https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

AGGREGATION, ALLOCATING SCARCE RESOURCES, AND THE DISABLED*

By F. M. Kamm

I. Introduction

In this essay, I first re-present certain claims that I have made in past work in which I distinguish between the moral significance of intrapersonal and interpersonal aggregation of (1) life years, (2) need, and (3) quality of life, for the purpose of making decisions about the allocation of scarce lifesaving resources among different people. (I call the latter *interpersonal allocation*.) Then I consider some partially contrasting views of Peter Singer on these matters. Finally, I present some new views—which contrast with both my earlier views and Singer's—on the role of quality and quantity of life in nondiscriminatorily allocating lifesaving and non-lifesaving resources between disabled and nondisabled individuals. I also examine how ex ante reasoning behind a veil of ignorance bears on these issues.

The issues I discuss in this essay, often using hypothetical cases, have relevance to many real-life cases in which lifesaving resources are scarce. Often, we cannot help everyone who needs and could benefit from these resources, and so we must choose whom to help. For example, flu vaccine, ICU beds, and organs for transplantation are scarce. The focus of most of this essay is on theoretical issues, especially in relation to aggregation, rather than on their application to these particular cases. However, the theoretical issues do bear on the real-life cases. For example, we shall consider whether how old someone is, or how good an outcome she will have, or how her outcome comes about, matters to whom we should help.

II. My Views in Past Work on the Significance of Intrapersonal Versus Interpersonal Aggregation for Allocating Scarce Resources among Different People

A. The Principle of Irrelevant Goods

In the course of earlier work on moral issues in the allocation of scarce lifesaving resources,[1] I sometimes distinguished between interpersonal

* This essay is based on three lectures given in May 2006, February 2007, and March 2007 in the Department of Philosophy and at the Program in Ethics and Health, Harvard University. I am grateful for comments from Ruth Chang, Shelly Kagan, Jeff McMahan, the other contributors to this volume, and its editors.

[1] For example, in my *Morality, Mortality, Vol. I* (New York: Oxford University Press, 1993). All references in the text to my "earlier work" will refer to this volume unless otherwise noted.

doi:10.1017/S0265052509090074

© 2008 Social Philosophy & Policy Foundation. Printed in the USA.

and intrapersonal aggregation. I shall here try to make this distinction salient and bring together, in a way I have not before, cases that show the difference between interpersonal and intrapersonal aggregation. By aggregation, I mean any way of combining goods, bads, values, and rights. In interpersonal aggregation, these things are distributed over different individuals and must be combined, while in intrapersonal aggregation, these things are concentrated within the life of one individual. In both types of cases, what is to be aggregated (let us say goods and bads, for short) may occur at one time or over a temporal span. As an example of how the difference between interpersonal and intrapersonal aggregation might affect decisions about the interpersonal allocation of scarce resources, I have in the past employed cases such as the following:

A		B	E		F		G
↓	vs.	↓	↓	vs.	↓	+	↓
10 yrs		15 yrs	10 yrs		10 yrs		5 yrs
Case 1			*Case 2*				

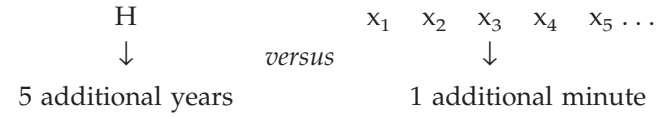
Case 1 represents the choice of saving person A for ten years or saving B for fifteen years, all other things being equal between them. Case 2 represents the choice of saving E for ten years or saving F for ten years plus saving G for five years, all other things being equal between them. My sense was that in Case 1, the extra five years for B, aggregated with the ten years that he as well as A could achieve, might reasonably *not* make a difference to whether we decide to save A or B. That is, perhaps we should continue to give each an equal chance. If so, this means that the extra five years for B is what I call a *morally irrelevant good* in this interpersonal allocation. The decision-making principle illustrated here is what I call the *Principle of Irrelevant Goods,* which says that sometimes the fact that we can produce an additional good if we choose to perform one act rather than another is morally irrelevant. This is not because that extra good is always irrelevant. It is possible that a good which is irrelevant to a choice in one context could be relevant to a choice in a different context. Hence, the Principle of Irrelevant Goods is consistent with its being preferable for B to live fifteen years rather than ten years; five years is not an irrelevant good from his intrapersonal point of view. And in the context where the only question is whether to (1) save B to live another ten years or (2) save B to live another fifteen years, we should choose the latter. Also, when producing the extra five years involves saving a third person (G), as in Case 2, then the five years, I believe, should determine that we save F and G rather than E. The fact that we could significantly affect another person (G), equal in all other respects to others,

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

counts morally in a way that benefiting one person (B) to an additional degree does not.[2]

I believe that this is true even independent of any effects of diminishing marginal utility that might be thought to account for the difference in the significance of an extra five years for B rather than for G. (That is, if G would die if he did not receive our aid and live another five years, then five years can mean more to him than it does to B, who would be saved for ten years anyway.) For suppose that G, unlike the others, will live for ten years regardless of what we do, but we can also extend his life for another five years if we save F rather than E. I suggest that this might determine that we should save F and G rather than E, even if we should give A and B equal chances. This is because we can give an additional five years to a second person if we help F and G rather than E, thus interpersonally aggregating the five years. This is in contrast to giving an additional five years to B, the same person who already will be getting ten, if we help B rather than A, thus intrapersonally aggregating the five years.

There could also be cases where intrapersonal aggregation affects interpersonal allocation decisions but interpersonal aggregation does not. Consider Case 3:



Case 3

In Case 3, we have to go either to one island to allocate a resource to H (who will then live fifteen more years instead of the ten more years he would live without the resource) or to another island to allocate the resource to an enormous number of people, equal in other respects to H, who will be saved for one extra minute of life over and above the ten years they would each live without the resource. I do not think that the interpersonal aggregation of the small benefit to each person, though it amounts to five years or more in total, should weigh against the significant benefit we can produce for H in allocating to him in this case of interpersonal choice.

[2] I have said that all candidates are to be imagined as equal in respects other than expected life years due to the resource. It is still possible that if they are equally young, the value to each candidate of the additional years will vary from what it would be if they were all equally old. At minimum, I only require one case in which the claim I have just described in the text holds—that is, the claim about the difference between the effect on interpersonal allocation of interpersonally aggregating (by giving five years to G) versus intrapersonally aggregating (by giving five years to B). Notice that even if we thought that B should be given a somewhat higher chance of getting the resource than A, this would still differentiate the effect of the extra five years in Case 1 from its effect in Case 2, where I think it could reasonably determine the outcome in which we help F and G.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

B. Need

In earlier work, I also argued that the need for a resource that provides additional life could be a function of how much (adequate conscious life) someone will have had if he is not given the resource. How much he will have had without aid is a function of how much he has had up to the point at which he needs aid and how much he will have in the future, even if he is not aided. I said that the person who would die having had the least amount of adequate conscious life was neediest for additional life, and he would be the person who would be worst off without additional life, other things equal. This means that intrapersonal aggregation of past years of life can help diminish the need for additional life. Note that if, at the time we must distribute a scarce resource, one person has had only twenty years of life and another has had forty, this will not necessarily mean that the first person is neediest. For if he would live on for twenty-five years even without the resource, but the second person would die in a year without it, it is the second person who is the neediest candidate. This is because the second person will die having had forty-one years of life while the first will die having had forty-five. (The second person is not neediest merely because he will die in a year, whereas the other person will die after many more years. For if the twenty-year-old would live for only ten years without the medicine and the forty-year-old would die in one year without the medicine, the first person would still be neediest—as I have defined the term.)

C. Outcome

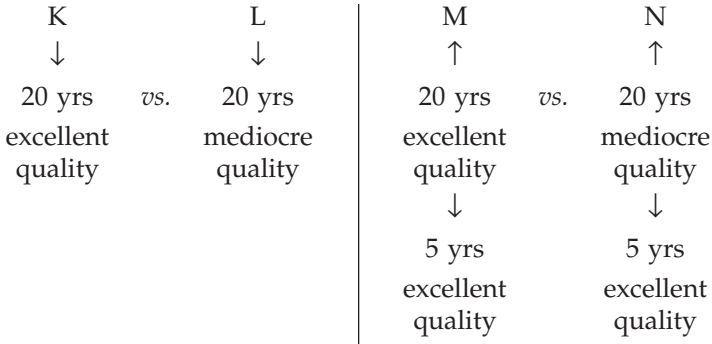
Expected outcome, in terms of future life years, also seemed morally relevant to me in deciding whom to help. If one person could only live another day with the resource and another could live five years, and other things were equal between them, then it seemed right to give the resource to the second person. Because both (1) expected future life years and (2) need (as explicated above) could matter in making interpersonal allocation decisions, we need to aggregate (i.e., combine) these two dimensions. I thought that sometimes a person with greater need should get a resource even though his expected outcome in life years would be worse, when all other factors were equal (including how good the additional years would be). For example, in Case 4, person I is someone who, if he is not saved, will die, having had twenty years of life, and person J is someone who, if not saved, will die, having had fifty years of life. If person I is saved, however, he will get only ten years of life, while if J is saved, he will get twenty, and all other things are equal between them (e.g., quality of past and future life). My view

was that person I should get the resource because his need is great (in absolute terms) and significantly greater than J's, and he would have a significant outcome (even if not as good as J's). This means that being older can be relevant to being denied a resource, even if it is not linked to a worse outcome.

A way to capture this conclusion is to aggregate need and outcome by assigning a multiplicative factor greater than 1 to need above a certain amount, varying the multiplicative factor according to degree of need, and multiplying the absolute value of outcome (in years, for example) by the multiplicative factors.

D. Quality of life

I also claimed that in making life-and-death allocation decisions, while differential lengths of future expected life in different candidates could sometimes be morally relevant, the expected quality of this life was not relevant, so long as it was minimally adequate conscious life (versus, e.g., a coma) and was desired by the person who would live it. The same would hold true for past life used to evaluate degree of need. To illustrate this, consider Cases 5 and 6.



Case 5 *Case 6*

In Case 5, the fact that K will have excellent quality of life in his expected future twenty years of life and L will not, when all other things (including their pasts) are equal, should not, I think, lead us to deny each of them an equal chance for the lifesaving resource, given that they each want to live on. In Case 6, the fact that M has had twenty past years that were better than N's should not, I think, lead us to favor N for the future five years of life, when all other things are equal between them. These views could be the result of just considering how many years of adequate conscious life someone will have had if we do not aid him and how many additional years of such life (that he wants) our resource can provide.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

E. DALYs and QALYs

This result conflicts with the use of QALYs (Quality Adjusted Life Years) and DALYs (Disability Adjusted Life Years), at least for purposes of deciding how to allocate a lifesaving resource among different people. The principle behind QALYs is to aggregate years of life and degree of good quality of life by multiplying a number for quantity times a number for quality in order to determine aggregate good, either in past years or in future years. The principle behind DALYs is to aggregate years of life and degree of disability of life by multiplying a number for quantity times a number for degree of disability in order to determine aggregate bad, either in past years or in future years.[3] If we were to use these devices to decide how to allocate a scarce resource between people, we could choose whoever will maximize future QALYs or have minimum future DALYs. Alternatively, we might give to the person who will be neediest if not given the resource, either in terms of lowest QALYs or highest DALYs. Another alternative would be to combine consideration of need with consideration of expected outcome. I took my conclusions about cases like 5 and 6 to imply that we should not use calculations of QALYs and DALYs in deciding who lives and dies. This meant that certain sorts of intrapersonal aggregation (e.g., the aggregate of years and quality) should not matter in the interpersonal allocation of lifesaving resources.

III. Peter Singer's Views on the Significance of Intrapersonal and Interpersonal Aggregation for Allocating Scarce Resources among Different People

A. Future quality of life

Now let us consider some views of Peter Singer. Singer has argued that in making life-and-death allocation decisions, there is no injustice or unfairness in deciding on the basis of future quality of life.[4] For example, whether someone will be disabled if he survives could make a difference

[3] Notice that the DALY does not seem to be a well-being measure but rather a record of a physical defect, whether or not compensation makes it have little impact on well-being. If this is true, it raises concern about the use of the measure. For example, suppose one society decides to spend x amount to compensate for a disability and another society spends the same amount to cure the disability. Well-being levels might be the same in the two societies, but one's DALY rating could be much higher than the other's. In such a case, should the DALY rating be at all relevant for deciding whom to help? I shall ignore the issue in this essay and assume (for the most part) that higher DALY ratings reflect differences in well-being or other forms of goodness in a life.

[4] For example, in Peter Singer, John McKie, Helga Kuhse, and Jeff Richardson, "Double Jeopardy and the Use of QALYs in Health Care Allocation," in Helga Kuhse, ed., *Unsanctifying Human Life: Essays on Ethics* (Oxford: Blackwell, 2002); first published in the *Journal of Medical Ethics* 21 (1995): 144–50. All references to Singer are to this essay.

to whom we choose to aid, without this involving unfairness.[5] I shall focus on three types of arguments that can be discerned in Singer's work for this conclusion, though I shall present them in my own way.

(S1) Someone's quality of life is worse if he is disabled than if he is not, other things equal. Evidence for this is that we try to cure or prevent disabilities if we can. We should maximize the good, and this implies saving the nondisabled person, other things equal. Suppose that one of two people (who would have the same quality of life) would live much longer than the other if he received a scarce treatment. It would not be problematic to choose to aid him. "For those who count only QALYs, it makes no difference whether the smaller number of QALYs gained comes from a lower quality of life or from a shorter expected life span."[6]

(S2) One would prefer to live a nondisabled life of a certain length rather than a disabled life of the same length, given the choice. We can reach this result by multiplying years of life to be had times the quality of life (i.e., by using QALYs). We can find out how to rank the quality of the disabled life by assigning 1 to a normal life and then seeing how many years of a life with the disability would be traded for normal years. People do, after all, sometimes trade length of life for quality of life. If a person would trade two disabled years for one normal one, then the quality of the disabled year is .5. The application of QALYs calculation to one person's life tells us which life for himself a person can reasonably prefer.[7] If someone can live either ten years at 1 or ten years at .5, it is only rational for him to prefer an outcome with a QALY rating of 10 (10 × 1) to one with a QALY rating of 5 (10 × .5). Indeed, as shown by the time trade-off test, it is also intrapersonally rational for someone to be indifferent as between one year without a disability and two years with such a disability that reduces the value of a year to .5.

