

Paper presented to HESG, Nottingham, 10-12 July 2000

**To reduce inequalities in health or to maximise health?
Determining the social welfare function using stated preference data**

Paul Dolan, Universities of Sheffield and Oslo (P.Dolan@shef.ac.uk)

Rebecca Shaw, University of York (rls5@york.ac.uk)

Aki Tsuchiya, University of York (at7@york.ac.uk)

Abstract

The maximisation of population health and the reduction of inequalities in health are two of the main objectives of the NHS. There are likely to be occasions when these twin objectives of efficiency and equity come into conflict with one another. The study reported in this paper looks at the extent to which 130 members of the general public are willing to sacrifice health gain in order to reduce inequalities that exist between people of different groups. The questionnaire was designed to allow the relevant parameters in a social welfare function to be determined. The results suggest that the curvature of the social welfare function is highly sensitive to *what* inequalities exist and *where* they exist. Respondents appear to be more concerned about reducing inequalities in life expectancy than they are about reducing differences in rates of long-term illness. Depending on the context, a health gain to social class V appears to be weighted between four and seven times as highly as compared to an equivalent gain to social class I. In contrast, respondents appear to be health maximisers when it comes to differences that exist between men and women and between smokers and non-smokers.

Key words: social welfare function, preference elicitation, equity-efficiency trade-off, inequalities in health

JEL Classification: I10

Acknowledgements

We would like to thank Roger Burrows, Atsushi Higuchi, Jan Abel Olsen, Peter Smith, Alan Williams, and all the respondents who agreed to take part. Rebecca Shaw is supported by the Economic and Social Research Council (Award No: 020R00456). Aki Tsuchiya is an overseas post-doctoral Research Fellow, Japan Society for the Promotion of Science, and a visiting Research Fellow, Centre for Health Economics, University of York. The usual disclaimer applies.

1. Introduction

The two main purposes of the NHS, and of most other publicly funded health care systems, are to maximise population health – an efficiency objective – and to minimise avoidable inequalities in health – an equity objective. A problem for any system with multiple objectives is what the balance between them ought to be. A policy that maximises population health might be one that makes relatively healthy groups of people even healthier whilst at the same time being of no benefit to less healthy groups, thereby increasing inequalities in health. Or, a policy that reduces inequalities might do so by foregoing the opportunity to improve the health of the relatively healthy, thereby compromising overall population health.

One way in which to determine how best to balance the competing objectives of efficiency and equity is to specify a social welfare function (SWF). Once the type of inequalities and the groups across which these inequalities exist have both been defined, the shape of the SWF will represent particular trade-offs between efficiency and equity. Thus, a SWF will allow the derivation of a set of relative weights to be given to a unit health gain to people from different population subgroups. In order for the SWF approach to be of any use to decision-makers, three questions need to be answered: 1) what is the space over which the SWF is to be specified; 2) which particular form of SWF is to be employed; and 3) how is the shape of the chosen SWF to be determined? Sections 2-4 deal with each of these questions in turn. Section 5 presents the results from an empirical study which elicited the public's preferences over the efficiency-equity trade-off in health. Section 6 shows how these data can be used to estimate the shape of the SWF. Section 7 discusses the implications of the results.

2. What is the space over which the SWF is to be specified?

Most economic models, including those that allow for the laundering of certain preferences, are welfarist in the sense that they rely on an individual's subjective assessment of her own well-being. As such, they are concerned with the distribution of *individual utility*. The focus on utility has been the subject of considerable criticism, much of which relates to the fact that welfarism is silent on how utility can be derived (see Sen and Williams 1982). This has led to alternative conceptions of welfare which, in various ways, use objective criteria to define well-being. For example, Rawls (1971) is concerned with the distribution of *primary goods*, not with what these goods do for people. By focusing on primary goods rather than utility, his theory avoids some of the interpersonal comparability problems (such as giving more income to people who are hard to please) associated with welfarism.

Sen (1987) agrees with Rawls' rejection of welfarism but criticises his focus on primary goods. If it is argued that resources should be devoted to people in poor health, despite there being no utility argument (because they might be contented) and despite there being no primary goods deprivation (because they have the goods that others have), then the basis for such an argument must lie elsewhere. Sen believes that what is at issue is the interpretation of needs in the form of *basic capabilities*. He suggests that focusing on basic capabilities is a natural extension of Rawls' concern with primary goods in that it shifts attention away from goods and towards what goods do for people.

Sen's work has been influential in the debate about why health care is considered to be more important than many other goods and services. For example, Culyer (1989) has argued that health care, through its impact on health, enables an individual to 'flourish'. This suggests that it may be appropriate to define a *health-related* SWF in terms of the different levels of health experienced by different groups (see Dolan 1998). Differences in health can then be represented as differences in length and/or quality of life. When the analytical objective is that of searching for a fair distribution of health (rather than utility), the general SWF approach is potentially more useful in a policy context since life expectancy and morbidity are more readily interpersonally comparable (see Olsen 1997).

For social welfare to be defined solely in terms of health, other attributes relating to the *causes* of ill health would have to be irrelevant. But society may not be neutral to the causes of ill health since there may be a range of 'lifestyle' characteristics by which the health gains received by one person might be weighted differently to those received by someone else. These could include the extent to which an individual is considered to be responsible for his own health, the degree to which he is informed about the health-related consequences of certain activities, the extent to which those activities are entered into voluntarily, and so on (see Olsen *et al* 2000)). Therefore, when the SWF is defined in terms of health, it might be important from a policy perspective to take account of these non-health attributes.

