THEORIA

THEORIA, 2018
doi:10.1111/theo.12167

Fairness and the Puzzle of Disability

by

GREG BOGNAR

Stockholm University, Stockholm Centre for Healthcare Ethics (CHE) and La Trobe University

Abstract: Consider two cases. In Case 1, you must decide whether you save the life of a disabled person or you save the life of a person with no disability. In Case 2, you must decide whether you save the life of a disabled person who would remain disabled, or you save the life of another disabled person who, in contrast, would also be cured as a result of your intervention. It seems that most people agree that you should give equal chances in Case 1: saving the life of the person with no disability would be unfair discrimination against the person with disability. Yet, in Case 2, it appears that many people believe that you are at least permitted to save straightaway the person who would have no disability after your intervention. There would be no unfair discrimination against the other person. I argue that these judgements present a puzzle for theories of resource allocation in normative ethics. The puzzle is straightforward for consequentialists: the two cases have the same outcomes, but the judgements are different. But the puzzle also presents a problem for nonconsequentialist views. After introducing the cases, I show this by reviewing a number of proposals for solving the puzzle. I argue that none of these proposals are successful. I then make my own proposal and conclude by spelling out its implications.

Keywords: disability, fairness, consequentialism

1. The Puzzle

LET ME BEGIN WITH THE sort of life-and-death example philosophers are so fond of. Suppose you are a physician in an emergency room and two accident victims are brought in. They will die without immediate surgery. Unfortunately, you are the only qualified surgeon there. No one else is available. You have to choose which patient to operate on. The other patient will die.

There is no morally relevant difference between the two patients. They are the same age and sex. They would survive for the same amount of time with the same quality of life. Neither of them is responsible for the accident, neither has greater family responsibilities or more friends, neither has rendered great services to society in the past or can be expected to do so in the future – they are equal in whatever factor you might believe can make a difference to what you ought to do.

A plausible thought is that you ought to select randomly the patient whose life you save. Consequentialists will point out that the outcomes of saving the first patient and saving the second patient are equally good. To decide whom to save,

you might as well toss a fair coin. Nonconsequentialists might argue that fairness requires that the patients are given an equal chance. Therefore, you should toss a fair coin, not merely to break the tie, but because it would be unfair not to give an equal chance to the two patients.

To make the case more interesting, here is a modification: suppose that the first patient has a disability. He has a permanent but manageable physical or mental impairment that has a substantive negative impact on at least one major life activity. The limitation caused by the impairment makes this patient's life, in one respect at least, worse than it would be otherwise. Since everything else remains equal, saving this patient's life would lead to a worse outcome than saving the life of the other patient, who could be returned to full health.

Note that setting up the example this way is not meant to suggest that the life of a disabled person is *necessarily* worse or worse all things considered than the life of anyone else. There is no doubt that people with disabilities can have lives that are just as good as the lives of others. Health is only one component of well-being. But in the example, everything is equal, except for the disability. Even if not all disabilities make life worse, there are at least some that, other things being equal, make most people worse off. I assume the patient has one of these.[1]

	Case 1	Case 2
First patient	disabled → disabled	disabled → disabled
Second patient	nondisabled → nondisabled	disabled → nondisabled

The example is illustrated as Case 1 on the figure (ignore Case 2 for the moment). The first row shows what happens to the first patient, *given that he is saved*; the second row shows what happens to the second patient, given that he is saved. The descriptions separated by the arrows indicate the health states of the patients before and after your intervention; the arrow represents the transition from a patient's *ex ante* to his *ex post* health state. That is, the first patient in Case 1 is disabled before the life-saving surgery and remains disabled after the surgery. The second patient in Case 1 has no disability before the accident and will remain fully healthy after your life-saving intervention. Remember also that only one of the patients can be saved; hence whether an *ex post* health state obtains depends on what you do. For instance, if you save the first patient in Case 1, he will remain disabled, while the second patient, rather than remaining nondisabled, will die. (This is not illustrated on the figure.)

1 In Bognar (2016), I examine in more detail the issue of whether disability can be considered, other things being equal, a harm.

Some people might have the following moral intuition about Case 1: it would be unfair to save the life of the second patient just because he does not have a disability. It should not count against the first patient that he would remain disabled after the life-saving surgery. Of course, you cannot treat the two patients equally (you cannot save both of them). So discrimination, in the sense of unequal treatment, is unavoidable. But discrimination is not always unfair. In particular, you can give the two patients equal chances. This is what fairness requires: it is not unfair if you discriminate against the disabled patient by saving the second patient, provided that both of them had an equal chance to be saved. But saving the second patient straightaway would constitute *unfair discrimination* against the person with disability.[2]

But consider Case 2 now. Here both of the patients who are brought into the emergency room are disabled (assume they have the exact same disability). You can operate on the first patient of this case with the same result as in Case 1: his life will be saved, but he will remain disabled. For the second patient, however, you can do more: you can not only save this patient's life, but also remove his disability. He will not remain disabled after the surgery. (Apart from this, I assume that everything else is equal: for instance, the surgical procedure used for the second patient in Case 2 is no more risky, complex or expensive than the procedure used for the other patients.)