If these are the outcomes that should be chosen intrapersonally, Singer believes that they should be chosen interpersonally as well, other things equal. For example, we should save someone with a higher QALY outcome, at least when other things are equal. We should give equal chances to live to someone who will live for five normal years and someone who will live for ten years with a .5 disability, or twenty years with a disability rated .25. Support for this conclusion, Singer thinks, comes from (S3).

(S3) Suppose one is behind a veil of ignorance and does not know whether one will be disabled or nondisabled. Singer describes the person who will be nondisabled as having "a stronger interest in continuing to live," compared to the disabled person, and as someone for whom continued life holds greater value. He says, "To maximize the satisfaction of

[5] Singer grants that other things—including the long-term effects on society of such a policy—may mitigate against choosing a candidate who promises to have a better outcome. I shall return to this point below.

[6] Singer et al., "Double Jeopardy," 284–85.

[7] I have modified the figures, but a discussion like this is present in ibid., 286.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

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."[8] Another way to put his point is as follows: Ex ante, one maximizes one's expected average utility by agreeing to a policy that will save the life of the nondisabled person, rather than giving each an equal chance of survival, if one assumes one has an equal chance of being either person. For example, if we assign a probability of 1 (certainty) to ten years ranked at quality 1 and a probability of zero to ten years ranked at quality .5 (i.e., (1)(10)(1) + (0)(10)(.5) = 10), we get a higher sum than if we assign a probability of .5 to both ten years at quality 1 and ten years at quality .5 (i.e., (.5)(10)(1) + (.5)(10)(.5) = 7.5). Singer concludes that if the just and fair principle of allocation between people is the one that they would choose in conditions of ignorance about who they will be while trying to maximize self-interest, then the principle of saving the nondisabled person is fair and just.

As Carlos Soto points out,[9] these arguments from Singer, presumably, should apply when disability, strictly speaking, is not in question, but when the interpersonal choice is between saving a life of intelligence and happiness versus a life of boredom and dullness. Singer's arguments should also make every small difference in QALYs relevant to interpersonal choice; there should be no irrelevant goods or bads.

B. Aggregation

Singer says that he is a straightforward aggregationist.[10] That is, whether goods or bads are aggregated intrapersonally or interpersonally is irrelevant to a moral decision in his view. Apart from cases where diminishing marginal utility is at issue (when we give one person more of a good rather than spreading it around), the distribution of goods and bads between people is irrelevant. Hence, he would disagree with my conclusion concerning Case 3. His view, however, also implies that if many disabled people can be saved rather than just one nondisabled person, the interpersonal aggregation of their QALYs could exceed *one* individual's better outcome.

Singer is known as a supporter of famine relief. Yet as a strict aggregationist, if he had a choice between saving a few thousand people dying of starvation or else curing the headaches of each of an enormous number of people, he should do the latter. Hence, despite his fame as an advocate of famine relief, Singer's theoretical position does not offer as strong a defense of such aid as does a position that emphasizes giving priority to

[8] Ibid., 289–90.

[9] Carlos Soto, "Choosing Whom to Aid," unpublished paper.

[10] He says this in his response to my "Faminine Ethics: Peter Singer's Ethical Theory," in D. Jamieson, ed., *Singer and His Critics* (Oxford: Blackwell, 1999).

satisfying the needs of the worst off even if this interferes with maximal aggregate benefits. It is not clear, however, that Singer consistently adheres to his aggregationist doctrine. For example, he claims that he supports the following principle governing aid: Unless something of *comparable* moral significance is at stake, we should relieve suffering.[11]

Singer thinks that this principle implies that one must bring oneself and one's family down to a level such that if one did any more in order to aid, one would be worse off than those whom one is trying to help. But this does not seem to be the correct limit, for it calls for something like a pairwise comparison of myself and any given individual I might help, to see if anyone I might help is worse off than I would be if I helped him. But Singer supposedly supports aggregation, not pairwise comparison, and it is possible that making oneself and the few people in one's family *worse off* as individuals than those whom one is trying to help might still prevent great suffering in each of many individuals and might prevent more suffering *in aggregate* when we total even minor losses in the many people whom our sacrifice helps. (This is especially likely to be true, given the relative costs to feed people in different countries; the money that comes from my depriving myself of food even if I let myself starve can be used to buy meals for hundreds of people in Africa.) For example, I think that, considered impartially, the death of me and my family at a young age is not of comparable moral significance to the avoidable deaths of thousands of others at even a slightly greater age. Hence, Singer's principle, when combined with a commitment to aggregation of equal or even lesser losses, demands more than he says it does.[12]

C. The worst off

Singer also gives no priority to helping those who will have been worse off if not aided (the neediest). Hence, aside from the issue of diminishing marginal utility, the fact that someone will have had many more years of life even if he is not aided does not (in Singer's view) weigh against him in competition with a person who will die at a much younger age if not aided, so long as the first person offers more additional QALYs in the future if he is aided than the second person would if he is aided.

Singer discusses selecting for a lifesaving procedure either someone who has had a poor quality past life through being paralyzed or someone who has been normal, when either could live for forty more years. He objects to the view that the disabled person should be saved as a way of compensating her for her worse past. This is another way of objecting to focusing on who will have been the worst off if not aided, when both

[11] Peter Singer, *Practical Ethics*, 2d ed. (New York: Cambridge University Press, 1993), 230–31.

[12] I made these points in my "Faminine Ethics: Peter Singer's Ethical Theory," 162–208.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

quality and quantity judgments are used in deciding who will have been worst off.[13] He claims, "To the extent that the purpose of health care is to lessen pain and suffering, and to the extent that nothing humanly possible can be done to lessen past pain and suffering, it would seem more reasonable to treat the person who is presently suffering more, since at least this person's suffering can be lessened. Nothing can be done about past suffering, whereas (often) something can be done about present and future suffering."[14] One might object to this that, if one allocates with an eye to the past, something can still be done about how much someone will have suffered in his life overall, and this seems important. Focusing on past pain, furthermore, may be misleading in several ways. When pain occurs, we think it is the intrinsic quality of the experience that is bad, rather than the fact that it interferes with having other good things. But if someone was disabled, he may not have been able to do as much as a normal person in the same period of time, and the things he did not do in the past (that the nondisabled person did) could be done in the future if he stays alive.[15] Similarly, if someone will die not having lived very long, this will interfere with his having more life in the future, and this future can still be given to him.

By contrast to Singer, even someone who thought it right to consider every difference in quality and quantity in life-and-death choices and equated interpersonal and intrapersonal aggregation might still give weight to need. Thus, she might choose to save a person who will have had a short life if he is not aided, even though the difference the aid makes to his future outcome will not be as good as the difference it could make to someone who will have lived a long life even without aid. She would do this because the first person will have been the worst off if not aided.

D. Further effects on overall good

Despite his views supporting the allocation of scarce lifesaving resources to the nondisabled rather than the disabled, Singer suggests a reason for treating them equally. He says, "We have rejected the claim that QALYs are unjust or unfair. . . . This does not necessarily mean that we think that health care should always be distributed so as to produce the largest possible number of health-related QALYs. There is more to overall utility

[13] Singer et al., "Double Jeopardy," 286. I have already noted that even someone who would take need into account might object to counting differences in past quality of life in deciding on need, if the past has been adequate conscious life.

[14] Ibid., 288.

[15] The impulse to compensate the disabled person would be especially strong, if quality of life were generally relevant to allocation decisions, when the disabled person would also be cured and live a normal life in the future. Yet Singer's forward-looking view also implies not giving preference to such a person in a contest with someone who had always led a normal life and would continue to do so. (But Carlos Soto raises other objections to "compensating" the nonrecently disabled in his "Choosing Whom to Aid.")

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

than health-related QALYs, and it is plausible to suppose that tilting the balance of health care towards the more disadvantaged members of society will reinforce feelings of concern and sympathy and lead to a more compassionate society. This, in turn, may be a society with more community feeling and therefore one that provides a higher level of general welfare than a less compassionate society."[16] The point is that Singer thinks the right acts are the ones that maximize utility overall. For someone who thinks this, it may make sense to perform an act that produces less than maximal utility directly because it leads to greater utility down the road than any other act. He suggests that one could recommend a policy allowing compassion for the worse off to override maximizing good outcomes, if this leads to more compassion in general and if more compassion in general leads to maximizing welfare.

One problem with this proposal from the point of view of someone trying to maximize welfare is that more compassion may just continue to lead to choices that do not produce maximal welfare. And compassion can produce no more welfare than continually doing what maximizes welfare. Hence, it seems hard for a maximizer to opt for a reliance on compassion as an indirect way to increase welfare rather than encouraging more direct maximization continually.

Singer not only considers whether compassion for the worse off could increase utility; he also suggests that people's decisions, as shown by surveys, not to maximize good outcomes in particular interpersonal allocation cases might be explained by either (1) their "concern for the effects that a direct maximization approach has on the kind of society we are,"[17] or (2) their view that it is important to "go beyond abstract justice or fairness," favoring "those who would otherwise feel themselves arbitrarily disadvantaged."[18]

Singer would no doubt prefer to interpret people's concern with a society of direct maximizers in (1) as a concern that the absence of compassion in some cases will reduce overall social welfare. However, there is another interpretation available: People may be concerned with a society of direct maximizers even if it leads to maximizing welfare, because they do not think that maximizing welfare is the correct standard for right action. With regard to proposal (2), Singer is suggesting that candidates for a scarce resource who will not have the best outcomes may incorrectly believe ("feel") themselves arbitrarily disadvantaged by a rule that selects according to best outcome, and people surveyed may believe that one way to maximize welfare in the face of such an incorrect belief is to go beyond justice and fairness. But an alternative interpretation is that people surveyed think that always deciding on the basis of maximizing

[16] Singer et al., "Double Jeopardy," 292.
[17] Ibid., 292–93.
[18] Ibid., 293.

welfare is not what justice or fairness permits or requires. For ordinarily, when beliefs are incorrect, we seek to change them rather than cater to them.

IV. The Allocation of Lifesaving Resources between the Disabled and the Nondisabled

Having laid out in Sections II and III some general positions (my own in past work, and Singer's) on aggregation and allocation, the rest of this essay will further examine the issues of intrapersonal versus interpersonal aggregation, the quality and quantity distinction, and the relevance of ex ante decision-making to the allocation of scarce resources between disabled and nondisabled persons.[19] I have elsewhere argued against Singer that the fact that we try to prevent disabilities or save people from them in an emergency room is not a sufficient ground to support the claim that disability lowers quality of life.[20] (I shall not repeat this argument here.) Nevertheless, I do agree with Singer that, in itself, and (generally) overall, disability does lower the quality of one's life, and I shall argue on this assumption.

A. Arguments for treating nondisabled and disabled alike

The Major Part Argument. Consider the following scenario: One person is on island A, and another is on island B. They share all the same properties, except that one recently lost a hand and the other did not. Because the loss to one person is so recent, the two people share equally good pasts. We can save the life of either one but not both. Arguably, each will be as badly off as the other if we do not help him, for each will be dead, having lived the same length and quality of life. But if we help the person who lost a hand, we do not produce an outcome with as few DALYs as we would produce if we saved the other person. (Call this the *Islands Case.*) I think it is morally wrong to decide whom to aid on this ground. We cannot rely on the principle of giving priority to the worse off to account for this conclusion, since each would, arguably, be as badly off as the other if not aided.

The Principle of Irrelevant Goods, however, might account for the right decision, for the following reason: In the Islands Case, what each person is capable of achieving is the part of the outcome about which each reasonably cares most, especially given the potential loss that is at stake

[19] Material in this section is based on, but adds new arguments to, a section of my essay "Deciding Whom to Help, Health-Adjusted Life Years, and Disabilities," in Sudhir Anand, Fabienne Peter, and Amartya Sen, *Public Health, Ethics, and Equity* (New York: Oxford University Press, 2005).

[20] See my "Disability, Discrimination, and the Principle of Irrelevant Goods," unpublished paper, from which material in this section also draws.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

for each. Put differently, what is reasonably held to be most important to each person can be had by either—long life saved with good quality of life. Furthermore, we should take seriously from an impartial point of view the fact that each person, from his subjective perspective, wants to be the one to survive. Given these facts, the additional benefit of a hand in one person's case can be seen as an irrelevant good in the competition to be saved. Fairness may require, therefore, that we not deprive either person of his equal chance for the great good of lengthy survival for the sake of the additional benefit of a hand in one person's case. This is especially true when that one person who would get the additional benefit is someone who would already be getting the other great benefit of additional life. That is, it is a case of an additional good concentrated in one of the two rather than dispersed over a third person.[21] I shall call this the *Major Part Argument,* for the irrelevance of some good means that either person can get the major part of what stands to be gotten, and this is all that is relevant to interpersonal allocation.

An analogy to this case is one in which only one of two people can be chosen to avoid extreme poverty. Each could reach a high income, but one could be somewhat richer than the other because he has more talent. One might argue that avoiding very bad poverty by achieving considerable wealth is the major part of what is at stake for each, and each person understandably wants to be the one to be helped; thus, the additional wealth that only one can achieve becomes an irrelevant good for purposes of choosing whom to help, at least when we are only concerned with these two people.

This case, of course, is disanalogous in that it does not involve a life-and-death choice. However, it reminds us that the Principle of Irrelevant Goods could also be applied to a *non-life-and-death* case in which we must choose whether to treat a non-life-threatening illness, such as gastritis, in a disabled person or a nondisabled person. (We are supposing that the drug to treat gastritis is scarce and we cannot treat both patients.) Suppose we could cure this illness as well in the person who lacks a hand as in the person who has two hands. One way to apply the Major Part Argument to this case leads to the conclusion that so long as each will have a life with the major part of what it is reasonable for each to want in life, each should get an equal chance for the treatment.

