



# Morality and the limits of societal values in health care allocation

Rebecca L. Walker<sup>a,\*</sup> and Andrew W. Siegel<sup>b</sup>

<sup>a</sup>*Life Sciences Values and Society Program and Department of Philosophy, University of Michigan, USA*

<sup>b</sup>*Bioethics Institute and School of Public Health, Johns Hopkins University, USA*

## Summary

In this paper, we consider whether there is a clear moral justification for the proposal that societal value preferences (SVPs) should be included in Cost Effectiveness Analyses (CEA) of health care resource allocations. We argue, first, that proponents of the use of SVPs need to be clear about the relationship between these values and moral principles. In particular, once moral principles are accepted as ruling out some SVPs (such as those that are irrational or revealing prejudice), an account is required of why we need to appeal to SVPs rather than moral principles to determine a just division of health care resources. Secondly, we consider whether an independent moral justification might underwrite the use of SVPs. In various places in the literature the notions of representation, presumed consent and democratic decision making appear to be invoked as candidates for fulfilling this justificatory role. We discuss some problems with each of these justifications in the hope of eliciting a more comprehensive proposal from the proponents of SVPs. We conclude that, although a number of interesting proposals have been made, no compelling justification for including SVPs in CEA has yet been systematically articulated. Copyright © 2002 John Wiley & Sons, Ltd.

**Keywords** cost utility analysis; cost effectiveness analysis; societal value preferences; justice; health care allocation

## Introduction

An important question in health economics and ethics is how to justly allocate health care resources under conditions of scarcity. One possible answer, offered by health economists, is to employ a version of the formal allocation method of Cost Effectiveness Analysis (CEA) using Quality Adjusted Life Years (QALYs). This form of CEA is generally referred to as Cost Utility Analysis (CUA). QALYs give a quality-weighted value between 0 and 1 to each year of life under particular health conditions. The most cost-effective distribution of resources is the one that

results in the most QALYs gained for the least dollars spent.

Implicit in this method of allocation is a form of utilitarianism, as it is assumed that the best allocation scheme is the one that maximizes the number of healthy life years in the population as a whole. In recent literature, this implicit utilitarian value standard has been challenged because people's actual preferences for allocation schemes seem to diverge from utility maximization. Instead of simply aggregating individual utility assessments, it is argued, societal value preferences (SVPs) should be taken into account.

In this paper, we consider whether there is a clear moral justification for the claim that SVPs

\*Correspondence to: Department of Philosophy, University of Michigan, 2215 Angell Hall, 435 S. State St., Ann Arbor, MI 48109-1003, USA. Tel.: +734-615-9725; fax: +734-763-8071; e-mail: walkerrl@umich.edu

should play a role in health care allocation. We begin by pointing out that proponents of the use of SVPs recognize that not just any such preferences should be taken into account. More specifically, they exclude morally objectionable preferences. Given the implicit priority granted to ethical principles, we argue that proponents of the use of SVPs have the burden of showing that SVPs are not superfluous to such principles. We then consider whether the use of SVPs may nevertheless be independently morally justified. Although a number of interesting options seem initially available, we argue that no compelling justification has yet been systematically articulated.

### The problem with CUA

To determine individual health-related utilities for the purpose of CUA, people are asked about the value they would place on a year of life under specified health conditions. Two of the standard methods for determining these values are time trade off (TTO) and standard gamble (SG). In TTO, respondents are asked how many years of their life in a specific health condition they would forfeit to live the other years of their life in perfect health. In SG, respondents are asked what risk of death they would accept to be cured of a specific health condition. For example, if a respondent would accept either a 20% risk of death or forfeit 2 of 10 years of life to be cured of blindness, then under a QALY analysis, one year of life with blindness is worth 0.8 (or 80% of a year of perfect capacity).

Serious ethical problems arise, however, when these individual health utility assessments are used in social policy decisions about health care through CUA. While it might be worth a 20% risk of death to one to be cured of blindness, one does not thereby agree that it is better to give priority to a program that will save the lives of ten sighted persons rather than ten blind persons. Yet the straightforward CUA seems to imply that this would be the case. For this reason, 'recommendations emerging from [CUA] can stand starkly at odds with the values of justice and non-discrimination' [1].

