

Rationing and Rationality

THE COST OF AVOIDING DISCRIMINATION

Nick Beckstead and Toby Ord

Since we cannot fund all of the health interventions that we would like to, we must use some method to decide which interventions to fund and which not to fund. Over the past 40 years, it has become standard to rank publicly funded health interventions in terms of how many quality-adjusted life years (QALYs) they produce for a fixed amount of money. A QALY is a unit that takes account of both the amount of life that a health intervention adds, as well as the intervention's effect on health-related quality of life. Length is measured in years, and quality is measured on scale of 0–1, with 1 being perfect health and 0 being equivalent to death. One determines total QALYs by multiplying years of life by their quality weight. Thus, 10 years of good health is equivalent to 20 years at a quality rating of 0.5 (a weight commonly assigned to blindness).

The QALY value of a health intervention is the difference between the total number of QALYs with the health intervention and the total number of QALYs without the health intervention. Some proponents of this metric favor ranking health interventions in terms of how many QALYs they produce for a fixed sum of money, then funding the most cost-effective interventions first. The rationale for this approach is clear: it produces the greatest amount of health for a fixed sum of money, provided the quality weights are assigned appropriately.¹ This approach has the potential to produce great gains in terms of efficiency of health interventions funded and transparency of decision making.

Despite its benefits, the QALY-maximizing approach has been attacked by disability rights advocates, policy makers, and ethicists on the grounds that

The main ideas of this paper were developed during the 2010 Brocher Summer Academy. We would like to thank the participants at Brocher Summer Academy for many helpful discussions on these issues and the Brocher Foundation for making the event possible. We are especially grateful to Erik Nord and Will Crouch for their insightful comments on late drafts of this paper.

¹ Of course, it only maximizes health gains if different health states are given appropriate weightings. We will assume that we are dealing with weightings that have been chosen appropriately because problems of choosing an appropriate weighting are orthogonal to our concerns here.

it unjustly discriminates against the disabled. The main complaint is that the QALY-maximizing approach implies a seemingly unsatisfactory conclusion: other things being equal, we should direct lifesaving treatment to the healthy rather than the disabled. This objection has been forcibly made by John Harris (who described it as “double jeopardy” for the disabled),² and it rose to national prominence in the United States after attempts to use the QALY approach in the state of Oregon were overturned on antidiscrimination grounds (Hadorn 1999). This argument pays insufficient attention to the downsides of the potential alternatives. We show that this sort of discrimination is one of four unpalatable consequences that *any* approach to priority setting in health care must face. Given the alternatives, it is far from clear that we should revise the QALY-maximizing approach in response to this objection.

Bubbles Under the Wallpaper

Attempts to avoid unequal treatment for the disabled have not met with great success. A solution favored by Erik Nord and others involves ignoring quality weights when deciding to whom we should give a lifesaving treatment, provided the people to be saved regard their lives as worth living (Nord et al. 1999). As Magnus Johannesson has pointed out, Nord’s proposal would sometimes conflict with individual preferences: it would sometimes rank one treatment higher than another, although this would be worse for someone and better for no one (Johannesson 2001). Johannesson offers his own proposal, which also faces devastating objections (Nord et al. 2003). In looking at such proposals, one gets the feeling that the task may be like trying to get a bubble out from behind the wallpaper; pushing down in one place simply moves the bubble elsewhere.

Now for our main point: discriminating against the disabled is one of four highly counterintuitive consequences that any system of priority setting must face. To see why, consider the following case:³

² Harris (1987). See Singer et al. (1995) for a response to Harris.

³ Note that, in this case, there is no preexisting disability. The case therefore differs from standard examples of “double jeopardy.” This makes the case stronger, not weaker. In a case of preexisting disability, prioritarians and egalitarians can claim that it makes sense to choose *X* over *Y* on the grounds that *X* has had less health in life and therefore should be given priority. However, in our case, Alice and Beth had equal health prior to the disease. Therefore, prioritarian and egalitarian adjustments to the QALY framework cannot justify choosing *X* over *Y*.

Prioritarians and egalitarians might argue that, in this kind of case, there is nothing wrong with favoring the healthy over the blind, although priority/equality considerations tell against it in cases of double jeopardy. We could address this point by changing the case so that Alice became blind a couple of years ago. In this context, equality and priority would have insignificant weight, so this defense would not apply. Yet it would still seem objectionable to choose *Z* over *X* on the grounds that Alice is disabled. The analysis of the revised case would proceed, more or less, as it does here.

