

## Impartiality and Disability Discrimination

**ABSTRACT.** Cost-effectiveness analysis is the standard analytical tool for evaluating the aggregate health benefits of treatments and health programs. According to a common objection, however, its use may lead to unfair discrimination against people with disabilities. Since the disability discrimination objection is seldom articulated in a precise way, I first provide a formulation that avoids some implausible implications. Then I turn to the standard defense of cost-effectiveness analysis and argue that it does not succeed. But this does not settle the question of whether the use of cost-effectiveness analysis leads to unfair discrimination. Rather, it shows that the controversy should be approached in a different way. Thus, I conclude by outlining an alternative strategy for answering the question.

Cost-effectiveness analysis is the standard analytical tool for evaluating the aggregate health benefits of treatments, interventions, or health programs. It works by comparing the ratio of costs and benefits of different alternatives. The lower the ratio, the more effective the treatment, intervention, or program. The use of cost-effectiveness analysis can ensure that scarce health care resources are allocated in a way that maximizes the satisfaction of health needs. According to a common objection, however, the use of cost-effectiveness analysis for setting priorities in the allocation of health care resources may lead to unfair discrimination against people with disabilities.

The aim of the first part of this paper is to clarify the disability discrimination objection. To set the scene, I provide a brief background to cost-effectiveness analysis in section 2. Since the discrimination objection is seldom articulated in a precise way, I provide a formulation that I believe captures the main worry behind it in section 3. Section 4 distinguishes this formulation from some alternatives.

In the rest of the paper, I address the moral basis of the discrimination objection—that is, the question of whether the discrimination against

people with disabilities that the use of cost-effectiveness analysis might lead to is *unfair*. Arguments on both sides of the controversy appeal to the requirements of *impartiality*. In particular, defenders of cost-effectiveness analysis argue that the measures of health benefit they use are impartial and hence disability discrimination is not unfair. In section 5, I show that their arguments do not succeed. This, however, does not settle the question. Instead, it shows that the controversy should be approached in a different way—namely, by providing an alternative account of the connection between impartiality and fairness. Thus, rather than directly answering the question whether disability discrimination is unfair, I conclude by outlining this alternative strategy for evaluating the objection.

## 2. QUALITY-ADJUSTED MEASURES

Cost-effectiveness analysis ranks alternatives in terms of their cost per unit of outcome. The lower the cost per unit, the more effective the treatment or program is. For a simple illustration, consider two health programs. Program 1 would save 1,000 lives for a total cost of \$5 million; program 2 would save 1,200 lives for a total cost of \$6,120,000. On the first program, the cost of one life saved is \$5,000; on the second program, the cost of one life saved is \$5,100. Therefore, even though program 1 would overall save fewer lives, it is more cost effective than program 2.

In most cases, saved life is not an adequate unit of outcome, since many treatments and programs improve the health of their beneficiaries rather than extend their lives, or they result in some combination of life extension and health improvement. In order to evaluate their effectiveness, a measure is needed that can compare treatments and programs whose outcomes are different. For instance, one treatment may restore mobility, while another may cure respiratory problems. An outcome measure must be able to compare improvements in mobility and the alleviation of respiratory symptoms.

In order to be able to make these comparisons, an outcome measure has to appeal to the value of health states. By “health state,” I mean an outcome that can be characterized by a set of symptoms, disease stage, treatment status, impairment, level of functioning, or any combination of these. In short, a health state is a description of the person’s health, and the required outcome measure must *evaluate* the person’s health. The most common—although not the only—way to do that is to appeal to the contribution of health states to the quality of life of the people who are in those health states. If curing respiratory problems contributes more than

improvements in mobility to the quality of life, this should be reflected in cost-effectiveness analysis.

Moreover, treatments and health programs often alter people's health state for an extended period of time, or they affect not only the quality of their life but its duration as well. It is thus necessary to include in an outcome measure effects on the quantity of life in addition to effects on its quality. Such combined measures are called "quality-adjusted measures."

The two most commonly used quality-adjusted measures are quality-adjusted life years (QALYs) and disability-adjusted life years (DALYs). On these measures, the value of a year of life spent in some health state is determined by the quality of life associated with that health state. Both measures assign values between 0 and 1 to health states. On the QALY measure, full health is represented by the value of 1 and death is represented by the value of 0; on the DALY measure, 1 represents death and 0 represents full health. Health states that are ranked lower than full health have values between 1 and 0 in decreasing order in QALYs and in increasing order in DALYs. For instance, if the quality-adjustment factor associated with a health state is 0.8, one year spent in that health state has a QALY value of 0.8, two years have a value of 1.6, and so on. On the DALY measure, if a health state has a disability weight of 0.2, then the burden of a year spent in that health state is 0.2, the burden of two years is 0.4, and so on.<sup>1</sup> The difference in representation is due to the interpretation of QALYs and DALYs: QALYs are intended to be a measure of *health benefit*, whereas DALYs are intended to be a measure of *disease burden*. Hence greater QALY and smaller DALY scores represent better outcomes.

Quality-adjusted life years are used extensively as an outcome measure in cost-effectiveness analysis.<sup>2</sup> Disability-adjusted life years have been used as a summary measure of population health in global burden-of-disease studies. These studies attempt to report the harm from premature mortality and disability from disease and injury for different regions and countries in terms of the DALYs associated with different conditions. Disability-adjusted life years are also recommended as an outcome measure for priority setting.<sup>3</sup>

### 3. DISABILITY DISCRIMINATION

Cost-effectiveness analysis has always been controversial. One of the deepest controversies concerns its appropriateness in relation to disabilities and chronic health conditions: it has often been argued that the use of

cost-effectiveness analysis unfairly discriminates against people with disabilities or chronic health conditions (henceforth simply “disabilities”).<sup>4</sup>

For the discussion of some of the putative forms of unfair discrimination, it will be helpful to distinguish settings in which resources can be used to save or extend lives and settings in which resources can be used to improve health. This distinction is drawn for analytical purposes only, because it is often impossible to make with any precision in practice. In many cases, saving the lives of some patients must be balanced with health improvements for other patients. More importantly, the distinction ignores uncertainty: many conditions have mortality risks, and thus it is frequently in principle impossible to determine which setting we are facing.

