



Utilitarianism and the Measurement and Aggregation of Quality – Adjusted Life Years

PAUL DOLAN

Professor in Health Economics, Sheffield Health Economics Group and Department of Economics, University of Sheffield, 30 Regent Street, Sheffield S1 4DA (E-mail: P.Dolan@sheffield.ac.uk)

Abstract. It is widely accepted that one of the main objectives of government expenditure on health care is to generate health. Since health is a function of both length of life and quality of life, the quality-adjusted life-year (QALY) has been developed in an attempt to combine the value of these attributes into a single index number. The QALY approach – and particularly the decision rule that health care resources should be allocated so as to maximise the number of QALYs generated – has often been equated with the utilitarian philosophy of maximising ‘the greatest happiness of the greatest number’. This paper considers the extent to which the measurement and aggregation of QALYs really is utilitarian by developing a new taxonomy in order to classify utilitarianism and the different aspects of the QALY approach. It is shown that the measurement of QALYs is consistent with a number of different moral positions and that QALYs do not have to be aggregated according to the maximisation rule. Therefore it is inappropriate to *necessarily* equate QALYs with utilitarianism. It is shown that much turns on what *in principle* the QALY represents and how *in practice* it can be operationalised. The paper highlights the category confusion that is often present here and suggests possible avenues for future theoretical and empirical research.

Key words: health policy, moral philosophy, priority-setting, quality-adjusted life-years, utilitarianism

1. Background

It is widely accepted that a principal objective of government expenditure on health care is to generate health. Since health is a function of both length of life and quality of life, the quality-adjusted life-year (QALY) has been developed in an attempt to combine the value of these attributes into a single index number. If QALYs represent health benefits and if the objective of health policy is to maximise health, then it follows that resources should be allocated so as to maximise the number of QALYs generated.

There has been considerable ethical debate about the morality of allocating scarce health care resources according to the number of quality-adjusted life-

years (QALYs) that they generate. Whilst this debate has been wide-ranging, much of it has focused on the role that health gain should play in prioritisation decisions across different groups or individuals. There have been those who have questioned the moral acceptability of health gain being given *any* weight (Harris, 1996) whilst many others have argued that health gain is only one of many different considerations (Broome, 1994).

Because the measurement and aggregation of QALYs has typically been equated with the principle that resources should be allocated so as to bring about the greatest gains in health, it has often been equated with the utilitarian philosophy of maximising ‘the greatest happiness of the greatest number’ (Bentham, 1789). The perception that the whole QALY approach is a utilitarian one is possibly reinforced by the fact that its chief proponents are typically health economists whose discipline arguably has utilitarianism at the heart of its orthodoxy (Hausman and McPherson, 1993). As a result, the QALY is often criticised or defended precisely because it is seen as adopting the utilitarian philosophy.

One of the aims of this paper is to consider the extent to which this is true. To do this, a new taxonomy is developed by which to classify utilitarianism and to distinguish it from some of its more prominent alternatives. In so doing, it helps to clarify precisely what is meant by the term ‘utilitarianism’ and, perhaps more importantly, it draws attention to different ethical foundations which are so commonly ignored by economists. Where the measurement and aggregation of QALYs fits into this taxonomy is then discussed. This means that the QALY can then be exposed to a range of moral arguments relating to each category within the taxonomy. In discussing the QALY in this way, the paper also aims to highlight those areas where future theoretical and empirical research might be conducted into the whole approach.

2. Five Steps to Utilitarianism

To distinguish utilitarianism from other moral philosophies, a taxonomy has been developed in which each of the criteria is defined according to a binary category; one for utilitarianism and one for ‘non-utilitarian’ philosophies. Of course, there are many moral philosophies besides utilitarianism and it would be impossible to discuss them all here. Rather, the most prominent alternatives for each criterion are discussed briefly. Note also that there are many philosophies which adopt the utilitarian position on one or more of the criteria below.

