = 5).

[7] For communication purposes, two graphs with different inequality aversion parameters (one low and one high) may also be produced.

Note In this analysis, net health impact and net equity benefit of Universal NRT and Proportional Universal NRT is compared to No NRT (baseline)

Step 6: Perform ethical sensitivity analysis by using different values of ε

Finally, calculate EDEH for the options in question by using a wide range of reasonable values (e.g. between 0 and 15 or higher) for the inequality aversion parameter and plot results in a graph (as in figure 13.11).

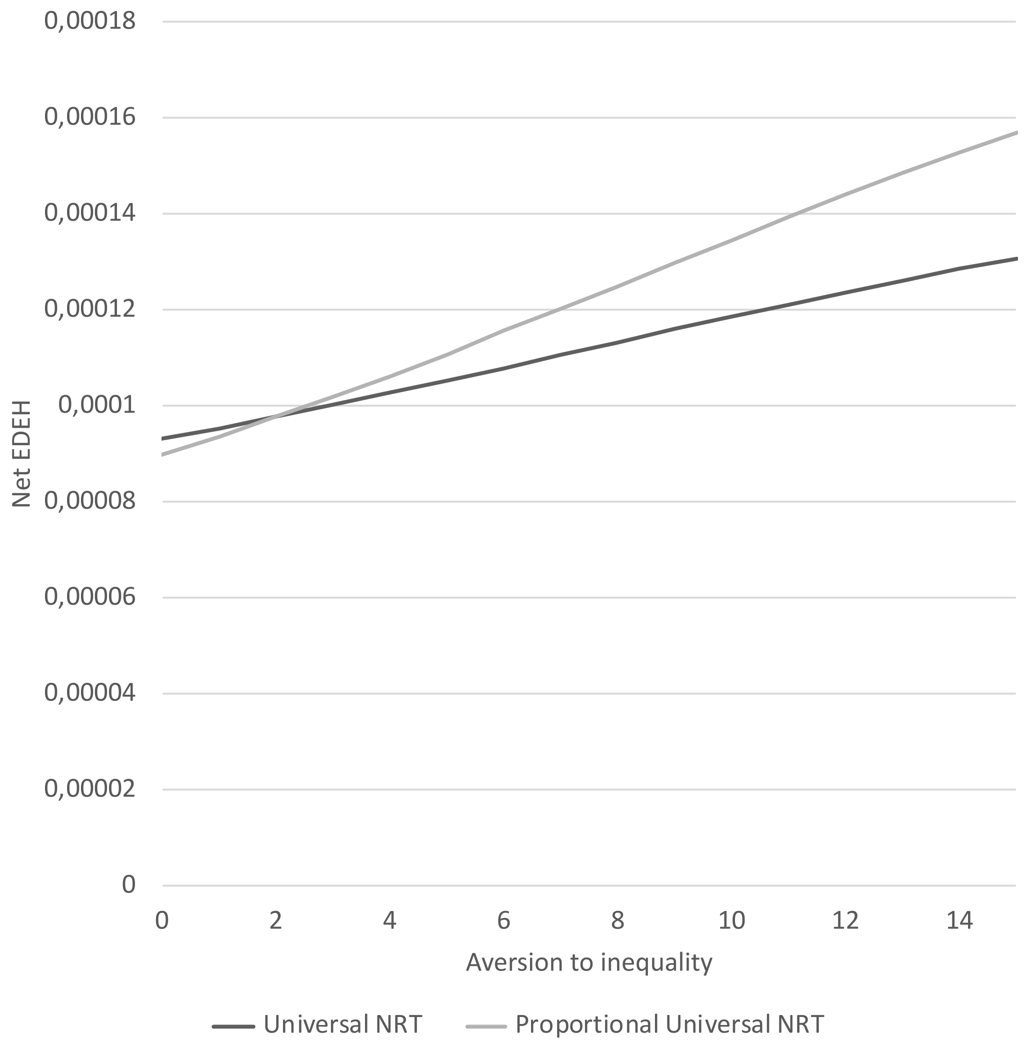

Figure 13.11 Net EDEH for the Universal and Proportional Universal NRT programme compared to no NRT as a function of different values of inequality aversion (ε = 0 to 15)

We see from the graph that net EDEH is highest for the Proportional Universal NRT programme if aversion to inequality is higher than about 2.2. Other forms of sensitivity analysis for other types of input where there is uncertainty is also possible and often necessary.

13.5 Comparison with rank-dependent SWFs

There are many similarities between level- and rank-dependent SWFs. Differences in practical implications may be few and the two types of analysis will often provide similar results. With respect to dominance ordering, there is little difference.

Both level- and rank-dependent SWFs are consistent with Generalized Lorenz (GL) dominance. In cases where GL dominance provides only a partial ordering of health distributions, both level- and rank-dependent SWFs are ways of completing the GL partial ordering. Both types of SWFs provide complete orderings that rank one distribution over another when the GL partial ordering leaves them non-comparable. The key point is that both types of SWFs are within the broader class of Paretian, Pigou-Dalton respecting SWFs (Adler, 2019; Bleichrodt & van Doorslaer, 2006).