TABLE 15.1 } The Alice-Beth dilemma

	Option X	Option Y	Option Z
Alice	45 years (blind)	–	–
Beth	–	60 years (blind)	35 years (full health)

The Alice-Beth Dilemma: Alice and Beth were both perfectly healthy 20-year-olds, but have recently contracted an unusual disease. This disease will kill them very soon unless treated, and even then they will suffer from serious complications, such as blindness and/or a reduced lifespan. To make matters worse, there are not enough resources to treat them both. There are, however three possible treatment options. outlined in Table 15.1.

In X, Alice is treated and will live for 45 years, but will lose her sight. Because Beth was infected by a slightly weaker strain, there are two treatment options available to her: in Y, she will live for 60 years but will lose her sight, in Z she will live for only 35 years, but will retain her sight. After a due course of reflection, Beth finds that she strongly prefers 35 years of life with full health over 60 years of life with blindness, and Alice would, given the hypothetical choice between X and Z, have the same preference (a result in line with most people’s preferences and with the commonly used QALY ratings).

However we decide to choose among these treatments, we will face severe difficulties.⁴ To see this, compare X and Y, then Y and Z, and then X and Z. There are three problems we’ll want to avoid here.

1. *Preference for Early Death:* X gives fewer years of life for Alice than Y does for Beth. This is the only difference between them. Since there is no reason to favor Alice over Beth, choosing X over Y would reveal a perverse preference for saving the person who would live for a shorter period of time.
2. *Pointless Violation of Autonomy:* Y is worse than Z for Beth, Beth has requested this treatment after a due course of reflection, and Z is not better for Alice. Choosing Y over Z thus violates Beth’s autonomy and benefits no one.

⁴ As a reviewer reminded us, the QALY approach is not typically used for bedside rationing, but for deciding which treatments should be available for the population as a whole. The case we have described is therefore not the typical setting for applying the framework.

Our example could easily be modified so that it fits a more typical application of the QALY-maximizing approach. One need only replace Alice and Beth with two groups of people with two diseases, A and B, that correspond to Alice and Beth’s conditions in terms of prognosis and treatment options. We can suppose that the government only has enough funds to authorize one of three treatments, X, Y, or Z, as here, and our argument could be repeated, in analogous form, in that context. However, the example we have given is a bit simpler and makes all the same theoretical points, so we have decided to leave it as it is.

3. *Disability Discrimination*: Z provides fewer years of life for Beth than X provides for Alice. The only reason to choose Z over X could be that X would leave Alice disabled. Thus, choosing Z over X involves the very kind of discrimination derided by those who try to revise the QALY approach to priority setting. (This problem is also discussed by Ole Norheim [2013, Chapter 14, this volume].)

To avoid these problems, we must choose Z over Y, Y over X, and X over Z. This leads us directly to the classical problem of *cyclic preferences*.⁵ Cyclic preferences violate the conditions of rational choice theory and open one up to being “money pumped.” If it is important to choose Z over Y, then one should pay a penny to switch from Z to Y. For the same reason, one should pay a penny to switch from Y to X, and another penny to switch from X to Z, leaving one back where one started, but with less money. Worse, there seems to be nothing keeping one from going around in a circle until one runs out of money.⁶ So, the fourth problem then, is this:

4. *Cyclic Preferences*: Choosing Y over X, Z over Y, and X over Z, involves cyclic preferences, which violate the conditions of rational choice theory and leave one open to irrational behavior, such as money pumping.

Since any way of setting priorities in health care will face one of these problems, we are forced to choose the least among these four evils. It is far from clear that the best way of resolving this problem will involve rejecting disability discrimination. Indeed, we find it doubtful that anyone will seriously contemplate *Preference for Early Death* or *pointless violation of autonomy*. For these reasons, we will consider other ways out in the next sections.

A Rights-Based Approach?

There is no way of choosing between these three options does not involve one of the four unappealing consequences just described. Some people respond by endorsing a certain rights-based approach that involves cyclic preferences, but avoids the other problems. On this approach, when we have equally expensive treatments but can save only one person’s life, the person who stands to gain the most life years (ignoring any quality adjustments) is awarded the right to treatment. The person may then select the treatment that she most

⁵ Note that we are using the term “cyclic preferences” in a broad sense, referring both to cycles of preferences within a given set of option and cycles of preferences across sets of options (such as the present case).