Some people might have the intuition that it is *not unfair* to save the second patient straightaway in Case 2.[3] After all, the surgery that you can perform on this patient will not only save his life, but also remove his disability. Surely, we should care about curing disabilities. It should be permissible to give priority to interventions that lead to better outcomes. So there is no unfairness in selecting

2 In each of the cases that I discuss, I assume that the disability is sufficiently severe so that, for instance, the two patients in Case 1 should have equal chances even though saving the second patient would lead to a substantially better outcome. The two patients should have equal chances despite this and not because the difference between the values of the outcomes would be so trivial that giving equal chances is merely dictated by the need of breaking what is essentially a tie. I am also assuming that compensation for past disability is not at issue. It might be that a person should be given priority (either in the form of being saved straightaway or given a greater chance of being saved) as a way of compensating her for past disability, chronic illness or suffering. I take it that the moral intuition that the patient with the disability should have an equal chance persists even when it is specified that the past does not matter (for instance, because the first patient's disability is too recent to have caused sufficient suffering or disadvantage up to this point). Finally, it should go without saying that the talk of "saving a life" is elliptical: you can never literally save a life; the best you can do is to extend it. In my examples, it is always assumed that you can prolong the lives of different patients by exactly the same amount of time.

3 Empirical surveys seem to suggest that many people have these intuitions. For reviews of empirical surveys of people's judgements in resource allocation cases, including cases like Cases 1 and 2, see Nord (1999), Green (2009) and Shah (2009). For philosophers who consider the moral intuitions in Case 1 and Case 2, see, for example, Brock (2009), Harris (1987) and Kamm (2004, 2009, 2013).

the second patient without giving an equal chance to the first patient. There is no unfair discrimination in Case 2.[4]

Or perhaps the two cases should be analysed differently. Perhaps it is not the case that saving the second patient straightaway in Case 2 is not unfair. It might be thought that even though it would be *pro tanto* unfair to save the second patient straightaway in Case 2, however, when all things are considered, there is some other consideration, absent from Case 1, that overrides or justifies this unfairness. So it may be, on one interpretation, that there is no unfairness in Case 2 at all, or it may be, on another interpretation, that the unfairness is justified all things considered. I will take into account both possibilities below. In either case, the concept of fairness plays a central role in these cases.

Suppose we accept the judgements that it would be unfair if the two patients were not given an equal chance in Case 1, but, all things considered, it would not be unfair (or impermissible) if the second patient were saved straightaway in Case 2. But now we seem to have a problem. In Case 1, it is impermissible to choose the better outcome and save the second patient straightaway; you should not allow his disability to "count against" the first patient, *despite* the fact that saving the second patient would lead to a better outcome. In Case 2, however, it is permissible (and perhaps even required) to save the second patient straightaway, apparently *because* this would lead to a better outcome, even though it seems that this means allowing the disability of the first patient to count against him. The problem is to explain how these two judgements can be compatible with one another.

The problem seems especially acute for consequentialists. They can agree that the second patient in Case 2 should be saved straightaway, since the outcome in which a person exists without a disability is better than the outcome in which a person exists with a disability (and everything else is equal). But their view implies the same for Case 1, since there is no difference between the two cases in terms of outcomes. On consequentialist grounds, you ought to save the life of the patient who will have no disability in both cases. But this conflicts with the moral intuitions that we have considered. Many people believe that not to give equal chances to the two patients in Case 1 is unfair.

---

4 It might also be thought that it is unfair to select the second patient straightaway, but it is not unfair not to give an equal chance to the first patient: instead, what fairness requires is a *weighted lottery* where the first patient has a less than 50% chance of being saved. To keep the discussion manageable, I am going to set aside the possibility of weighted lotteries until Section 3. Hence for now, I will say that it is permissible (and perhaps even required) to save the second patient straightaway. As far as I can tell, none of the arguments below would change if I said instead that it is permissible (and perhaps even required) to give the second patient a greater-than-equal chance – but it would be more cumbersome. In any case, none of the proposals that I consider below include weighted lotteries.

Consequentialists, of course, can reject these moral intuitions. In particular, they can argue that the intuition in Case 1 should not be accepted. When you have to allocate scarce resources, the only thing that matters is how much good you can do with those resources. But this reply is not entirely satisfactory. It does not explain why there is no unfairness in Case 1. Consequentialists should be able to say more in defence of their view. They must address the problem of fairness. I will look at one attempt in the next section.[5]

The cases raise a problem for nonconsequentialists too. Nonconsequentialists believe that the goodness of outcomes is not the only morally relevant factor in determining the rightness of actions. Some nonconsequentialists even believe that the goodness of outcomes is not a relevant factor at all. Both types of view seem to explain the judgement in Case 1 well. Since disability leads to lower quality of life (given that everything else is equal), saving the second patient in Case 1 would lead to a better outcome, all things considered. But goodness is not the only relevant consideration (or it is not relevant at all). What matters is that the two patients would be treated unequally on account of the first patient's disability. This kind of unequal treatment is unfair. Of course, you cannot treat the two patients strictly equally, since you can only save one of them. But you can treat them equally in the sense of giving them equal chances. This is what fairness requires you to do.