Categorising different philosophical perspectives is an incredibly complicated task. There are problems associated with the definition of the categories

Table 1. The defining characteristics of utilitarianism and some of its alternatives

Utilitarianism	Alternative viewpoint
<i>Consequentialist</i> evaluates only the consequences of actions	<i>Non-consequentialist</i> (e.g. Nozick, 1974) considers rights to be of intrinsic value
<i>Monist</i> utility is a homogeneous measure of benefit	<i>Pluralist</i> (e.g. Rawls, 1971) different (incommensurable) dimensions
<i>Welfarist</i> goodness depends only on individual utility	<i>Non-welfarist</i> (e.g. Sen, 1987) non-utility information may matter
<i>Preference-satisfaction</i> Actual preferences are what matter	<i>Non-preference based</i> (e.g. Scanlon, 1975) distinguishes between 'needs' and 'wants'
<i>Aggregation according to sum-ranking</i> social welfare is the sum total of utilities	<i>Relative advantage</i> (e.g. Rawls, 1971) social welfare depends on relative position

themselves and many of the concepts do not neatly fit into one category or another. Therefore, whilst these criteria would ideally be mutually exclusive, it is inevitable that there will be some overlap between them. However, the defining feature of utilitarianism according to the taxonomy developed in this paper is that it is a consequentialist, monist, welfarist, preference-based philosophy in which advantage is aggregated according to sum-ranking. The taxonomy is summarised in Table 1.

2.1 *Consequentialism or non-consequentialism?*

Utilitarianism is a consequentialist philosophy which evaluates every action according to the consequent states of affairs. This means that in the context of the debates about the importance of rights and the role of duties, it judges any right or duty according to what happens as a result of the exercising of that right or performing of that duty. In this way, it is possible to trade-off different rights against one another in different contexts and, as such, no one right is considered of more intrinsic value than any other.

This is not true of non-consequentialist (or deontological) theories which judge some rights to have greater moral value than others. Possibly the most famous deontological philosopher of them all, Immanuel Kant (1785), viewed autonomy (the right to rational self-legislation) as perhaps the central moral value. More recently, Nozick (1974) has argued that individuals have certain inviolable rights, such as the right to the fruits of their own labour, which must be respected no matter what their consequences in terms of the distribution of income.

In the context of how to evaluate different social states, consequentialism (and hence utilitarianism) will look only at the outcomes associated with those states. This is criticised by Hahn (1982) on the grounds that “My utility may not only depend on what I get but on the manner in which I get it. That is, my utility may not only depend on the consequences of policy but on the policy itself.” The suggestion, then, is that the process by which a decision is made may have an effect independent of the outcome of that decision. In principle, however, any attribute that can be embedded in the description of a state can be the object of consequential moral evaluation. Therefore, it is important to distinguish between the conceptual possibility of including certain considerations in a consequentialist view from the moral question of what ought to be included.

2.2 *Monist or pluralistic morality?*

Utility is the focus of interest in utilitarian philosophy. This makes it a monist morality since all interests, ideals and desires are regarded as commensurable. The idea that there exists one homogeneous magnitude offers a standard of consistency by which to rank different states of the world. Pluralist morality, on the other hand, recognises that there may be different incommensurable dimensions across which social states could be compared. The implication of this could be that the social states are incommensurable with one another but pluralist morality does not necessarily admit incompleteness. For example, Rawls’ Theory of Justice (1971) provides a hierarchy of criteria (basic liberties, freedom of choice of occupation, powers and prerogatives of office, income and wealth, and social bases of self-respect) by which to judge social states. Therefore, the contrast between monist and pluralistic moralities does not rest on the issue of completeness *per se*, but on the way completeness is achieved when it is achieved.

2.3 *Welfarism or non-welfarism?*

Utilitarianism is welfarist in that it considers the goodness of a state to be a function only of the utility information regarding that state and views persons as locations of their respective utilities. Sen and Williams (1982) are critical of this approach, arguing that “Persons do not count as individuals in this any more than individual petrol tanks do in the analysis of the national consumption of petroleum.” Sen (1987) argues that judging the well-being of a person exclusively in the metric of happiness has damaging limitations in the context of interpersonal comparisons of well-being. For example, one implication is that if a poor person gets half the utility from a given amount of income than

a rich person, then the income should be given to the latter rather than the former.

Rawls (1971) argues that the most important principle of justice is that each person should have an equal right to the most extensive scheme of basic liberties. Thus, by arguing for the priority of liberty (i.e. non-utility information), his Theory of Justice lies outside of welfarism. Moreover, by focusing on primary social goods rather than utility, his theory avoids the problems associated with welfarism in the context of interpersonal comparisons of well-being. However, if it is argued that resources should be devoted to the worst-off, 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 (1987) 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.