⁶ The status of money pump arguments is somewhat controversial. This much, however, is not: if policy makers want to set health care priorities using a ranking system that they can follow in general, the ranking system will be susceptible to a money pump if the ranking is cyclic.

prefers to receive, under the advisement of her doctor. (This is the simplest modification of Nord's proposal that avoids *Pointless Violation of Autonomy*.) Although it sounds very sensible, we will now argue that this approach has serious problems.

Before discussing these problems, let's take a minute to understand how this approach works. When the options are *X* and *Y*, Beth gets the right to treatment and will choose *Y*; so this approach avoids preference for early death. If the options are *Y* and *Z*, Beth is again awarded the right to treatment. Since she prefers *Z* to *Y*, she will choose *Z*; this approach thus avoids pointless violation of autonomy. If the options are *X* and *Z*, Alice will be awarded the right to treatment because she stands to gain more years of life; thus, this approach avoids *Disability Discrimination*. Of course, this puts us back at cyclic preferences, violating rational choice theory and opening us to money pumps.

Although some people will regard *Cyclic Preferences* as enough to reject this kind of view outright, others may find this less objectionable, especially since it arises naturally out of the informed choices of individual people. Because we are unlikely to actually get money-pumped in practice, they may find that accepting cyclic preferences is preferable to the other alternatives.

The point of the money pump argument is not so much that we are afraid of bankrupting the public health budget if the situation ever arises, but that the possibility of being money pumped reveals a certain unreasonableness inherent in the policy. So, ultimately, this complaint that money pumping situations will not arise in practice is unpersuasive. Still, we think it is worthwhile to illustrate some additional problems faced by the rights-based approach.

Notice that the rights-based approach will deliver different results when seemingly irrelevant options appear or disappear. If the choices are *X*, *Y*, and *Z*, then treatment *Z* will be chosen. If the choices are *X* and *Z*, then *X* will be chosen. Since *Y* will not be chosen either way, it seems strange that adding this irrelevant option should change anything about which options are worth choosing. This strangeness can manifest itself in disturbing ways.

Consider, for instance, the following embellished version of the Alice and Beth case depicted in Table 15.1. Suppose that for each treatment (*X*, *Y*, or *Z*) that could be delivered, there is a corresponding vial of medicine that must be administered. After this, the patient must receive a very uncommon medicine, of which the clinic has only a single dose. Following the rights-based approach, the doctor decides on option *Z*, so he walks over to the table and selects vial *Z*, then fills a syringe with it. Just as he is about to inject Beth with this medicine, he hears a small crash: vial *Y* has just fallen off the table and shattered, making treatment *Y* unavailable. The doctor then realizes that it would now be wrong to give treatment *Z*, as it has become a choice between only *X* and *Z*, so he goes back to the table and fills a syringe from vial *X* to give to Alice instead.

This kind of behavior seems bizarre. The problem could arise in more ordinary contexts as well. A doctor might know that X and Z are available but not know whether Y is available. If Y is available, choosing Z would involve a pointless violation of autonomy and should be avoided. If Y is not available, choosing Z would involve discriminating against the disabled. Not knowing what to do until he knows whether Y is available, the doctor might be required to search the clinic or telephone suppliers, even though he knows he won't use the drug he's looking for. One could imagine more extreme versions of the case in which the doctor runs expensive tests to determine whether it is possible to use treatment Y . In still more disturbing versions of this example, the doctor would decide who to treat on the basis of how Y and T compare, where T is some other treatment that no one wants anyway (perhaps T and Y involve saving the most life years, but offer very low quality of life, so patients aren't interested in them anyway).

These objections suggest that, ultimately, the rights-based approach is untenable. Since problems of this kind arise as a result of *Cyclic Preferences*, they illustrate the unreasonableness of resolving the problem by embracing cyclic preferences.

A Lottery Solution?

In the section in which we presented the original problem, we assumed that the only options were X , Y , and Z . Someone might point out that, in practice, we would have many more options, corresponding to all of the lotteries over X , Y , and Z . Since many philosophers believe that fairness sometimes requires the use of lotteries when indivisible goods must be distributed, it is natural to wonder whether a lottery-based solution might avoid all of the problems we've been worried about.⁷ In this section, we argue that any lottery solution—any method of allocating health-care resources on the basis of chance—will face problems analogous to those presented above.