The moral problems that are generated by using intra-personal utility assessments in inter-personal allocation decisions drive some theorists to try to develop methods of directly assessing individual's

preferences for inter-personal health care allocation. One method designed to elicit these preferences is the person trade off (PTO) question. In a PTO question respondents might be asked how many people in a specific health condition would need to be cured before the funding of this treatment is more socially valuable than saving the lives of ten other people. In general, the empirical evidence seems to show that values other than utility maximization come into play when people are specifically asked about inter-personal allocation preferences. For example, people generally prefer to help the least well off first, and to distribute goods equitably, other things being equal. Proponents of SVPs argue that these kinds of preferences should be taken into account in CEA in order to escape the moral problems posed by the straightforward use of CUA.

### Moral principles and societal value preferences

One argument advanced by those who advocate taking SVPs seriously is that the values are often supported by independent moral arguments, and thus have a greater status than mere 'preferences.' For example, principles of fairness, which require that we give priority to helping 'the worst off,' support the preferences elicited in PTO studies attaching significance to the severity of illness. Those calling for the inclusion of SVPs in CEA maintain that it is 'ethically objectionable' not to consider such values where there exist independent arguments underwriting them:

Some aspects of treatment are such that omitting them from allocative decisionmaking is ethically objectionable. This is a more demanding notion than the simple claim that such aspects involve overlooked 'preferences' or unexamined 'values.' For the omission of a factor to be ethically objectionable, one has to be able to articulate some argument for the preferences it generates, not merely point out that people hold those preferences [1].

The claim here is not that it is necessary and sufficient for justifying the inclusion of a societal value into CEA that a moral argument can be generated which supports the value. For we *can* generate moral arguments on behalf of SVPs that proponents of the use of SVPs would seek to exclude from allocation decisions. Imagine that the

PTO questions revealed that most individuals are prepared to sacrifice the lives of persons with disabilities to save the lives of persons without disabilities. While it might be possible to produce a version of a utilitarian argument in support of these values, those advocating the use of SVPs would not want such values to figure into health policy decisions. On their view, preferences count only 'assuming that those preferences are not irrational or ethically objectionable, as when they reflect discriminatory attitudes.' [1]

But while it seems reasonable to condition the use of SVPs on their being ethically acceptable, this proviso gives rise to the question of why it is necessary to consider SVPs at all. If we cannot justify an appeal to SVPs without first establishing that they are not ethically objectionable, then SVPs will be superfluous to the extent that we can resolve resource allocation issues through the application of ethical principles.

One might respond that we are not able to provide decisive moral arguments for many of the allocation questions we are concerned with. This is not necessarily to say that decisive arguments do not exist. One could instead hold that unless we discover such arguments, it is legitimate to accord SVPs the status otherwise reserved for moral values. Peter Ubel expresses sentiments along these lines:

Many allocation dilemmas have no simple solutions, and highly trained, intellectually rigorous philosophers would completely disagree with each other about the best solution, for example, about the extent to which severely ill patients deserve treatment priority. In such situations, the public deserves a role [2].

Although Ubel seems skeptical about the possibility of consensus on many allocation issues, it appears implicit in his remarks that we ought to follow the lead of philosophers and other intellectuals to the extent they *are* able to achieve consensus. But then the claim seems to be that we should give priority to ethical principles over preferences, at least where such consensus exists. While proponents of the use of SVPs in CEA would no doubt decry the creation of philosopher kings, their own arguments may ultimately provide the crowning.

Of course, there are often deep disagreements amongst moral and political philosophers about the values that should inform health policy. But consensus on policy need not elude us in the face

of such disagreement. It is sometimes possible to construct policy around points of convergence between theories that oppose one another in their more comprehensive forms [3]. Those advocating the use of SVPs need to better establish the precise limits of moral and political theory for settling allocation issues. In the absence of such an account, it is not clear whether employing SVPs is warranted. On the other hand, even if we suppose that ethical analysis cannot resolve most allocation problems, it is not clear that there is an adequate justification for appealing to SVPs. We turn to this question in the next section.