One pragmatic advantage with a level-dependent SWF is that it relies only on available data from standard economic evaluations: baseline levels of health and incremental health. Analysis is quite straightforward and does not require additional data on personal characteristics like income or area of living. However, if the study needs only to deal with bivariate or multivariate equity analysis (e.g. health inequality by personal characteristics such as income, or a combination of income and education), the rank-dependent SWFs could be chosen (Wolfson & Rowe, 2001). The level-dependent SWFs (Atkinson and Kolm-Pollack) are inherently univariate.

Rank-dependent SWFs do not satisfy the axiom of separability. Suppose the health of separate individuals in a population can be represented by a vector. Separability means that the ranking of the vectors is independent of the well-being levels of the unaffected individuals. SWFs that capture a concern for relative rank violate the separability axiom, while level-dependent SWFs satisfy the separability axiom. If circumstances indicate that to assume separability is unjustified, evaluations using rank-dependent SWFs must include data on the distribution of health over the whole population. Since rank may affect also members of the population that are unaffected by the intervention in question, weights may also change and hence evaluation of distributions is not independent of non-affected individuals. This type of analysis can be quite data hungry.

13.6 Conclusion

In this chapter we have described how to conduct equity-efficiency trade-off analysis using an inequality aversion parameter within a level-dependent health-related social welfare function. We have also noted that standard CEA does not avoid these thorny value judgements: it simply makes a very specific value judgment that decision makers do not care about health inequality at all. If a decision maker, or the analysts, care about health inequality, this can be captured by standard measures of health (in)equity, such as the Gini or the Concentration Index and a rank-dependent SWF. We have discussed level-dependent SWFs that are compatible with the ethical framework of prioritarianism. The basic prioritarian principle is compatible with different value judgements about how much priority to give to the worse off. Level-dependent SWFs can be used to order different distributions of health following from different decisions. Even if the analyst does not know how much weight to assign to the worse off, we have shown how equity parameters can be used consistently and systematically to perform a sensitivity analysis of how different degrees of concern for the worse off imply different sets of equity weights for different people and can have different policy implications. Equity-efficiency trade-off analysis does not replace the need for decision makers to exercise their own judgement and wrestle with thorny ethical issues, but it can provide a helpful framework for thinking about the issues and useful analytical insight to inform their choices.

13.7 Further Reading

Adler provides a useful introduction to level-dependent social welfare functions (Adler, 2019). Asaria and colleagues provide an example of how level-dependent equity weights can be used in practice in distributional cost-effectiveness analysis (Asaria et al., 2015). Parfit is the classic statement of the general prioritarian view in ethics (Parfit, 1997). Ottersen sets out the prioritarian view as applied to lifetime health (Ottersen, 2013).

References

Adler (2012). *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. New York: Oxford University Press.

Adler. (2019). Measuring Social Welfare: An Introduction. New York: Oxford University Press.

Ali, Tsuchiya, Asaria, & Cookson (2017). How Robust Are Value Judgments of Health Inequality Aversion? Testing for Framing and Cognitive Effects. *Med Decis Making*, 37, 635-646.

Anand (2002). The concern for equity in health. *J Epidemiol Community Health*, 56, 485-487.

Anand, Diderichsen, Evans, Shkolnikov, & Wirth. (2001). Measuring disparities in health: methods and indicators. In T. Evans, M. Whitehead, F. Diderichsen, A. Bhuiya, & M. Wirth (Eds.), *Challenging inequities in health: from ethics to action* pp. 49-67). Oxford: Oxford University Press.

Asaria, Griffin, & Cookson (2016). Distributional Cost-Effectiveness Analysis: A Tutorial. *Med Decis Making*, 36, 8-19.

Asaria, Griffin, Cookson, Whyte, & Tappenden (2015). Distributional cost-effectiveness analysis of health care programmes--a methodological case study of the UK Bowel Cancer Screening Programme. *Health Econ*, 24, 742-754.

Atkinson. (2013). Health inequality, health inequity and health spending. In N. Eyal, S. Hurst, O.F. Norheim, & D. Wikler (Eds.), *Inequalities in health: ethics and measurement*. New York: Oxford University Press.

Baker, Bateman, Donaldson, Jones-Lee, Lancsar, Loomes, et al. (2010). Weighting and valuing quality-adjusted life-years using stated preference methods: preliminary results from the Social Value of a QALY Project. *Health Technol Assess*, 14, 1-162.

Bansback, Brazier, Tsuchiya, & Anis (2012). Using a discrete choice experiment to estimate health state utility values. *J Health Econ*, 31, 306-318.

Bleichrodt, & van Doorslaer (2006). A welfare economics foundation for health inequality measurement. *Journal of Health Economics*, 25, 945-957.

Cookson, Ali, Tsuchiya, & Asaria (2018). E-learning and health inequality aversion: A questionnaire experiment. *Health Econ*.