Nevertheless, just as consequentialists have a problem with Case 1, nonconsequentialists have a problem with Case 2. If fairness requires you to ignore the disability of the first patient in Case 1, shouldn't it also require you to ignore the disability of the first patient in Case 2? Shouldn't the two patients have equal chances? Yet it seems it is morally relevant that you can remove the disability of the second patient in Case 2 in addition to saving his life. Consequently, nonconsequentialists who reject that the goodness of outcomes is morally relevant need to explain why the common moral intuition in Case 2 should be rejected; and nonconsequentialists who accept the moral relevance of the goodness of outcomes need to explain why it is decisive in Case 2 but not in Case 1.

These cases show that disability presents a puzzle for the ethics of resource allocation – for consequentialist and nonconsequentialist views alike. Few people would deny that it is unfair if people with disabilities suffer disadvantages merely on account of their disability. At the same time, it is important that scarce resources are used in ways that bring about the best outcomes. Since (other things

5 Of course, it might be possible to formulate more sophisticated versions of consequentialism on which the value of an outcome is partly a function of whether it was reached in a fair way. I am going to set such versions aside to keep the discussion manageable. Note that any view like that would first have to solve the puzzle in order to determine the (fairness-included) value of the outcomes in Cases 1 and 2.

being equal) disabilities make life worse, there is a general moral requirement to minimize their prevalence. This is why the prevention, cure and rehabilitation of disability are so important. But sometimes these aims conflict with the idea of nondiscrimination. Disability presents a puzzle for our views on fairness.

I will proceed as follows. In Section 2, I discuss an argument that aims to defend the consequentialist position by rejecting the common moral judgement in Case 1. In Section 3, I apply the currently most influential account of fairness to the puzzle. In Section 4, I consider two further proposals. I will conclude that none of these arguments and proposals succeeds. Then, in Section 5, I outline a new proposal. I argue that it provides a solution to the puzzle, but it also has some implications that some people might find hard to accept.

2. The Veil of Ignorance Argument

I said the puzzle poses a problem for consequentialists. It is not difficult to see why. Suppose that we represent the values of the outcomes in our cases with numbers between 0 and 1. 0 represents a health state that is not better than death, 1 represents full health, and numbers from 1 to 0 represent health states with increasingly severe illness or disability. Suppose, for the sake of simplicity, that the badness of the disability in our cases can be represented by the value of 0.5.

With respect to the value of their outcomes, Case 1 and Case 2 are identical. In both cases, if you decide to save the life of the first patient, he will remain disabled, and the second patient will die. The outcome can be represented by (0.5;0), where the first number is the value of the outcome for the first patient, and the second number is the value of the outcome for the second patient. If, in contrast, you save the life of the second patient, he will survive with no disability, and the first patient will die. The outcome is (0;1). Again, this is true of both cases. It does not make a difference that the second patient in Case 2 is disabled *ex ante* – what matters is that he will have no disability *ex post.*[6]

Hence consequentialists cannot accept that the right action differs across the two cases. For them, only the value of the consequences matter. Since the two choices are identical with respect to the value of their outcomes, the right action

6 A scale like this can provide the quality-adjustment factors for *quality-adjusted life years*, or QALYs. The QALY is a measure of health-related quality of life that combines the value of health states with the time spent in those health states. Thus, 1 QALY can represent 1 year in full health, 2 years in a health state whose value is 0.5, 4 years in a health state whose value is 0.25, and so on. If the numbers represent QALYs in the examples, then their interpretation might be that all patients who are saved survive for 1 year only. If that seems too short, you could multiply all the numbers by 10, for instance, to represent 10 years. (For more detail on health state evaluation and ethical issues in the use of QALYs in the allocation of health care resources, see Bognar and Hirose, 2014.)

must be the same. In both cases, it is to save the life of the second patient. This accords with the common moral intuition in Case 2, but not in Case 1. Therefore, consequentialists must provide an argument for rejecting the common moral judgement in Case 1.

One argument they have made appeals to the device of the *veil of ignorance.* It is proposed by Peter Singer in collaboration with others.[7] Here it is in rough outline. Suppose that the patients in our examples are rational and they pursue only their own interests. Suppose that they have to decide how to allocate a scarce life-saving resource (for instance, the surgery that can save their life when they are brought into the emergency room). Suppose also that neither of them knows whether he is the first or the second patient. Their identity is concealed by the veil. What moral principle would their choice reveal?

The patients know that other things being equal, their life is better if they do not have a disability. Hence they will have a *stronger interest in continued life* when they are not disabled. Given this, they would not agree to give equal chances to a patient who would end up disabled and a patient who would have no disability. As Singer and his collaborators put it, "to maximise the satisfaction of their own interests, rational egoists would have to choose a system that gives preference to saving life when it is most in the interests of the person whose life is saved" (Singer et al., 1995, p. 148).

Thus, suppose you are one of the patients in Case 2, but you do not know whether you are the first or the second patient. According to the argument, you should choose a moral principle (or "system") that gives priority to the second patient, because he has a stronger interest in continued life. Going on living without disability is a greater benefit. This is the rational choice that maximizes the expected satisfaction of your interests, given that you might end up as either patient when the veil is lifted.