Consider the life-saving setting first. Suppose that a scarce medical resource, enough to save the life of only one patient, must be allocated between patients A, B, and C, who would all otherwise die. The allocation decision is made on the basis of some quality-adjusted measure. I assume for the moment that the treatment costs are invariant for the three patients and also that the patients are alike in all relevant respects, including their age and sex. The only difference between them is that two of the patients, B and C, have prior disabilities that influence their prospects. Thus, patient A has a high quality of life and normal life expectancy; patient B has the same life expectancy but a lower quality of life; and patient C, has a lower life expectancy than A but the same quality of life. More benefit would be produced by saving the life of A rather than that of B or C on account of her higher quality and greater expectation of life. But if she is chosen, patients B and C are discriminated against. Since this is exclusively due to their lower *ex ante* quality or expectation of life, the discrimination is systematic.

For illustration, consider table (i) on the left in the figure. To keep matters simple, only the prospects of two patients are included. The numbers are the scores on some quality-adjusted measure that represents the value of the outcomes for the patients; greater numbers indicate better outcomes. You can give the resource to A, in which case A will survive for a year in full health, as shown by the number 1 in A’s row in the first column; in this case B will die, as shown by the number 0 in B’s row in the first column. Or you can give a resource to B, in which case A will die and B will survive, as shown in the second column. But since B has a prior disability due to which her quality of life, even if she survives for a full year, will only be 0.8, more benefit would be created if A were saved. Because the lower score for B is the product of combining the quality and

TABLE 1.

(i) <i>Life-saving setting</i>			(ii) <i>Health-improvement setting</i>		
	Give resource to			Provide treatment to	
	A	B		A	B
A	1	0	A	1	0.8
B	0	0.8	B	0.4	0.7

expectation of life, the same result would obtain if B's disability produced a lower life expectancy rather than a lower quality of life. The prior disability diminishes the patient's capacity to benefit—therefore the overall benefit of saving the nondisabled will always be greater.

A parallel example can be constructed for the more general setting of health improvement. Suppose that A, B, and C suffer from a condition that can be cured, but not all of them can be treated with the available resources. Although they have the same health care needs and they are identical in all relevant respects, their prospects are different. A can be restored to full health. B, however, has a prior disability that makes her relative quality-of-life gain from the treatment smaller—she cannot be restored to full health. C has another disability that shortens her life expectancy even though the treatment restores her to full health. The relative improvement from the treatment is lower in this case as well. If the allocation decision is made on the basis of quality-adjusted measures, A will be guaranteed treatment, while B or C (or both) will have to be left untreated. This sort of discrimination, however, is systematic *only if* the prior disability affects the improvement that can result from the treatment in such a way that it is systematically less than the improvement that can result from the same treatment for those without the disability.

To see this point, consider table (ii) on the right in the figure. If you treat A, the person without the prior disability, she can be restored to full health—from a score of 0.8 to 1. B will remain at 0.4—a very bad health state because of the combined burdens of the prior disability and the present condition. Her health burden could be greatly reduced by treating her instead, restoring her to a score of 0.7. So, in this example, the relative improvement, as well as the overall benefit, would be greater by giving the treatment to B. This choice would be more cost effective: it would result in a greater health improvement for the given cost. Opponents of the use of cost-effectiveness analysis, however, argue that such cases are atypical:

prior disabilities tend to make treatments more complex, expensive, or uncertain.<sup>5</sup>

It is important to point out that opponents of the use of cost-effectiveness analysis need not deny the relevance of effectiveness considerations in all cases. They can concede that it is unobjectionable if it is used for choosing between alternative treatments or health programs for the same patient or the same group of patients. As a matter of fact, the majority of cost-effectiveness studies are conducted for these purposes. What the opponents find problematic, however, is the use of cost-effectiveness analysis for choosing *which* patient, group of patients, or condition to treat. But even with this qualification, the objection poses a serious challenge for the use of cost-effectiveness analysis in health care resource allocation, since macro-level allocation decisions cannot avoid prioritizing condition-treatment pairs, which unavoidably gives precedence to some groups of patients. Even though cost-effectiveness analysis purports to be simply a measure of aggregate health benefits, it unavoidably disadvantages people with a diminished capacity to benefit.

Still, what is wrong with that? After all, it is desirable to allocate scarce resources in a way that can do the most good. It is not sufficient for the proponents of the discrimination objection to point out that the use of cost-effectiveness analysis may disadvantage those who have a diminished capacity to benefit: they must show that these disadvantages are *unfair*. For instance, they may argue that people with a diminished capacity to benefit belong to vulnerable groups. Thus, the use of cost-effectiveness analysis (at least if no other considerations are taken into account) systematically disadvantages vulnerable groups of the population.<sup>6</sup>

The difficulty for the proponents of this argument is to formulate the discrimination objection in a way that does not imply that the use of effectiveness considerations will be unfairly discriminatory even with regard to groups that do not seem to be disadvantaged and that we would not normally think are vulnerable in any relevant way.