However, the non-life-and-death cases also raise other issues. For example, the person who recently lost a hand might be worse off without treatment for gastritis (given that he will then have to deal with *two* problems) than the person with two hands would be.[22] This, it might be

[21] For discussion of the moral difference that distributing rather than concentrating the benefit can make, see the discussion of Cases 1 and 2 above.

[22] I first discussed cases of this sort in arguing against Thomas Scanlon's view that it is only how the treatment we have to give bears on the problem it is meant for in each candidate that is relevant to its allocation, in my essay "Owing, Justifying, and Rejecting,"

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

suggested, could give one a reason to favor treating the disabled person, even though we could cure the gastritis equally well in both.[23] However, as the disabled person without treatment might still have the major part of what each person wants in life, favoring the worse-off person seems to go beyond the Major Part Argument per se.

Most importantly, when we can treat non-life-threatening illness equally well in a disabled and a nondisabled person, the overall outcome with treatment will differ by just as much, relative to what the outcome would be without treatment, whomever we treat. It is true that if we treat the nondisabled person, a person with a 1 QALY ranking may result, and if we treat the disabled person instead, a person with no more than a .9 may result. But if we had not treated the nondisabled person, he would have been alive with close to a 1 ranking anyway.[24] And if we do not treat the disabled person, we are left with someone alive with the lower QALY-rated life (at an even lower level than it would be with gastritis treatment). Only if treating the nondisabled person, for example, led to other good effects (healthwise or nonhealthwise) that would not occur in the life of the disabled person, would there be a difference in overall outcome relative to what would occur without treatment.

At least, in the Islands Case, we can acknowledge that not having a hand makes life worse, other things equal, yet this acknowledgment is compatible with the view that relative to the question of whose life we should save, the missing hand could be a morally irrelevant consideration. Hence, targeting funds to replace a missing hand in one person because life without it is worse (e.g., harder) than life with it, other things equal, is not inconsistent with giving equal weight to saving the lives of the disabled and the nondisabled. This is contrary to what a simple use of QALYs in distributing scarce resources would predict, and it is a reason for rejecting the latter approach.

The Moral Importance Argument. Consider a case involving a larger disability. We must choose between saving the life of someone who has recently become paraplegic (where paraplegia is assumed, for the sake of argument, to be rated at .5 on a QALY scale, on the basis of a time

Mind 111, no. 442 (2002): 323–54. Different issues about whether to give priority to helping the worse-off person would arise if the person without the hand had lived without it for a long time in the past as well (or even instead). I shall not discuss these different issues in this essay.

[23] However, defending such a reason would require showing that it was consistent (a) not to hold someone's disability against him because it is irrelevant and (b) to hold his disability in his favor because it is a relevant burden. For if (a) is correct, why is it not also correct that the nondisabled person has an extra good that is irrelevant when it comes to choosing against him in the matter of preventing gastritis? I was prompted to think of the issue by Carlos Soto's work on the tension between compensating for past paraplegia but ignoring it in future outcomes in life-and-death cases. See his "Choosing Whom to Aid."

[24] Presumably, we are not interested merely in creating perfect specimens, as this would commit us to curing minor conditions in the almost well rather than making big differences to those who are very badly off though they will never be perfect.

trade-off argument) and a person who would be saved to live a perfectly healthy life (a QALY rating of 1). They have identical pasts and can live an equal number of future years. It might be said that when the prospect each faces is to fall to 0 on a QALY scale (death), it is a significant good merely to achieve .5, and a person should not be deprived of the equal chance he wants to get that good merely because someone else could achieve that good plus an additional benefit that brings him to a QALY rating of 1.

Why might this be so? We cannot say, in this case, as we could in the Islands Case, that what each would get is at least the major part of what both stand to get and the difference is frosting on the cake. If one person can be saved to a life of perfect health (rated 1), this is assumed to be equivalent to an extra .5 beyond the .5 level of a paraplegic life. We can conceive of each individual's getting at least up to .5 on a QALY scale but only one person's getting the .5 increase to 1 on the scale. But reaching .5 on the scale, achievable by each, is *not* the major part of 1; it is half of 1.

In response, it might be said that it may be morally more important to give someone the basic goods that help him avoid the worst evil and *make his life one worth living* than to give him the goods which admittedly double the overall value of his life. Another way of putting this point might be that it is having a life worth living that is of crucial significance, and if paraplegics can have this, they have what is relevant to deciding to give equal chances for life, given that each individual (the paraplegic and the nondisabled person), from his personal point of view, wants to be the one to have a life worth living. I call this argument for giving the disabled and the nondisabled equal chances the *Moral Importance Argument,* because it emphasizes the moral importance of giving people equal chances at what is most important in life. The analogous argument in the case of income would claim that avoiding extreme poverty and having a reasonable income is what is most important. Hence, we should give equal chances to escape extreme poverty to someone who will achieve middle-class income and to someone who will (with a combination of what we provide and his superior talent) become a millionaire. This assumes that each wants to be the person chosen and that we are only concerned with these two people.

As this is not a life-and-death case, it reminds us of the case of choosing whether to treat a non-life-threatening illness, such as gastritis, in someone who has recently become paraplegic or treating the same illness equally successfully in a nondisabled person. One way to apply the Moral Importance Argument to this case implies that as the paraplegic whose illness is treated would have those things that it is most important to have as much as a nondisabled person whose illness is treated, each should be given an equal chance.

It is beyond the scope of the Moral Importance Argument to take account of the possibility that someone who will be paraplegic should be given

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

priority in avoiding the additional burden of the gastritis over someone who will not have the additional problem of paraplegia with which to cope. Taking account of who will be worse off, however, might well be a relevant consideration in deciding whom to aid in this non-life-and-death case. Again, a very important point in non-life-and-death situations is that if we can treat the gastritis as successfully and there are no further differential effects of nongastritis in the disabled and the nondisabled, the overall outcome with treatment would differ by as much relative to what the outcome would have been without treatment, whomever we treat.

The Sufficiently Good Only Option Argument. Another argument for sometimes ignoring the difference between a QALY rating of .5 and a rating of 1 in lifesaving cases is as follows. Suppose one can only have a life rated at .5 and not 1, and the alternative is 0 (death), which is very bad. One may reasonably want .5 as much as one would want 1 if one could have it. So, for example, given that .5 is all that one can have and 0 is very bad, one might reasonably do as much to achieve .5 (e.g., spend as much money, suffer as much) as one would do to achieve 1 if one could have it. This is consistent with the willingness even to risk losing .5 and falling to 0 for a chance at 1, as would happen if someone who is a paraplegic decides to have surgery that could cure his paraplegia, even though he might die in the surgery. (What is known as the *standard gamble test* examines how bad paraplegia is by considering how big a risk of death one would take to avoid it or cure it.) That is, the fact that one would risk falling to 0 in order to get 1 instead of .5 does not show that one would be more likely to risk falling to 0 in other pursuits (e.g., risky leisure activities) simply because one would only be losing a life rated .5 instead of 1. Hence, a paraplegic might reasonably choose to risk death in order to get a better life for himself as a nonparalyzed person because he *cares more for* (in the sense of *prefers*) the nonparalyzed life, though he *will reasonably not care more about* this nonparalyzed life, if he gets it, than he cares about the life he already has. This shows that "if one can have only *x*, one cares about it as much as one would care about *y* if one had it" is not equivalent to "one cares to have *x* as much as one cares to have *y*." It suggests that the only external measure for 1 being worth more than .5 to someone is the willingness to exchange .5 for 1 but not vice versa, or to risk falling below .5 to go from .5 to 1 but not vice versa. The difference in worth between 1 and .5 need not show up in a difference in the other goods that one would give up to get or keep 1 or .5, if either were one's only option. All this may seem paradoxical, yet I think it is true.[25]

[25] I thank Susan Wolf, David Sussman, and other members of Philamore for discussion of this point. It is worth pointing out, as an analogy to someone's concern for his own life, a parent's concern for the life of his child. One child may have many more good traits than another. A parent could prefer that his child have better traits. He might even allow the child to take certain risks in order to become better for the child's own good. But that does not mean that a parent would make greater sacrifices to save the life of the child when he has

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

What I have said conflicts with Singer's view, described in (S2) and (S3) in Section III above, that the nondisabled person has a stronger interest in going on living than the disabled person. Singer concludes this on the basis of the time trade-off argument (rather than the standard gamble argument). He describes someone with a severe disability who would (let us assume, reasonably) exchange forty years with the disability for four years without it.[26] He concludes that her interest in living forty years with the disability is less than a nondisabled person's interest in living forty years. But what I have said implies that the fact that someone reasonably *cares for* a shorter life with higher quality need not imply that she reasonably *cares about* her only option of the longer life with lower quality less than a nondisabled person cares about the same length of life.

It also supports my view that willingness to take an *intrapersonal* risk of death in order to achieve a better life for oneself does *not* translate into the permissibility of imposing the same risk of death interpersonally on someone whose own life cannot be improved by it, if imposing that risk will result in a better life for another person.[27] It supports my view that an intrapersonal preference for a life rated at 1 over a life rated at .5, and thus the relevance of the aggregation of .5 and .5 intrapersonally, does not translate into the importance for interpersonal allocation of a difference between 1 and .5.

These results are contrary to the results implied by Singer's arguments (S1), (S2), and (S3). For arguments (S2) and (S3) imply that whatever chance of death it would be reasonable to take to avoid one's own disability is the same chance of death it is reasonable to impose on one person so that someone else will live a nondisabled life.[28]

To make my point even more graphically, suppose that someone could undergo surgery that has an equal chance of (a) curing his disability or (b) definitely killing him if he remains disabled. In this case, if the person cannot be rendered nondisabled, he will definitely die. The fact that it

better traits than when he has less-good traits. And this may be due to love for the actual child, rather than from duty.

[26] Singer et al., "Double Jeopardy," 286.

[27] Though this view might be supported in other ways as well.

[28] However, it is worth emphasizing a difference in the causal route in the intrapersonal and interpersonal cases. Intrapersonally, a person may give up, or risk giving up, something of his in order to get something better for himself. The exact parallel in the interpersonal context would be that someone gives up something, or risks giving it up, in order to cause someone else to get something. Here the "in order" involves a causal relation. Someone might object to a principle permitting the latter and yet distinguish it morally from what happens when we outright give a scarce resource to one person with the result that another person loses his life or loses a chance to keep it. In this case, the loss, or risk of it, for the second person does not cause the gain to the first person. Without denying the moral significance of this difference, what I am discussing throughout this essay is the possibility that one person could object to our concluding, on the basis of the reasonableness of his intrapersonal sacrifice of something, that it is permissible to outright give a scarce resource to another person when it results in the first person's losing, or risking the loss of, that same type of thing.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

could be reasonable for someone to choose such surgery does not imply that it is reasonable to definitely let a person die who will be disabled in order to be able to save *another person* who will be nondisabled. This is true even though, ex ante, the probability of being a disabled person or a nondisabled person is the same in the intrapersonal and interpersonal scenarios, and the ex ante chance of death is also the same. It could even be true when the intrapersonal chance of death that it is not unreasonable to take in order to be nondisabled is greater than the ex ante chance of death one would run if the nondisabled person is automatically helped relative to the different disabled person.

The failure to appreciate the difference between "caring for" and "caring about" may be the result of not distinguishing two different notions of the worth of life. First is the notion that involves evaluating lives as better or worse, where all properties of the life are included in this evaluation. (When Singer speaks of those whose continued life holds greater value, he may have this in mind.) Second is the notion of the worth of life, or going on living, to someone. On this second notion, the quality of the person's life—thought of as a set of synchronic properties that modify any period of her life—is treated as a background condition, and we ask whether going on living for a certain quantity of time—thought of as a diachronic property—is worth as much to her as it would be to someone who had different synchronic properties. The view I am proposing is that going on living could reasonably be worth just as much to someone who has a less favorable set of synchronic/quality properties as to someone who has a more favorable set.[29] (If so, we cannot conclude that a life that holds greater value is also, as Singer claims, the life that a person has a "greater interest in continuing to live.") This implies that going on living, at least as a self-conscious person with a certain amount of good in his life, is a separate good for someone, to be distinguished from other quality-of-life goods that his life may be instrumental to achieving.

Nonetheless, there are such quality-of-life goods, and someone who lacks them might exchange some quantity of life (or risk losing all quantity of life) in order to get the other goods. This does not mean that if he is reasonable, he should care less for continuing life per se if he cannot have these quality-of-life goods, as compared to someone who can have these quality-of-life goods. (Still, there are also some conditions, such as extreme unending suffering, that negate or override the value of continuing to live.) The person with life of higher quality has what it is more in any person's interest to have, by comparison with a lower quality of life. But the fact that one person will get more of what it is in anyone's interest to have, if he goes on living, need not imply, I suggest, that it is more in

[29] This claim may be in some tension with the view that the badness of death for someone is a function of the goods of which it deprives him.

his interest to go on living than it is in the interest of another person to go on living.

Notice that some lives are worth living (e.g., an additional three happy months of life), but it might not be reasonable for the person, whose only alternative it was, to sacrifice as much to gain this additional three months of life as it would be to gain a much longer period of life with a QALY rating of 1.[30] Hence, it is not always reasonable to do as much to get one type of worthwhile life, even if it is all one can get, as it would be reasonable to do to get a better type of worthwhile life. Hence, unlike what is true of the Moral Importance Argument, in this additional argument, it is not merely having a life worth living that is crucial. On account of this, I shall call the additional argument I have now given the *Sufficiently Good Only Option Argument*.[31]

If either the Sufficiently Good Only Option Argument or the Moral Importance Argument is correct, it helps expand the reach of the Principle of Irrelevant Goods. This is because these arguments let us see how the Principle of Irrelevant Goods can be used to argue for treating equally those with large disabilities, even when the Major Part Argument cannot be used to support the Principle of Irrelevant Goods.