The very same argument applies to Case 1. Because of the veil, you do not know whether you are the disabled or the nondisabled patient. But you do know that you have an equal chance of being either one, and that the patient without the disability has a stronger interest in continuing to live. Therefore, rather than giving them equal chances, you should choose a principle that enjoins you to save the second patient straightaway. That principle is benefit maximization, which tells you to choose the action that leads to the best expected outcome. Therefore, it is permissible (and perhaps even required) to save the life of the patient with

7 See Singer et al. (1995) and McKie et al. (1998). I have provided a detailed criticism of their argument in Bognar (2011). I will not repeat my objections here. I will consider only whether their case for giving up the judgement in Case 1 is persuasive.

no disability without giving a chance to the disabled patient. But this conclusion conflicts with the judgement that many people have.

Should we accept the veil of ignorance argument? Should we revise our judgement in Case 1?

I think not. Here is why. Behind the veil, your choice, by hypothesis, is self-interested and rational. Since the veil conceals your identity, it is also impartial. Singer and his collaborators believe that the combination of self-interested rationality and impartiality ensures fairness. Evidently, they could not otherwise claim to have dealt with the objection that saving the second patient is a case of *unfair* discrimination. The principle for resource allocation chosen behind the veil is supposed to be fair.[8]

Consider the crucial move in the argument: the claim that the person who has no disability has a *stronger interest* in continuing to live. This claim is not self-evident.[9] Singer and his collaborators identify a person's interest in continued life with the *goodness of that person's life* – more precisely, with how good it would be for the person whose life it is to continue living it. Given this view on what determines a person's interest in continued life, it is undoubtedly true that it is rational to choose to save the patient with no disability. Continuing that life would be better.

But this just tells us what we already knew. We already knew that (other things being equal) a life with no disability is better, and from the perspective of benefit maximization, it would be better to save the patient with no disability. You do not need the veil of ignorance argument to show this. Those who insist that the two patients should have equal chances in Case 1 do not (or should not) dispute that it would lead to a better outcome if the second patient is saved. But they can complain that the argument does not show that this would not be unfair.

What is more, the opponents of the veil of ignorance argument can point out that this is not the only relevant understanding of the concept of a person's interest in continued life. In another sense, a person's interest in continued life can be determined by how the *person herself values her life*, or how much she cares about it. There is no reason to think that this is an incoherent concept. Thus, the opponents can suggest that the question that rational and self-interested people would ask themselves behind the veil of ignorance might instead be this: "How much would I value, or care for, my life if I turned out to be any of these persons, some of whom have less good lives than others?"

8 Singer and his collaborators follow Rawls (1971) in combining rationality and impartiality to ensure fairness.

9 Both Singer et al. (1995) and McKie et al. (1998) spend surprisingly little time defending it.

Often, people who face severe or terminal illness or unmanageable pain value their lives less – in some cases, they even wish to end it. Many people believe that we should give weight to the patients' values and wishes in those cases. They are morally relevant. But then there is no reason to reject their moral relevance in general – and in the veil of ignorance argument in particular. How people value their lives from their own subjective point of view matters.

Moreover, there is no reason to believe that a person with a disability values or cares for her life less than others. There is certainly no reason to attribute this belief to the decision makers behind the veil of ignorance. On the contrary, it seems that they should be aware that people with disability might value and care for their lives just as much as others. If this is the most important consideration, then it is not irrational to select a principle that gives equal chances. On this interpretation of a person's interest in continuing to live, the two patients should have equal chances in Case 1, since (by assumption) they do not differ in the way they value or care for their lives.

At the end of the day, however, I do not think we should accept this argument. For one thing, it goes astray in Case 2: it entails that the two patients should have equal chances even though the second patient would be rid of his disability if he were chosen. But the argument is not obviously wrong. Therefore, it can be used to make a more general point – to illustrate the main problem with veil of ignorance arguments. That problem is that the conclusion of these arguments depends crucially on how you interpret the "input". In the case of disability discrimination, the conclusion depends on which of two interpretations of the concept of a person's interest in continued life is accepted.[10]

3. Fairness and Goodness

In this section, I turn to what is probably the currently most influential account of fairness, and apply it to the puzzle. The account was proposed by John Broome (1994, 1999).

First, I need to explain some background. When resources must be distributed between different people, there are several kinds of reason why any person should have a particular good. For instance, there are *reasons of benefit*: if a person can benefit from getting a good, there is a reason for giving it to her. This kind of

10 As Jeff McMahan pointed out to me, both of these interpretations have implausible implications in any case. The first implies that the worst death for someone is one that occurs immediately after she begins to exist (for instance, if one begins to exist at conception, then the worst death is that which takes place right after conception). The second interpretation implies that those who are not self-conscious and hence cannot value their lives do not have any interest in continuing to live. Thus, neither animals nor infants have an interest in going on living.

reason is determined by the person's capacity to benefit: how much good it would do if she got it. The greater her capacity to benefit, the stronger the reason for her getting the good.

Of course, it is often the case that other persons can also benefit from a good. There are reasons of benefit on the side of more than one person. In addition, these reasons may well differ in strength, since some people could benefit more than others from getting the good. Sometimes, however, the good can be divided in a way that maximizes benefits. The person with the greatest capacity to benefit gets shares of the good until her capacity to benefit from further shares falls below the capacity of the second person to benefit from those shares; then the second person gets the additional shares until her capacity to benefit falls below the capacity of the next person to benefit from additional shares; and so on until all of the good is distributed. This way, each unit of the good is given to the person who would benefit the most. As a result, the overall benefit derived from the good is maximized.

