

Bursting Bubbles? QALYs and Discrimination

Ben Davies

The Oxford Uehiro Centre for Practical Ethics, University of Oxford

This is a pre-publication draft – please cite the published version.

Abstract

The use of Quality Adjusted Life Years (QALYs) in healthcare allocation has been criticized as discriminatory against people with disabilities. This paper considers a response to this criticism from Nick Beckstead and Toby Ord. They say that even if QALYs are discriminatory, attempting to avoid discrimination – when coupled with other central principles that an allocation system should favour – sometimes leads to irrationality in the form of cyclic preferences. I suggest that while Beckstead and Ord have identified a problem, it is a misdiagnosis to lay it at the feet of an anti-discrimination principle. The problem in fact comes from a basic tension between respecting reasonable patient preferences and other ways of ranking treatment options. As such, adopting a QALY system does not solve the problem they identify.

1. INTRODUCTION

The Quality-Adjusted Life Year (QALY) is a measurement that takes additional life-years gained from a health intervention and weights them by quality. We can use QALYs to decide between individual patients. We can extend Rikesh's or Agnieszka's life by ten years. Agnieszka will be healthy while Rikesh will be blind. If life in full health is better than life while blind, choosing Agnieszka produces more good. We can also use QALYs to set thresholds for cost-effectiveness by refusing to fund treatments that cost more than a certain amount of money per additional QALY.

In calculating QALY weightings, typically members of the public imagine themselves as having a health condition, and their preferences regarding that condition generate a decimal score for one year of life in full health with that condition, with '1' equivalent to full health.¹ For instance, they might say how many years with the condition they would trade for one year of full health. If people standardly would trade two years while blind for one year in full health, blindness receives a QALY weighting of 0.5.

This has led to criticism² that QALYs discriminate against disability. Persad et al summarize one version of this complaint: 'Even if a life-year [while disabled] ... is worse than a healthy life-year, someone adapted to [the disability] ... might reasonably value an additional life-year ... as much as a non-disabled person'.³ The issue here is how far judgements of social value should take priority over personal utility judgements. Others worry that QALYs put disabled patients in

¹ D. Franklin, 'Calibrating QALYs to respect equality of persons', *Utilitas* 29 (2017), pp.65-87, at 66.

² J. Harris, 'QALYfying the value of life', *Journal of Medical Ethics* 13 (1987), pp. 117-123; P. Menzel, P. Dolan, and J. Richardson 'The role of adaptation to disability and disease in health state valuation: A preliminary normative analysis', *Social Science and Medicine* 55 (2002), pp. 2149-2588; F. Kamm, 'Deciding whom to help, health-adjusted life years and disability', *Public Health, Ethics, and Equity* ed. S. Anand, F. Peter & A. Sen (Oxford, 2004), pp. 225-242; D. Brock, 'Cost-effectiveness and disability discrimination', *Economics and Philosophy* 25 (2009), pp. 27-47.

³ G. Persad, A. Wertheimer and E.J. Emanuel, 'Principles for Allocation of Scarce Medical Interventions', *The Lancet* 373 (2009) pp. 423-31, at 427.

‘double jeopardy’:⁴ if their lives really are worse, they have already been disadvantaged, and so it seems unfair to disadvantage them again. Finally, QALY judgements may routinely misrepresent disability. For instance, Barnes⁵ argues that people’s disabilities are less intrinsically harmful for them than most people assume. This would mean that the social judgements embodied in QALYs are themselves discriminatory, since they reflect inaccurate bias about the quality of disabled people’s lives.

Defenders of QALYs often reject these charges outright.⁶ This paper considers a different response, from Nick Beckstead and Toby Ord.⁷ They maintain that even if QALYs are discriminatory, no healthcare allocation system can avoid all problems; trying to solve one problem generates others, like pushing down ‘bubbles under the wallpaper’ only to find new

⁴ Harris, ‘QALYfying’, pp. 119-20.

⁵ E. Barnes, *The Minority Body: A Theory of Disability*. (Oxford, 2016).