3. Which form of SWF is to be employed?

The specification of SWF used in this paper is one which assumes a constant elasticity of substitution (CES):

$$W = [aH_a^{-r} + bH_b^{-r}]^{\frac{1}{r}}, \quad H_a, H_b \geq 0, \quad a + b = 1, \quad r \geq -1, \quad r \neq 0,$$

where W is the level of overall population health and H_a and H_b are the level of health of groups a and b . The nature of the SWF, and the resulting iso-welfare curves, is determined by r and a . Further details of this type of SWF can be found in Wagstaff (1994) (see also Williams 1997 and Tsuchiya 2000).

The objective of this study is to identify the degree of people's aversion to inequalities in health, and to derive the implied weights to be given to a unit health gain to one group relative to another. This is represented as the marginal rate of substitution (MRS) along the relevant iso-welfare curve. As such, reference is made only to contours of the SWF, and not to the *level* of social welfare implied by these contours. In this respect, the CES SWF is equivalent to the so-called Atkinson SWF (Atkinson 1970) which was first proposed to address income distributions and has recently been used in empirical studies in the health context (see Dolan 1998; Dolan and Robinson 1999). We have chosen the CES specification rather than the Atkinson one since the interpretation of the parameters is more straightforward.

The parameter r measures the degree of aversion to inequality, as represented by the convexity of the iso-welfare curves. If $r = -1$, social welfare is equal to the sum of individual health and thus there is no aversion to inequality. This utilitarian-type SWF results in iso-welfare curves that are parallel straight lines with a gradient of $-a/b$. In

Figure 1, curve (a) represents the case where $\mathbf{a} = \mathbf{b}$. If $r > -1$, then there is aversion to inequality which exhibits a diminishing MRS i.e. the greater the inequalities in health between the two groups, the greater is the weight given to the worst-off group relative to the better-off group. This results in iso-welfare curves that are convex to the origin, and is represented as curve (b). In the extreme, the worst-off group are all that matters and $r = \infty$. This will result in a Rawlsian-type SWF with right-angled iso-welfare curves, as illustrated as curve (c).

The parameters \mathbf{a} and \mathbf{b} reflect the weight given to one group relative to the other, as reflected in the steepness of the iso-welfare curves. If $\mathbf{a} + \mathbf{b} = 1$ then social welfare is measurable on the same scale as health (which is a convenient but not essential property). If both groups are considered to be equal in every respect other than health, then $\mathbf{a} = \mathbf{b} = 1/2$, thus resulting in contours that are symmetric around the 45° line. This is represented by curve (d) in Figure 2. There, the gradient of the curve where it crosses with the 45 degree line is equal to -1 . If the groups are considered to be different (such as if one group is felt to be more responsible for their good or bad health than the other), then the gradient of the contour at the 45° line will be equal to $-\mathbf{a} / \mathbf{b}$. Curve (e) represents the case where $\mathbf{a} < \mathbf{b}$, and therefore the tangent is less steep than that for curve (d).

4. How is the shape of the SWF to be determined?

The question now is how do we measure the values of r and \mathbf{a} ? Williams (Williams 1997) suggests that respondents could be presented with the current unequal distribution of health and then asked to think about an equal distribution of health that makes them indifferent between the two distributions. However, this would require the health of the better off group to decline and this seems implausible in a policy context. This particular framing would also be susceptible to respondents giving greater weight to the losses of one group as compared to an equivalent gain to the other group. This loss aversion has been used to explain the results from preference elicitation studies across a range of different contexts (see Casey 1995 and Schweitzer 1995). In addition, it is implausible to consider health programmes which would result in a completely equal distribution of health between two groups. More realistically, the choice will be between programmes which maximise health and those which reduce (but do not eliminate) inequalities.

All of this suggests that questions of the general kind illustrated in Figure 3 would be best suited to estimating the parameter r in the SWF. The initial situation (I) is presented to respondents together with a programme (A) that will benefit both groups by the same amount. They are then presented with an alternative programme (B) that will be of greater benefit to the worse-off group. The aim then is to determine how much programme B would have to benefit the worse-off group in order to be considered equally as valuable as programme A. Since it has been shown that loss aversion may also be present when considering *potential* as well as actual losses from a particular reference point (see Dolan and Robinson 1999), it is important that neither programme is seen to represent the status quo. Once indifference between programmes A and B has been established, the value of r can be calculated. From this, the weight given to a unit health gain to the worse-off group vis-à-vis the better-off group can be estimated.

The parameter a can be determined by a different kind of question which asks respondents to consider two programmes which result in the same health benefit to each individual who receives them. Programme A is targeted on a group of people with a particular set of characteristics and Programme B is targeted on a different group with a different set of characteristics. Respondents are then asked to set the number of people in each group that would make them indifferent between the two programmes. The relative weight of the two groups is then calculated from the ratio of the number of people in each group. Person trade-off (PTO) questions of this kind have been used to address other distributional considerations in health care (see Nord 1995). Figure 4 represents a question of this kind, where potential recipients of programme A are given less weight than potential recipients of programme B.