### Independent moral justification for societal value preferences

One approach to justifying the use of SVPs would be to show that their use is independently morally justified, rather than being relevant only when moral arguments run out. If there were such a justification, then there would seem to be good *prima facie* reason to believe that SVPs should play a role in health care allocation.

To give an example of how an independent moral justification would function, we return to individual utility preferences. There is a strong moral argument to be made for letting competent adults make decisions about their own well being. This claim is stronger than the idea that individuals are usually the best judges in these matters. That would be a merely practical, and therefore contingent, claim. There are moral reasons to respect persons' choices about their own utility even when the choices that they make are not the ones that we think are best for them. Of course we do not always let people make these decisions for themselves. When the cost to their well being is great we may turn paternalistic in the hope that they will realize later that they were in error about their own best interests. The important point is that we would need to show in any particular case that we are justified in ignoring the *prima facie* moral value of respecting a person's own decision about his or her well being.

What is needed in the case of SVPs, then, is an argument for their independent moral justification. Like the case of respect for individual's decisions about their own utility, this justification would register on the ethical scales even when the

decision generated by the use of SVPs does not maximize QALYs for the society in question (as would the straightforward use of CUA). At this juncture, one might be tempted to claim that the justification for the use of SVPs in the inter-personal case is a rather obvious extension of the justification for the use of an individual's own CUA in the intra-personal case. Just as personal values are relevant to decisions about the distribution of goods at the intra-personal level, social values are relevant to decisions about the distribution of goods at the inter-personal level. After all, it seems trivially true that social values are relevant to decisions that are grounded in questions of such value. As it seems that social values are at least partly at issue in questions about the distribution of scarce health care resources, one might naturally conclude that SVPs are to that same extent relevant in answering these questions.

This conclusion leaves open the extent to which social values should factor into decisions about the distribution of health care resources. Another critic might be tempted to argue that these values per se bear no direct relation to the real issue at stake: namely, the just distribution of scarce health care resources. But we do not want to argue along these lines. We do not reject out of hand that social values are relevant to questions about the distribution of shared but scarce social goods. While we would dispute the claim that these values are the *only* such relevant factor (indeed, we think that the fact that the proponents of SVPs acknowledge their potential moral limitations [1] shows that they would also disagree with such a conclusion), we allow that they may play *an* important role. However, the move from aggregating statistically representative individual preferences about social values to the claim of generating useful social values needs justification. In other words, what is needed is a non-circular way of showing that social values are indeed represented by SVPs.

A clue as to what this justification might be can be found by looking at a remark by Paul Menzel:

One of the strongest responses that a society can make to particular parties who are disadvantaged by a social policy is to note that those parties themselves (or persons relevantly similar to them) have had a major role in the decision-making process that led to the policy in question. . . . CEA's *moral and political* future will be brighter if it can note precisely where in its process influential values have been contributed

by representatives of those who are affected by final decisions [4, our italics].

Following Menzel, one might argue strategically that political support would be stronger for allocation schemes that are perceived as taking into account the preferences of those affected by the outcomes. In so far as this claim is about increased political acceptability, it is an empirical claim. It *might* turn out that the political future of CEA will be 'brighter' if people feel that moral theorists, economists, and health care delivery experts spent considerable time and effort determining which moral principles should modify cost-utility considerations. In any case, the political point would not, without further substantive argument, provide a moral justification for taking SVPs into account in allocating social goods.

However, Menzel means to make a moral claim as well as a political point. In the context of this quote the argument for the moral claim is not entirely apparent, but we might glean some insight from another part of the passage cited above:

If in the process of discerning the critical values involved, CEA has queried people who can reasonably be construed to represent the patients in a category that ends up disadvantaged by resulting allocation, a ready response to critics is available [4].