What does the Sufficiently Good Only Option Argument imply for cases in which we must decide whether to treat non-life-threatening gastritis in someone who is recently paraplegic or in someone who is nondisabled, other things equal? One way to apply the argument implies that if life as a paraplegic without gastritis is a sufficiently good only option, then we should give equal chances to this person and the nondisabled person. It might also be argued that since the paraplegic would have to cope with two problems if he also had gastritis, but the nondisabled person would have to deal only with gastritis, we should favor the disabled person when deciding whom to treat. This view, however, goes beyond what the Sufficiently Good Only Option Argument itself implies. Again, in non-life-and-death cases, a very important point is that if we can treat the gastritis as successfully and there are no further differential effects of nongastritis in the disabled person and the nondisabled person, the overall outcome would differ by just as much, relative to what the outcome would be without treatment, whomever we treat.

[30] I owe this point to David Sussman.

[31] This argument implies that when .5 by itself is compared with 0, it is worth some maximal sacrifice *x,* and when 1 by itself is compared with 0, it too is worth *x.* This suggests that by the measure of sacrifice, and relative to 0, .5 equals 1. But that does not mean that .5 is equivalent to 1 per se on all measures. For example, we could be willing to give up .5 for 1 and also willing to risk going from .5 to 0 to get 1, but (obviously) not willing to risk going from 1 to 0 to get .5. This argument could even be taken to generate an intransitivity: that is, .5 equals *x,* and 1 equals *x,* but .5 does not equal 1. But the supposed intransitivity is explicable because of the effects of different contexts, where different alternatives are available: .5 is worth *x* in a context where 0 is the only alternative; 1 is always worth *x;* .5 is not worth as much as 1 when both are alternatives to 0 for a given person.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

Being on a par. It is worth noting that the three arguments just discussed, which support the irrelevance for interpersonal choice of a good that is relevant for intrapersonal choice, seem to be similar to arguments about values that are "on a par." It has been said that if x and y are of equal worth, and $x + n$ is of greater worth than x, it should be of greater worth than y. But if x and y are merely "on a par," then $x + n$ may not be worth more than y even if it is worth more than x. An example that can be used to illustrate this point is as follows: Medium good baroque music is on a par with medium good romantic music; slightly better baroque music is clearly better than medium good baroque music, but it is not necessarily better than medium good romantic music; the two can remain on a par.[32]

This might be analogous to saying that while in the life of person A it is better to have a QALY rating of 1 than a rating of .5, a rating of .5 in the life of person A can remain on a par with a rating of 1 in the life of person B. A clearly preferable outcome, which should be selected in an intrapersonal choice, remains on a par in an interpersonal choice. What could make them be on a par is that they are being considered relative to death as the alternative, and each person wants to be the one to live. Each person functions like a separate category, the way baroque and romantic music do, and those categories provide insulation from intrapersonal (intracategory) changes in interpersonal (intercategory) rankings of outcomes. On this view, the separateness of persons makes changes that are relevant intrapersonally irrelevant to the interpersonal choice.[33]

B. A problem with the arguments: Differences in quantity of time

Now consider what I think is a problematic implication of the Major Part Argument, the Moral Importance Argument, and the Sufficiently Good Only Option Argument in life-and-death cases. They would seem to imply that we should treat sufficiently good only options that involve *quantity* of life in the same way as we treat sufficiently good only options that involve *quality* of life. That is, large differences in how long someone can live if we save him should make no more difference to whom we save than large differences in quality of life, other things equal. (This is the

[32] This issue is discussed by Derek Parfit in his *Reasons and Persons* (New York: Oxford University Press, 1981), and by Ruth Chang and others in *Incommensurability, Incomparability, and Practical Reasoning,* ed. Ruth Chang (Cambridge, MA: Harvard University Press, 1997).

[33] Ruth Chang has suggested (in conversation) that the separateness of persons serves as insulation in interpersonal allocation even when it is not reasonable for one person to care about his sufficiently good only option as much as another person cares about his better option. This is certainly true when we would have to take away from someone something that is his, in order to help someone else. I believe she is suggesting that it may be true even when we deny him help, as in the cases we are discussing. This would mean that the explanation of the fact that an additional good that matters in the intrapersonal context does not matter in an interpersonal context can be independent of the distinction between "caring for" and "caring about." The possibility of such an alternative explanation was suggested in note 27 above.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

concern raised by Singer that I describe in (S1): How can we distinguish quality and quantity of life in allocation decisions?) For example, suppose one person can be saved to live for five years and another for fifty years, and everything else is equal between them. Five years is a very significant good, and given that it is someone's only option, she might well reasonably do everything to get it that someone who could live for fifty years would do to get that. If there is a moral difference between taking account of expected length of life (both in the future and in how long someone will have lived if not aided) and taking account of disability or some other quality-of-life factor (both in the future and in the life someone will have had if not aided), then another argument apart from those we have considered is needed to justify this.

Here is one possibility. Each person is entitled to equal respect and (at least for purposes of an impartial distribution of scarce resources) equal concern. That may mean that (at least certain) synchronic properties, such as whether one is or will be paralyzed, even assuming that they significantly affect quality of life, should not bear on selection for scarce resources. If these synchronic properties are appropriately thought of as determining one's identity, one might say that equal respect makes identity irrelevant for purposes of allocation. (Call this the *Principle of Irrelevant Identity*.) However, taking into account, for example, how long a person can live if he gets a scarce resource is not treating someone differently because of the type of person he is or will be qualitatively; the latter (it is being suggested) is done only if we consider someone's synchronic properties (properties that determine the character of his time alive). It is, theoretically, compatible with each synchronic type that a person could be, that he could be that type for longer or shorter amounts of time.

However, this does not rule out that having a certain synchronic property could cause longer or shorter life. For example, having a disability might make impossible doing exercises necessary for longevity. This is what I call linkage—a causal relation between a disability (or any other property) and other effects that are morally relevant to allocation. One does not, I think, hold against someone his synchronic property per se, if one takes account of its causal links. Hence, it may not violate equal respect and concern for different types of persons to consider how long they will have lived or will live.[34]

This argument does not focus on the worth of living to a person; rather, it offers an interpretation of meeting the requirements of equal respect and concern, so I shall call it the *Equal Respect Argument*. In another way,

[34] The analogous contrast in the case of income is between the following two choice-sets: (1) One person can avoid extreme poverty by having $50,000 a year and another by having a million a year; (2) One person can avoid extreme poverty for five years and another for fifty years (whether at $50,000 a year or a million). The analogous conclusion would be that it could be reasonable to help the second person in (2), even if both should have an equal chance in (1).

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

it attempts to extend the Principle of Irrelevant Goods while distinguishing quality and quantity considerations. It suggests that we may take account of large differences in outcome when this does not deny equal respect and concern to people in virtue of synchronic properties they have or will have. Thus, if we had a choice between saving a paralyzed person (with a quality rating of .25) who could live for twenty years, or an unparalyzed person (with a rating of 1) who could live for six years, the Equal Respect Argument could recommend saving the paralyzed person. This is so even though, in an *intra*personal choice, the aggregation of quality and quantity could make it reasonable for someone to choose six normal years over twenty years with a quality as low as .25.[35]

Does the Equal Respect Argument have implications for cases where we must decide whose non-life-threatening illness should be treated? One possibility is that equal respect and concern for different types of people implies giving equal chances for treatment of gastritis to the disabled and the nondisabled person. However, equal respect and concern need not mean equal treatment, but rather treatment as an equal. Suppose the disabled person who is to be treated as an equal would suffer more healthwise overall than a disabled person if he has gastritis because he also must deal with a disability. Possibly equal respect and concern *itself* implies that the disabled person should be given priority for treatment. Again, it is also important to remember that in non-life-and-death cases, the difference in outcome with treatment *relative* to outcome without treatment can be the same, whomever we successfully treat.

C. The Treatment Aim Principle

Are the results of the Major Part, Moral Importance, Sufficiently Good Only Option, and Equal Respect arguments, and the Principle of Irrelevant Goods that they support, consistent with what is standardly understood to be the correct account of a nondiscriminatory policy on the treatment of the disabled and the nondisabled? This account says that if a treatment aims to correct a particular problem (e.g., gastritis, or upcoming death due to kidney failure) and is equally successful in achieving that aim in a disabled and a nondisabled person, the difference in outcome represented by the continuing presence of the disability in one person is irrelevant. I shall call this the *Treatment Aim Principle.* It is one way to capture the result noted above, that in non-life-and-death cases,

[35] How long someone will live is not a quality of a life such that someone might say that we are discriminating against those who will not live long (or have the property of short-livedness), at least when we are deciding so as to determine whether there will be a long or a short life. A context in which we would be discriminating against those who will not live long is when we refuse them admission to a park, for example, more than we refuse admission to those who will live a long time, even though both can make equally good use of the park.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

the overall outcome with treatment would differ by just as much, relative to the outcome without treatment, whomever we treat.

However, for various reasons, the Treatment Aim Principle differs from the other principles so far considered. First, the Principle of Irrelevant Goods could imply treating candidates equally when the Treatment Aim Principle would distinguish between them. This is because even if the aim of treatment is to correct a particular problem, the Principle of Irrelevant Goods would imply that some differences in the degree to which *that* particular problem is successfully treated could also be morally irrelevant in deciding whom to treat. So, if one person's gastritis could be treated slightly less successfully than another's, this might be morally irrelevant to who gets the treatment. Alternatively, another way to apply the Moral Importance Argument and the Sufficiently Good Only Option Argument is to argue that if the most important part of the good of treatment is possible in either candidate, or if each could get a sufficiently good improvement due to treatment, we should not select on grounds of different success in treating. If this were true of a difference in treatment outcome between two nondisabled candidates, it should also be true that the difference is not what determines a selection between a disabled and a nondisabled candidate. A similar conclusion could be drawn about life-and-death cases where one candidate would go on to live somewhat longer than another.

What if we could treat a fatal or nonfatal illness *much more* successfully in one person than in another? The Treatment Aim Principle should favor the person in whom treatment is much more successful. For example, suppose that someone's disability interfered with his doing exercises that should accompany a drug that treats gastritis, and so his treatment would not be successful. The principle takes this as a conclusive reason to help the nondisabled person. The other principles I have considered need not lead to this result, if they are combined with giving priority to the worse-off person. But the Treatment Aim Principle is inherently in tension with an argument based on the importance of need in deciding whom to aid (i.e., priority to the worse off) when treatment is less successful. In the case of lifesaving treatment, giving priority to the worse off could direct us to give a successful lifesaving treatment, or even one that is significantly less successful, to the person who would be overall worse off if he did not get it (in terms of years lived) rather than to someone else (who would die at an advanced age, for example). By contrast, the Treatment Aim Principle by itself not only ignores this other consideration but seems to exclude it. It implies that we should just consider what our treatment itself can do for people, rather than considering additional problems (e.g., dying much younger or having an additional problem with which to cope) that one person would have rather than another.

Putting to one side the issue of helping the neediest (and continuing our assumption that the disability is recent), let us consider other cases in

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

which treatment of fatal or nonfatal illnesses will be much more successful in the nondisabled person than in the disabled person. Then both the Treatment Aim Principle and the other principles I have considered could favor the person in whom treatment is most successful, even if it is due to his disability that the treatment is much less successful in the disabled person. This is an example of linkage—a causal relation between a disability and other effects that are relevant to allocation decisions. Consider, first, life-and-death cases. Suppose that if we give a scarce, lifesaving organ to a nondisabled person, he will live twenty years. If we give it to a disabled person, he will live five years because his disability makes it impossible to do as good surgery on him and/or because he cannot do certain exercises subsequent to surgery. In this case, the Treatment Aim Principle, as well as the Equal Respect Principle, can recommend giving the organ to the nondisabled person. In this sort of linkage case, we are not holding against the disabled person his disability as a lived component of his life but rather considering its causal effects on how many years someone will live if he is treated. This is no different from considering the causal effects of being *non*disabled—a state which is considered preferable to disability in itself—if it were imagined to directly cause a much poorer length-of-life outcome of a lifesaving surgery. In the latter case, according to the Treatment Aim and Equal Respect Principles, the surgery should be performed on the disabled person.

Consider a nonfatal illness. Suppose that someone's disability prevents a successful surgery to correct gastritis. Then the Treatment Aim Principle could recommend allocating the scarce surgery to the nondisabled person. The same conclusion might be yielded by one way of applying the Sufficiently Good Only Option Argument if the outcome of treatment did not make a sufficiently good improvement in the disabled person due to his disability.

However, there is another way that linkage can occur and be relevant to the allocation of scarce resources and that the Treatment Aim Principle does not capture. Suppose that a disabled person and a nondisabled person each have a fatal disease *X* and can be treated equally well for it, so that *X* does not lead to a shorter life in one person than in another. However, the disability on its own will result in death shortly after disease *X* is cured. Strictly speaking, treatment of disease *X* is equally good in either patient, and so the Treatment Aim Principle should be neutral as to whom we treat. But, presumably, it is wrong not to take account of the fact that the very same negative factor that our treatment seeks to avoid—namely, death soon—will occur anyway, though from a different cause. The Treatment Aim Principle by itself does not distinguish between helping someone who will succumb in a few weeks to another illness and helping someone whose successful treatment implies that he will live for twenty years. By contrast, the other principles that I have discussed can recommend that the nondisabled person get the scarce resource in such a case.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

What if a scarce treatment for a non-life-threatening illness that causes a lot of pain can work equally well in a disabled and a nondisabled person, but the disabled person will wind up with the same degree of pain caused by his disability alone. Again, the Treatment Aim Principle, strictly speaking, would recommend not distinguishing between the two candidates, but the principles that I have discussed need not do this.

To avoid these problems, we would have to modify the Treatment Aim Principle to what I call a *Treatment Similarity Principle,* which would focus on whether we could cure a fatal or nonfatal illness without there being conditions similar to the ones we aimed to cure occurring in the patient in any case. (The Treatment Similarity Principle, however, will not correctly deal with a case in which we could successfully treat a *nonfatal* disease equally well in a disabled or a nondisabled candidate, but his disability will soon cause the disabled person *to die.* This is because while the soon-to-occur death should undercut the candidacy of the disabled person, this is not because a negative condition *like the one* we are trying to treat will occur in any case from a different cause. Further modification, which I shall forgo, could correct this problem.)