2.4 Preferences or non-preferences?

In general terms, attention can be focused on an individual's manifest (actual) preferences or on what might be regarded as her 'true' preferences which will not necessarily coincide with her actual preferences if she is ill-informed or does not know what is 'best' for her. Whilst the utilitarian view of preferences has evolved over time, it has consistently focused attention on actual preferences. The original hedonistic perspective of Bentham (1789), according to which people do things simply to attain pleasure and to avoid pain, and the more 'refined' ethics of Mill (1859), which distinguishes between 'higher' and 'lower' pleasures, as well as the plethora of more recent derivatives of these models, are concerned only with manifest preferences.

However, Broome (1991) argues that the preference-satisfaction model "can be refuted by pointing out the undeniable fact that people sometimes prefer, of two alternatives, the one that is worse for them, because they are misinformed about the merits of the alternatives". Many utilitarian economists would require that preferences are suitably 'corrected' for mistaken beliefs, and sometimes even 'laundered' to exclude certain antisocial preferences, such as envy and malice (Harsanyi, 1982). But most would argue that other-regarding preferences should not be included in a conception of individual well-being. In keeping with this, Broome (1991) proposes a model which would represent an individual's preferences if she were rational, well-informed and self-interested. In contrast, Sen (1987) argues that defining preferences solely in terms of self-interest is too narrow and claims that a

mixture of selfish and selfless behaviour can be seen in a wide variety of situations.

Of great relevance to public sector resource allocation decisions, Scanlon (1975) has criticised preference-satisfaction on the grounds that it fails to distinguish between needs and wants. He suggests that there are occasions (for example, when deciding whether to give money to a beggar) where we will make our decision using objective criteria relating to what we think the person 'needs' rather than subjective criteria relating to what they want. In this way, we may consider it our moral obligation to feed a beggar, but not to finance his gambling, say.

2.5 Aggregation according to sum-ranking or not?

The utilitarian philosophy aggregates utility information according to sum-ranking i.e. by looking only at the sum-total of utilities, and this is clearly related to absolute positions. Therefore, even the tiniest gain in the total sum would be taken to outweigh distributional inequalities of the most blatant kind (provided, of course, that the inequalities themselves did not directly affect utility). Rawls (1971, 1982) argues that social states should be evaluated according to the difference principle relating to relative positions. He argues that social and economic inequalities must be to the greatest benefit of the least advantaged but 'maximin' can be criticised for its lack of interest in the magnitude of utility gains and losses. Of course, deontological philosophies will not aggregate advantage at all since all individuals have certain inviolable rights and hence aggregation across individuals is meaningless.

3. Utilitarianism and QALYs

This section discusses where the theoretical and empirical work into the QALY fits into this taxonomy and considers the appropriateness (or otherwise) of the positions adopted. The discussion is not intended to cover every viewpoint but rather to provide an overview of the philosophical position adopted by most health economists who have undertaken research into QALYs and to highlight those areas where most disagreement amongst health economists more generally can be found. In so doing, areas for future research are identified.

It is also worth noting at this point that the maximisation of QALYs can be defended or criticised on many grounds besides the extent to which the decision rule conforms to the utilitarian philosophy. For example, Menzel (1990) provides a contractarian justification for basing resource allocation decisions on the number of QALYs generated which turns not on the

consequences of this rule but on the possibility that a QALY-maximising rule may reflect a population's *prior consent* to a particular allocation of resources.

3.1 *Consequentialism or non-consequentialism?*

That the QALY is defined as a measure of health *outcome* means that it is firmly located within a consequentialist framework. One right is rarely considered to be morally superior to another. It is, of course, possible (and entirely consistent) for some rights to act as 'side-constraints' on public policy. For example, the QALY approach could be used to evaluate paediatric services subject to the constraint that every woman has the right to an abortion. Using QALYs, then, does not mean that rights need to be ignored. But the important point in the context of this discussion is that the QALY concept is only operational once these constraints have been imposed and in this sense the approach is regarded as consequentialist.

Mooney and others have been critical of the emphasis that QALYs have placed on health outcomes and have suggested that the utility from process should also be taken into account when making resource allocation decisions. However, it would seem that some of the examples of 'process utility' that are given, such as the fear of dentists (Mooney, 1994: 74), are better defined as the (outcome) utility associated with the treatment itself. In principle, such information (at least about one's current or future health status) could be taken account of within an appropriately defined 'mental health' dimension of an instrument designed to measure changes in an individual's utility (or well-being). And this, of course, would be entirely consistent with all but the narrowest of consequentialist frameworks. Other examples, such as autonomy in the decision-making process (Mooney, 1994: 19), might represent genuine process utility and would be difficult to capture within the QALY approach.