The moral claim, then, seems to be that patients who are disadvantaged by the allocation have nevertheless had their views represented. Since their views have been adequately represented, they have no strong basis for disagreement with the resulting outcome. But it seems that no one is reasonably construed to *represent* the choices that we would have made if their choices differ from ours. In fact, whether or not we were part of the sampled group, if the aggregated preferences that result are not the ones that we would have chosen, they are not reasonably construed as representative of *our* values.

One might object here, that in the case of collective decision making an aggregated outcome may differ from what we would have chosen, and yet still 'represent' our choice because of the procedure by which the collective decision was determined. Democratic procedures may offer one example of this type of collective 'representation'. In the next section, we will consider whether there is a democratic procedural justification available to proponents of SVPs. Before we consider this

alternative, however, we will need to see what is wrong with the claim that others may represent our choices when they are asked their opinions in our stead.

The problem with Menzel's claim seems to stem from an essential yet subtle ambiguity in the terms 'represent' and 'representative.' There is an important difference between being represented by one's lawyer and being representative of a certain group (say generation X-ers). We think of people as representative of a certain class or group if they share particular characteristics. This is the sense of 'represent' that Menzel is justified in using in the passage above. The problem is that this sense of 'represent' is not the morally weighty sense that he needs to make the argument work. In order for someone to represent someone else in a morally (or legally) significant manner, they must either be that person, or be appointed in some way to represent their interests (as in some cases of artificial persons). When and how a person (or institution) can be truly thought to represent another is not our concern here, although it is a difficult and important question. Our point is simply that 'representing' in a statistical sense (in this case by sharing certain identifiable characteristics), is not 'representing' in any sense that would justify the conclusion offered above.

Menzel acknowledges that the defense he offers *may* not be as strong as the consent of those disadvantaged. Since Menzel claims that the strongest form of moral defense *would be* the consent of the disadvantaged, we should see whether this argument might go through. To consider this claim, we need to turn to a passage from his book, *Strong Medicine*:

Seldom will the actual individuals who get disadvantaged by a quality adjustment have been previously consulted. . . . To most people, though, the fact of the sampling itself is probably not very bothersome. . . . With a good enough initial sample, therefore, it is acceptable to presume people's consent to the trade-offs others choose [5].

So it seems that while Menzel does not want to claim that those disadvantaged by the quality adjustment have given actual consent (unless, of course, they have), there are nevertheless conditions under which we can presume their consent. This claim is initially surprising, since there are well-known theoretical problems with the notion of presumed consent. It seems *prima facie* true that the only way to get consent from person A, is for

A to give his or her consent or for a person who represents (in the morally robust sense) the interests of A, to give consent for A. If an allocation scheme disadvantages A and he or she does not agree to this disadvantage (through his or her own consent or the consent of a proper representative), then it seems that he or she has not given consent, presumed or otherwise.

Menzel attempts to justify the claim that presumed consent, while not as morally weighty as actual consent, nevertheless may serve as sufficient moral justification for a health care allocation plan under some conditions [5]. Of particular relevance for our discussion is the following condition on the use of presumed consent: actual consent would be impossible or prohibitively expensive (where 'prohibitively expensive' is defined by what the person him or herself would (presumably?) view as too costly) [5]. Menzel also points out that any justification relying on presumed consent bears the burden of showing that what is presumed is what people would have actually agreed to (although the difficulty of getting reasonably good evidence about what people would have consented to is not to stand in the way of the general appropriateness of presuming consent) [5].

But do these qualifications have any real bearing on the question of whether presumed consent itself can bear the weight of moral justification? Menzel thinks that the answer to this question is clearly 'yes'. He writes,

If others are careful and do not bypass one's actual consent where it is feasible to consult one. . . . attempts to ascertain one's hypothetical will do not disrespect one's autonomy. We should put the matter even more strongly: if others ignore the possibility of presuming one's prior consent, they assault an important dimension of one's freedom amid the imperfect realities of historical social life [5].

So it seems that, for Menzel, in some cases discerning presumed consent is not only morally justified, it is morally required.