5. The empirical study

In the present study, the health inequalities were defined in terms of average life expectancy and rates of limiting long-term illness. This enables us to compare whether the SWF has a different shape for length of life than for quality of life. It may have been possible to present people with variations in quality-adjusted life expectancy but such data is not as reliable as that relating to differences in life expectancy and long-term illness. In addition, explaining such information to respondents would have increased the cognitive burden on them.

There are a number of groups across which inequalities in health exist. The most obvious in a UK context relate to social class. On average, people in social class I (such as doctors) live five years longer than those in social class V (such as cleaners). For males aged 45-64, 12% of those in social class I report limiting long-term illness compared to 40% of those in social class V. Interestingly, differences of the same magnitude in average life expectancy exist between women and men and the same differences in long-term illness exist between non-smokers and smokers. To test whether the value of r is determined by the groups across which the inequalities exist, different samples of respondents were presented with data relating to different groups. To serve as a control case, some respondents were presented with groups which were defined simply as 'healthy' and 'unhealthy'.

In order to test whether the cause of ill health was a relevant consideration, respondents were presented with a question which asked them to trade-off a group of people who had cared for their own health (in terms of a good diet, regular exercise, not smoking or drinking to excess) against a group of people who had not. In this way, the parameters α and β in the SWF can be determined for groups defined according to the extent to which they are responsible for their own health.

5.1. The questionnaire

The questionnaire was constructed during a period of qualitative research involving in-depth interviews and focus groups and was refined during extensive piloting (the details of which are to be discussed in a separate paper). Since inequalities in health are not very prominent in lay presentations (see Blaxter 1997), it was decided to present the questions to respondents in the form of graphical representations. See the Appendix for details of the questionnaire and Table 1 for a summary of the study design.

The interviews were carried out in two rounds. Respondents in each round were further divided into two groups. In the first round of interviews, half of the respondents were presented with inequalities in length and quality of life by social class and half received inequalities in length of life by sex and quality of life by smoking status. In the second round of interviews, half of the respondents were presented with inequalities by social class and half received the healthiest 20% and the unhealthiest 20% of the population as the two comparison groups. To counter claims that people may be following a particular pattern of responses, the response categories were altered in the second round of interviews.

Respondents were asked a third equity-efficiency trade-off question in the context of rates of childhood mortality across the various population sub-groups used in the first two questions. However, the responses to this question do not allow the parameters in a SWF to be determined and so they are not reported here. Finally, a PTO question involving two groups of 55 year-olds (one which have taken care of their health and one which have not) was asked in the second round of interviews. All respondents were asked how difficult they found the questions and the interviewer coded how much help was needed and the time taken to complete the questions.

5.2 The sample

Every 8th person appearing on the electoral register in 3 wards in York was contacted and invited to participate. Out of a total of 1,500 letters initially distributed, 467 people (31%) agreed to participate. To ensure that the sample was broadly representative of the wider population, respondents were selected for interview based on information on a broad range of characteristics obtained from their reply slips. In total, 130 individuals were interviewed out of a total of 140 who were invited to attend. All interviews were carried out by one of the authors and two other researchers. Each interview took place at the University of York and lasted for approximately one hour, of which ten to fifteen minutes was spent on the questions analysed in this paper. Table 2 shows the characteristics of the sample. The characteristics of those who attended the interviews were similar to those of the population of Yorkshire and Humberside.

5.3 Results

Columns 2-4 in Table 3 present the distribution of responses in the context of average life expectancy. For those respondents who initially chose to target on programme B but then switched at some point to programme A, their point of indifference has been taken to be half-way between the last point at which they chose B and the first point at which they chose A. The precise trade-offs made by those who choose not to target and by those who always choose to target are, of course, indeterminate. The median respondent is indifferent between people in social classes I and V living on average to be 80 and 75, respectively, and them living to be 78 and 75.5, respectively. This is also the median response when the sub-groups are defined in terms of the healthiest and unhealthiest quintiles of the population. However, when the sub-groups are defined by sex, the median preference is not to target on men at all.

Columns 2-4 in Table 4 show the results from the long-term illness question. The median response to the social class question implies indifference between a decrease in the rate of long-term illness of 7% for both social classes and a 2% reduction for class I with a 8.5 reduction for class V. When the sub-groups used are the healthiest and unhealthiest, the median respondent is indifferent between a 7% reduction for both groups and a 2% reduction for the healthy with a 10% reduction for the unhealthy. However, when the sub-groups are defined according to smoking status, the median preference is not to target on smokers at all. Table 5 shows the results from the PTO question. Column 2 shows that the median response is indifference between 100 people who haven't taken care of their health and 45 people who have.

The analysis of the effect of various background characteristics on people's stated preferences is at a preliminary stage and will be reported more fully at a later stage. However, preliminary analysis suggests that differences in responses cannot be explained by differences in personal characteristics. For example, better educated people do not seem to be any less willing to target social class V than their less educated counterparts.