Therefore, the debate about the importance (or otherwise) of 'genuine' process utility needs first to be placed in a coherent conceptual framework that is, so far as possible, agreed upon by all health economists. This will enable us to move on from the debate about precisely what process utility is and will highlight those contexts (if any) in which a genuine trade-off exists between process utility and outcome utility. This will then provide a meaningful background for empirical research into the extent to which people are willing to make a trade-off between processes and outcomes in the allocation of scarce health care resources.

3.2 *Monist or pluralistic morality?*

Proponents of the QALY adopt a monist approach in which they attempt to collapse the benefits from health care into a homogeneous magnitude. The

logic is that, since all resource allocation decisions make an implicit trade-off between different attributes, consistency requires that those trade-offs be made explicit. Although it is recognised that individual behaviour may be shot through with incommensurable values arising out of pluralistic considerations, the argument is that public policy requires consistency in a way that personal decisions may not (see Nelson, 1994).

However, there is the unresolved empirical question as to whether people give such a great weight to consistent decision-making. They may recognise that all policy decisions imply some trade-off between attributes (which at some level implies commensurability) but it is conceivable that they would prefer for these trade-offs to vary across decision contexts in ways which could not be captured by even the most sophisticated QALY measure. For example, they might wish for policy-makers to have some discretion over the trade-offs they make so that context-specific public participation (or ‘voice’; see Anand, 1999) can have a real impact on particular decisions at the margin.

3.3 *Welfarism or non-welfarism?*

Whether the measurement of QALYs is consistent with welfarism or not depends initially on what QALYs are supposed to represent. If they represent an individual’s cardinal utility index, then they are clearly within a welfarist framework. However, this requires a number of very stringent assumptions, including that the health component of an individual’s utility function is additively separable from the wealth component (see Bleichrodt and Quiggin, 1999). Because of this, most economists define the QALY as a “cardinal utility index of health” (Culyer 1991, emphasis added). Thus, since there is more to utility than health, the QALY approach is not truly welfarist.

Although QALYs leave open the question of how health state values are determined (see Broome, 1993), most of the empirical work has been welfarist insofar as individuals are asked to assign a utility value to a state of health that they are experiencing, or are asked to imagine themselves experiencing. Where the approach deviates greatly from the welfarist one is that these values (particularly if they are elicited from the general public) are often then attached to health states that are experienced by *other* people; for example, those in a clinical trial.

This assumes, then, that a particular level of dysfunction has the same effect on all individuals. This is done to avoid one of the problems associated with welfarism; namely, adaptive preference formulation, which is important in health care since people are known to adapt to illness. But it does, of course, raise interpersonal comparability problems of its own; for example, being a wheelchair user will have a much smaller effect on an academic’s professional life than on a builder’s. Therefore, there is the need for much

more theoretical and empirical research into which sources of difference in health state values across respondents are considered to be legitimate ones and which are considered to be illegitimate ones.

Much of the discussion about those in poor health (or with fewer QALYs) has been influenced by Sen's (non-welfarist) capabilities approach. For example, Culyer and Wagstaff (1993) suggest that good health is necessary for an individual to 'flourish' as a human being, and therefore propose that a just distribution of health (or QALYs) is an equal distribution. Of course, many authors have argued that equality of health is not desirable because some people may choose to be less healthy than others (see Mooney and Jan, 1997), and have consequently focused attention on equality of access. But again the emphasis on equality of opportunity is clear (and arguably more explicit). A possible avenue for future research, then, would be one that focused attention on which particular definition(s) of equality of opportunity were considered most appropriate in which contexts.

3.4 Preferences or non-preferences?

Much of the research undertaken into QALYs lies within the preference-based framework but it is important to remember that QALYs do not equate benefit with happiness or utility since they are restricted to the health domain only (see section 3.3 above). And those who focus on equality of QALYs typically adopt a non-preferences approach in which it is possible, at the margin, to override individual preferences in the interests of improvements in the overall level or distribution of QALYs. For example, Culyer and Wagstaff (1993) argue that "[t]here is a danger in straining out the gnat of offending personal liberty that one swallows the camel of enduring and outrageous inequalities of health." Those, like Mooney (1994), who focus more on access to health care adopt a more preferences-based approach in which the exercising of individual choice is a legitimate reason for differences in health. It is important, then, to consider those circumstances where it is justifiable to override individual preferences in order to enhance individual well-being. To date, this is an issue which has received little attention from health economists.