⁶ J. Cubbon, ‘The principle of QALY maximisation as the basis for allocating health care resources’, *Journal of Medical Ethics* 17 (1991), pp.181-184; P. Singer, J. McKie, H. Kuhse and J. Richardson, ‘Double jeopardy and the use of QALYs in healthcare allocation’, *Journal of Medical Ethics* 21 (1995), pp. 144-150; M. Rawlins and A. Dillon, ‘NICE discrimination’, *Journal of Medical Ethics* 31 (2005), pp. 683-684.

⁷ N. Beckstead and T. Ord, ‘Bubbles under the wallpaper: Healthcare rationing and discrimination’, *Bioethics: An Anthology* 3rd edn., ed. H. Kuhse, U. Schüklenk, and P. Singer (Oxford, 2016), pp. 406-412. See also N. Beckstead and T. Ord, ‘Rationing and rationality: The cost of avoiding discrimination’, *Inequalities in Health: Concepts, Measures, and Ethics*, ed. N. Eyal, S. Hurst, O. Norheim, and D. Wikler (Oxford, 2013), pp. 232-239.

bubbles elsewhere.⁸ After outlining and refining their argument (Sections 2 and 3), I suggest that they have misdiagnosed the problem (Section 4), which comes not from a commitment to avoiding discrimination, but from a clash between two ways of ranking outcomes: by patient preference, or by some method external to patient preference. Section 5 considers possible responses and argues that a solution need not discriminate against disabled patients.

2. NECESSARY COSTS

Even if QALYs discriminate unjustly, we cannot automatically conclude that we should abandon them; they may be the best available option. This is Beckstead and Ord's line on disability discrimination. They argue that as well as disability discrimination ('Disability'), allocation systems should avoid two other problems.

The first is preference for smaller benefits ('Smaller'): where there is 'nothing else to choose' between two patients, we should not prefer 'producing smaller benefits rather than larger ones'. To think otherwise, they suggest, is 'morally perverse'. We must also avoid 'pointless' violations of autonomy ('Autonomy'). When a patient prefers one treatment, we should not insist that she receives a second treatment 'for no gain at all'.⁹ I discuss each principle below. First, though, I will complete their argument.

⁸ Beckstead and Ord, 'Bubbles', p. 407.

⁹ Beckstead and Ord, 'Bubbles', p. 407.

We cannot respect all three of these claims, fully, in all cases. Consider Case 1 where X, Y and Z represent incompatible treatment options, and their prospects for extending Alice’s or Beth’s life (at the same cost). Like many people, Beth prefers Z to Y.¹⁰

Case 1

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

Principles deriving from an attempt to avoid each problem collectively lead to cyclic preferences. Cyclic preferences are preferences where each option is preferred to another, in a way that gives no overall hierarchy of preference (e.g. $Y > X > Z > Y$). This means we can never alight on a final decision. For any possible option, some principle will recommend another choice as preferable.

‘Disability’ (together with ‘Smaller’) tells us to prefer X to Z: avoiding discrimination requires that we judge benefit size only according to the number of life-years gained. As such, the only relevant difference is that X gives Alice more additional years than Z gives Beth. So, $X > Z$.

‘Smaller’ tells us to prefer Y to X. Even if we don’t ignore disability, Beth benefits more from Y than Alice does from X. So, $Y > X$.

The apparent transitivity of preferability suggests that if $Y > X$ and $X > Z$, then $Y > Z$. But given Beth’s reasonable preferences, Beckstead and Ord say it would violate ‘Autonomy’ to insist that

¹⁰ Beckstead and Ord, ‘Bubbles’, p. 407.

she receive a treatment she does not want, when the treatment she does want is available for the same cost. As such, $Z > Y$. I discuss this principle further in the next subsection, but as an initial interpretation, let us say that it is acceptable for patients to reject the most beneficial treatment, so long as their preference for a different treatment is reasonable. If *only* Beth were being treated, and she simply had to choose between two options, it would be perverse to deny her the option she (reasonably) prefers.

The preferences under consideration are the preferences of those making the final decision about which treatment option should be chosen. We are not directly comparing Beth's intrapersonal preference for Z over Y with the interpersonal preferences (e.g. embodied in judgements of social value) that rank Y over X and (if we ignore disability) X over Z. If we were directly comparing these preferences, one might object that we have no reason to *expect* them to avoid cyclicity. But we *should* expect the overall preferences of those making distributive decisions to avoid cyclicity.