Consider, for instance, the differences in life expectancy between the sexes. Men usually have a lower life expectancy than women, which is partly a result of their riskier lifestyles and partly a result of their biologically determined lower survival potential. Yet most people would probably not consider cost-effectiveness analysis unfairly discriminatory with regard to men, even if, other things being equal, fewer QALYs (or more DALYs) are produced by choosing treatments and programs that save the lives or improve the health of men. To put it more cautiously: the lack of

discussion of this form of discrimination in health economics and public health ethics may be an indication that sex discrimination is usually not considered unfair.<sup>7</sup>

Another illustration is the elderly. Older people usually have lower life expectancy or worse quality of life (and often both) than members of younger generations. Although the elderly may be a vulnerable group in many—perhaps most—societies, numerous philosophers, health economists, and respondents in empirical studies agree that they are not unfairly discriminated against if they receive lower priority due to their worse prospects.<sup>8</sup>

Age and sex are personal characteristics that affect a person's *ex post* quality and expectation of life.<sup>9</sup> Other things being equal, outcomes are worse on quality-adjusted measures for men and the elderly even if their *ex ante* health states are as good (or bad) as those of others. Yet to most people it does not seem unfair if cost-effectiveness analysis assigns different priorities due to these characteristics. Therefore the discrimination objection must be formulated in a way that can distinguish between relevant and irrelevant personal characteristics—otherwise any disadvantage owing to the use of effectiveness considerations would be regarded as the result of unfair discrimination. This conclusion is surely too strong: even though effectiveness is not the only relevant consideration in the allocation of health care resources, it is one of the relevant considerations on any plausible view.

This suggests that an adequate formulation of the discrimination objection should focus on the moral relevance of personal characteristics. On such a formulation, cost-effectiveness analysis is morally problematic because the measures of benefit it employs include characteristics that are irrelevant. That is, quality-adjusted measures violate certain *moral constraints* on the information that should be taken into account in evaluating benefits and burdens. A proponent of the disability discrimination objection has to make the case that while age and sex may be relevant personal characteristics, prior disability is not.

The root of the problem is that people with disabilities have a characteristic that tends to make the benefit they can get from a resource smaller than the benefit that those who do not have the characteristic can get from that same resource. Prior disabilities tend to reduce a person's capacity to benefit. Hence cost-effectiveness analysis discriminates between persons: people with disabilities who are already unfortunate may systematically be disadvantaged by receiving lower priority than their more fortunate peers.<sup>10</sup>

## 4. THE VALUE OF LIFE AND THE WORTH OF A PERSON

Those who worry about disabilities and the use of cost-effectiveness analysis sometimes seem to make a stronger point than that people with prior disabilities are unfairly disadvantaged by it. Consider the following:

A severely disabled person will have a much lower QALY ranking than a person in full health and therefore each year they live will have a lower (normative) quality of life ranking. But does this mean that the former person's life is less worth living than the latter's; is it thus *worth less*? This goes against a profound belief, both spiritual and secular, that all lives are equally valuable. (Phillips 2006, 47, his emphasis)

How should we understand the idea that all lives are equally valuable? There are at least two senses in which the phrase "the value of life" is used. In one sense, the value of life is the value of a person's life *for that person*. The value of life in this interpretation is determined by the contents and features of that life. It depends on how well life goes for the person whose life it is—that is, on the well-being of the person. Since each life has different content, the lives of different people go differently well for them—depending on what happens in that life and on the characteristics of the person that affect what can happen in that life. The value of life in this sense varies from person to person; not all lives are equally valuable.

In another sense, the value of a life does not depend on how well it goes for the person whose life it is. Rather, life has value in virtue of certain psychological capacities. The effects of these capacities for the content of life, however, are irrelevant to this sort of value. What matters is their presence. The value of life in this sense is the *worth of the person*. Since it is generally agreed that the capacities in question are necessary for personhood, and their value does not depend on the degree to which they characterize a person, all persons have equal value or worth in this sense.<sup>11</sup>

Perhaps the argument of the opponents of cost-effectiveness analysis should be interpreted in these terms. What is morally important for just resource allocation is not the value of people's lives, but their worth as persons. The appropriate attitude toward beings with this worth is *respect*, and since each person has equal worth, each person is owed equal respect. It is impermissible, then, to differentiate between persons on the basis of their personal characteristics.

Thus, in the present context, the argument would be that quality-adjusted measures are objectionable because they disregard the worth of persons: when scarce resources must be distributed, setting priorities on

the basis of capacity to benefit fails to respect persons. This is because although equal respect for persons does not imply treating persons equally, it does imply *treating persons as equals*. Distinguishing between persons by their personal characteristics and hence by their capacity to benefit, however, is incompatible with treating persons as equals. Therefore, the use of cost-effectiveness analysis for setting priorities is morally impermissible. The discrimination that results is a consequence of failing to respect persons, and it is unfair.

This argument, however, faces a dilemma. On the one hand, if treating people as equals indeed implies that *no* personal characteristic can be taken into account in health care resource allocation, then the objection prohibits taking into account *any* effectiveness considerations. If no personal characteristic is relevant, then a person's capacity to benefit is irrelevant. Consequently, effectiveness is not a relevant consideration in priority setting.

On the other hand, if treating people as equals is not incompatible with taking into account *some* personal characteristics, then a person's capacity to benefit becomes relevant insofar as those characteristics affect her capacity to benefit. And then it remains an open question which characteristics should be included and which should be excluded.

Rather than trying to resolve the dilemma, I set this argument aside. I focus instead on the formulation of the objection that relies on the value of life in the former sense and accepts that taking certain personal characteristics into account is not incompatible with treating persons as equals. Thus, it avoids the implausible implication that all effectiveness considerations are excluded from resource allocation decisions. This formulation agrees that it is morally acceptable to consider capacity to benefit but rejects prior disability as a relevant characteristic. In my view, this is the most persuasive formulation of the discrimination objection.<sup>12</sup>

##### 5. DISABILITY AND THE VEIL OF IGNORANCE

The debate about cost-effectiveness analysis and disability discrimination has focused on the use of quality-adjusted measures like QALYs and DALYs. The arguments on both sides are couched in terms of *equality* and *impartiality*. Defenders of quality-adjusted measures claim that these measures are "egalitarian" because all units of benefit are counted equally regardless of differences in wealth, race, or other morally irrelevant factors. For instance, it is often pointed out that a QALY gain has the same value regardless of who receives it. Thus, everyone's interests are given

equal weight.<sup>13</sup> This is considered a requirement of impartiality. Thus, John Harris (1988, 83) argues that “the whole burden of the defence of QALYs comes down to whether or not they operate impartially.” His opponents agree:

The QALY approach is often criticized for being inegalitarian, but . . . allocating resources in a way that maximizes human well-being, giving equal weight to each person’s interests in the process, is an expression of our refusal to allocate resources in a biased way, and reflects our commitment to impartiality. (McKie et al. 1998, 90–91)

This quote is from a joint work of two health economists (John McKie and Jeff Richardson) and two philosophers (Peter Singer and Helga Kuhse) who offer a defense of quality-adjusted measures—in particular, the QALY—by an argument based on the device of the veil of ignorance. They ask how rational, self-interested persons would choose to allocate health care resources behind the veil, where they do not know their own identity. One way the decision makers might reason is this. They imagine a situation in which there is enough life-saving resource for only one of two patients. The patient who does not get the resource will die. They also suppose that one of the patients has a stronger interest in continuing to live, because she can benefit more from being saved. A person’s interest in continuing to live, on this view, is determined by that person’s capacity to benefit.

The decision makers behind the veil know that they may end up being either one of these patients once the veil is lifted. Since they are motivated by self-interest, they want to maximize their potential gain, and thus it would not be unreasonable for them to ask what distribution method would best promote their interests *given that they do not end up the person who is not saved*. As McKie and his coauthors put it, “To maximize the satisfaction of their own interests, it would seem, rational egoists would choose a system that gives preference to saving life when it is most in the interests of the person whose life is saved” (McKie et al. 1998, 105). Since quality-adjusted measures reflect the strength of a person’s interest in continuing to live, the decision makers would accept some such measure as the correct outcome measure for the purposes of resource allocation. Consequently, the use of cost-effectiveness analysis can be defended by an appeal to the rational choice in a situation that satisfies the requirements of impartiality.

Of course, it is not uncontroversial which principle of resource allocation rational decision makers would choose behind the veil of ignorance.

Rawls (1971) himself famously argued that the decision makers should make their choice using the maximin rule, arriving at a principle of resource allocation according to which inequalities are justified only if they work to the advantage of the worst off. McKie and his coauthors disagree. More importantly, however, they point out that in the present example, once the veil is lifted, there is no person in the worst-off position, since the person who is not saved will die. The maximin rule cannot be applied. Even if you would prefer using maximin when the alternative is a benefit-maximizing rule, there seems to be no choice in these circumstances but to accept the benefit-maximizing rule for making the choice behind the veil.<sup>14</sup>

There is, however, a problem with this argument. McKie and his coauthors only discuss the life-saving setting; they do not extend the veil of ignorance argument to the distribution of resources for health improvements in general.

Consider a similar choice behind the veil of ignorance except that a scarce resource is to be used for treatment that results in health improvement rather than life saving. There are two patients who suffer from the same condition, but only one of them can be treated. One of the patients can be restored to full health; the other, because of a prior disability, has less potential for health improvement—treating her would lead to a smaller QALY gain or DALY reduction. The decision makers behind the veil do not know which patient they might be. Would it be reasonable for them to focus on the question of what distribution method would best promote their interests *given that they do not end up the person who is not treated*? It does not seem so. What made the parallel question (arguably) not unreasonable in the life-saving setting was that if you end up the patient who is not saved, no further harm can come to you: giving weight to the benefits of the patient who is not saved would not make a difference to that patient.

In the life-saving setting, it is perhaps not unreasonable to focus on the interests of the patient with the greater capacity to benefit. But this rationale is not available in other settings. In the case of health improvements, there does not seem to be any reason to disregard the quality of life prospects of the patient who has a diminished capacity to benefit—especially if she would still have a reasonably good *ex post* quality and expectation of life.

The point can be put differently. Suppose you think that maximin is an attractive rule to use behind the veil of ignorance. Nonetheless, you might agree that it cannot be applied in the life-saving setting, since there is no one in the worst-off position. But that argument will not move you in the

health-improvement setting, because there is a person who *will* be in the worst-off position. The maximin rule can be applied.

To be sure, in certain circumstances the argument given by McKie and his coauthors might apply in the health-improvement setting as well. Consider again table (i) on the left in the figure from section 3. The table illustrates two alternatives: giving the resource to A and giving the resource to B. In the first case, the resulting state of affairs will be  $x = (1, 0)$ , where the first number represents the outcome for A and the second number represents the outcome for B. If the resource is given to B, the resulting state of affairs will be  $y = (0, 0.8)$ . By convention, a QALY score of 0 stands for death. B dies in the state of affairs represented by  $x$ , and A dies in  $y$ .

Now one might reconstruct the argument the following way. Since the decision makers do not know whether they will occupy the position of A or B once the veil is lifted, they hold that  $y = (0, 0.8)$  is just as good as another state of affairs,  $y' = (0.8, 0)$ . This might be interpreted as an expression of impartiality. Since  $x$  is better than  $y'$ —there is no one for whom it is worse, and there is someone for whom it is better—and  $y'$  is just as good as  $y$ , it follows that  $x$  is better than  $y$ .<sup>15</sup>

If this is an adequate reconstruction of the argument, it seems to work in some cases when the choice behind the veil does not concern life saving. For suppose that  $x = (1, 0.5)$  and  $y = (0.5, 0.8)$ . In this case,  $y$  is just as good as  $y' = (0.8, 0.5)$ . Again,  $x$  is better than  $y'$ , and therefore it is better than  $y$ . The decision makers behind the veil would prefer to give the resource to the first person, A, who has the greater capacity to benefit. But B would not die. She would simply be left untreated.