At first glance, this solution is unresponsive to the philosophical problem. If the only options were X , Y , and Z (perhaps because there is no time to roll dice, or we are working in a community that forbids leaving important decisions up to chance), we would still like to know what to do. If an approach gives unsatisfactory results even in imaginary cases, that counts against using that kind of approach.

A second problem is that even if lotteries seem sensible in particular clinical cases, they will seem less reasonable in other contexts. The QALY approach can, in principle, to be used to rank many health options, including research

⁷ See, for instance (Broome 1984).

options. To suggest that the public health organizations should sometimes roll dice to determine which research to fund seems rather implausible, although analogues of our problems would arise in this context. Therefore, it is doubtful that a lottery solution could have fully general scope.

But the most serious objection to a lottery solution is that it falls prey to a more general version of the problem that we've been worrying about. To see this problem, think about what probabilities we would assign to the different alternatives when the choices were X and Y , Y and Z , and X and Z . There are three problems we'll want to avoid:

- 1*. *Preference for Early Death**: When the potential treatments are X and Y , the lottery allows Alice at least as great a chance as Beth even though Alice and Beth are equally healthy and Beth would live for an additional 15 years.⁸
- 2*. *Pointless Violation of Autonomy**: When the potential treatments include Y and Z , we force Beth to gamble over Y and Z , rather than just letting her choose Z , even though she prefers Z and giving it to her would be worse for no one.
- 3*. *Disability Discrimination**: When the potential treatments are X and Z , the policy demands a lottery over X and Z that gives Beth at least as great a chance as Alice, even though (a) Beth stands to gain fewer years of life, and (b) we would favor Alice if she were not disabled.

If we avoid these problems, we introduce a fourth. Since Z merits more probability than Y , Y merits more probability than X , and X merits more probability than Z , we are back to another kind of cyclic ranking. Intuitively, one alternative should only get more probability than another if choosing that alternative would be preferable to choosing the other (if chosen deterministically). This would imply that the relation "is preferable to" is cyclic. If "is preferable to" is cyclic, that brings back all of the problems of *Cyclic Preferences*, which we were trying to avoid.

Conclusion

We have shown that it is impossible for a policy to provide guidance in interpersonal tradeoffs between length of life and quality of life without facing one of four very challenging conclusions. Of these, we think that those most likely

⁸ Some people may be tempted to think that it isn't so bad to embrace *Preference for Smaller Benefits** on the grounds that a fair coin toss is the appropriate response to this case. A difference of 15 years should be enough to make this implausible. We could adjust the case by choosing a more debilitating condition and allowing an even larger gap in years. For this solution to work in general, one must be willing to do fair coin tosses even when the difference in benefits could be very great. This idea is absurd; it is anathema to the very idea of priority setting in health care.

to be accepted are *Disability Discrimination* and *Cyclic Preferences*. We argued that approaches involving *Cyclic Preferences* face severe problems and should be rejected.

Although we have not directly argued in favor of the QALY approach's treatment of the disabled, we have shown it to be substantially more plausible in light of the challenges faced by all of its competitors. This makes it a lot less clear that we should change the QALY system and thereby throw away the great health gains it has achieved.

References

- Broome, J. (1984). Selecting people randomly. *Ethics*, 95, 38–55.
- Hadorn, D. (1999). The Oregon priority-setting exercise: Quality of life and public policy. *The Hastings Center Report*, 21(3), 11–16.
- Harris, J. (1987). QALYfying the value of life. *Journal of Medical Ethics*, 13, 117–123.
- Johannesson, M. (2001). Should we aggregate relative or absolute changes in QALYs? *Health Economics*, 10, 573–577.
- Nord, E., Pinto J. L., Richardson, J., Menzel, P., and Ubel, P. (1999). Incorporating societal concerns for fairness in numerical valuations of health programs. *Health Economics*, 8, 25–39.
- Nord, E., Menzel, P., and Richardson, J. (2003). The value of life: Individual preferences and social choice: A comment to Magnus Johannesson. *Health Economics*, 12, 873–877.
- Norheim, O. (2011). Atkinson's measure of inequality: Can measures of economic inequality help us understand tradeoffs in healthcare priority setting? In N. Eyal, S. A. Hurst, O. F. Norheim, and D. Wikler (Eds.), *Inequalities in health: Concepts, measures, and ethics* (Chapter 14). New York: Oxford University Press.
- Singer, P., McKie, J., Kuhse, H., and Richardson, J. (1995). Double jeopardy and the use of QALYs in health care allocation. *Journal of Medical Ethics*, 21, 144–150.