Beckstead and Ord suggest that the unappealing conclusions 'most likely to be accepted' are Disability, or cyclicity.¹¹ Since they clearly regard cyclic preferences as irrational, it is obvious which they prefer to jettison. Disability discrimination is unfortunate, but unavoidable.

However, their treatment of the principles is unclear. So, before continuing with the argument, I offer some clarification and, it turns out, revision of the principles.

2.1 The principles: 'Disability'

¹¹ Beckstead and Ord, 'Bubbles', p. 410.

Beckstead and Ord assume that avoiding disability discrimination means requires disability altogether. While that might be challenged, I will grant this assumption. Even so, a decision to ignore disability, and to measure benefit size only in life-years, doesn't require a belief that disability *cannot* make a person's life worse. As mentioned in Section 1, some who call QALYs discriminatory agree that some disabilities make life worse, but deny that this should influence healthcare allocations.

In such cases, the decision to ignore disability is an artificial restriction, made for the sake of justice. This is similar to the way many factors that affect welfare are treated by a QALY system. The fact that someone is very poor may affect how much a treatment benefits them. But a person's QALY prognosis typically does not directly include this information. This isn't because those who apply QALYs deny that poverty makes one's life worse, but because it would be unfair to include. I return to this in my discussion of 'Smaller'.

Others deny that (certain) disabilities makes people's lives intrinsically worse; in that case, the judgement of benefit size by life years is less artificial and reflects the denial that disabled years are intrinsically worse than full-health years.

2.2 The principles: 'Smaller' and 'Autonomy'

Beckstead and Ord say surprisingly little about what gives violation of autonomy a 'point', but as we will see, the most likely interpretation requires considering Autonomy and Smaller together. First, take Smaller. Beckstead and Ord think we should prefer the greater benefit when there is

‘nothing else to choose’ between patients. This is important because one might otherwise protest that there are often *very* good reasons for preferring smaller benefits.¹²

However, ‘larger benefits’ could have two meanings. To see why, recall the point from the previous section, that any allocation system used by an institution for measuring benefit size may ignore some facts about people that affect how much an intervention will benefit them. For instance, QALYs consider only *health* benefits. But people who allocate by using QALYs might know that other things will determine how much a treatment will benefit a person, or how much she values it. I already mentioned the role of poverty in how much a treatment might benefit someone. Nonetheless, we may avoid allowing this knowledge to influence our allocation decision because this would be discriminatory. Similarly, someone may prefer to live longer, even at a lower quality of life, because he wants to be in his children’s lives for as long as possible. And yet we presumably do not want to have our evaluation criteria formally assume that people with children will automatically prefer longer, low-quality lives.

This means that those making allocation decisions might make judgements in two capacities. In their *institutional* capacity, they may judge that treatment T is more beneficial than treatment S for a patient, because T is better with respect to factors their allocation system explicitly considers. Yet in their *personal* capacity they may agree with the patient that, all things

¹² For instance, there may be egalitarian reasons to prioritize worse off patients. See T. John., J. Millum and D. Wasserman, ‘How to allocate scarce health resources without discriminating against people with disabilities’, *Economics and Philosophy* 33 (2016), pp. 161-186, at 167-8.

considered, S is a reasonable preference. This will occur unless their institution decides to include all possible factors in its calculation of benefit.

When Beckstead and Ord advocate larger benefits, they may mean *as measured by the factors our healthcare system has decided to explicitly consider*. Alternatively, they may mean that we should prefer larger benefits, *as measured when considering all possible factors*, including those that we have decided to have our system ignore for other reasons. The second option seems perverse: if an institution has good reason to ignore some factors, insisting that we must nonetheless take them into account in a comprehensive fashion undermines that institutional decision. Furthermore, we must be consistent in our approach: either we ignore such factors in both cases, or neither. This means that, in a system that explicitly ignores disability (as does the one in which Alice and Beth are treated), we cannot implicitly appeal to the idea that a shorter life without disability is ‘really’ better than a longer one with disability. So, in the above scenario, Smaller must tell us that $Y > X > Z$, which is a consistent ordering.