So far, I have discussed the role of linkage in making it permissible to deny a scarce resource to a disabled person, whether we are using the Treatment Aim Principle or other principles I have discussed. But linkage involves a disability having a causal relation to some other property relevant to allocation. The Treatment Similarity Principle shows us a way in which the disability property itself (i.e., the synchronic property) could sometimes be morally relevant to the allocation decision. Suppose a nonfatal disease will cause paraplegia and also pain; and we are interested in treating it in order to prevent both of these bad states. We can prevent the pain equally effectively in either a paraplegic or a nonparaplegic candidate, but the fact that we can successfully prevent the disease from attacking the nerves in either candidate results in the prevention of paraplegia only in the candidate who is not already a paraplegic. If avoiding paraplegia is the most important part of our aim, it could be morally correct that the disabled person not be a candidate for the treatment. Furthermore, this is not because of the further causal effects of his disability but simply because of (what I call) the role of disability as a *component of his life.* We have previously seen that the Treatment Similarity Principle extends the limits of the Treatment Aim Principle. Now, we see that it also sets a limit to the Principle of Irrelevant Identity (that I introduced in connection with the Equal Respect Argument). It accounts for why it may sometimes be permissible to deny a scarce resource to someone who is disabled because he is disabled per se.

Consider again the Treatment Aim Principle. Arguably, a treatment can be considered less successful even if it achieves its aim, if it has bad enough side effects. Suppose that giving a scarce non-lifesaving treatment to a disabled person will result in a bad side effect caused by the inter-

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

action of the treatment with his disability. The bad effect is not so bad as to totally outweigh the good the treatment does him. A nondisabled person will get the same good effect without the bad side effects. The Treatment Aim Principle, if it were modified in another way, might imply that we should help the nondisabled person. The other principles I have discussed could imply the same if the side effects were significant enough.

In the next section, I shall consider further the possible role of side effects of treatments on the allocation of scarce resources to the disabled versus the nondisabled.

D. Further grounds for sometimes not ignoring disability when allocating scarce resources

Switch Cases and the Causative Principle. Suppose that the Principle of Irrelevant Goods implies that we should not prefer saving an unparalyzed person to saving a paraplegic person (other things equal). Then it would also imply that we should not prefer saving a paraplegic who, as a side effect of lifesaving treatment, will also be able to walk again to saving a paraplegic who will remain a paraplegic. That is, it would imply *no role* for such a quality difference in expected outcome among equally disabled candidates. It would also imply no role for such a difference in outcome between equally nondisabled candidates—so if one unparalyzed person would become paraplegic as a side effect of lifesaving treatment but another would not, this should make no difference in whom we choose to save. I call these cases in which the disability condition of a person changes as a result of lifesaving treatment the *Switch Cases.* Are these results correct?

These Switch Cases raise the following possibility: (1) Sometimes a sizable extra synchronic good (or bad) *that we can produce* in the outcome, if we treat one person rather than another, should be morally relevant in deciding whom to help with a lifesaving resource.[36] (2) Yet, if candidates for treatment who present themselves have this difference in good (or bad) between them, and this is why it shows up in the outcome, the extra good (or bad) should be morally irrelevant in deciding whom we help with a lifesaving resource. The Principle of Irrelevant Goods cannot account for the simultaneous truth of (1) and (2). I call this the *Asymmetry Problem.*

Proponents of the view involving both claims (1) and (2) need a principle that will explain why the fact that a person is and will be disabled to a certain degree should sometimes be irrelevant in deciding whose life to save, but the presence of the same disability in an outcome can sometimes be morally relevant and used in a nondiscriminatory fashion in deciding whose life to save. That principle would imply that a QALY evaluation could sometimes be relevant in making life-and-death choices,

[36] This revises the view I presented in my *Morality, Mortality, Vol. 1.*

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

depending on how the difference in quality comes about. This implication differs from my position described in Section II, which excluded a role for most quality-of-life evaluations in allocating scarce resources.[37] It is still different from Singer's position, described in Section III, which would generally require quality-of-life evaluation, other things equal.

There could also be Switch Cases and an Asymmetry Problem in non-life-and-death situations. For example, suppose that a scarce drug that prevents impending blindness will also cure paraplegia in one candidate but not the other. Alternatively, the same scarce drug could be imagined to prevent blindness in either nondisabled candidate but cause paraplegia as a bad side effect in only one candidate. These Switch Cases are meant to contrast with one in which we can prevent blindness in either someone who has recently become paraplegic and will remain so or in someone who will remain nondisabled. In these cases, the principles suggested earlier imply that either equal chances should be given or we should take into account that someone will have a much harder life if he will be both paraplegic and blind rather than just blind. The Principle of Irrelevant Goods and some other principles discussed earlier might be understood to imply that we should ignore the good and bad side effects in allocating the drug in non-life-and-death Switch Cases, at least on the supposition that blindness is a significantly worse fate than paraplegia.[38] Those who support claims (1) and (2) in non-life-and-death situations will also need a principle that makes these compatible.

Even someone who thought a difference in outcomes such as disability should not matter, regardless of how it comes about, might be interested in seeing if we can distinguish a *discrimination objection* to counting the difference from a more general objection to counting such a difference in outcome based on the Principle of Irrelevant Goods. One claim about these Switch Cases might be that sometimes, even if we incorrectly violate the Principle of Irrelevant Goods in taking account of certain differences in outcomes, we would not be engaged in (invidious) discriminatory conduct. Our conduct might be wrong, but not because it involved discrimination.

It will be useful to present some cases and the judgments in each for which the new principle would try to account.[39] Let "P" stand for paraplegia and "U" for unparaplegia in the following cases in which we must decide to whom to give a scarce lifesaving procedure.

[37] Hence, this implication departs from the view I defended in *Morality, Mortality, Vol. 1,* that in life-or-death situations, quality of life (above a certain minimal level) should not affect our decision about who gets a scarce lifesaving resource, given that each person wants the resource.

[38] If this were not true, we might imagine that the drug prevents impending blindness *and* deafness, on the supposition that these together are worse than paraplegia.

[39] I will here deal only with cases in which treatment also causes the side effect of a cure for a disability in a recently disabled person. A fuller range of cases would also include ones in which a treatment also causes the side effect of a disability in a previously nondisabled person. I discuss the fuller range of cases in "Deciding Whom to Help, Health Adjusted Life-Years, and Disabilities," and in "Disability, Discrimination, and Irrelevant Goods."

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

	Case 7		*Case 8*	
Persons	A	B	E	F
	P	U	P	P
	↓	↓	↓	↓
	P	U	U	P

In Case 7, a lifesaving procedure will save A or B, but not alter *their* initial status as paraplegic or not. I shall assume at this stage (given what was said in Section IV.B) that we should not favor B because a better outcome will thereby result (one with more QALYs).

In Case 8, we select between two paraplegics, but only in the case of E will the lifesaving procedure have the additional good effect of curing her paraplegia. In this case, I believe some might think it permissible to choose to save E rather than F on the basis of the fact that we thereby get a better outcome by producing a cure of paralysis, as well as saving a life.

Notice that Case 8 differs from Case 9:

	Case 9	
Persons	E*	F*
	P	P
	↓↓	↓
	U	P

In Case 9, the double arrow leading from P to U is intended to symbolize the fact that E*'s unparalyzed state is not caused by the lifesaving procedure or anything else we do, but rather by an independent cause that would have resulted in an end to E*'s paralysis so long as he lived on. While the state of E* when we treat him is paraplegia, he is the sort of person who, independent of anything we do, will become nonparalyzed. I believe, in this sort of case, that those whose judgments I am considering might say we should *not* decide to save E* on the basis of his expected better outcome. That is, Cases 7 and 9 are morally alike. There should be no difference in how we decide to allocate a lifesaving resource just because one person is permanently paraplegic and the other is only temporarily paraplegic *in this way.* However, to be "temporarily paraplegic" because our lifesaving cure can also sometimes cure paraplegia has, it might be argued, different moral significance.[40]

[40] I am using paraplegia in discussing the Switch Cases. But it is possible that only if we could cure a worse disability (e.g., quadriplegia) would it be appropriate to respond differently to cases in which a better outcome occurs as a result of our curing a disability as opposed to our saving a person who is not disabled. I am concerned only with whether such causation ever matters morally, not so much when it does.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

A principle that can account for the responses I have described in the preceding cases can be referred to as the *Causative Principle.* It tells us that we may decide whom to help based, in part, on the synchronic difference we can make to a person's situation, not on the difference he brings to the situation. (By "we," I do not mean doctors in particular, but the health-care intervention system generally.) More particularly, the Causative Principle is concerned with the differential effectiveness of our treatment in producing nondisability and our being entitled (though not obligated) to bring about a better outcome by using our skills in this way. (The principle can also be concerned with not causing disability. So it could distinguish between someone who was and remains P and someone whom we cause to be P. I shall not discuss that aspect here.[41] It could also be concerned with causing someone to remain U who would otherwise change from U to P.) The Causative Principle tells us to ignore this difference in outcome when it arises in any other way, whether because the disability inheres in the person or will arise because of what inheres in him, or even will arise from causes outside of him other than our treatment. (For example, if we know a criminal will do something to one nonparalyzed person to make him paraplegic if he lives, we should ignore this knowledge.) This is all against a background in which the outcome that results is (i) still one that is a life worth living, and (ii) reasonable for someone to do a great deal to retain. I also continue to assume that one can permissibly compare outcomes regarding length of life for purposes of choosing whom to save, even when the difference in quantity of life is *not* caused by what we do (though it may follow on what we do). (The difference we *make* does not include every change that will occur in people's lives through other causes that follow on what we do to them.)

For convenience, I will just say that the Causative Principle tells us to "ignore who the candidates are, as evidenced by the synchronic properties they have and will have, and look to what we do." The Causative Principle can be combined with a limited use of the Principle of Irrelevant Goods, in that some differences we cause are still morally irrelevant if they are relatively too small.

In Case 7, we save life in the paraplegic person as much—by hypothesis, no more than and no less than in terms of length—as we save life in the nonparaplegic person, and we do not cause nonparaplegia in the nonparaplegic person. By contrast, in Case 8, we cause the additional large good of nonparaplegia in one patient. In doing this, we produce a nonparalyzed life *by saving a life and by making it unparalyzed.* By contrast, when we save the life of a nonparalyzed person, we produce a nonparalyzed life *by saving a life that is nonparalyzed;* the outcome *piggybacks* on a property the person brings with him. We might say that this is

[41] I discussed it in my "Deciding Whom to Help, Health-Adjusted Life Years, and Disabilities," and in my "Disability, Discrimination, and Irrelevant Goods."

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

the difference between doing what makes people unparalyzed and doing what makes unparalyzed people. (There is a comparable distinction between causing a paralysis, and a paralysis piggybacking on a property someone brings with him.)

In non-life-and-death cases, there is a comparable difference, for example, between (a) an outcome in which a person is nonparalyzed and not blind because we piggyback by preventing blindness in a nonparalyzed person, and (b) an outcome in which a person is nonparalyzed and not blind because we cure paralysis as well as prevent blindness.

One proposed justification of the Causative Principle is that when outcomes are affected by who a person is and/or by what we do, counting only what we do is consistent with the account of respect for different types of persons given above (in Section IV.B), at least so long as what we do is significant. Counting the qualitative differences that the people themselves bring, I have suggested, is not consistent with respect for different types of persons, except when the difference results in a life below a certain minimum (e.g., a life not worth living, or not worth doing a lot to save), or there is significant conflict with the Treatment Aim Principle or the Treatment Similarity Principle. However, we can reasonably value paraplegia less than nonparaplegia as a state, without this in itself showing disrespect for paraplegic people, and without treating a person who is or will be (through causes independent of us) in the state less well just because he is or will be in the state. But if we accept the Causative Principle, valuing the nonparaplegic state more than the paraplegic state can have a worse consequence for one person than for another. For if we use these values to choose whom to aid when we can, for example, do more good for one paralyzed person by making him nonparalyzed, the person who will not be cured of paraplegia will not be given the scarce lifesaving resource (and a person who will not be cured of paraplegia will not get the scarce blindness-preventing resource).

Notice that the Equal Respect Argument said that choosing whom to save based on how long each of the people will live is consistent with respect for different types of people, and it drew no distinction between extra life being (1) something we produced (e.g., by being able to do a certain type of procedure on one person but not another), and (2) something that results only from what the person brings with him (e.g., extra genetic hardiness). It is only in dealing with the synchronic properties that characterize the time a person lives and the type of person he is that, I suggest, we should distinguish between producing and piggybacking in order to act consistently with equal respect for persons.

One concern about the Causative Principle is how it will deal with other quality-of-life differences. That is, suppose that of two people who have always been mentally dull, if we save the life of one person but not another, we can also make him intelligent. Should the fact that we can

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

cause this difference make it permissible for us to choose to save one person rather than another, when all else is equal between them and each wants to live for an extra ten years each could have?

Treatment Aim Principle Modified. As a preface to introducing the Switch Cases, I noted that it is possible that a treatment is often considered more or less successful depending on whether it has bad side effects, even if it achieves the aim of treatment per se. This raises the possibility that the Treatment Aim Principle might also be able to deal with the Asymmetry Problem raised by the Switch Cases.

Suppose that if we use a certain treatment for a deadly heart problem in one paralyzed person, but not if we use it in another, it also cures his paralysis. This was not our original aim. For this reason, the Treatment Aim Principle might seem to be equally satisfied whichever paralyzed person is treated. However, it might be argued that once we know that the drug can treat two conditions (heart failure and paralysis), we could *aim* at treating them both. Since we cannot achieve both our aims by treating the person who remains paralyzed, this would imply that we do not violate the Treatment Aim Principle in not giving him a chance for treatment of his heart condition. (Another way of looking at this case involves the Treatment Similarity Principle: the person who remains paralyzed has a condition that is like the one we are trying to treat. This could be true if the cause of paralysis in the two people differs.)