This procedure involves the weighing of reasons against each other at each step. The relative strengths of the reasons of benefit determine the final distribution. To be sure, this sounds more complicated than it is. Often, there is no need for a step-by-step procedure and shares of the good can be allocated by weighing the reasons of benefit against each other all at once. For instance, if a good is indivisible, it should go to the person who has the strongest reason of benefit on her side.

Thus, Broome argues that if only reasons of benefit matter, the person with the greatest capacity to benefit should get the good. This is the consequence of the weighing of reasons of benefit.

In Cases 1 and 2, the good that is to be distributed is having one's life saved. This good is indivisible. Now, if only reasons of benefit mattered, you should save the second patient in both of these cases. In each case, the second patient would go on living without disability. Other things being equal, living with no disability is a greater benefit than living with a disability. The reason to save the second patient is stronger than the reason to save the first patient.

But there is another kind of reason that is relevant in resource allocation. Sometimes, a person should get a good as a matter of a *duty owed to that person herself* – that is to say, the person has a *moral claim* to the good. Moral claims provide a different kind of reason for why a person should have a good. Evidently, it is one thing to say that you ought to have a good because you would benefit from it, and quite another to say you ought to have a good because you have a moral claim to it. In Broome's view, fairness is concerned with the latter only. The distinction between reasons of benefit and reasons of moral claims is important: a reason of benefit is not owed to you the way a claim is owed to you.

It is not unfair if you do not get a good when the reason for your having it is merely that you would benefit from it. (Perhaps someone else would benefit more from it.) It is, however, unfair if you do not get the good when the reason for your having it was that you had a claim to it. Only in this case is the good owed to you. Fairness applies only to this kind of case.

How do moral claims determine the allocation of resources? On Broome's account, claims should be satisfied in proportion to their strength. If two people have equal claims to a divisible good, they should have equal shares. If one person has a stronger claim, then the good should be shared proportionally. The stronger a person's claim, the greater her share. But as long as a person has a claim to a good, no matter how weak, she should have some of it. Note how different this is from the way reasons of benefit are weighed against each other: in their case, the good is not divided proportionally, but each share is given to the person with the strongest reason until the resource is exhausted. If your capacity to benefit from all units of a divisible good is greater than those of others, you should have all of the good. But if you have the strongest claim to a good, you should still not have all of it – others with weaker claims should have some of it too. Claims are not weighed against one another the way reasons of benefit are.

Consider an indivisible good. If only reasons of benefit are present, the good should go to the person whose capacity to benefit is the greatest. Others will have none of the good. In the case of moral claims, in contrast, the good should be shared proportionally. But this is not possible with indivisible goods. So, Broome suggests, a lottery should be held among all of the people with a claim to the good. A lottery provides a sort of "surrogate satisfaction" of people's moral claims. It is a way to achieve fairness when not everyone's claim can be satisfied.

This account of fairness applies directly to our puzzle. Can it solve it?

Helpfully for my purposes, Broome (1994, pp. 38–39) discusses Case 1. Because of his disability, the first patient has a diminished capacity to benefit from your life-saving intervention. If only reasons of benefit mattered, what you should do would be clear: you should save the life of the second patient. The reason to save the first patient is weaker than the reason to save the second patient, since he can benefit less from being saved. But most people believe that saving the second patient straightaway would be unfair, and we can now explain why: both of the patients brought to the emergency room have an equal claim to the life-saving intervention. (Remember that all things apart from the disability are equal, and few people would argue that people with disabilities have weaker moral claims to life-saving resources *just because* they have disabilities.)

Unfortunately, you cannot save both patients. Hence some unfairness is inevitable. What you can do instead is to give them equal chances to the good of having their life saved. This does not completely remove the unfairness, but at least

minimizes it by providing a "surrogate satisfaction" of their equal moral claims. It is a second-best solution – but a way of treating them equally.

Fairness requires that the two patients get an equal chance of being saved. On grounds of benefit only, the second patient should be saved. These are different considerations, and they work differently. How can they be compared to determine what you ought to do, all things considered? Broome suggests that when the considerations of fairness and benefit maximization are both relevant, you should return to weighing: fairness should get priority when the sacrifice in overall benefit would be relatively small, and benefit maximization should get priority when the sacrifice in overall benefit would be too great. In Case 1, fairness is the weightier consideration, since the sacrifice in overall benefit (the difference in quality of life between being disabled and having no disability) is relatively small compared to loss of life. Consequently, the patients should have equal chances, and the reason of benefit on the side of the second patient should be overridden.

So far, so good. But let us now apply the foregoing analysis to Case 2. Here, both patients are initially disabled. The only difference is that the second patient can be returned to full health, if his life is saved. On the one hand, it should be uncontroversial that this does not make a difference to the strength of their claims. Just as in Case 1, their claims to the good of being saved are equally strong. Since only one of them can be saved, fairness requires that they are given equal chances. On the other hand, the reasons of benefit on their sides are different. The second patient has a greater capacity to benefit. In fact, the outcomes would be just as good as in Case 1, both if you save the life of the first patient and if you save the life of the second patient. Since the sacrifice in overall benefit is just as small, it follows that the two patients should have equal chances.