6. Estimating the SWF

The questions relating to average life expectancy and long-term illness ask respondents to identify two points, A and B , on an iso-welfare curve. If the notation $H_a(B)$ is to stand for the value of health of subgroup a at point B , then the corresponding gradient at the midpoint of A and B is:

$$\frac{dH_b}{dH_a} = - \left(\frac{H_b}{H_a} \right)^{(1+r)} = - \left(\frac{(H_b(A) + H_b(B))/2}{(H_a(A) + H_a(B))/2} \right)^{(1+r)}.$$

Furthermore, by definition, the same gradient can be approximated as:

$$\frac{dH_b}{dH_a} \approx \frac{H_b(B) - H_b(A)}{H_a(B) - H_a(A)}.$$

Therefore, by taking the logarithms of these and solving for r :

$$r \approx \frac{\log(H_b(A) - H_b(B)/H_a(B) - H_a(A))}{\log(H_b(A) + H_b(B)/H_a(A) + H_a(B))} - 1.$$

In other words, given $A(H_a, H_b)$ and $B(H_a, H_b)$, it is possible to approximate the value for r that identifies the contour through these two points.

Strictly speaking, this is only possible for those respondents who switch from programme B to programme A at some point. For those who chose A in the initial pairwise comparison, it is tempting to assume that $r = -1$ (i.e. utilitarian-type preferences). However, for this to be the case, such respondents would need to be indifferent between A and B when they bring about the same overall benefit. Since respondents were not given this as a response option, it may be reasonable to assume that they were indeed health maximisers but we cannot rule out the possibility that some respondents may have

been inequality prone. It is also impossible to determine a precise value for r for those who always chose B. It would seem reasonable here to assume that these respondents are concerned only with equality and hence $r = \infty$ (i.e. Rawlsian-type preferences) but again we cannot be sure.

The 5th column in Tables 3 and 4 represent the implied value of the r parameter for all question options. Since respondents who always choose option B are assumed to have an infinitely large r , this makes the calculation of the group mean impossible for most questions. Further, for those questions where the median respondent chooses option A in the initial pairwise comparison, it is not completely straightforward to assume that the value of r that represents the median preference is -1 . The Tables show that the value of r increases disproportionately with the extent of the equity-efficiency trade-off. This is due to the convex shape of the iso-welfare curve. Note that since the SWF is defined in terms of resulting levels of health, the value of r is independent of the level of health at the initial point.

Once the value of r is obtained, the weight to be given at the margin to the health of subgroup a relative to the health of subgroup b at the initial point is given by the negative MRS:

$$-\frac{dH_b}{dH_a} = \frac{\mathbf{a}}{\mathbf{b}} \left[\frac{H_b}{H_a} \right]^{(1+r)},$$

and since $\mathbf{a} = \mathbf{b}$, this becomes

$$\left[\frac{H_b}{H_a} \right]^{(1+r)}.$$

The implied weights for each of the question options are given in the 6th column of Tables 3 and 4. For more details, see Tsuchiya (2000).

The 7th column in Tables 3 and 4 indicate the implied equity-efficiency trade-offs in terms of health at the initial point. This is calculated as the difference between average health and the “equally-distributed equivalent health”. The (negative) values indicate how much of a decrease in overall population health people would be willing to accept for there to be an equal distribution of health between the two groups. So, in the case of average life expectancy by social class, the median respondent would be willing to move from the initial point (where social classes 1 and V live to be 78 and 73, respectively) to a point where both groups live to be 74.5 i.e. they would be willing for average health to be one year less if it were distributed equally.

The PTO question assumes that the two groups are exactly the same in all relevant respects except for whether or not they have taken care of their own health. Again, two points, A and B, that lie on the same iso-welfare curve are identified. This time, since the

two groups are identical, it is reasonable to assume that $r = -1$, implying a straight-line contour. The relevant function to solve here is:

$$aH_a(A) + (1-a)H_b(A) = aH_a(B) + (1-a)H_b(B).$$

All values except for a are known, and this can easily be calculated. Since $r = -1$, the implied weights is:

$$\frac{a}{b} = \frac{a}{1-a}.$$

Values for a and the implied weights are presented in columns 3 and 4 in Table 5. The weight given to someone who has not cared for her health is about half as much as that given to someone who has cared for the health. Note that, since the relevant iso-welfare curve is a straight line, the weights apply not only to the initial state point but also to *any* other point in the relevant space.

7. Discussion

The study reported in this paper attempted to determine the extent to which the general public are concerned about reducing inequalities in health as compared to maximising overall population health. Preferences over various equity-efficiency trade-offs were used to determine the shape of the SWF in the context of average life expectancy and limiting long-term illness. The results generated should not be regarded as definitive numbers. Rather, by asking respondents to think about the sacrifices in overall population health that they would be willing to make, they allow us to draw some general conclusions about the extent to which real-world inequalities in health really do matter to them.

In this regard, the results presented here make some intuitive sense. It would appear that differences in average life expectancy by social class are more important to people than differences in rates of long-term illness by social class, perhaps reflecting the view that to die younger is more of an injustice than to be in poorer health whilst alive. In both cases, though, there is a willingness to trade-off efficiency for equity and the resulting values of r in the SWF were both found to lie between -1 and ∞ which appeals to ‘the common-sense conception of justice’ (Elster 1992). Similar results to social class were found when the groups were defined in terms of the health ‘healthiest’ and ‘unhealthiest’ quintiles in society.

Differences between the average life expectancy of men and women did not seem to matter so much. This might be explained by the fact that it is not obvious whether it is better to be a man or a woman. In fact, with a few exceptions, most people would seem to be happy with the sex they were given at birth. Therefore, both men and women might be happy to accept the health differences that exist between them. In contrast, most people would prefer to be in social class I rather than in social class V and, as such, it may be seen as more of an injustice for there to be health differences between the classes. Such an explanation is akin to Varian’s notion of envy-free justice (see Varian 1974,1975).