However, the idea that we could ever adequately ascertain a person's 'hypothetical will' in a way that would warrant the application of SVPs to allocation decisions seems problematic. Responses to SVP questions are, after all, not uniform by any stretch of the imagination. To take just one example, in a study cited by Menzel *et al.* of 150 Norwegian politicians 'accountable for health policy at the county level' there was no consensus

on how resources should be divided [1]. When asked to choose an allocation of resources between treatments that would offer a little help for a group with severe illness or treatments that would offer considerable help for a group with moderate illness, 45 per cent chose equal division of resources, 37 per cent gave priority to those with the more severe illness, and 11 per cent gave priority to those with moderate illness [1]. While this is most assuredly evidence that this group generally favours helping the least well off first (even when this will not maximize utility) it could not possibly count as evidence for the presumed consent of any particular politician to such an allocation scheme. Similarly, for any given SVP, one can easily assume that the individual value preferences that go into the overall calculation of the 'societal value' may diverge from one another by wide margins. If this is the case, it is hard to see how the project of accurately presuming consent for the application of SVPs can even get off the ground.

While rationing is necessary, presumed consent based on an appeal to SVPs does not seem to offer adequate moral justification for such rationing. To solidify the argument on this point, we should consider the likely bases for the claim of presumed consent to a health care allocation plan. The two most plausible candidates are rationality (or reasonableness) and values. In other words, we might argue that we could presume consent to an allocation scheme because that scheme is the most rational and hence that all reasonable people would accept it. On the other hand, we might presume consent on the grounds that the allocation accurately reflects the values expressed by those persons affected by the allocation.

Menzel's claim is of the second sort. He argues against idealized versions of presumed consent (such as that represented by an appeal to the 'most rational' allocation) writing, 'However few or many irrationalities of real people are acceptable in actual consent, exactly those could also be acceptable in presumed consent- no more, no fewer' [5]. Moreover, rationing plans that appeal to what reasonable persons would accept would have no need for bothering with determining SVPs. We shall assume, then, that the underlying foundation for Menzel's appeal to presumed consent is the (non-idealized) SVPs of those affected by the allocation.

The crucial difference between using idealized and non-idealized versions of presumed consent is that any actual person disagreeing with the allocation plan cannot be presumed to consent to it in the non-idealized case. In the idealized case, however, such consent might still be presumed since it depends for its validity, not on the real wishes of actual persons, but on the idealized wishes of model persons. Hence it seems that an appeal to non-idealized presumed consent could not justify the application of SVPs in the denial of health care to patients, as we cannot expect that the values of the persons disadvantaged will be uniform.

## Procedural justice

Perhaps the justification we are looking for can be found in a claim about the right procedure to use in allocating health care. Although moral principles might be the best guide in determining what counts as a just distribution of health care resources, it might also be true that the right procedure to use in determining that allocation is reliance on SVPs. To illustrate the point about right procedure, imagine that you have a friend who really does know you better than you know yourself (this probably is not too hard for most of us to imagine). In fact, she is perfectly reliable in judging your best interests, whereas you are only about 80 per cent reliable. If this is true, should you not defer to her opinion in all choices that impact your well being? In fact would it not be irrational to act otherwise? You may disagree; and not necessarily out of irrational self-indulgence. Even though you agree that she is a great person to go to for advice about difficult decisions, you might still think that you should be your own ultimate guide in these matters. You might rightly claim, 'Even if someone else is better than I am at judging where my interests lie, it is better in general that I myself make decisions that impact my own well-being. This claim about the right procedure for making these choices is independent of the goodness of the outcome.

If a claim about just procedure is available, then proponents of the use of SVPs will find an independent reason for incorporating these preferences in health care resource distribution. Perhaps this is the justification that Menzel has

in mind when he claims that taking into account the interests of the disadvantaged (through persons thought to be representative) gives society a response to the objections of the disadvantaged. Perhaps he means to argue that since the interests of the disadvantaged parties were taken into account, the procedure of incorporating SVPs into health care allocation decisions is just. Since the procedure is just, the outcome is justified to that extent (whether or not it represents the most just distribution on its face).