But this conflicts with the judgements that we have accepted. Broome's account of fairness is incompatible with the view of unfair discrimination that many people seem to accept. I take this to be a serious, and perhaps fatal, problem for this account of fairness.

In truth, Broome could reject the judgement in Case 2. If fairness can outweigh the balance of reasons of benefit in Case 1, then it should do so in Case 2. Sometimes we should let our best theory revise our judgements.

But I find it hard to let go of the judgement in Case 2. Here is an example to illustrate the theoretical costs of this. Suppose there are two types of cancer drug. The first leads to 5 years of remission but does not improve quality of life (the effects of the cancer, such as difficulties in breathing, persist). The second also ensures survival for 5 years, but it also removes all of the cancer symptoms. For those 5 years, patients are completely healthy. Suppose also that after 5 years,

the cancer returns in both cases and leads to quick and painless death. Now, it seems clear which of these drugs the health care system should prefer if it can offer only one of them. This seems true even if you assume, in addition, that only half of the cancer patients would benefit from the second drug. The other half would die, even though they could have benefited from the first kind of drug.

Notice that you do not need to deny that there is unfairness against the first patient in Case 2 (or the patients who would not benefit from the second drug in the cancer case). All that needs to be pointed out is that when fairness and benefit maximization are weighed against each other, the result in Case 2 changes compared to Case 1. But it is puzzling why it should. Broome's account of fairness cannot explain it. The puzzle remains unresolved.

Incidentally, my discussion in this section also shows why a weighted lottery cannot solve the puzzle. One proposal might be that in Case 2 you could strike a balance between fairness (requiring equal chances for the two patients) and benefit maximization (suggesting that the second patient should be saved straightaway) by holding a lottery in which the first patient has a smaller chance than the second patient. This way, the moral intuition that the value of the outcomes matter is taken into account by the greater chance of selecting the second patient; at the same time, the first patient's moral claim is not ignored, since he is given at least some chance of being selected. The precise weights might depend on the strengths of the moral claims of the patients and the difference in value of the outcomes – that is, the relative benefits.

Once again, the problem with this proposal is that the very same argument seems to apply to Case 1. If the second patient's greater capacity to benefit in Case 2 justifies giving him a greater chance of being saved, then it should justify a greater chance to the patient with no disability in Case 1. But this remains unfair towards the disabled patient. Moreover, if it is accepted that the two patients have equal moral claims to the life-saving intervention, it becomes unclear why they should nevertheless have unequal chances.

4. A Pluralist Proposal

Perhaps a solution to the puzzle can be provided by a *pluralist proposal.* The values within the pluralist proposal that I am going to consider briefly are benefit maximization and equality.[11]

Consider Case 1 first. In this case, benefit maximization recommends saving the life of the second patient. This course of action, however, would lead to

11 This proposal was made by Jeff McMahan on an earlier draft of this article.

greater lifetime inequality between the two patients: the second patient would survive for the same amount of time with a higher quality of life than the first patient. If you saved the first patient, there would be less inequality between the two patients. Now suppose that these two considerations are roughly equally strong. Giving equal chances seems the natural response to different values that are on a par. Hence the two patients should have equal chances.

Consider Case 2 now. Once again, benefit maximization recommends saving the second patient; equality recommends selecting the first patient. But, one may argue, the reason deriving from equality is weaker in this case. This is because greater equality would be achieved by denying the second patient an *additional* benefit: the second patient would receive the same benefit as the first (having his life saved) and also the benefit of getting rid of his disability. So, in this case, the reasons provided by the values of benefit maximization and equality are not on a par; the former is stronger. Thus, it is permissible to save the second patient straightaway.

I can think of at least three objections to the pluralist proposal. First, in Case 1, the proposal justifies equal chances as a *mere tie-breaker*. But most people would say that equal chances are required in this case for more substantive reasons. Even if the size of the benefit (or the size of inequality) changed, they would insist that equal chances should continue to be given.

Second, there is something slightly paradoxical about the argument that saving the first patient in the two cases would make the *outcome* more equal. In the outcome of both of these cases, only one patient exists. There is no inequality between the patients after your intervention. It is true that from a broader perspective the two lives are more unequal if you save the second patient. But given that no patient would ever experience the effects of the inequality, letting the second patient die for the sake of increasing the equality between the lives of the two patients does not seem to be the right kind of reason. On the pluralist proposal, the badness of inequality is not located in its effects.

The final, decisive objection is this. The second patient in Case 2 would get no greater benefit than the second patient in Case 1. The benefit – having his life saved and continuing to live without a disability – is the same as the benefit the second patient would get in Case 1. Although it can be considered an *additional* benefit, the overall increase in well-being is the same. And it should not matter how benefits are individuated. Why should the *number* of benefits matter instead of their total value?

Consequently, there should be no difference between the two cases on the pluralist proposal. If the values of equality and benefit maximization are on a par in Case 1, then they are on a par in Case 2. The reason deriving from equality is no

weaker in Case 2. The pluralist proposal does not provide a solution to the puzzle.[12]

5. A New Proposal

I have argued that disability presents a puzzle for theories of resource allocation in normative ethics. I have looked at a number of proposals to solve this puzzle, but found that none of them was successful.