On average, respondents were unwilling to reduce inequalities in health that were due to smoking status. And the responses to the PTO question suggest that people who have not cared for their own health are given half as much weight as those who have cared for their health. These findings are not surprising since there is now evidence that many people wish to give less priority to those who are considered to be in some way responsible for their ill health (see Dolan *et al* 1999).

It is also encouraging that the distribution of responses remains more or less constant across the two rounds of interviews. Moreover, the interviewers felt that respondents were able to provide meaningful responses to the trade-off questions and suggested that these questions may be amenable to a postal questionnaire. Overall, then, it would appear that this study has demonstrated that the theoretical concept of a SWF can be combined with stated preference data to generate intuitively appealing results.

Of course, there are many theoretical and methodological questions which a study of this kind raises. To begin with, there are normative questions relating to the space over which the SWF ought to be specified. It appears that health is accorded special status in most countries, and so it would seem reasonable to specify the SWF over the health space rather than in terms of individual utility (see Evans and Wolfson 1982). Doing so certainly makes the questions presented to respondents more manageable and facilitates interpersonal comparisons of welfare.

The particular form of SWF employed in this study has a constant elasticity of substitution. This, of course, represents only one specification and it may be that the data generated from this study are amenable to other functional forms. In addition, there are other ways in which the parameter r in the SWF might be calculated. In this paper, r has been calculated according to the resulting *levels* of health, rather than in terms of the *gains* in health. Since the actual questions given to the respondents were phrased in terms of gains from the initial point, many may have interpreted the equity-efficiency trade-off in terms of these gains. Therefore, it would be interesting to see how the implied weights given to particular groups vary according to how the SWF is specified and how r is calculated.

If the results generated from studies of this kind are going to be used to inform policy decisions, then it is important to represent group preferences in a meaningful way. Given the highly skewed nature of the values for r , it is not possible to represent group preferences using the mean value. It is difficult, therefore, to account for the strength of each individual's preferences in the overall views of the group. Based on arguments around the median voter hypothesis, the median is considered by many to be the most appropriate measure of central tendency in a policy context. The median might mitigate against some of the problems of a skewed distribution but it certainly does not eliminate them. In this study, the median respondent is not willing to target some disadvantaged groups (such as men) at all and so even the median value of r is indeterminate for some questions. The issue of how to aggregate preferences in order to represent a SWF has received considerable attention in the literature but this highlights the need for yet more work here.

There are also questions concerning the implications for policy of some of the qualitative findings reported in this study. For example, there is a willingness to reduce inequalities in rates of long-term illness that exist between social classes I and V but an unwillingness to do so when the same inequalities exist between non-smokers and smokers. But, of course, the prevalence of smoking is higher in social class V and so to target this group *would* be to target smokers. Since respondents were not asked to consider health inequalities by social class *and* smoking status, there is no internal inconsistency in the data reported here. However, there will be an inconsistency so far as practical policy is concerned if some of the reason for the higher rates of long-term illness in social class V is the higher prevalence of smoking in this group.

From earlier qualitative work (to be reported elsewhere), many people consider social class to be a factor largely outside of an individual's control and this could certainly explain the desire of most respondents to reduce social class differences in health. Since the preferences of those presented with differences between groups of 'healthy' and 'unhealthy' people were similar to those presented with differences by social class, it would seem that the extent to which someone is 'healthy' is also seen to be outside an individual's control. This is interesting since that same qualitative work revealed that people considered general 'lifestyle' to be the most important determinant of health.

Beyond the policy usefulness of the findings, there are questions relating to the general robustness of the results. It is now recognised that people's stated preferences across a range of decision contexts are to some extent shaped by the processes that are used to elicit them (see Rabin 1998). This study was designed to minimise the effects of certain well-established framing effects, such as reference point effects, the status quo bias, omission bias and loss aversion. But, of course, it is impossible to remove every potential bias. In particular, respondents were always presented with response categories in the same order; that is, programmes A and B start out being equally effective and then B (the targeted programme) becomes incrementally less effective. This ordering was chosen to highlight the equity-efficiency trade-off and to eliminate any confusion that might result from respondents being asked to 'ping-pong' backwards and forwards between response categories. However, there is the possibility of a start-point bias whereby some respondents get 'locked into' choosing B throughout.

In addition, there are questions relating to the reliability of stated preference data. Economists are brought up to believe that preferences that are not motivated by any degree of self-interest cannot be trusted (see Johannesson 1999). This certainly explains their scepticism of stated preference studies that do not provide respondents with any (financial) incentives to reveal their 'true' preferences. However, this scepticism follows from the assumption that social welfare is primarily a function of the utility levels of self-interested individuals, and this is certainly contestable. Although self-interest exists, it does not necessarily follow that it must be the basis for social welfare since society may adopt any objective or set of objectives that it desires. If an ethically defensible set of society-regarding preferences can be derived from moral intuition or forms of ethical reasoning which are separate from mere self-interest, then the social perspective might be a legitimate one (see Menzel 1999).