Turning now to Autonomy, one thing that isn’t clear in Beckstead and Ord’s argument is the briefly mentioned role of experts. In characterising the choice of Y as a pointless violation of Beth’s autonomy, they claim that ‘Y is worse for Beth than Z’, and that choosing Y would produce ‘what she and experts both regard as a worse outcome’.¹³ My interpretation of this is that Z is supposedly better because it offers a larger benefit (if we do not ignore disability), and the evidence for this is that experts prefer it. This suggests that if Beth preferred Y, which offers fewer QALYs, violating her autonomous preferences *would not* be pointless: the point to violating her

¹³ Beckstead and Ord, ‘Bubbles’, p.407.

autonomy would be to maximize benefits. On this view, we respect Autonomy only when it aligns with Smaller, where what is smaller is determined always by 'experts' (who accept the QALY system!).

That causes significant problems. As Beckstead and Ord say in defending the claim that a discrimination-avoiding principle should prefer X to Z, 'the only thing Z has in its favour is that Beth would be at full health, whereas X would leave Alice with a disability. Thus, choosing Z over X involves discriminating on the grounds of disability'.¹⁴ However, this seems inconsistent with the thought that we should only attend to Beth's preference because she prefers what is 'better' for her, as determined by a measurement system that discriminates on grounds of disability. The same principle – whether it ignores disability or not – must apply in all pairwise comparisons. Effectively, Beckstead and Ord 'smuggle in' the claim that we ought not ignore disability when judging benefit size, even though the hypothetical system they are judging does just that.

I have drawn a distinction between factors that an institution explicitly allows to affect measurements of benefit size, and factors which we might acknowledge as reasonable for patients to use in their personal judgements. This seems a more reasonable ground on which to make the appeal to Autonomy: while Beth prefers the lesser benefit (as measured by the disability-ignoring allocation system), she does so reasonably and thus has a (defeasible) claim to that preference being satisfied. This is not an absolute commitment to satisfying patient preferences. As one anonymous reviewer notes, we should not allow Beth to refuse treatment and demand the equivalent cash value. Some may think that this applies more generally: once

¹⁴ Beckstead and Ord, 'Bubbles', p.407.

we have selected the best treatment, patients should not have the right to opt for a different treatment. But insofar as she disagrees *reasonably* with the institutional assessment of her treatment options, it seems excessive to force her to take an option she does not prefer, when a preferred alternative is available.

So, I suggest that Autonomy is plausible only if it has some independence from judgements about which option is best, contrary to my reading of Beckstead and Ord's version of the principle. We should try to avoid violating patient's autonomous choices, even if they choose the option that we think is worse for them, or which comes out as worse according to our way of measuring benefit.¹⁵ Even with this revision, though, the problem Beckstead and Ord identify remains: the three principles give rise to cyclic preferences. Section 4 argues that this is not for the reason they suggest.

3. RIGHTS

First, however, I will briefly consider an attempt to rescue Case 1 from cyclicity, while respecting Disability, that Beckstead and Ord explicitly reject. This proposal appeals to rights. First, when only one patient can receive treatment, the patient who has the best available outcome, ignoring disability (in this case, Beth with option Y) has the *right* to receive treatment. But patients have

¹⁵ This aligns to some extent with public views on parents' rights to choose substandard treatments for children, with a recent survey of public attitudes showing a majority tolerating substandard treatments that increased the risk of death by up to 5%. See T. Nair, J. Savulescu, J. Everett, R. Tonkens and D. Wilkinson 'Settling for second best: when should doctors agree to parental demands for suboptimal medical treatment?' *Journal of Medical Ethics* 43 (2017), pp. 831-40.

separate rights, once they qualify for treatment, to choose among reasonable options. Once Beth is selected, she may choose Z instead of Y, even though Y qualified her for treatment. Their response¹⁶ is that this violates a further plausible constraint on decisions, ‘independence of irrelevant alternatives’, which says that if we know an option will not be chosen, it should not affect our decision.

Given this response, though, it is unclear why we should consider Option Y in the original case, since it will not be chosen. Rather than assigning the right to treatment to the patient who *could* benefit most, why not assign the right to treatment to the patient who *will* benefit most? Beth prefers Z to Y, so this means comparing Z with X, and ignoring Y. In this case, since we are ignoring disability, $X > Z$.