Merely taking into account the interests of those disadvantaged by the outcome does not seem to be a sufficient foundation for a just procedure. To establish a just procedure, we need to consider the way in which these interests are taken into account. In this case, we need a more general account of why it is that an aggregation of statistically representative individual's preferences can itself offer a fair method for the distribution of a morally significant good such as health care. Initially, it might seem that an obvious response is to defer to the justifications of democratic procedures in deciding issues that impact the general society. In fact, Menzel and co-authors seem to rely on a claim like this at one point writing, 'A good moral argument can be made that it is paternalistic or antidemocratic to ignore *any* of a population's preferences about allocation of health care' [1]. Although the only explicit claim here is that it would be antidemocratic to ignore a population's preferences, it seems implied that it is democratic to respect these preferences. In this case, an appeal to democratic procedure is an implied moral justification for the use of SVPs in health care allocation. In what follows, however, we will show that the procedural justification used to support democratic decision-making does not carry over into a support of the use of SVPs.

Why is a democratic procedure a just procedure? One argument is based on the notion of the fundamental equality of all members of a political community. When a political community encompasses members with very diverse sets of values, it is often best to allow each to flourish as they see fit (where this does not infringe on other's flourishing). However, choices must often be made for the political community as a whole. Allocation of resources controlled by the government offers one such example. Since the democratic procedure of one vote for each member of the community gives everyone an equal opportunity to influence the outcome, one could argue that this offers a fair

method for deciding such controversial issues. Of course, for any given vote, no one can guarantee that things will go their way. The outcome is a collective decision that may end up pleasing no one in particular. The point, however, is that if the procedure is itself just, the outcome is justified to that extent. So, in the case of a vote that does not go my way, I must still abide by the outcome once I have accepted the legitimacy of the procedure.

SVPs are established by polling statistically representative samples of the relevant (that is, either directly or potentially affected) population. There is much concern in the literature about 'whom' to ask in establishing SVPs. In particular, researchers debate whether to sample members of the general public (as potential patients) or members of the particular diagnostic classes that would be affected by any particular resource allocation decision. As one can imagine, SVPs may end up looking very different depending on what group one chooses as the 'relevant' group for the purpose of preference solicitation. Eric Nord, Peter Ubel, Paul Menzel and others recommend using a two-step model. In the first step actual patients are asked to determine utility measures for their particular conditions. The second step is to use these utility measures in asking a 'representative sample' of the general public to determine preferences for hypothetical allocation schemes. These preferences could then be used as preference (or as Nord prefers, 'fairness') weights in QALY analyses, or would be otherwise incorporated into CEA [6,7].

While it is certainly important methodologically to determine whom to ask in assessing societal preferences for distributive schemes, this debate relies on an assumption that is not itself fundamentally democratic. That assumption is that it is possible to substitute a statistically representative sample of preferences for each potentially affected person's equal opportunity to register their opinion by voting. This substitution, however, is a centrepiece of the methodology involved in eliciting public preferences. As Menzel writes, 'In reality we question only a sample of people to establish our basic map of proportional quality ratings, and we question only a relatively small sample of patients to place them on the spectrum of health states' [5]. While this method may lead to the same outcome as a general vote in terms of 'winning' preferences, the democratic procedural justification is lost.

The following example might help in clarifying this point. Important elections and ballot referendums are usually accompanied by extensive public opinion polling. As a matter of empirical fact, these polls are almost unerring in the accuracy with which they predict the results of the actual voting process. So what would be wrong with substituting such polls for the tedious (and expensive) practice of holding actual votes? The answer should be obvious. Even if the results are comparable, only an actual vote carries with it the procedural justification that is essential to democracy. This would be the case even if, for example, the actual ballots were lost and so it was impossible to determine the outcome of the vote. In other words, even if sample polling were the only way to determine a winner, this method would not carry with it the democratic procedural justification.