A related criticism of studies of this kind is that some respondents may have given what Miller (1992) refers to as 'Sunday Best' responses; that is, "the views that people think

they ought to hold according to some imbibed theory as opposed to the operational beliefs that would guide them in a practical situation.” However, there is a strong argument that only those preferences and social values that people are prepared to air publicly should be used to inform social policies which are designed to incorporate the public’s views on social justice. Gauthier (1986) has argued that the preferences people reveal in behaviour should match the attitudes they enunciate verbally and it is arguably the latter that should be used in priority-setting in publicly-funded national health care schemes.

Overall, the fact that responses to the questions reported here appear to make intuitive sense leads us to conclude that it is possible to use stated preference to say something about the shape of the *health-related* SWF. We make no claims to have provided definitive answers but hope that we have shown how empirical data can be used in conjunction with theoretical constructs to provide policy relevant information.

References

Atkinson, A B. 1970. "On the measurement of inequality." *Journal of Economic Theory* 2:244-263.

Blaxter, M. 1997. "Whose fault is it? People's own conceptions of the reasons for health inequalities." *Social Science and Medicine* 44:747-756.

Casey, J T. 1995. "Predicting buyer-seller pricing disparities." *Management Science* 41:979-999.

Culyer, A J. 1989. "The normative economics of health care finance and provision." *Oxford Review of Economic Policy* 5:34-58.

Dolan, P. 1998. "The measurement of individual utility and social welfare." *Journal of Health Economics* 17:39-52.

Dolan, P., Cookson, R. and Ferguson, B. 1999. "Effect of discussion and deliberation on the public's views of priority setting in health care: focus group study" *British Medical Journal* 318:916-919.

Dolan, P and Robinson, A. 1999. "The measurement of preferences over the distribution of benefits: the importance of reference points." *European Economic Review* (forthcoming).

Elster, J. 1992. *Local Justice: How institutions allocate scarce goods and necessary burdens.* New York: Russel Sage.

Evans and Wolfson, 1982

Gauthier, D. 1986. *Morals by agreement.* Oxford: Oxford University Press.

Johannesson, M. 1999 "On aggregating QALYs: a comment on Dolan" *Journal of Health Economics* 18 5:381-386.

Menzel, P. 1999 "The allocation of health care resources: an ethical evaluation of the 'QALY' approach - McKie, J., Richardson, J., Singer, P. And Kuhse, H." *Bioethics* 13 5:436-440.


Miller, D. 1992. "Distributive justice: what do people think?" *Ethics* 102:555-593.

Nord, E. 1995. "The person trade-off approach to valuing health care programmes." *Medical Decision Making* 15:201-208.

Olsen, J A. 1997. "Theories of justice and their implications for priority setting in health care." *Journal of Health Economics* :625-639.

- Olsen, J A, Dolan,P., Menzel, P. and Richardson, J.** 2000. “*The moral relevance of personal characteristics in setting health care priorities.*” Presented to the Nordic Health Economics Study Group.
- Rabin, M.** 1998 “Psychology and Economics” *Journal of Economic Literature* 36 1:11-46.
- Rawls, J.** 1971. *A Theory of Justice*. London: Harvard University Press.
- Schweitzer, M.** 1995. “Multiple reference points, framing and the status quo bias in health care financing decisions.” *Organisational Behaviour and Human Decision Processes* 63:69-72.
- Sen, A.** 1987. *On Ethics and Economics*. London: Blackwell.
- Sen, A and Williams, B.** 1982. “*Utilitarianism and Beyond.*” . Cambridge: Cambridge University Press.
- Tsuchiya, A.** 2000. “*The design and calculation of the fair innings weights.*” Paper presented at the HESG meeting, Newcastle .
- Varian, H.** 1974. “Equity, envy and efficiency.” *Journal of Economic Theory* 9:63-91.
- Varian, H.** 1975. “Distributive justice, welfare economics and the theory of fairness.” *Philosophy and Public Affairs* 4:223-47.
- Wagstaff, A.** 1994. “QALYs and the equity-efficiency trade-off.” pp. 428-447 in *Cost-Benefit Analysis (2nd ed.)*, edited by A. Laylard and S. Glaister. Cambridge: Cambridge University Press (reprinted from *Journal of Health Economics*, 10:21-41, 1991).
- Williams, A.** 1997. “Intergenerational equity: an exploration of the 'fair innings' argument.” *Health Economics* 6:117-132.

Figure 1: The effect of change in value of r



H_a, H_b : health of sub-population a and b


I : the initial state

I' : the point where the health of the two groups are exchanged

Three different types of social welfare contours:

- (a): $r = -1$... cf. the classical utilitarian
- (b): $-1 < r < \infty$... cf. efficiency-equity trade-off
- (c) $r = \infty$... cf. the Rawlsian

Figure 2: The relationship between \mathbf{a} and \mathbf{b} .