A concern with this proposal, to which I return in section 5, is that if Beth *knows* she may be in competition with others, she has a perverse incentive to misreport her preferences. If she claims to prefer Y, she will be picked for treatment. Though she prefers Z to Y, she presumably prefers Y to X; but on the proposed system, it is X that will be chosen if she reports her true preference. It is not entirely clear how serious this flaw is. It is unlikely that most people will be aware of the system that is used or would engage in such manipulation. On the other hand, it is problematic if administrators have reason to hide, or poorly publicize, details about the system in use: transparency is a virtue of public decision-making.

The remainder of this paper shows that Beckstead and Ord’s preferred allocation system leads either to cyclic preferences, or to one of the two problems they have not given up on

¹⁶ Beckstead and Ord, ‘Bubbles’, pp. 408-9.

(‘Smaller’; ‘Autonomy’). This is because the problem they identify comes from a more basic tension between respect for reasonable patient autonomy, and any system for calculating benefit size that is somewhat independent of patients’ subjective preferences. After considering several ways to resolve this tension, I conclude that maintaining a commitment to disability rights has no greater costs than a pure QALY system; if avoiding disability discrimination is of independent value, we should prefer an allocation system that does so.

4. BURSTING BUBBLES?

The cyclic preferences in Case 1 derive from a conflict between respecting patients’ preferences, and a notion of benefit size that is not fully governed by those preferences. As such, cyclicity can occur for any system that adopts an objective measure of benefit size. Beth’s preference for QALY-maximization conflicts with the method for determining benefit size operating in Case 1 (pure length of life). But this is not because the method ignores disability. All we need for a clash is that:

- A patient, P, has two treatment options
- Their (reasonably) preferred treatment option is worse (according to our method of measuring benefits) than the only option available to a second patient, Q
- P has another option that is better (according to the same method) than Q’s only option, but which P reasonably does not prefer.

The same tension can come from a commitment to any method of measuring benefits, if patients might reasonably prefer a ‘smaller’ benefit. Consider:

Case 2

	Option A	Option B	Option C
Chris	60 years (blind)	Death	Death
Don	Death	55 years (blind)	35 years (full health)

Assume Don prefers a longer life while blind over a shorter life in full health. Let’s treat this case as Beckstead and Ord prefer, and apply an autonomy-constrained QALY system. This means assigning lower weighting to life-years that involve disability, aiming to maximize benefit size, and respecting reasonable patient preferences. To begin with, Chris stands to benefit more from A than Don does from B. So, per ‘Smaller’, $A > B$.

If we do not ignore disability, ‘Smaller’ might say that Don’s gain from C is greater than Chris’s gain from A. This will occur if blindness receives a QALY score of less than 0.58. Assume this is true, and that $C > A$.

To avoid cyclicalities, we should find that $C > B$. Indeed, applying QALYs alone gives us this transitive ordering. But Beckstead and Ord argue for a QALY system that avoids pointless violations of autonomy. I argued in Section 2.2 that, if this principle is not empty, it cannot say that violating autonomy is non-pointless whenever it would lead to a greater benefit as judged by QALYs. So, given Don’s preference, ‘Autonomy’ (as I interpret it) tells us that $B > C$.

So, we again have cyclic preferences, generated without avoiding ‘Disability’. In other words, the tension can be generated by adhering to the two principles Beckstead and Ord think we should keep.

5. POSSIBLE RESPONSES

I have argued that although Beckstead and Ord have identified a genuine problem, once we understand Autonomy in the way that is required for it to be an independent principle (i.e. as not entirely subservient to benefit size), the identified problem can occur even in allocation systems that don't ignore disability. I will now consider some alternative proposals to avoid or resolve this issue and note that each has problems. However, my more general conclusion is this: since a commitment to ignoring disability does not make the problem worse; sometimes mitigates it (see Section 5.2); and is by Beckstead and Ord's own lights a desideratum of healthcare allocation, the tension identified gives us no reason to abandon a commitment to ignoring disability.