But how would the Treatment Aim Principle explain, in a different case, treating the unparalyzed person who will not become paralyzed rather than the one we will paralyze? It might simply be said that successfully treating someone is a function not only of achieving the treatment's aim but also of avoiding bad side effects of treatment. I am not certain that this modification is consistent with the spirit of the Treatment Aim Principle. This is because, in this case, the bad side effects in one person will not be as bad as the death she will otherwise face, and ordinarily one would not refuse to treat the more serious problem just because a less serious side effect will occur. It is only the fact that someone else will not have the bad side effect that tempts us not to treat.

However, if these modifications were successful, then what I shall call the *Treatment Aim Principle Modified* would result. It may overlap, at least in part, with the Causative Principle. And both have to be constrained by the Principle of Irrelevant Goods, I believe; for minor differences that we cause in treating or achieving the aims of treatment may be morally irrelevant when what is at stake is a chance for a great good that could be had by either patient.

The Treatment Aim Principle Modified could also apply in non-life-and-death cases. Either we turn the additional large good side effect into a further aim, or we judge a treatment to be more successful if it avoids a significant bad side effect.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

Switch-and-Reduce Cases and how causation can affect the role of intrapersonal quality and quantity trade-offs in interpersonal allocation. Now consider a further modification of the Switch Cases. Suppose that in a modified version of Case 8, G and H have a disability giving them a QALY rating of .25. We can save G and also render him nondisabled by using a scarce resource in a particular way that would reduce the number of years he can live to five. He has the option of going through such a procedure or of going through a different lifesaving procedure. The latter is the only procedure open to H, and it uses the scarce resource differently so that either G or H could live for twenty years but remain disabled (.25). If G chooses the latter nonswitch option, we should give him and H equal chances for the procedure using the scarce resource. However, it is not unreasonable for G to trade some extra years of life for improved quality of life. If it were true that the disability reduced the QALY rating of a life to .25, then it would not be unreasonable for G to take the switch option with five years of life. Suppose he takes (what I shall call) this *Switch-and-Reduce* option. Should we then continue to give him an equal chance to receive the lifesaving resource because it makes it possible to cause in him as much good (though in a different form, combining quality and quantity in a different way) as will be the result of helping H? I believe so.

Suppose the Causative Principle is correct for some quality properties, and it is also permissible to choose between people on the basis of a big difference in quantity (such as the fact that one will live five years but another will live twenty years), when other things are equal. Then, giving G and H equal chances when G takes the Switch-and-Reduce option is consistent with still *not* giving equal chances to H and B′ in Case 10 (which is a modified version of Case 7), when H will be paralyzed but live for twenty years and B′, who is and continues to be *un*paralyzed, will live for only five years. Hence, even though G, when he chooses the Switch-and-Reduce option, has the same outcome as B′, G is given an equal chance to be saved relative to H, but B′ would not be given an equal chance relative to H.

Case 10

Persons		G		H		B′
		P		P		U
	↙	= ↓	=	↓	>	↓
	20P	5U		20P		5U

This difference in the relation of B′ and G to H implies that there is a way, after all, in which an aggregation of quality and quantity that makes sense in an intrapersonal choice bears on interpersonal choice. In the case we have considered, when a person chooses a certain quality/quantity

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

trade-off that we can produce, the higher quality is allowed to compensate for the reduction in quantity. By contrast, the higher quality is not allowed to compensate for the lower quantity when the quality is due to piggybacking, at least according to principles for allocation that I have so far discussed.

What if G had no choice to make between five and twenty years because we could only save him by using the scarce resource in a way that makes him nondisabled (i.e., he goes from .25 to 1) and gives him five years of life? I think that, if causation should matter, then this combination of goods, whose quality part we at least cause, should result in G's having an equal chance relative to H, even if B′ should be disfavored relative to H. This would mean that it is not just respect for someone's choosing one option over another and giving up longer life that justifies allowing intrapersonal aggregation of quality and quantity to affect interpersonal allocation of lifesaving resources.[42]

There are also Switch-and-Reduce Cases in non-life-and-death contexts. For example, suppose that one of two paraplegic candidates for a scarce resource that could permanently prevent blindness also has a further option: If we deliver the drug in a certain way, it will also cure his paralysis but at the cost of less successful retention of vision (e.g., he will have only half the years of sightedness of the other candidate). It might well make sense for someone to trade off some vision for some years of free mobility. Each of the candidates should have had an equal chance for the resource if both would remain paralyzed and have as good prevention of blindness. It seems that if a candidate chooses to produce a balance between two aims that makes intrapersonal sense, he should not lose his equal chance for the resource. However, if someone who was all along not paralyzed could be saved from blindness only half as long as a paralyzed candidate for the blindness cure, the principles we considered earlier suggest that he should not get an equal chance for the scarce resource. An outcome that is a result of piggybacking on nonparalysis is treated differently from an outcome one produces, if the Causative Principle is correct.

The Treatment Aim Principle Modified might yield the same conclusions, if producing nonparalysis become a second aim of the treatment. It would be considered invidiously discriminatory to give a blindness treat-

[42] We shall consider G and H from Case 10 again below. The Switch-and-Reduce Cases are derived from my initially considering a scenario in which a paralyzed person who will remain paralyzed and will live twenty years gets a lifesaving resource rather than someone who will remain unparalyzed but live for only five years because we abide by one of the principles discussed in Section IV.B. After getting the resource, the paralyzed person opts for a separate surgery that he knows will cure his paralysis but reduce his life expectancy to five years. Presumably, we would not want to rule out his possibly reasonable choice of the surgery simply on the ground that he is opting for an end state identical to the one that resulted in someone else's being deprived of an equal chance for the scarce resource relative to him. Yet the case seemed problematic to me.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

ment to a nondisabled person that turns out much better in a disabled person, other things equal. Yet it seems not to be invidiously discriminatory to give a treatment to a disabled person because we can count the change in his disability status as an aim of treatment that compensates for the reduced effectiveness of treatment for blindness. (If the Treatment Aim Principle were not modified, and it focused only on how we achieve the original aims of blindness-prevention or lifesaving, it would not make the same prediction as the Causative Principle.)

The Switch Cases and the Causative Principle opened the way for QALYs to be used in allocating scarce lifesaving and non-lifesaving resources interpersonally, even though candidates' expected life years (and achievement of the original treatment aim) were the same, because we could cause a significantly improved quality of life in one candidate. The Switch-and-Reduce Cases open the way for QALYs to be used in allocating lifesaving and non-lifesaving scarce resources interpersonally when candidates' expected life years (and achievement of the original treatment aim) are different, if each candidate's outcome is equivalent by the intrapersonal quantity/quality (that-we-cause) trade-off test.

Switch-and-Reduce Cases and the Nonconstant Role Principle. Consideration of the Switch-and-Reduce Cases also suggests another way of dealing with large quality differences that does not depend on the Causative Principle and that, perhaps, may better capture some intuitions about cases. Suppose that someone who will continue to be U (unparalyzed) can be saved for only five years, whereas someone who will continue to be P (paralyzed) can be saved for twenty years. (We play no causative role in their quality of life per se.) It might be suggested that someone's piggybacking as P or U should make no difference when other significant factors (such as length of life) are the same, but that even piggyback U can compensate for a deficiency in other factors in a way that P cannot. U can do this to the degree represented by an intrapersonal equivalence point between U with fewer years of life and P with more years of life. The intrapersonal equivalence of two outcomes implies that a given person would get as much out of U for five years as he would get out of P (rated at .25) for twenty. The interpersonal implication of this, according to the new way of dealing with quality differences that I am now exploring, is that even if someone will piggyback on U for five more years, he should get an equal chance with another person who is P (.25) and will live for twenty years. This conclusion differs from what the Causative Principle implies for such a case (as we saw above) because the Causative Principle provides no reason not to favor a person who is P and will live for twenty years, over someone who piggybacks as U and will live for five years.

In a sense, this new proposal envisions *a nonconstant role for U,* at least as a piggyback property: It cannot add any positive weight for interpersonal allocative purposes in order to move us away from equal chances when all other relevant factors besides piggyback properties U and P are

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

equal between candidates; it can only sometimes lead us to "forgive" the absence of other positive factors and so maintain equal chances. The absence of comparable positive factors would not be "forgiven" if P were present instead of U. Hence, if a candidate with P could live for only five years, he should not be given an equal chance with a candidate with U, who can live for twenty years, according to this proposal. If the U candidate can live for twenty-two years, *intra*personally this is preferable to P (.25) for even eighty years. That is, (1)(22) is greater than (.25)(80). Nevertheless, *inter*personally, a person with piggyback U for twenty-two years shows up as having only a two-year difference by comparison with a person with piggyback P who can live for twenty years. This is because each can have at least twenty years, and U's positive weight is not being called on to compensate for any deficiency in respect of years of life expected in the U person. Because the proposal envisions a nonconstant role for at least piggyback U (and perhaps certain other positive properties), I shall call it the *Nonconstant Role Principle.*

Here is an analogy (though it is not intended to support the permissibility of using the Nonconstant Role Principle): If a man and a woman have the same skills, they should be treated alike. If men lack a certain skill, this should not be held against them in a contest with a more skilled woman, but if a woman lacks the skill, this should be held against her. In this case, if we assume that there is nothing intrinsically better about being a man or a woman, we would say that the nonconstant role for the property of maleness is wrong. But if we assume that U (being unparalyzed) is truly preferable to P (being paralyzed), the question then is whether it can be coherent and morally permissible to have such a nonconstant role for U (whether we cause U or not). We would certainly be favoring piggyback U over P, in the sense that a U person gets to retain equal chances in circumstances where a P person would not; a P person is held to a higher standard than a piggyback U. However, the latter is never allowed to win in a contest merely because he is U; the most he gets on account of U per se is sometimes retaining an equal chance and so avoiding an immediate loss.[43] (And, of course, in an intratype context, a U person who will live for five years would lose to a U person who will live for twenty, even though relative to a P person who will live for twenty years, the two will have equal chances.[44] This is explained by the different factors present in the three pairwise comparisons.)

Consideration of the Switch-and-Reduce Cases raised the possibility of the Nonconstant Role Principle. I shall, however, set it aside for now and

[43] Notice that the fact that piggyback U adds positive weight when the candidate will live five years cannot be accounted for by the U property's standing out in this particular context but not when a candidate can gain twenty years. For if a P candidate could also gain only five years, the Nonconstant Role Principle would recommend equal chances for a P candidate and a U candidate, with U adding no additional positive weight.

[44] This was pointed out to me by Shelly Kagan.

focus on principles that emphasize our causative role. To recap before we proceed: We have now considered principles that suggest how calculating QALYs should sometimes play a role in life-and-death and non-life-and-death decisions (as shown in the Switch Cases, where we cause a change in QALYs) and how intrapersonal aggregation of quality and quantity can bear on interpersonal allocation choices (as shown in the Switch-and-Reduce Cases).

A problem for the Causative Principle. Now consider a problem for the Causative Principle.

Case 11

Persons	I	J
	P	U
	↓	↓
	U	U

Case 11 is one in which if we save person I, we will also cure his paralysis, whereas if we save J, we will just save his life, there being no paralysis to cure. Here the outcomes are assumed to be the same in quality and quantity, but *we* produce a significantly larger difference if we treat person I than if we treat J. Our treatment is, in that sense, more effective with the paraplegic person. Therefore, the Causative Principle, as presented above, says that it is permissible to save the paralyzed person rather than the unparalyzed. However, I think that this is the wrong conclusion and that there is no good reason for favoring one over the other.

When the same outcome (U) will come about in the case of the other candidate (J), albeit not because of our causative power, the benefit of having U as an outcome can be achieved no matter which person we choose. This is what seems morally important, not whether we causally contribute to U. What motivates the Causative Principle is not pride in our causing more good, but permission to seek a better outcome for someone when we can cause U, without violating moral side-constraints.[45]

[45] Here is a way in which we should *not* explain our judgments about Case 11: we should not disaggregate interpersonally and turn the greater difference we make intrapersonally into a greater difference we make interpersonally. If we disaggregate, we claim that the two-person choice case is like the three-person choice case. That is, Case 11 is treated as though it were equivalent to the following case:

U		P1		P2
↓	*vs.*	↓	+	↓
				(cure of paralysis only)

In this disaggregated case, we choose between saving the life of an unparalyzed person or saving the life of a paraplegic (who remains such) and also curing paraplegia in a third person who does not need his life saved. In this case, where the additional good is distrib-

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

Though we achieve the lifesaving aim of treatment equally in both candidates, suppose we acquire two aims once we learn of the additional effect that can be achieved in one patient. Do we then also satisfy the Treatment Aim Principle Modified better if we treat the paralyzed person? If so, the Treatment Aim Principle Modified faces the same problem with Case 11 as the Causative Principle because, strictly speaking, our treatment does achieve more in the paralyzed than in the unparalyzed person.

Similar results hold for a non-life-and-death case. If we can prevent someone from going blind and also cure his paralysis, the Causative Principle and the Treatment Aim Principle Modified seem to imply that we should help him rather than help an unparalyzed person from going blind. However, I think there is no good reason to favor one over the other. (In addition, if we can prevent a paralyzed person or an unparalyzed person from going blind, but the cure will cause paralysis in the unparalyzed person, the two principles suggest that we should avoid the treatment that has the worse side effect. Yet given that both will wind up paralyzed, I think there is no good reason to favor the paralyzed person over the unparalyzed one.)

The Principle of Irrelevant Identity. Suppose we accept the implications that the Causative Principle has for some cases (such as Case 8) but reject its implications for others (such as Case 11). Then we will have to find a principle to account for both sets of judgments. Here is a suggestion: The important point of the Causative Principle is that what the person is (consisting not only in what he is now but what he would become due to causes independent of us if he survives) should not determine whether he is helped, so long as the outcome that can come about through helping him is significant (e.g., a life for which it is reasonable to take important risks, rather than saving someone already in a permanent coma). This does not mean that the quality of outcome does not matter per se, only that we do not let the differences in what the person is (independent of the change we make) make a difference to whom we choose to help. This has two implications. When the synchronic outcomes expected in different individuals are different, to pay attention to anything but the causative difference we make would be to make the difference in them affect our decision of whom to aid. Hence, we can abstract from who they are by attending to the causative component. Or, alternatively, we can imaginatively add the good

uted over a third person, I believe the extra good done could permissibly determine our choice of whom to help. But if this conclusion carried over to Case 11, it would imply that we should save I (who is paralyzed) rather than J (who is not), and I think this is wrong. (We already saw in discussing Case 1 and Case 2 that turning a case of a concentrated aggregate of goods into a case of a distributed aggregate of goods need not lead to the same interpersonal allocation decision. In that set of cases, the outcome in the two-person case differed from what it was in the three-person case: A would live ten years and B fifteen years. In Case 11, the outcomes of I and J are the same.)