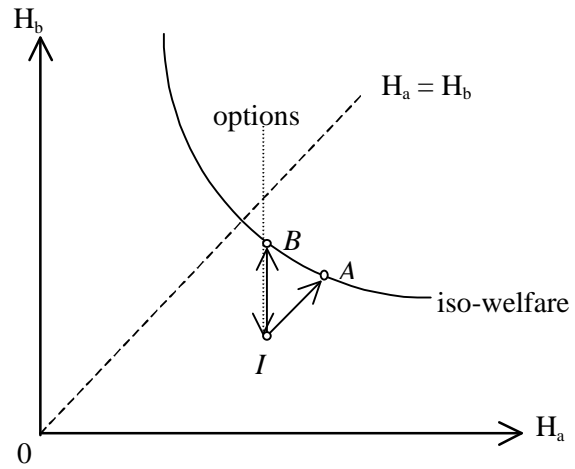


I : the initial state

Two different types of social welfare contours:

- (d): $\mathbf{a} = \mathbf{b}$... the gradient of the tangent to the contour at the 45 degree line is -1
- (e): $\mathbf{a} < \mathbf{b}$... the same gradient ($=\mathbf{a} / \mathbf{b}$) is less steep

Figure 3: Question 1



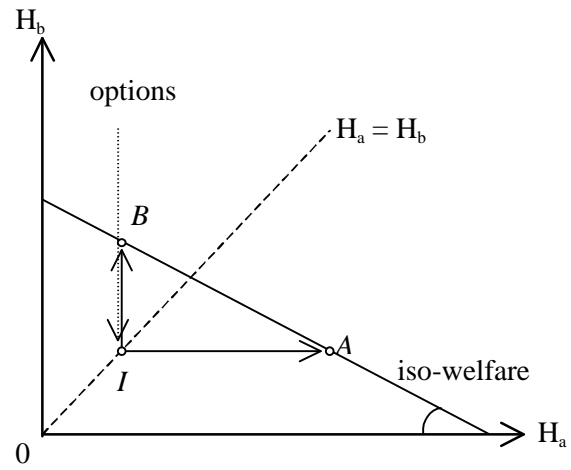
I : initial point

A : outcome offered by programme A

options: the set of options offered by the alternative programme

B : the point at which a given respondent is indifferent between the two programmes

Figure 4: Question 4



I : initial point

A : outcome offered by programme A

options: the set of options offered by the alternative programme

B : the point at which a given respondent is indifferent between the two programmes

Table 1: Questionnaire variants

Question	1 st round of interviews		2 nd round of interviews	
	Average Life Expectancy	Social Class	Sex	Social Class
Long-term illness	Social Class	Smoking status	Social Class	Healthy/ Unhealthy
Person trade-off	Life-style	Life-style	Life-style	Life-style
N	29	31	37	33
Total	60 respondents		70 respondents	

Table 2: Respondent characteristics

Category	Interviews 1	Interviews 2	Yorkshire and Humberside*
Sex:			
Male	48%	49%	47%
Female	52%	51%	53%
Age:			
<44 years	48%	51%	49%
>45 years	52%	49%	51%
Mean age	50	48	
Dependants:			
Children	57%	76%	66%
No-children	43%	24%	34%
Smoking status:			
Smoker	15%	17%	27%
Ex-smoker	-	33%	
Non-smoker	85%	50%	73%
Economic status:			
Employed	60%	49%	56%
Other	40%	51%	44%
N	60	70	4014

* The Annual Survey of English Housing 1997/1998 and The British Household Panel Survey 1997.

Table 3: Average life expectancy questions

80 and 75 = 78 and ..	Social class n = 66	Sex n = 31	Healthy/un n = 33	Implied r	Implied weight at initial point	Implied trade-off
77 or more	24	20	12	-1.0	1.0	0.0
76.75	0	1	-	2.3	1.3	-0.1
76.50	5	-	2	5.9	1.6	-0.3
76.25	0	0	-	9.8	2.1	-0.4
75.75	2	1	-	20.1	4.0	-0.8
75.50	11	-	14	27.9	6.8	-1.0
75.25	11	5	-	41.2	16.4	-1.4
75 or less	13	4	5	Infinity	Infinity	

median response in bold

“-” indicates where a response category was not offered to respondents

Table 4: Limiting long-term illness questions

95 and 67 = 90 and ..	Social class n = 66	Smoking n = 31	Healthy/un n = 33	Implied r	Implied weight at initial point	Implied trade-off
72 or more	14	27	9	-1.0	1.0	0.0
71.75	1	0	-	-0.8	1.1	-0.2
71.5	3	-	1	-0.6	1.2	-0.5
71.25	0	0	-	-0.4	1.2	-0.7
70.5	1	0	-	0.2	1.6	-1.6
70.0	5	-	9	0.7	2.0	-2.3
69.5	6	2	-	1.3	2.4	-3.0
68.5	7	1	-	2.9	4.4	-4.9
68.0	13	-	14	4.1	7.1	-6.2
67.5	5	0	-	6.2	16.0	-7.9
67 or less	11	1	0	Infinity	Infinity	

median response in bold

“-” indicates where a response category was not offered to respondents

Table 5: Person trade-off question

100 and...	Life-style n = 70	Implied a	Implied weight At initial point
100 or more	1	0.5	1.00
95.5	4	0.49	0.96
77.5	21	0.44	0.78
45.0	24	0.31	0.45
17.5	1	0.15	0.18
		0.05	0.05
5 or less	5		

median response in bold

Appendix A – Example of the questions.

Question 1 – Average life expectancy by social class.

Question 1

As you might know, average life expectancy differs by social class.

Whilst actual life expectancy varies between individuals, on average, people in social class 1 live to be **78** and in social class 5 they live to be **73**.


Imagine that you are asked to choose between two programmes which will increase average life expectancy. Both programmes cost the same.

In the two graphs below the light grey part shows average life expectancy, and the dark grey part shows the increase in life expectancy. There is a separate graph for each of the programmes.