We do not seem to have a similar moral puzzle when it comes to other forms of discrimination. For instance, it would be clearly unfair if some people were given priority in resource allocation because of their race or sex. Why can we not say the same about disability?

Race and sex are morally irrelevant personal characteristics. It would be unfair discrimination if they were used to decide who gets life-saving surgery when they are the only difference between two patients. So it might be tempting to think about racial or sexual discrimination as providing a model for disability discrimination. But disability is unlike these characteristics. In the case of race or sex there is no analogue to Case 2. Race and sex are not characteristics that have any impact on a person's capacity to benefit; they are not conditions that (other things being equal) make a life worse. Since they lack this feature, it is easy to see why they are morally irrelevant.

Disability, in contrast, does have an impact on capacity to benefit. It is a condition that (other things being equal) makes a life worse. This factor is normally morally relevant. Hence, it needs to be explained why it should be irrelevant when it comes to disability. Simply treating disability like race or sex would be begging the question.

12 Some may think that prioritarianism can handle the cases in the "correct" way. (I thank an anonymous reviewer for pressing this point.) Prioritarianism is the view that it is morally more important to benefit those who are worse off, and the worse off these people are, the more important the benefit. Thus, in Case 1, the smaller benefit that you can bestow on the first patient who is worse off might be morally just as important as the greater benefit that you can give to the second patient who is better off. In Case 2, the smaller benefit that you can provide to the first patient may be morally less important than the greater benefit that you can give to the second patient, given that the two patients are equally badly off. Now, it is possible to construct a prioritarian weighting function according to which saving the first patient and the second patient in Case 1 have equal moral value (hence you should be indifferent between the two courses of action), while saving the second patient in Case 2 has greater moral value than saving the first patient. But by choosing some different weighting function, saving the second patient has greater moral value in both of the cases. That is, whether prioritarianism can handle the cases in the "correct" way depends entirely on the choice of the function along with the numbers to represent the value of outcomes. If the numbers are different (being disabled is slightly better or worse compared to being in full health, say), the same function will deliver different results. Thus, if prioritarianism can be considered to "solve" the puzzle, it does it in an *ad hoc* way.

It might be tempting to try another argument. One of the ways the two cases that I have discussed differ is that in Case 1, a disabled patient stood in competition for the life-saving intervention with a nondisabled patient. In Case 2, in contrast, both patients had a disability. The candidates were equal in this respect. Now in Case 1, the common moral intuition is that the two patients should have equal chances; in Case 2, the common moral intuition is that it is permissible to give priority to the patient whose disability can be removed. This suggests that when a disabled person stands in competition with a person without disability, it is unfair to discriminate by taking into account their capacity to benefit; but when only people with disabilities are candidates for a good, it is not unfair to discriminate by their capacity to benefit.

The problem with this idea is that it provides no explanation why we should accept this distinction. And this is precisely what is at issue. We reject the distinction, for instance, when it comes to illness and disease. Why should disability be any different?

Perhaps the distinction can be justified by the consideration that people with disabilities are often a *vulnerable group*: they are frequently discriminated against, they are often subject to prejudice, they tend to be socially disadvantaged, and their disadvantages are often compounded. These are usually (though not always) not true of those who are ill. So the impermissibility of discriminating against people with disabilities might ultimately be rooted in contingent historical, social and political conditions. The reasons for the common judgements in the cases we have considered reflect social and political experience. Just like other vulnerable groups, people with disabilities are owed special consideration, given the social and political conditions that they have to face. This is what is reflected in our ordinary moral judgements.

Appealing to vulnerability is a "political" solution to a moral problem. What I mean by that is that the solution points to extrinsic considerations. It explains the moral intuitions that we have examined, but does not justify them – at least not if discriminating against people with disabilities is thought to be *intrinsically* wrong. In that case, it is no help with the moral puzzle. In conditions of perfect background justice – that is, when all social and political forms of disadvantage and prejudice against people with disabilities have been eliminated – the problem raised by Case 1 and Case 2 would remain. It would then be permissible or fair to discriminate against the disabled patient in Case 1 by saving the second patient.

To be sure, the view that people with disabilities are owed special concern due to contingent social, historical and political conditions is common. It is held by many people. But the argument suggests that it is *only* due to

contingent conditions that discrimination against people with disabilities is wrong. If those conditions did not obtain, disability discrimination would not be unfair. It is not clear that those who hold the view that the disadvantages of disability are due to contingent conditions would welcome this conclusion.

Here is a tentative way to take these considerations into account in order to provide a solution to the puzzle. It is inspired by *rule consequentialism.* Perhaps what you should consider is the consequences of adopting a particular rule or policy in resource allocation with regard to the role of disability. For instance, if you accept the general policy of saving the nondisabled patient in Case 1, it is likely that there will be further negative consequences: distrust in the medical system by those who need it most, an erosion of solidarity between disabled and nondisabled members of the community, increasing intolerance, disrespect and prejudice against people with disabilities. In order to avoid such negative consequences, it is better to adopt a policy of giving equal chances.

There are likely to be fewer negative consequences if the policy adopted in Case 2 is to save the patient whose disability can be removed. It might increase distrust in the medical system by those whose disability cannot be removed, but it is unlikely to erode the solidarity between people with disabilities and other members of the community, or to increase intolerance, disrespect or prejudice against people with disabilities. Therefore, for these sorts of cases, it is permissible to adopt a policy of giving priority to those who have a greater capacity to benefit.