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

property that one party is missing into his outcome (or imaginatively subtract the good property the other party has from his outcome). But when the synchronic outcome we expect in different individuals is the same, the decision to attend only to the causative difference we make (from how they were or would have been independent of us) results in the differences in who they are (independent of our action) playing a role in our decision regarding whom to treat.

That is, when the bottom line is the same, paying attention to the difference we make is an indirect way of treating people differently on the basis of who they are. Showing that we need to do less to reach a given bottom line (outcome) implies that the person had more to begin with, and thus the Causative Principle would imply, for example, that we should not give an equal chance to someone who had more good to begin with.[46] The principle would hold against the person who is nondisabled the fact that he is nondisabled, since this is what accounts for our making less of a difference in his being alive and nondisabled. If we do not want who the person is to count against him (or for him) when outcomes are the same, we could imaginatively add (in our calculation) the disability condition to one person who actually lacks it or take it away from the other who actually has it. Then the causative difference we make would be the same in each person's case. But we achieve the same result by just attending to the fact that the outcomes are the same.

Notice, however, that our cases are consistent with "who the person is" still making a difference, for example, to how each of two people compare with a third. For example, in Case 7, person B, who remains unparalyzed, will not be favored for the scarce resource over A, who remains paralyzed; but in Case 8, person E, whom we switch from P to U, can be favored over someone identical to A. The person who remains U has the same outcome as the person who switches to U, but according to the Causative Principle, they have different prospects relative to someone who remains P. This means that it is only the identity that a person brings to the circumstances—whether someone starts as U or starts as P—that determines his differing chances relative to someone who remains paralyzed.[47] This has the result that in "within-type" contests, for example, between two people who are both paralyzed to begin with, certain factors

[46] It is very important to repeat that I continue to imagine that the disabled person has (only recently) become disabled, and also that the nondisabled person has always been nondisabled. I do this in order to avoid the issue of one candidate's being needier than another because, if not aided, he will have lived a worse life, other things equal, in having lived for a long time as disabled. If we were to give some priority to helping the person who would be worse off if not aided, this might be a reason to hold someone's having lived a long time nondisabled against him in a choice between him and the disabled person (though not if only quantity- and not quality-of-life considerations should play a role in the evaluation of need in life-and-death choices). But I have constructed the cases so that this is not a factor.

[47] I owe this point to Gideon Yaffe.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

(such as being nonparalyzed in the outcome) can count in favor of a candidate that cannot count in "between-type" contests, for example, between a paralyzed and an unparalyzed person. Hence, if we say that "we do not want who the person is to count against him (or for him) when the outcomes would be the same," we must mean "when the outcomes would be the same for the contestants for the resource in a pairwise contest."

Similarly, recall that in the Switch-and-Reduce Case 10, person B′ had the same outcome (U for five years) as G (U for five years). They fare differently relative to H (with an outcome of P for twenty years), with B′ disfavored and G treated equally, if we use the Causative Principle. This is because of their starting points. Yet, if what I have said about Case 11 is correct, if we had to decide whether B′ or G gets the scarce resource, each should be given an equal chance. One difference that Case 10 introduces is that when the quality property is something we cause, it can aggregate with quantity to determine the outcome that is relevant to an interpersonal allocation. This is what accounts for equal chances between G and H, even though their outcomes are not alike except in QALY terms.

What if the three types of people found in Cases 7 and 8 are all present at once? For example, consider Case 12, where each person we might save would live equally long in the future.

Case 12

Persons	B	F	E
	U	P	P
	↓	↓	↓
	U	P	U

It would seem that E may be chosen over F on grounds that we cause a better outcome without either one's identity counting for or against him. Then it is a toss-up between B and E, because we would hold the identity of B against him if we were to favor E. If B wins this fair toss-up, F cannot complain about the lack of an equal chance with B, because there is sufficient reason why he is eliminated from the contest by E. Why do we choose this way of proceeding rather than the following way—start with a toss-up between B and E, which B might win, and then have a toss-up between F and B that F might win? I believe the answer to this question is important. We may decide on the basis of which path leads to the best outcome so long as doing so does not hold anyone's identity for or against him. (I shall return to this point below.)

Hence, the dominant point of the discussion in this subsection is that we should treat persons so that who they are (their type-identity, but also,

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

more generally, what type they will be independent of the change we make) does not count for or against them. This only sometimes commits us to the subsidiary Causative Principle. To capture this dominant point in a principle, I shall use the *Principle of Irrelevant Identity,* which I originally introduced in connection with the Equal Respect Argument. Note that we are concerned with abstracting from what people are independent of the change we make, though we are still concerned with achieving the best outcome. Usually, abstracting from characteristics people have and treating them equally is associated with theories of individual rights; such theories try to capture the idea that it is certain characteristics that all and only persons have which are sufficient for certain forms of treatment (whatever other characteristics a given person may have). But these theories of rights are also associated with ignoring potential outcomes entirely when deliberating about what to do. The limited form of abstraction from outcomes (i.e., abstraction from what people are, independent of what we do) that we are considering is interesting, in part, because of its contrast with these rights theories.

Indeed, the fact that we are still concerned with outcomes can help us place the point of the Principle of Irrelevant Identity in its proper context. One may put this as follows: (1) We aim at doing what results in the best outcome, (2) on the condition that we not hold someone's identity for or against him. Failing in (2) will defeat our pursuing the aim in (1). What's more, (3) abiding by the Causative Principle is important when it defeats the charge that we have not met the condition in (2). That is, the Causative Principle defeats the defeator of (1), thereby allowing us to proceed with (1). However, (4) if we focus on the Causative Principle more than as a way of meeting (2), we will also violate (2). (Points 1 through 4 are reflected in the path that we took in deciding what to do in Case 12.) This all also implies (roughly) that rather than focusing on whether we *cause* a good or bad outcome, we should focus on whether we *cause* a (significantly) better (or worse) outcome than would otherwise exist among any of the candidates. If not, we should give all the candidates equal chances. (Determining whether we cause a better outcome can involve aggregation of quality and quantity, as evidenced by the Switch-and-Reduce Cases.)

I have used life-and-death cases in trying to deal with problems raised by the Causative Principle, but the same points apply to non-life-and-death cases.

The causal and componential role of identity in relation to the Principle of Irrelevant Identity. When I originally introduced the Principle of Irrelevant Identity in connection with the Equal Respect Argument, I pointed out that the principle did not rule out *linkage;* that is, it permitted counting relevant differences causally due to a person's identity. This is quite independent of our causing differential effects. I also said that it was a permissible limit on the principle to take account of the noncausal, component role of who the person is in accordance with the Treatment Similarity

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

Principle. Let us review these points again, now that we have discussed our producing the nondisabled or disabled conditions.

Distinguishing between (a) what the person is and (b) what results we will produce (abstracted from what he is) should not, I believe, prevent our attending to results that occur *due to what he is.* We should, if we wish not to discriminate invidiously, abstract from the disability a person brings to the treatment situation as a *component of his life,* not from its *causal* role. Thus, for example, suppose that if we save a paraplegic, he will remain a paraplegic and live for two years, but if we save a nonparaplegic, he will remain a nonparaplegic and live for ten years. Assume the two candidates for treatment are alike in all other morally relevant respects. Even if the fewer years of life we can expect for the paraplegic are entirely due to his being a paraplegic, it need not be discriminatory to take them into account in the allocation decision. This is *linkage.* It takes account of paraplegia's effects (its causal role), for example, on quantity of life, but it is consistent with not attending to the synchronic properties of paraplegia in deciding how to allocate resources. (Such properties would include, e.g., not being able to walk, being in more pain, and whatever else are components of the life of the paraplegic in virtue of paraplegia.) The causal effects could also include other disabilities. For example, suppose that the drug that could save life causes blindness in the disabled person, but not in the nondisabled person, due to an interaction between the disability and the treatment. This difference in the outcome that we can produce in the two candidates might justify giving the scarce resource to one candidate rather than another, and so might not constitute invidious discrimination. This is so even though it involves making use of the value judgment that it is worse to be blind than not to be blind when judging the outcome we produce. (The Treatment Aim Principle Modified could also take account of bad side effects in deciding whom to treat.) Obviously, to make this account work, more would have to be done to distinguish the characteristics of a condition from its effects.

The distinction between the purely causal role of paraplegia and its component properties is also crucial in answering one potential objection to the Principle of Irrelevant Identity.[48] It may be said that if our treatment cures paraplegia in addition to saving a life only in one person and not in another, this must be because of some difference between the two people, for example, an allergic reaction in the second person that blocks a cure. Therefore, to let the difference in outcome count makes the difference between people count, and is that not contrary to the Principle of Irrelevant Identity? I would argue, however, that there can be a moral difference between counting against someone his allergy's synchronic properties as a *component* in his life and counting the allergy's *causal* effect. Only the former involves treating people

[48] The objection was raised by Douglas MacLean and John Broome.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

differently because we dislike the allergy they have, independently of our causing its presence.

That we are not here holding the disability that someone would have (independently of what we do) against him as a component of his life is further reinforced, I argued earlier, by considering something that improves a life when it is considered as a component. For example, suppose painfreeness adds positively to a life, and yet painfreeness is caused by a protein that interacts badly with our lifesaving treatment and so causes paraplegia or much shorter life. Ignoring painfreeness as a good component and thus not discriminating in favor of the people who have it would be consistent with attending to its effect. Deciding not to help the person with painfreeness because of its causal effects would not constitute discrimination against him on grounds of his painfreeness.

I also suggested that there could be cases in which identity as a component feature may permissibly count in selection because its properties are like ones that it is the aim of the treatment in question to deal with. When identity factors are used as a basis for selection for this reason, there may also be no invidious discrimination. Suppose that there are two patients who must be treated for fatal kidney disease—a patient who is weak and one who is hardy. The first patient's weakness is not due to the kidney disease. It has another cause that cannot be affected even by a kidney cure that reverses past effects of kidney disease. We can cure the kidney disease equally effectively in both, and this means that we can prevent the weakness *due to kidney disease* equally in both persons. The Principle of Irrelevant Identity would tell us to ignore the weakness that preexists from another cause in one of the patients. The Treatment Aim Principle and Treatment Aim Principle Modified agree; the first because we cure kidney disease as well in each, and the second because there are no differential good or bad side effects of our treatment in each.

However, the weakness that has another cause is *like* one of the effects of the disease that we are trying to treat. After all, it may be said, one of the reasons we treat kidney disease is to prevent not only death but continual weakness; we can prevent that state (by stopping the kidney disease) in one patient, but we cannot prevent that state in the other patient because it is overdetermined. (That is, we can stop weakness from kidney disease but not from the other cause.) The *Treatment Similarity Principle* says that when we cannot prevent in a person a state that is like one that gives us a significant reason to try to treat a disease, it is not wrong to ignore the fact that the state is an identity factor.

If this principle has any role to play, it should still be constrained by the Principle of Irrelevant Goods, so that at least small differences do not make a moral difference. It remains open that something like curing weakness in one person but not another is too small a difference to make a moral difference relative to avoiding death.

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

Now suppose, counterfactually, that paraplegia was one effect of kidney disease that we were interested in preventing. If one candidate were a paraplegic as a result of some other event in his life besides his kidney disease, then our treatment of his kidney disease would not cure or prevent his paraplegia the way it would cure or prevent the paraplegia of a second candidate. Suppose that, in our previous case, it would be permissible to treat the hardy candidate as opposed to a candidate who is weak due to causes other than his kidney disease, despite equal prospects for success in treating the effects of his kidney problem. Then it would also be permissible to treat the nonparaplegic patient rather than the paraplegic patient in the present example. The Treatment Similarity Principle is here overriding the Treatment Aim Principle and the Treatment Aim Principle Modified. However, we cause the unparalyzed person to remain unparalyzed—he is not piggybacking on his initial unparalyzed state—and this is consistent with the Principle of Irrelevant Identity.

Views of discrimination and a decision procedure. The previous discussion implicitly gestures toward *three possible understandings of (invidious) discriminatory conduct*. The *first* possible understanding tells us that acting in any way on the differential value attributed to being nondisabled or disabled is discriminatory. But this would imply that common surgeries undertaken to cure people of paraplegia are discriminatory conduct if we perform them because we think it is better to be unparalyzed than paralyzed. I believe we should reject this first suggestion. Rejecting it helps us see how the claim of discrimination against the disabled differs from the claim of racial discrimination. For a claim of racial discrimination could be supported just by showing that our act was undertaken because we believe it is better to be white than to be black per se (i.e., independent of any other factors, such as having to live with negative social attitudes toward blacks). Notice that a discriminatory attitude against blacks relative to whites could be present even if we were deciding *only between blacks* who should get a scarce resource. This would be true, for example, if we decided to give a scarce lifesaving resource to one black person rather than another because it had the side effect of turning him white. (This implies that the fact that we are only choosing between paraplegic people in Case 8 does not, by itself, settle the question of whether we would be engaged in discrimination by deciding to help the person who will stop being paraplegic.)

The *second* possible understanding of discriminatory conduct tells us that acting on the differential value attributed to being nondisabled or disabled when this makes the person who will be disabled worse off than he might otherwise have been (for example, losing a scarce resource he might have gotten on a coin toss) is discriminatory. This understanding of discrimination would rule out the Causative Principle and the Principle of Irrelevant Identity. If we think this is wrong, we could endorse a *third*

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

possible understanding of discrimination: It is discriminatory to act on the differential value attributed to being able or disabled as a component of someone's life, if that is not like what we aim to treat, when this makes the disabled person worse off than he might have been because of what he is or would be, independent of what we do. It leaves it open that the value of a component can sometimes count for or against someone if we cause that component (even if this further implies that those in whom we cause a component, for example, are favored relative to someone else when those who have the same component independently of what we do do not get favored in the same way).