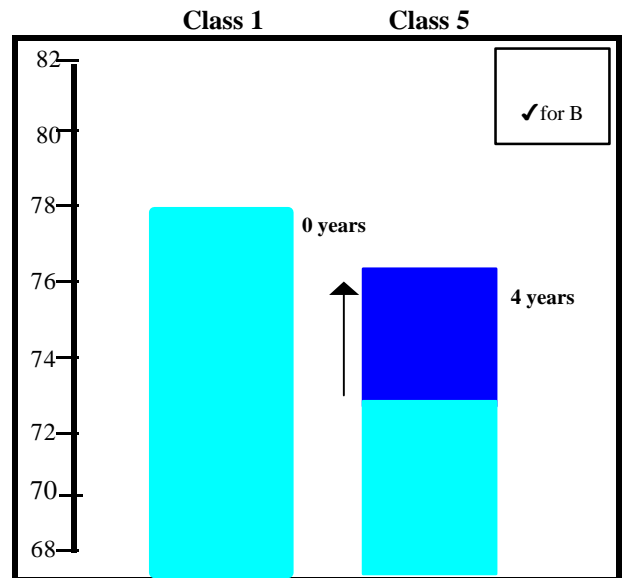
As you can see, Programme A is aimed at both social classes equally and Programme B is aimed more at social class 5.

Please indicate whether you would choose A or B by ticking one box.

Programme A



Programme B



Respondents who chose A were directed to the next question. Respondents who chose B were asked to turn over the page.


**Respondents in the second round of interviews were presented with four options.
Respondents in the first round of interviews were presented with six options.**

Choosing Programme B might mean that the increase in life expectancy is less overall. For each of the **four** choices below, please tick one box to indicate whether you would still choose B, or whether you would now choose A.



Programme A

Programme B



1.





2.



3.




4.




Programme A

Programme B


1.




2.




3.




4.



5.



6.



Question 4 – person trade-off

Question 4

Imagine that you are asked to choose between two programmes which could benefit two different groups of people. Both programmes cost the same.

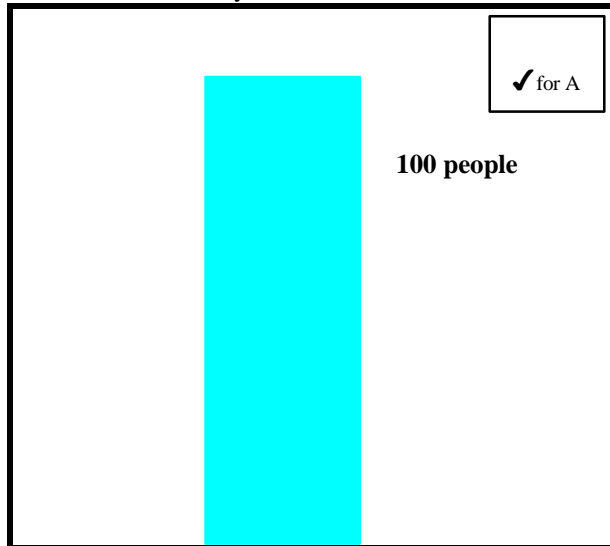
Without the intervention, patients will die within a few days but with the intervention they will live for another 10 years in good health and then die.

People in both groups are 55 years old. They are similar to one another except that those in Programme A have not taken care of their health, whilst those in Programme B have.

Please indicate whether you would choose A or B by ticking one box.

Programme A

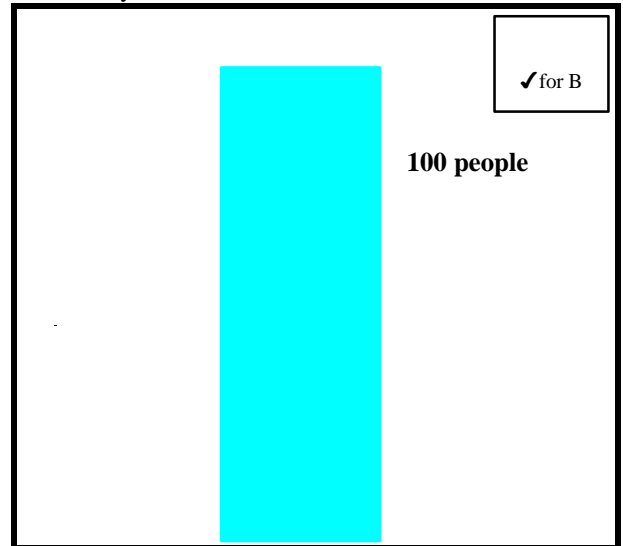
Number of 55 year olds who **haven't** taken care of their health, who will live for ten more years, rather than die in a few days.



A rectangular box containing a cyan vertical bar on the left side. To the right of the bar, the text "100 people" is written. In the top right corner of the box, there is a small square checkbox containing a checkmark and the text "for A".

Programme B

Number of 55 year olds who **have** taken care of their health, who will live for ten more years, rather than die in a few days.




A rectangular box containing a cyan vertical bar on the left side. To the right of the bar, the text "100 people" is written. In the top right corner of the box, there is a small square checkbox containing a checkmark and the text "for B".

Choosing Programme B might mean that fewer people can be treated. For each of the **four** choices below, please tick one box to indicate whether you would still choose B, or whether you would now choose A.


Programme A

1. People who **haven't** taken care of their health





Programme B



People who **have** taken care of their health



2 .



3.



4.