The proposal takes us back to the point I made above about having a political solution to the moral problem. Such a solution explains the moral judgement in Case 1 by appealing to the extrinsic, contingent effects of a choice or policy. It cannot demonstrate that the choice or policy is wrong for intrinsic reasons. Thus, discrimination against the patient with disability in Case 1 turns out to be merely extrinsically, rather than intrinsically, wrong. Therefore, discrimination against people with disabilities is *not intrinsically wrong*, even though it is often wrong for extrinsic reasons.

Even though some people may not find anything objectionable in this conclusion, others may find it unsatisfactory. It entails that there is a fundamental difference between discrimination by disability on the one hand, and discrimination by race or sex on the other. The latter kind of discrimination is intrinsically wrong, or unfair for intrinsic reasons. Disability discrimination, in contrast, is not intrinsically wrong. When it is unfair, that is for extrinsic reasons. In a world in which conditions of perfect background justice obtain, it would not be wrong to save the life of the second patient in Case 1.

6. Conclusion

I have argued that disability presents a puzzle for theories of resource allocation in ethics. The puzzle arises from the conflict of the intuitive moral judgements in Cases 1 and 2. For the sake of the argument, I did not question these intuitive moral judgements. I looked at a number of proposals, consequentialist and non-consequentialist alike, that might be applied to the problem. I argued that none of them is successful. I ended by introducing another proposal, inspired by rule consequentialism. I tried to show that it is able to justify the intuitive moral judgements. However, it also implies that discrimination against people with disabilities is wrong, or unfair, only for extrinsic reasons. Disability discrimination is not intrinsically wrong. This is an implication that some people might find hard to accept. They might therefore prefer, rather than accepting my proposal, to contend that the puzzle remains open.

Acknowledgements

This article grew out from a presentation I gave at the SPAWN 2014 Conference on the philosophy of disability at Syracuse University. I would like to thank the participants of the conference for their comments. I owe special gratitude to Jeff McMahan, who was my commentator at the conference, for his detailed, written comments and our discussions. In addition, I would like to thank audiences at the University of Melbourne, at the Center for Applied Philosophy and Public Ethics (CAPPE) at Charles Sturt University in Canberra, and at Stockholm University; I am also grateful to Gustaf Arrhenius, Lisa Furberg, Markus Schrenk, Carlos Soto, Gerard Vong, an anonymous reviewer and the editor for their comments. The first draft of this article was written while I was a visiting researcher at the Fondation Brocher in Geneva, Switzerland; I thank the foundation for its support and hospitality. The completion of the article was also supported by an ARWU Initiative research grant at La Trobe University in Melbourne.

References

Bognar, G. (2011) “Impartiality and Disability Discrimination.” *Kennedy Institute of Ethics Journal* 21: 1–23.

Bognar, G. (2016) “Is Disability Mere Difference?” *Journal of Medical Ethics* 42: 46–49.

Bognar, G. and Hirose, I. (2014) *The Ethics of Health Care Rationing: An Introduction*. New York, NY: Routledge.

Brock, D. W. (2009) “Cost-Effectiveness and Disability Discrimination.” *Economics and Philosophy* 25: 27–47.

Broome, J. (1994) "Fairness Versus Doing the Most Good." *Hastings Center Report* 24: 36–39.

Broome, J. (1999) "Fairness." *Ethics Out of Economics*, pp. 111–122. Cambridge: Cambridge University Press.

Green, C. (2009) "Investigating Public Preferences on 'Severity of Health' as a Relevant Condition for Setting Healthcare Priorities." *Social Science & Medicine* 68: 2247–2255.

Harris, J. (1987) "QALYfying the Value of Life." *Journal of Medical Ethics* 13: 117–123.

Kamm, F. M. (2004) "Deciding Whom to Help, Health-Adjusted Life Years and Disabilities." In S. Anand, F. Peter and A. Sen (eds), *Public Health, Ethics, and Equity*, pp. 225–242. Oxford: Oxford University Press.

Kamm, F. M. (2009) "Aggregation, Allocating Scarce Resources, and the Disabled." *Social Philosophy & Policy* 26: 148–197.

Kamm, F. M. (2013) "Rationing and the Disabled: Several Proposals." In N. Eyal, S. Hurst, O. F. Norheim and D. Wikler (eds), *Inequalities in Health: Concepts, Measures, and Ethics*, pp. 240–259. New York, NY: Oxford University Press.

McKie, J., Richardson, J., Singer, P., and Kuhse, H. (1998) *The Allocation of Health Care Resources: An Ethical Evaluation of the 'QALY' Approach*. Aldershot: Ashgate.

Nord, E. (1999) *Cost-Value Analysis in Health Care: Making Sense out of QALYs*. Cambridge: Cambridge University Press.

Rawls, J. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Shah, K. K. (2009) "Severity of Illness and Priority Setting in Healthcare: A Review of the Literature." *Health Policy* 93: 77–84.

Singer, P., McKie, J., Kuhse, H., and Richardson, J. (1995) "Double Jeopardy and the Use of QALYs in Health Care Allocation." *Journal of Medical Ethics* 21: 144–150.