But even on this reconstruction the problem remains: the argument cannot be generalized. It works only as long as either the worst-off positions of any two states of affairs are equally good or the best-off position in one of them is not better than the worst-off position in the other. In other cases, the argument is inconclusive. For instance, consider table (ii) on the right in the figure from section 3. In this situation, the two states of affairs are  $x = (1, 0.4)$  and  $y = (0.8, 0.7)$ . Neither of these states of affairs can be shown to be better by the procedure I've just described. For consider another state of affairs,  $y' = (0.7, 0.8)$ .  $y$  is just as good as  $y'$ ; but it is not true that there is no one for whom  $x$  is worse, since even though it is better for the first person, it is worse for the second person. In this example, a person using this procedure would not choose  $x$  behind the veil. In fact, she would not be able to compare these alternatives.<sup>16</sup>

It is crucial to the success of the veil of ignorance argument that it be generalizable. The allocation of health care resources is concerned with all health burdens and benefits. No general conclusions can be drawn from pure life and death examples (or other special cases). As I said in section 3, the distinction between settings in which resources can be used to save lives and settings in which resources can be used to improve health is helpful only for analytical purposes. Often, resources must be allocated between treatments and health programs whose outcomes include some combination of health improvement and life extension. In other cases, because of the uncertainties involved in utilizing resources, it is unclear which setting we are facing. Furthermore, in the case of large-scale resource allocation decisions for which cost-effectiveness analysis is suitable, the conditions that characterize the situations in which the resources will be used are usually not known in advance. It is a mistake to choose an example with special features as the paradigm case of health care resource allocation and then proceed on the basis of arguments that depend on those features. By focusing on life-saving settings, McKie and his coauthors do just that.

#### 6. IMPARTIALITY AND FAIRNESS

The veil of ignorance argument does not succeed. But this does not show that the use of cost-effectiveness analysis for the allocation of health care resources leads to unfair discrimination against people with disabilities. It does not settle the matter either way. Indeed, I doubt that the approach taken in the debate can settle the matter at all. Let me explain.

Behind the veil of ignorance, the fairness of the choice is guaranteed by the rationality and impartiality of the decision makers. The choice is *rational*, since the decision makers are assumed to aim to secure the best outcome for themselves. The choice is *impartial*, since the veil of ignorance ensures that it is not distorted by the decision makers' own interests or biased toward those they care about. Impartiality is achieved by denying the decision makers knowledge of their identity, personal characteristics, and the circumstances they will find themselves in once the veil is lifted. This way, they are forced to take into account the interests of all affected parties. If they do not know the identity of the persons whom their choice affects, they cannot be influenced by which of them are harmed or benefited.

On this approach, impartiality is ensured by excluding information. Once this feature is in place, all you have to do is to figure out how to make the appropriate choice. Hence the problem reduces to deciding

which decision rule to use—that is, finding the uniquely rational decision-making rule. Famously, this is the approach taken by Rawls and some of his utilitarian opponents.<sup>17</sup> This is also the approach taken by McKie and his coauthors. For them, fairness is a matter of the rational and impartial *selection* of a rule rather than its application.

Not surprisingly, there has been little agreement on what fairness requires on this approach. Some philosophers have argued for rules that ensure that the interests of the worst off are given priority. Thus, the harms and benefits that go to these persons should have more weight. Others reject this view. They argue that impartiality consists in giving *equal* weight to the harms and benefits of all of those whose interests are considered. This is how fairness is achieved. Some philosophers even argue that any view that gives more weight to the interests of the worst off would violate impartiality and lead to unfairness. It is hard to see how these disagreements could be resolved.

Another problem with this approach is that it provides no criteria for deciding what pieces of information should be concealed by the veil and what pieces of information should be available to the decision makers. How “thick” should the veil be? For instance, in a choice like the one described by McKie and his coauthors, should information on the prevalence of disability be available to the decision makers? If that information is available, it might make a difference to the principles of health care resource allocation that they choose. Although the decision makers do not know whether they themselves are disabled, they might choose a rule that minimizes discrimination against people with disabilities if the prevalence of disability is high. But, if the prevalence is low, they might risk choosing a benefit-maximizing rule. Perhaps one might want to argue that information such as the prevalence of disability should not be available. But what arguments can be made to support this suggestion? After all, such information is compatible with not knowing your identity, and it is relevant to the decision makers behind the veil, since it can make a difference to their choice. Nonetheless, would the availability of such information be a source of bias? I simply don’t know how to decide whether having this information is compatible with impartiality. This approach does not tell us how to answer such questions.

There is, however, a different and more promising approach by which we can evaluate the discrimination objection. It begins from a different view of the connection between impartiality and fairness. On this view, impartiality is not a matter of ignorance, and fairness is not achieved by

identifying a uniquely rational decision-making rule. Rather, fairness requires the *impartial application* of a rule, which means applying the rule consistently in each relevant situation and with regard to each person who is affected, taking into account only the morally relevant features of the situation and the affected persons. These are, however, only necessary conditions of fair choice. For it is possible to apply an unfair rule impartially: a rule that tells you never to keep your promises if doing so causes inconvenience for you can be applied consistently and impartially, but this does not make the breaking of promises fair. A fair choice is based on the impartial application of an appropriate rule.<sup>18</sup>

Thus, on this approach, the fairness of a choice depends on at least two components. First, the rule that is used to arrive at the choice must itself be morally justified. Second, the rule must be applied impartially. Thus, fairness is a matter of the rational and impartial *application* of a rule rather than its selection.

Impartiality itself involves more than merely applying a rule consistently and without being influenced by your own interests or by considerations of who is harmed or benefited by a choice. It also involves avoiding being influenced by the morally irrelevant features of the situation and taking into account all of its morally relevant features. This is separate from applying the rule, in the sense that it raises questions about which features of the situation and which characteristics of the affected persons can be taken into account.<sup>19</sup>

Proponents of the disability discrimination objection argue that the use of cost-effectiveness analysis leads to unfair discrimination against people with disabilities. The way I have formulated their objection, they do not object to all use of cost-effectiveness considerations in health care resource allocation. They allow that it is morally acceptable to take into account capacity to benefit, and they allow that taking into account certain personal characteristics does not lead to unfair discrimination. What they do reject, however, is the moral relevance of prior disability. Prior disability is not a relevant personal characteristic. Its impact on a person's capacity to benefit should not be taken into account.