We are not, of course, claiming that a democratic procedural justification requires that literally everyone must have the option to register his or her opinion on every decision. Rather, such a justification requires only that everyone who is relevantly affected by the decision (and who meets other legally recognized criteria) is given the option to have his or her voice heard. So, to put this in the context of SVPs, only those who are potentially affected by the particular allocation issue at hand would need to be consulted. For example, in cases where the allocation scheme is limited to a particular area or group of patients, only members of that constituency would need to be consulted. The problem is just that the methodology in determining SVPs is to sample members of this morally relevant group of persons with the assumption that the preferences of *these* persons are statistically representative of the preferences of others in that same group. However accurate such a method may be, it is not to be mistaken for a democratic procedure.

We do not mean to claim that supporters of SVPs have somehow mistaken their methodology for the equivalence of a vote. In fact, they are quite specific on this point. Ubel writes, for instance, 'Because governments and insurers, who often play a large role in determining which services will be available, cannot rely solely on voting booths or free markets to tell them what patients want, they must have some other way to find out' [2]. This 'other way' of measuring what is wanted, is 'sample polling' SVPs.

We do, however, mean to argue that SVPs, determined as they are by a relatively small sample of the affected population, cannot be justified by appeal to democratic procedure. Since this aspect of the methodology is essential to the project, we are skeptical of the claims that using societal preferences is more democratic than not using such preferences. Menzel asks, 'Are the rights of individuals going to be represented better in the cumbersome political process than they are via the consent foundation of a sensitively elicited quality-adjustment map?' [5]. Contrary to Menzel's own answer, we think the answer may be 'yes.'

## Conclusion

In this paper we have argued, first, that proponents of the use of SVPs need to be clear about the relationship between these values and moral principles. In particular, once moral principles are accepted as ruling out some SVPs, an account is needed of why we need appeal to SVPs at all where moral principles might do the same work. Second, we considered whether an independent moral justification might underwrite the use of SVPs. In various places in the literature the notions of representation, presumed consent, and democratic decision making appear to be invoked as candidates for fulfilling this justificatory role. We have raised some problems with each of these justifications in the hopes of eliciting a more comprehensive proposal from the proponents of SVPs.

The move to incorporate SVPs into CEA appears to be inspired in part by a commitment to justice and fairness in the allocation of scarce health care resources. In so far as the authors we have discussed are motivated by the desire to find a practical way to incorporate such social values into health care allocation schemes, we have a great deal of sympathy for their work. While we do not want to deny that social values may play a legitimate role in formulating allocation policies, we stick fast to the claim that the use of SVPs needs to be justified. Without this justification, we worry that the move from CUAs to SVPs is just a move from an implicit reliance on a questionable utilitarian standard to an explicit reliance on the popularity of moral values.



## Acknowledgements

The authors would like to thank Ruth Faden, Marc Spindelman and PJ Ivanhoe for their extremely helpful comments on earlier drafts of this paper. We would also like to thank our anonymous reviewers whose insightful and challenging comments were invaluable.

## References

1. Menzel PT, Gold MR, Nord E, Pinto-Prades J, Richardson J, Ubel P. Toward a broader view of values in cost-effectiveness analysis of health. *Hastings Cent Rep* 1999; **29**: 7–15.
2. Ubel P. *Pricing Life*. MIT Press: Massachusetts, 2000; 94.
3. Gutmann A, Thompson D. *Democracy and Disagreement*. Harvard University Press: Massachusetts, 1996.
4. Menzel PT. How should what economists call social values be measured? *J Ethics* 1999; **3**: 249–273.
5. Menzel PT. *Strong Medicine*. Oxford University Press: New York, 1990; 22–35, 89–91.
6. Nord E, Pinto JL, Richardson J, Menzel P, Ubel P. Incorporating societal concerns for fairness in numerical valuations of health programmes. *Health Econ* 1999; **8**: 25–39.
7. Ubel P, Richardson J, Menzel P. Societal Value, The person trade-off, and the dilemma of whose values to measure for cost-effectiveness analysis. *Health Econ* 2000; **9**: 127–136.