Notice that the Treatment Aim and Treatment Similarity Principles can lead to someone's being worse off than he would have been because of his disability, but not *directly* because his disability *is disvalued per se* relative to nondisability. For example, it was pointed out that a positively valued state (e.g., hardiness), as much as a disability, could reduce the effectiveness of our ability to treat an illness. And if a disability is another cause of a condition that it is our aim to treat, this could also be true of a positively valued state (e.g., hardiness). Whether the state is negatively or positively valued, the Treatment Aim and Treatment Similarity Principles could prefer the individual without the trait that has these effects. However, the disvalue of the disability *does* come up in deciding what we should treat.

Can we summarize our discussion of principles as a decision procedure (for life-and-death and non-life-and-death cases) to decide if (the third view of) discrimination is involved? Here is one attempt:

(1) Check the level of well-being (including quantity and quality) to which you can bring someone relative to another;
(2) Check to see if, in reaching this level, counting his starting point (or factor independent of what we do), when it is not similar to a condition we are trying to treat, would make the starting point (or factor independent of what we do) work, as a component feature, in favor of or against him relative to others;
(3) If the answer to (2) is no, deciding whom to help by differences in level of well-being to which you can bring someone will not involve invidious discrimination (on the third view of discrimination). If the answer to (2) is yes, deciding whom to help by differences in level of well-being to which you can bring someone will involve such discrimination; and
(4) If avoiding such discrimination is all that should stand in the way of producing the best outcomes, then decide how to allocate by following (1), (2), and (3).

This leaves it open that differences in outcome should not matter because they are morally irrelevant goods, even if attending to the goods would

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

not involve the third view of discrimination. Indeed, my conclusion is that while it may be wrong to ignore the Principle of Irrelevant Goods, doing so need not involve inappropriate discrimination. Hence, some complaints on behalf of the disabled may have to appeal to the Principle of Irrelevant Goods, rather than to a claim of discrimination.

If we decide according to this four-step procedure, it is what we can do for someone (in the sense of the level of well-being at which our behavior will leave him) that will matter, at least so long as making this matter does not involve treating factors (that are unlike those we are trying to treat) beyond what we do, as component features in favor of or against some person relative to another.

One problem with this decision procedure is that it just attends to our treatment of one person relative to another in deciding whether there is invidious discrimination. However, suppose we treat B correctly relative to A, and B loses the scarce resource, but we do not treat C correctly relative to A (or D) in order that C not lose the scarce resource. Might not B complain that he is being discriminated against relative to C, for we are holding him to standards to which we do not hold another, and that is enough to support a claim of discrimination?[49] In order to avoid this problem, we might modify step (3) to: If the answer to (2) is no, and the procedure in steps (1) and (2) is applied pairwise generally (or impartially), deciding whom to help by differences in levels of well-being to which you can bring someone will not involve discrimination. If the answer to (2) is yes, or the procedures in steps (1) and (2) are not applied generally or impartially, deciding whom to help by differences in this level of well-being will involve discrimination.[50]

In sum, I have argued that there are several ways in which differences between people that we can produce (and sometimes even differences we do not produce) allow us to decide whom to treat without committing the wrong of invidious discrimination. But narrowing the ground on which one may complain of discrimination may only empower the Principle of Irrelevant Goods instead, for that principle may possibly be called on to prohibit actions that one might have thought could be ruled out on grounds of discrimination.

What if there are cases where we see a role for both the Principle of Irrelevant Identity, which has a causative component, and the Nonconstant Role Principle, which does not have a causative restriction on counting U? Then the Principle of Irrelevant Identity might be modified to allow identity—in the sense of quality properties that are present independent of what we do—to preserve equal chances for those whose shorter lives would be as good for them as a longer life of lower quality.

[49] I owe this point to James Nelson.

[50] I discuss how to manage apparent intransitivities that may result from using this procedure in my essay "Disability, Discrimination, and Irrelevant Goods."

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

E. An inadequate ground for considering disability when allocating scarce resources

Let us now consider the ex ante choice behind a veil of ignorance as Singer envisions it, which results in all quality and quantity distinctions being relevant to interpersonal allocation. The argument might be criticized as follows: We will get the wrong result for a moral principle if we think, as Singer does, that ex ante reasoning behind a veil of ignorance involves each person's thinking of himself as having a possibility of occupying each outcome-position, for example, disabled or nondisabled. Rather, the veil of ignorance should be understood as a device that leads each person to take seriously the fates of the separate persons who will actually occupy each of the outcome-positions, disabled and nondisabled. Thomas Scanlon makes this point in distinguishing his form of contractualism from that of John Harsanyi.[51]

Scanlon's basic objection to Harsanyi's approach is with its interpretation of impartiality—that is, what it means to choose principles that are not designed to favor one person over another and that one could accept independently of knowing what one's actual position is or will be. Scanlon agrees that it can be helpful in finding such principles for each person to imagine that he is in every other person's actual position in life (outside the veil of ignorance), having that person's perspective on things, in order to see whether each person in any actual position could reasonably approve (or not reasonably reject) a proposed principle governing relations between people. While Scanlon does not make use of a veil of ignorance, his interpretation of impartiality implies that (i) if no person in any position outside the veil could reasonably reject a principle on his own behalf, then (ii) any individual behind a veil of ignorance could agree to the principle. If someone from some position outside the veil could reasonably reject the principle, then it could not be agreed to by a person behind the veil of ignorance, because it is not a principle that one could agree to without knowing into which position one would fall. The reasonableness of rejection outside the veil is, according to Scanlon, a function of comparing the possible complaints of people in different positions.[52]

Scanlon rejects the alternative interpretation of impartiality that Harsanyi accepts and that Singer also seems to accept—that to be impartial is best

[51] See Thomas Scanlon, "Contractualism and Utilitarianism," in Amartya Sen and Bernard Williams, eds., *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982). Scanlon discusses John Harsanyi, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy* 63, no. 4 (1955): 309–21. Singer's approach to the veil of ignorance seems like Harsanyi's. My discussion in this subsection of the essay is based on my comments on Allan Gibbard's *Reconciling Our Aims: The Tanner Lectures in Human Values,* University of California, Berkeley, 2006 (New York: Oxford University Press, forthcoming).

[52] See Thomas Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1999), chap. 5.

understood as deciding behind a veil of ignorance on the assumption that one has an equal probability of being in any person's position outside the veil. The assumption of equal chances of being in any position is supposed to insure that one does not choose principles that favor one position over another. But there is still no requirement on this view that a decision-maker behind the veil first take account of whether each person actually occupying each position could reasonably reject a principle, before he, behind the veil, decides to accept it. Rather, the order is the reverse of what it should be in Scanlon's view: namely, the order is: (i)′ if any individual behind the veil of ignorance would choose a principle on his own behalf, then (ii)′ no person in any actual position outside the veil could reasonably reject the principle. It is this order reversal, as I refer to it, that lies at the heart of Scanlon's disagreement with Harsanyi's approach and the approach Singer takes.

Scanlon also argues by example against the conclusion that what an individual would choose behind a veil of ignorance in order to maximize his average expected utility determines whether someone outside the veil can reasonably raise a moral complaint against a principle. Suppose one maximizes one's average utility by accepting a principle that allows many people to receive small benefits when only a few will be very badly off. Scanlon thinks that people who are very badly off could reasonably reject the principle because no one of the many stands to gain much individually, while some others are very badly off, and there is an alternative that could greatly improve the worse off at a small cost to each of the many. It is the latter alternative that is preferable, but it reduces aggregate utility and hence is not what maximizes any individual's ex ante average utility. Scanlon considers a principle from the perspective of each type of position that will actually be occupied by someone and compares its effects on a position pairwise with every other position in order to decide whether the principle should be adopted.[53] I believe that Scanlon's example shows that attending to the maximal sum of utilities, regardless of what distribution of these is involved, would not yield the right principle.

I suspect that emphasizing the distinction I drew in Section IV.A between (1) caring more for (or to be in) one position rather than another and (2) caring equally about remaining alive in the less favorable position in which one winds up, when there is no possibility of moving to a better position, coincides (at least in part) with Scanlon's view. His approach, in telling us to focus on how someone will approach things from the actual position he will be in, is like (2), insofar as (2) implies what Singer denies, namely, that a person can have as strong an interest in going on living, even when the objective good in his life is less than in someone else's life.

[53] Scanlon follows Thomas Nagel, who argues for pairwise comparison as coming closest to being the correct way to combine people's different interests in an outcome. See Nagel, "Equality," in his *Mortal Questions* (Cambridge: Cambridge University Press, 1995).

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

It contrasts with considering someone's preferences when he thinks he has a chance to be in any of many positions, which is like (1).[54] Perspective (2), rather than (1), is the relevant one to take even in an ex ante thought experiment.[55] If we engage in policies that we know will leave some to die simply because we can achieve some additional goods for others, we must be able to give a justification to the individuals who actually occupy the position of those threatened with death. We must be able to show them that the additional good is worth depriving them of an equal chance at life. Saying to them, "Agreeing to such a policy would have maximized your ex ante chances of benefit because you might have been the nondisabled person" is not the right kind of justification. Let me expand on this last point.

The point is connected with the objection to Harsanyi's and Singer's use of the veil of ignorance—that it does not take the separateness of persons seriously. To make this objection, it is sometimes said that a person's willingness to risk *his* having a bad fate for the sake of maximizing *his own* expected average utility (intrapersonal risk) does not bear on whether he may endorse a principle that risks *someone else's* having a bad fate for the same goal (interpersonal risk). (My earlier discussion of the case in which someone will risk immediate death if a surgery does not render him nondisabled, in order to possibly be cured of a disability, was supposed to show this. However, it was framed in terms of the person not necessarily being willing to die if he is disabled as a result of our choosing to help someone else live a nondisabled life. We could also say that the fact that someone would exchange a certain quality and length of life for a better quality and length for himself does not mean that he would agree to a principle that involves his having the better life because someone else loses a chance to have the less good one, or vice versa.)

In one sense, this response seems misplaced. For suppose it is the case that *any* individual—A or B—deprived of knowledge that distinguishes his eventual position from that of others would take the small risk of being in a bad position in order to maximize his own expected average utility. Then when A suffers a bad fate while B has a good fate, this does not straightforwardly mean that A suffers because *B decided to risk A's having a bad fate,* for A would have taken this risk for himself.

What must be emphasized in order to make the objection that Singer's and Harsanyi's use of the veil of ignorance ignores the separateness of persons when considering whether risk-taking is appropriate is the difference between one's *actually* (beyond the veil, not behind it) losing out

[54] I shall have more to say below about what sort of "chances to be in any of many positions" are relevant to focusing on (1).

[55] Indeed, Scanlon believes that forcing people to identify in this way with each person is why Rawls uses a thick veil of ignorance (excluding probability calculations). See Scanlon's "Contractualism and Utilitarianism." Singer, like Harsanyi, does not object to probability calculations, as can be seen in his discussion of the morality of slavery (Singer et al., "Double Jeopardy," 291).

on something or running a risk of death in order to have some good, and being willing to have someone else actually (beyond the veil, not behind it) lose out on something or run a risk of death as a result of one's having one's good. In order for this difference to come to the fore, one must not interpret one's "being willing to run a great risk to have one's good" as including a decision behind the veil of ignorance to run a great risk that one will die (for example, if one turns out to be the person who is in the bad position) in order that one get to keep one's better life (if one turns out to be the *other* person who is in the good position). However, this is how reasoning behind the veil of ignorance in the Harsanyi/Singer manner works; hence, it tends to assimilate interpersonal to intrapersonal sacrifice, and one might object to it on this ground.

Thus, one could refuse, beyond the veil, to let something happen to another as a result of one's outright being chosen to benefit, although one is willing, beyond the veil, to have the same thing happen to oneself as a result of seeking one's own benefit. Notice that the difference between self and other that is retained here has nothing to do with whether each person behind the veil would or could reasonably agree to accept the same risks. They each may be willing to accept the same risks behind the veil and yet still distinguish between (1) the case where they actually (not behind the veil) run risks for which only they may have to pay (if they lose) and (2) the case where others run risks or suffer losses, including being left to certain death, so that one may have a good. This, I think, is the way to understand the point of the initial response that we considered to Harsanyi/Singer–style veil-of-ignorance reasoning—the response that said one might be willing to take on risks and suffer losses oneself for one's own benefit but not be willing to accept that others must outright do the same if one is to benefit.

V. Conclusion

I have been discussing factors relevant to the allocation of lifesaving and non-lifesaving resources. I have suggested that it is important to consider need and not only outcome in deciding how to allocate such resources. In evaluating outcomes, the difference a resource can make to quantity of life in one person rather than another may be relevant. A difference in quality of life in one person's outcome rather than another's may be relevant when a large difference in quality between persons would be produced by our efforts (rather than when we save a person whose much better quality of life is not due to us) or when someone's unchangeable quality is like, or causally linked to, what we aim to treat. This difference in how quality and quantity should be treated conflicts with a general use of QALYs in allocating resources.

I have also argued that the fact that a person may generally trade off quantity and quality and take risks intrapersonally in order to achieve a

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press

better outcome (1) may not mean that he has less of an interest in retaining a worse outcome than a better one, when the former is all he can have, and (2) does not imply that the same trade-offs and risks can be assigned interpersonally as intrapersonally. These two points conflict with premises in some arguments for the general use of QALYs in allocating resources. Nevertheless, if the use of QALYs is permissible when our efforts produce a large quality-of-life difference between persons, this could indicate that some use of QALYs in the allocation of lifesaving and non-lifesaving resources can be defended even if points (1) and (2) are correct.

On the basis of these conclusions, I distinguished different views of invidious discrimination against the disabled in allocating resources and suggested a decision procedure for allocating resources that would not involve invidious discrimination.

Philosophy and Public Policy, Kennedy School of Government and Philosophy Department, FAS, Harvard University

https://doi.org/10.1017/S0265052509090074 Published online by Cambridge University Press