3.5 Sum-ranking or something else?

QALYs are typically aggregated according to sum-ranking which could be justified on the grounds that it is morally preferable to confer greater benefit than less. However, in choosing how to prioritise between patients, the size of the benefit is not the only consideration. Broome (1991) suggests that we look for reasons why one person should be given priority over another, and proposes a class of reasons, referred to as 'claims'. He argues that fairness

is about mediating the claims of different people and requires that claims should be satisfied in proportion to their strength. Harris (1996) believes that each citizen is entitled to (or has a 'claim' to) an equal chance of having his health needs respected.

Whilst it is common practice to aggregate QALYs according to sum ranking, there is nothing in the measure *per se* which forces benefits to be aggregated in this way. Given the plethora of studies which assume that benefits should be weighted equally, regardless of the relative 'claims' of the recipients, it seems that this is one area where empirical investigation is urgently required. To incorporate distributional considerations into resource allocation decisions requires data on what the general public and others consider to be legitimate claims and on the relative weight they attach to these claims.

4. Discussion

This paper has attempted to develop a new taxonomy in order to classify different moral philosophies and, specifically, to distinguish utilitarianism from other theories. A more modest aim has been to consider where the QALY fits within the taxonomy in order to assess the claim that the QALY is a utilitarian concept. Figure 1 summarises the extent to which the QALY *in principle* satisfies the conditions required for utilitarianism. It can be seen that it strictly satisfies two of the five conditions; consequentialism and monism. Those for whom monism is not acceptable would be forced to reject the QALY approach. But whilst the QALY is incompatible with a non-consequentialist moral theory, there would seem to be considerable scope for moral argument about what ought to be included in the description of a health consequence and what *in practice* can be incorporated into the QALY model.

The welfarism/non-welfarism and preferences/non-preferences categories in the taxonomy have received considerable attention from health economists. This probably stems from the fact that health has certain characteristics which make it unlike most other commodities and which challenge many of the assumptions of standard welfare economics. By focusing attention on the QALY, this paper has served to show where future theoretical and empirical research might be conducted into such issues as whether there is more to social welfare than individual well-being and the role that individual preferences should play in determining priorities.

The reason why the whole QALY approach has been so closely linked to the utilitarian one is probably related to the point made in the introduction that the QALY concept is often seen to be synonymous with QALY maximisation. This stems partly from that fact that proponents of the QALY approach often

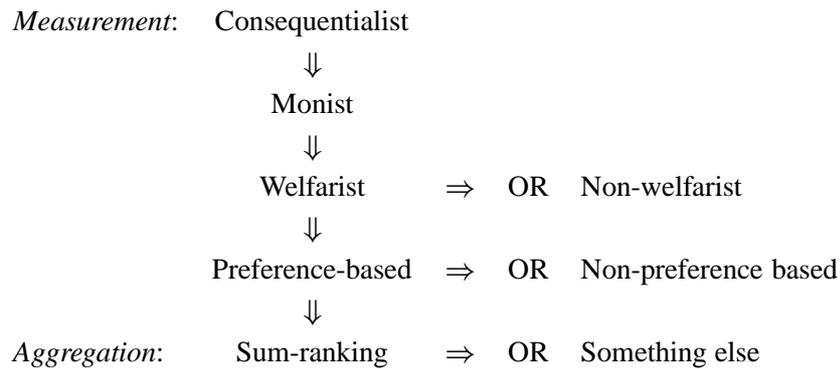


Figure 1. The defining characteristics of QALYs

endorse the *principle* of sum-ranking (Culyer, 1997) and partly from the fact that *in practice* most empirical studies have aggregated QALYs in this way. Rejecting the QALY on the grounds that QALY maximisation is unacceptable is analogous to throwing out the baby with the bathwater but there is certainly the need for much more empirical research into how QALYs can be weighted to take account of a whole range of distributional considerations.

Conceptually, then, the measurement and aggregation of QALYs is at least 2/5ths utilitarian (assuming that all dimensions are weighted equally, of course) and at most entirely utilitarian. Therefore, this paper has demonstrated that the debate about the QALY does not and should not be limited to a debate about the relative merits of utilitarianism. Ultimately, much turns on what *in principle* the QALY represents and how *in practice* it can be operationalised. To date, there would seem to have been some category confusion here.