Costa-Font, & Cowell (2019). Incorporating inequality aversion in health-care priority setting. *Social Justice Research*, 32, 172-185.

Dolan. (2003). Fairness in Health: What the Public Thinks. In A. Oliver (Ed.), *Health care priority setting. Implications for health inequalities. Proceedings from a meeting of the health equity network*. London: The Nuffield Trust.

Dolan, Edlin, & Tsuchiya. (2008). The relative societal value of health gains to different beneficiaries. Discussion Paper. HEDS Discussion Paper 08/12.

Dolan, & Tsuchiya (2009). The social welfare function and individual responsibility: some theoretical issues and empirical evidence. *J Health Econ*, 28, 210-220.

Dolan, & Tsuchiya (2012). It is the lifetime that matters: public preferences over maximising health and reducing inequalities in health. *J Med Ethics*, 38, 571-573.

Edlin, Tsuchiya, & Dolan (2012). Public preferences for responsibility versus public preferences for reducing inequalities. *Health Econ*, 21, 1416-1426.

Gakidou, Murray, & Frenk (2000). Defining and measuring health inequality: an approach based on the distribution of health expectancy. *Bull World Health Organ*, 78, 42-54.

Harper, King, Meersman, Reichman, Breen, & Lynch (2010). Implicit Value Judgments in the Measurement of Health Inequalities. *The Milbank Quarterly*, 88, 4-29.

Johansson, & Norheim (2011). Problems with prioritization: exploring ethical solutions to inequalities in HIV care. *Am J Bioeth*, 11, 32-40.

Kolm (1976). Unequal inequalities. *Journal of Economic Theory,* 12, 416-442.
Lancsar, Wildman, Donaldson, Ryan, & Baker (2011). Deriving distributional weights for QALYs through discrete choice experiments. *J Health Econ,* 30, 466-478.
Lindholm, & Rosen (1998). On the measurement of the nation's equity adjusted health. *Health Econ,* 7, 621-628.
Norheim. (2013). Atkinson's index applied to health: can measures of economic inequality help us understand trade-offs in health care priority setting? In N. Eyal, S. Hurst, O.F. Norheim, & D. Wikler (Eds.), *Inequalities in health: ethics and measurement* pp. 214-230). New York: Oxford University Press.
Norman, Hall, Street, & Viney (2013). Efficiency and equity: a stated preference approach. *Health Econ,* 22, 568-581.
Otsuka, & Voorhoeve. (2018). Equality versus priority. In S. Olsaretti (Ed.), *The Oxford Handbook of Distributive Justice* pp. 65-85).
Ottersen (2013). Lifetime QALY prioritarianism in priority setting. *Journal of Medical Ethics,* 39, 175-180.
Ottersen, Mæstad, & Norheim (2014). Lifetime QALY prioritarianism in priority setting: quantification of the inherent trade-off. *Cost Eff Resour Alloc,* 12, 2.
Parfit (1997). Equality and Priority. *Ratio*, 10, 202-221.
Parfit. (2002). Equality or Priority? In M. Clayton, & A. Williams (Eds.), *The ideal of equality* pp. 81-125). New York: Palgrave, Macmillan.
Petrou, Kandala, Robinson, & Baker (2013). A person trade-off study to estimate age-related weights for health gains in economic evaluation. *Pharmacoeconomics,* 31, 893-907.
Robson, Asaria, Cookson, Ali, & Tsuchiya. (2015). Estimating the Level of Health Inequality Aversion in England. HESG Leeds Discussion Paper.
Rowen, Brazier, Keetharuth, Tsuchiya, & Mukuria (2016a). Comparison of Modes of Administration and Alternative Formats for Eliciting Societal Preferences for Burden of Illness. *Appl Health Econ Health Policy,* 14, 89-104.
Rowen, Brazier, Mukuria, Keetharuth, Risa Hole, Tsuchiya, et al. (2016b). Eliciting Societal Preferences for Weighting QALYs for Burden of Illness and End of Life. *Med Decis Making,* 36, 210-222.
Sen (1997). *On economic inequality. Expanded edition with a substantial annexe by James E. Foster*. Oxford.: Clarendon Press.
Shah, Tsuchiya, & Wailoo (2014). Valuing health at the end of life: an empirical study of public preferences. *Eur J Health Econ,* 15, 389-399.
Shah, Tsuchiya, & Wailoo (2015). Valuing health at the end of life: a stated preference discrete choice experiment. *Soc Sci Med,* 124, 48-56.
Tsuchiya, & Dolan (2005). The QALY model and individual preferences for health states and health profiles over time: a systematic review of the literature. *Med Decis Making,* 25, 460-467.
Tsuchiya, & Dolan (2007). Do NHS clinicians and members of the public share the same views about reducing inequalities in health? *Soc Sci Med,* 64, 2499-2503.
Wolfson, & Rowe (2001). On measuring inequalities in health. *Bulletin of the World Health Organization,* 79.