In other words, the disagreement is not about which decision making rule should be used in health care resource allocation but about how that rule can be applied without leading to unfairness. Hence the controversy cannot be adjudicated on the basis of the veil of ignorance argument. Proponents of the discrimination objection should not object to using a benefit-maximizing rule. After all, as I showed with the example of table

(ii) in the figure from section 3, a disability often increases a person's capacity to benefit. Surely, no one should object to giving priority to people with disabilities in such cases. So proponents of the objection can agree that effectiveness is a relevant consideration in the allocation of health care resources, even if it is not the only relevant consideration. But they insist that fairness and impartiality matter too.

What about the systematic discrimination that results from the use of cost-effectiveness analysis in life-saving settings, as illustrated by table (i)? There are several proposals for avoiding unfairness in such situations. Perhaps this is a case in which fairness directly conflicts with benefit maximization, and there is some further moral principle that overrides benefit maximization. For instance, that principle may require equal chances of being saved for all affected persons, regardless of disability status. Perhaps in pure life-saving settings, effectiveness considerations are indeed irrelevant.

Alternatively, consider a proposal from health economics. It has been suggested that for cost-effectiveness calculations in life-saving settings, we should assume that a person with a prior disability would be restored to full health after treatment—that is, a year of life gained by a disabled or chronically ill patient should be always given a QALY score of 1, irrespective of the quality of life in that year (at least as long as it is judged to be worth living by the patient). This amounts to completely disregarding the impact of the prior disability on the person's capacity to benefit.<sup>20</sup>

One problem with this proposal is that it assumes that there is a clear-cut distinction between life-saving and health-improvement settings. As I have argued, the distinction is often unclear in practice. Another problem is that if health benefits are calculated this way, they will be systematically overestimated. Marginal benefits for some patients could outweigh substantial benefits for other patients, and that in itself may be a form of unfairness. If you have to decide whether to save the life of a chronically ill 75-year-old or to save the life of an otherwise healthy 20-year-old, you might be hesitant to ignore completely their different capacity to benefit. Also, the fact that a person's chronic condition has developed into a life-threatening illness, such that her life could only be extended for a short time with low quality, might be a morally relevant feature. At the same time, an unrelated disability that diminishes the person's capacity to benefit should not, for that reason, worsen a person's chances of being saved. It is not a morally relevant feature.

These examples bring us to the main problem with these proposals: they assume that the question of whether disability discrimination is unfair is an all-or-nothing matter, with a conclusive yes or no answer. This assumption is pervasive in the debate.<sup>21</sup> Arguments that use the device of the veil of ignorance cannot even avoid it: since they presuppose that the issue comes down to choosing the appropriate decision rule for the allocation of health care resources in an impartial situation, they cannot make distinctions between different disabilities and circumstances. Once, however, we recognize that the problem of disability discrimination is about deciding what it takes to apply a rule impartially by taking into account only morally relevant features, it becomes possible to ask more detailed questions about which disabilities, under which conditions, are morally relevant. This approach not only gives a better interpretation of the moral basis of the problem of disability discrimination; it may even make it possible for defenders and opponents of cost-effectiveness analysis to find common ground.<sup>22</sup>

## 7. CONCLUSION

Perhaps the most important objection to cost-effectiveness analysis concerns the discrimination against people with disabilities and chronic health conditions that the use of cost-effectiveness considerations may lead to in the allocation of health care resources. Here I have presented the discrimination objection in what seems to me its most persuasive formulation—as an objection to taking into account the effects of a prior disability on a person’s capacity to benefit in quality-adjusted measures. This formulation has the right scope: it avoids the implications that all effectiveness considerations are excluded from cost-effectiveness analysis and that some forms of discrimination that we would normally not object to turn out to be unfair.

I have argued that a well-known defense of quality-adjusted measures does not succeed. But it does not follow that the use of cost-effectiveness analysis does not lead to unfair discrimination. Rather, I have tried to show that opponents and defenders of cost-effectiveness analysis should adopt a different account of the connection of impartiality and fairness to settle their disagreement.

I am indebted to audiences at Queen’s University, New York University, and at the Ethical Issues in the Measurement of Health and the Global Burden of Disease conference organized by the Harvard University Program in Ethics and Health at Harvard University for valuable comments on an earlier draft of this paper, as well as to Francis Kamm (who was

my commentator at that conference), Dan Brock, John Broome, Norman Daniels, Iwao Hirose, Sam Kerstein, Bill Ruddick, Carlos Soto, and two anonymous referees.

## NOTES

1. This is not strictly speaking correct, since both QALYs and DALYs are commonly subject to discounting, and DALYs are also often age weighted: the value of a given health benefit steadily diminishes with age after young adulthood. I ignore these complications in the subsequent discussion, along with other methodological differences between the two measures that need not concern us here. For details, see Murray (1996), Nord (1999), and Bognar (2008).
2. Famously, they were used in the initial phases of the Oregon Medicaid reform. Subsequently, however, cost-effectiveness considerations were dropped from the reform proposal, primarily because of objections that they would lead to unfair discrimination against people with disabilities. For a detailed account, see Blumstein (1997); see also Bognar (2010).
3. See, for example, the contributions to Jamison et al. (2006). When cost-effectiveness analysis uses QALYs as an outcome measure, it is traditionally called cost-utility analysis. Given the ubiquity of QALYs and DALYs in economic evaluations of health care resource allocation today, I, however, use the broader term.
4. Many different conceptions of disability have been proposed. For the purposes of this paper, I accept the definition used by the Americans with Disabilities Act: a disability is “a physical or mental impairment that substantially limits one or more major life activities.” (See <http://www.ada.gov>.) This definition has at least two advantages: first, it is broad and inclusive, incorporating a wide range of physical and mental conditions; second, it is compatible with both medical and social models of disability—that is, it does not take a stand on whether the disadvantages associated with disability are primarily due to functional limitations or social factors.
5. Note that in my illustrations I am not assuming that a person with a disability *necessarily* has a lower quality of life. By putting the argument in terms of diminished capacity to benefit, I can focus on those cases in which discrimination would be morally the most worrisome: when the disability does make a quality-of-life improvement that we can attain for others impossible for those with the disability, or when the treatments that could in fact provide a comparable quality-of-life improvement for them are less cost effective due to the factors I mention. None of this implies that people with disabilities cannot have a high quality of life or could not attain substantive quality of life improvements by cost-effective means.