Acknowledgements

It would be impossible to thank everybody who has contributed to the development of this paper but I would particularly like to thank the following people: Richard Cookson, Tony Culyer, Paul Menzel, Jan Abel Olsen, Jeff Richardson, Jennifer Roberts and Alan Williams.

References

- Anand, P. (1999) Procedural fairness in economic and social choice: evidence from a survey of voters, *paper presented to the Royal Economic Society*, University of Nottingham.

- Bentham, J. (1789) *An introduction to the principles of morals and legislation*, W. Harrison (Ed.), Oxford: Blackwell, 1967.
- Bleichrodt, H. and Quiggin, J. (1999) Life-cycle preferences over consumption and health: when is cost-effectiveness analysis equivalent to cost-benefit analysis? *Journal of Health Economics* **18**(6), 681–708.
- Boadway, R.W. and Bruce, N. (1984) *Welfare Economics*, Ch. 6, Oxford: Blackwell.
- Broome, J. (1991) *Weighing goods*, Oxford: Blackwell.
- Broome, J. (1993) Qalys, *Journal of Public Economics* **50**, 149–167.
- Broome, J. (1994) Fairness versus doing the most good, Hastings Center Report, July – August, pp. 36–39.
- Culyer, A.J. (1991) The normative economics of health care financing and provision, in A. McGuire, P. Fenn and K. Mayhew (Eds.), *Providing health care*, Oxford: Oxford University Press.
- Culyer, A.J. (1997) The principal objective of the NHS should be to maximise health: the case for, in B. New (Ed.), *Rationing: Talk and Action in Health Care*, London: King's Fund and British Medical Journal Publishing.
- Culyer, A.J. and Wagstaff, A. (1993) Equity and equality in health and health care, *Journal of Health Economics* **12**, 431–457.
- Hahn, F. (1982) On the difficulties of the utilitarian economist, in Sen, A. and Williams, B. (Eds.), *Utilitarianism and beyond*, Cambridge: Cambridge University Press, pp. 187–199.
- Hammond, P. (1982) Utilitarianism, uncertainty and information, in Sen A and Williams B (Eds.), *Utilitarianism and beyond*, Cambridge: Cambridge University Press, pp. 85–102.
- Harris, J. (1996) What is the good of health care?, *Bioethics* **10**(4), 269–291.
- Harsanyi, J. (1982) Morality and the theory of rational behaviour, in A. Sen and B. Williams (Eds.), *Utilitarianism and beyond*, Cambridge: Cambridge University Press, pp. 39–62.
- Hausman, D. and McPherson, M. (1993) Taking ethics seriously: economics and contemporary moral philosophy, *Journal of Economic Literature* **31**: 671–731.
- Hughes, T. and Larson, L. (1991) Patient involvement in health care: a procedural justice viewpoint, *Medical Care* **29**, 3.
- Kant, I. (1785) *Groundwork of the metaphysics of morals*, translated by H. Paton, London: Harper and Row, 1948.
- Menzel, P. (1990) *Strong medicine: The ethical rationing of health care*, Oxford: Oxford University Press.
- Mill, J.S. (1859) *On liberty*, C.V. Shields (Ed.), London: Macmillan, 1985.
- Mooney, G. (1994) *Key issues in health economics*, New York: Harvester Wheatsheaf.
- Mooney, G. and Jan, S. (1997) Vertical equity: weighing outcomes? or establishing procedures?, *Health Policy* **39**, 79–87.
- Nelson, J. (1994) Publicity and pricelessness: grassroots decision-making and justice in rationing, *Journal of Medicine and Philosophy* **19**, 333–342.
- Nozick, R. (1974) *Anarchy, state and utopia*, New York: Basic Books.
- Rawls, J. (1971) *A Theory of Justice*.
- Rawls, J. (1982) Social unity and primary goods, in A. Sen and B. Williams (Eds.), *Utilitarianism and beyond*, Cambridge: Cambridge University Press, pp. 159–185.
- Scanlon, T.M. (1975) Preference and urgency, *Journal of Philosophy* **77**, 655–69.
- Sen, A. (1987) *On ethics and economics*, Oxford: Blackwell.
- Sen, A. (1987) Justice, in J. Eatwell, M. Milgate and P. Newman (Eds.), *The new polgrave: a dictionary of economics*, London: Macmillan.
- Sen, A. and Williams, B. (1982) Introduction, in A. Sen and B. Williams (Eds.), *Utilitarianism and beyond*, Cambridge: Cambridge University Press, pp. 1–21.