5.1 Proposal 1: Give 'Smaller' lexical priority

I criticized Beckstead and Ord for – it seemed to me – interpreting Autonomy in a way that gives it value only when it agrees with Smaller. But they might claim this as a feature rather than a bug: perhaps benefit size *should* take lexical priority over autonomy, so that we have reason to respect a patient's autonomous decisions only when they choose the objectively best outcome.

However, this considerably weakens their argument, which claims that QALYs require giving up the fewest fundamental commitments to avoid irrationality. This is because this proposal amounts to giving autonomy no independent value at all: if we only respect your decisions when you choose what is best, you have only one option and hence are not free to choose. So, this option involves rejecting two desiderata, not just one as Beckstead and Ord claim.

5.2 Proposal 2: Dismiss Case 2

Beckstead and Ord might acknowledge that Case 2 demonstrates a tension between ‘Autonomy’ and ‘Smaller’. However, they might insist that this tells us nothing about the how we should respond to the entirely separate conflict that Case 1 highlights, which is between ‘Autonomy’ and ‘Smaller’ on the one hand, and ‘Disability’ on the other. Perhaps one of ‘Autonomy’ and ‘Smaller’ is most important (see proposals 1 and 4), but both trump ‘Disability’.

Yet not only does a principle of ignoring disability not complicate Case 2: it helps. If we ignore disability, Chris will benefit most. So Smaller tells us to prefer A over B *and* C. What’s more, the problem in Case 1 comes not from Disability directly, but from its establishing opposition between autonomy and benefit size; the reverse effect is why it simplifies Case 2. So, it misrepresents things to present these as two separate tensions.

5.3 Proposal 3: ‘Subjective’ QALYs?

The identified tension emerges because patients can (reasonably) prefer the option that is rated lower by our allocation system. Can we resolve the tension by calculating QALYs subjectively, rather than by appeal to public opinion? On this view, B generates more QALYs than C for Don, though it wouldn’t for most people. For Beth, Z generates more QALYs than Y. By stipulation, Smaller points the same way as Autonomy.

Yet whatever else we think about the idea, Proposal 3 clearly has significant (and possibly insurmountable) practical disadvantages in eliciting preferences. A healthcare system cannot reasonably be expected to elicit not only patients’ preferences, but a reliable weighting system that could be used in interpersonal comparison, from every individual patient.

5.4 Proposal 4: Give Autonomy lexical priority

I suggested earlier that we might prioritize patient autonomy lexically by only considering options competing patients prefer, while respecting Disability. But it's also possible to also do this in a QALY-based system. We would ignore whichever choices would not be chosen by the relevant patient (i.e. Y and C) and assign treatment based on whoever has the larger QALY gain from their preferred treatment. In Case 1, that's Beth. In Case 2, it's Chris. Benefit size is still determined according to QALYs, not life-years. But benefit size is trumped by respect for patient preference within a reasonable range of options.

This option inevitably suffers from the problem raised earlier, of incentivising misreporting. If Don knows he is choosing the officially 'worse' option, he will also know that he is thereby increasing the chance of his life not being saved at all. For he will know that there is at least a possibility that he will be in competition with someone like Chris, i.e. someone who faces a better QALY score than Don's preferred option, but a worse QALY score than the option he doesn't prefer. As such, he has a perverse incentive to misreport his preferences, undermining our commitment to patient autonomy.

Moreover, following Beckstead and Ord's logic, if two systems are otherwise equivalent, but one meets an additional desideratum, we should opt for that system. So, if they choose Option 4, Beckstead and Ord cannot resist reinstating 'Disability', as per my version of this proposal.

6. CONCLUSION

Beckstead and Ord's argument aims to show that we cannot accept a commitment to avoiding disability discrimination, together with commitments to respecting patient autonomy and to preferring larger benefits, without generating irrational sets of preferences. I have shown that

the tension they identify arises independently of disability discrimination. It affects any system that aims for larger benefits and respects patient autonomy without defining one purely in terms of the other. Though no solution seems entirely attractive, there seems to be no reason to think that the best proposal will require disability discrimination.¹⁷

¹⁷ Thanks to participants in the King's College London summer seminar (2017), Dale Miller, and to two anonymous referees for feedback on this paper.