6. See, for example, Bickenbach (2005).
7. A few philosophers do take up the question of the fairness of sex discrimination. Amartya Sen (2004), for instance, argues that sex discrimination is not unfair. Shlomi Segall (2009), however, presents an argument for the view that differences in life expectancy constitute an unfair disadvantage for men. In this paper, I do not take a stand on this issue. In all of my examples, I assume that all candidates for a resource are of the same sex.
8. For discussion and empirical studies, see Nord et al. (1996) and Tsuchiya (1999). Some philosophers and health economists even propose that benefits to the elderly should be given lower weight than benefits to younger people, which would exacerbate their relative disadvantage on unweighted quality-adjusted measures. See, for example, Lockwood (1988). The formula for calculating DALYs also includes age weighting. I assessed some of the common justifications for age weighting in Bognar (2008).
9. “Personal characteristics” is a term used in health economics to refer to features of people’s lives that affect how much utility they can derive from different goods. It is intended to capture nonutility information that is important from a normative point of view but that would be neglected if we focused exclusively either on utilities or commodities and services. The term encompasses some of the external circumstances as well as the internal characteristics of the person. See, for instance, Culyer (1990).
10. One of the first critics of QALYs, John Harris (1987), called this the problem of “double jeopardy.” The architect of the DALY measure, Chris Murray (1996, pp. 32–33) regards this feature of quality-adjusted measures a “vexing moral problem.”
11. The distinction goes back at least to Kant. My formulation follows Jeff McMahan (2002, pp. 240–65). What these capacities are is a hotly debated question that can be set aside here. I also set aside the case of people with cognitive disabilities so severe that they lack the capacities in question.
12. This formulation of the objection is admittedly narrower than those that are usually offered by opponents of cost-effectiveness analysis. In an examination of their arguments, Dan Brock (2009) lists six different forms of disability discrimination. Some of these are general objections to the use of effectiveness considerations, and they are vulnerable to the problems I discuss in section 3. Other formulations have only a tenuous connection to discrimination against people with disabilities: one formulation, for instance, is concerned with disabilities that are the *result* of treatment when no disabilities are present in any of the patients before the treatment. I find it hard to see why this form of discrimination would be discrimination against people with disabilities.

13. Christopher Murray (1996, p. 6) proposes the following principle for governing the formulation of the DALY measure: “The non-health characteristics of the individual affected by a health outcome that should be considered in calculating the associated burden of disease should be restricted to age and sex.” The idea is that the discrimination by age and sex that results from the use of DALYs is not unfair.
14. As they also point out, the use of quality-adjusted measures, and cost-effectiveness analysis more generally, is part of a utilitarian account of distributive justice; it is, they state, “rather artificial to separate the two” (McKie et al. 1998, p. 54). Nevertheless, the present argument must be interpreted as a separate defense of one aspect of a utilitarian theory—namely, benefit maximization in health care resource allocation. If utilitarianism was independently justified, all McKie and his coauthors would have to show is that QALYs or some other similar measure is the most appropriate maximizandum, not that it is rational to allocate resources by maximizing prospective benefit.
15. This reconstruction of the argument is proposed by Iwao Hirose (2008).
16. This suggests that Hirose’s reconstruction of the argument of McKie and his coauthors is inaccurate, since they would argue that in this example *y* would be chosen behind the veil, given that it is the benefit-maximizing choice. The reconstruction does not allow for this conclusion.
17. See Rawls (1971) and Harsanyi (1982).
18. Bernard Gert (2005) offers an account of impartiality that develops this approach. (But note that what I say in this paper does not depend on Gert’s specific account of impartiality.) Brad Hooker (2005) distinguishes *formal* fairness—the impartial application of a rule—from *substantive* fairness, which is the impartial application of a rule that is in itself fair.
19. Gert recognizes this point but only addresses it in a note. There he gives this example: “That a judge is influenced by the guilt of the party before him is not incompatible with his acting impartially with regard to all those who come before him in the respect of applying the law” (1995, p. 123, n. 1). That is, the impartiality of the judge is not merely a matter of applying the law consistently, and avoiding being influenced by her own interests, but is also a matter of being influenced only by morally relevant features: the race of the defendant, for instance, is morally irrelevant. (Gert discusses morally relevant features in more detail in Gert 1999.) See also Barry (1995, p. 228) who, in the context of impartiality and health care resource allocation, says that “we can talk about equal treatment only after we have settled the prior issue of what should count as relevant and what should be excluded as irrelevant.”

20. For this proposal, see Nord et al. (1999). For the health-improvement setting, they suggest that we introduce “potential weights” to reduce the discrimination against patients with less capacity to benefit. Benefits to patients who have a smaller potential for health gain would be given greater weights. The objections I make below apply to this suggestion as well.
21. Empirical research on people’s judgments on the relevance of different personal characteristics has also been couched in all-or-nothing terms. Researchers have mostly asked people only in general terms about their judgments on the role of the severity of health conditions and the potential for health improvement in health care resource allocation without distinguishing between different kinds of disability or other personal characteristics. For a survey of empirical studies and their discussion, see Olsen et al. (2003) and Nord (1999).
22. It is, however, beyond the scope of this paper to present such a view. I undertake this task in Bognar (2010). There I argue that we need to distinguish between, on the one hand, prior disabilities that are unrelated to, or independent of, the condition whose cost effectiveness we are evaluating, and, on the other hand, disabilities that are genuine comorbidities. I argue that it is unfair to take into account the impact of unrelated disabilities but that it might not be unfair to take into account the impact of comorbidities. However, any account of the difference between unrelated disabilities and comorbidities is likely to remain controversial. In my view, this at least partly explains why reasonable people continue to have disagreements about particular cases of disability discrimination.

## REFERENCES

- Barry, Brian. 1995. *Justice as Impartiality*. Oxford: Clarendon Press.
- Bickenbach, Jerome. 2005. Disability and Health Systems Assessment. In *Quality of Life and Human Difference: Genetic Testing, Health Care, and Disability*, ed. David Wasserman, Jerome Bickenbach, and Robert Wachbroit, 237–66. Cambridge: Cambridge University Press.
- Blumstein, James F. 1997. The Oregon Experiment: The Role of Cost-Benefit Analysis in the Allocation of Medicaid Funds. *Social Science and Medicine* 45: 545–54.
- Bognar, Greg. 2008. Age-Weighting. *Economics and Philosophy* 24: 167–89.
- . 2010. Does Cost Effectiveness Analysis Unfairly Discriminate against People with Disabilities? *Journal of Applied Philosophy* 27: 394–408.
- Brock, Dan W. 2009. Cost-Effectiveness and Disability Discrimination. *Economics and Philosophy* 25: 27–47.

- Culyer, A. J. 1990. Commodities, Characteristics of Commodities, Characteristics of People, Utilities, and the Quality of Life. In *Quality of Life: Perspectives and Policies*, ed. Sally Baldwin, Christine Godfrey, and Carol Propper, pp. 9–27. London: Routledge.
- Gert, Bernard. 1995. Moral Impartiality. *Midwest Studies in Philosophy* 20: 102–28.
- . 1999. Morally Relevant Features. *Metaphilosophy* 30: 13–24.
- . 2005. *Morality: Its Nature and Justification*. Rev. ed.. New York: Oxford University Press.
- Harris, John. 1987. QALYfying the Value of Life. *Journal of Medical Ethics* 13: 117–23.
- . 1988. More and Better Justice. In *Philosophy and Medical Welfare*, ed. J. M. Bell and Susan Mendus, pp. 75–96. Cambridge: Cambridge University Press.
- Harsanyi, John C. 1982. Morality and the Theory of Rational Behaviour. In *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams, pp. 39–62. Cambridge: Cambridge University Press.
- Hirose, Iwao. 2008. Disability Discrimination in Health Care. Unpublished manuscript.
- Hooker, Brad. 2005. Fairness. *Ethical Theory and Moral Practice* 8: 329–52.
- Jamison, Dean T.; Breman, Joel G.; Measham, Anthony R.; Alleyne, George; Claeson, Mariam; Evans, David B.; Jha, Prabhat; Mills, Anne; and Musgrove, Philip, eds. 2006. *Disease Control Priorities in Developing Countries*. Washington, DC: Oxford University Press and the World Bank.
- Lockwood, Michael. 1988. Quality of Life and Resource Allocation. In *Philosophy and Medical Welfare*, ed. J. M. Bell and Susan Mendus, pp. 33–55. Cambridge: Cambridge University Press.
- McKie, John; Richardson, Jeff; Singer, Peter; and Kuhse, Helga. 1998. *The Allocation of Health Care Resources: An Ethical Evaluation of the 'QALY' Approach*. Aldershot, UK: Ashgate.
- McMahan, Jeff. 2002. *The Ethics of Killing: Problems at the Margins of Life*. Oxford: Oxford University Press.
- Murray, Christopher J. L. 1996. Rethinking DALYs. In *The Global Burden of Disease: A Comprehensive Assessment of Mortality and Disability from Diseases, Injuries, and Risk Factors in 1990 and Projected to 2020*, ed. Christopher J. L. Murray and Alan D. Lopez, pp. 1–98. Cambridge, MA: Harvard School of Public Health on behalf of the WHO and the World Bank.
- Nord, Erik. 1999. *Cost-Value Analysis in Health Care: Making Sense out of QALYs*. Cambridge: Cambridge University Press.

- Nord, Erik; Pinto-Prades, Jose-Luis; Richardson, Jeff; Menzel, Paul; and Ubel, Peter. 1999. Incorporating Societal Concerns for Fairness in Numerical Valuations of Health Programmes. *Health Economics* 8: 25–39.
- Nord, Erik; Street, Andrew; Richardson, Jeff; Kuhse, Helga; and Singer, Peter. 1996. The Significance of Age and Duration of Effect in Social Evaluation of Health Care. *Health Care Analysis* 4: 103–11.
- Olsen, Jan Abel; Richardson, Jeff; Dolan, Paul; and Menzel, Paul. 2003. The Moral Relevance of Personal Characteristics in Setting Health Care Priorities. *Social Science and Medicine* 57: 1163–72.
- Phillips, David. 2006. *Quality of Life: Concept, Policy and Practice*. New York: Routledge.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Segall, Shlomi. 2009. *Health, Luck, and Justice*. Princeton, NJ: Princeton University Press.
- Sen, Amartya. 2004. Why Health Equity? In *Public Health, Ethics, and Equity*, ed. Sudhir Anand, Fabienne Peter, and Amartya Sen, pp. 21–33. Oxford: Oxford University Press.
- Tsuchiya, Aki. 1999. Age-Related Preferences and Age Weighting Health Benefits. *Social Science and Medicine* 48: 267–76.