

Against “Saving Lives”: Equal Concern and Differential Impact

(Forthcoming in *Bioethics*)

Richard Yetter Chappell*
University of York

April 3, 2015

Abstract

Bioethicists often present “saving lives” as a goal distinct from, and competing with, that of extending lives by as much as possible. I argue that this usage of the term is misleading, and provides unwarranted rhetorical support for neglecting the magnitudes of the harms and benefits at stake in medical allocation decisions, often to the detriment of the young. Equal concern for all persons requires weighting equal interests equally, but not all individuals have an equal interest in “life-saving” treatment.

Introduction

Alice, Brian, and Cody lie in hospital, suffering from different diseases, but all in need of the same rare drug treatment if they are to survive. Our

*Thanks to Anne Barnhill, Ezekiel Emanuel, Alexander Guerrero, Govind Persad, Peter Singer, Dominic Sisti, Helen Yetter-Chappell, and the Practical Philosophy group at the University of York, for helpful discussion and comments.

stores of the drug would prove sufficient to save either Alice alone, or else Brian and Cody both. All else equal, the latter option—saving the most lives—seems like the way to go. But further suppose that Brian and Cody are cancer patients, with a life expectancy of only five years each (that is, if they survive their current illnesses), whereas Alice is in otherwise good health, with a remaining life expectancy of twenty years. Should this change our verdict? I will argue that this further information *should* lead us to changes our minds, rebutting arguments to the contrary that are extant in the bioethics literature. I will further argue, more controversially, that “saving lives” is not *in itself* a fundamental moral goal at all. To a first approximation, it is *life-years*, rather than *life-numbers*, that our decisions ought to track. And while there are indeed some respects in which the distribution of life-years matters, I will argue, it is not straightforwardly a matter of wanting them to be distributed over as large a number of people as possible, in the way that ascribing value to “saving lives” would suggest.

We often speak of “saving lives”, but such loose talk can be misleading. Strictly speaking, the best we can do is to save someone several *years* of life. Death is never banished; merely postponed. We know there will be exactly one for each life, and there’s no changing that (modulo immortality research). What we typically mean instead is that we aim to *extend* life. It’s worth being clear on this, since a rhetorical focus on “life-saving” can occlude the important fact that not all life-extensions are equal.

I illustrate this problem by way of two case studies. In section 1, I address

Harris (1987)’s ‘discrimination’ and ‘double jeopardy’ objections to the use of Quality-Adjusted Life Years (QALYs), showing how both ultimately rest on this rhetorically-induced conceptual confusion. Section 2 then argues against the common view, expounded most clearly by Bognar and Kerstein (2010), that we should treat person-*numbers* aided as an important moral goal, competing with that of simply aiding people (in the aggregate) as much as possible. To this end, I explore some important respects in which not all life-years are equal—a fact that can help explain our intuitions about cases that might otherwise be taken to suggest that person-numbers intrinsically matter. These conclusions have important implications for (*inter alia*) health policy, research priorities, and resource allocation—in short, whenever we might face a conflict between “saving” more lives or extending lives for longer.

1 Harris on the Value of Life

1.1 QALYs, Discrimination, and Equal Concern

Harris (1987) famously objects to the use of QALYs in guiding medical rationing decisions, on the grounds that: “Each life counts for one and that is why more count for more. For this reason we should give priority to saving as many lives as we can, not as many life-years” (p.120). This makes no sense if, as Harris goes on to admit, “talk of ‘life-saving’ [...] must always be understood as ‘death-postponing’.” The idea that “each life counts for one” should not be understood to mean that all life-extensions are equal. To treat

one saved year for me as equal in value to twenty saved years for you would, on the contrary, be to massively over-weight my interests in comparison to yours (cf. Broome 1984, 52).

To say that “each life counts for one”, or to affirm our equal worth as persons, is simply to say that we must “treat each person with the same concern and respect.” (Harris 1987, 121) The question in dispute is how we are to achieve this—that is, what *treating people with equal concern and respect* ultimately amounts to. I suggest that we do not achieve this by neglecting the comparative *weight* of the interests at stake. Harris himself seems to accept this general principle when he approvingly quotes Dworkin (1977, 227): “If I have two children, and one is dying from a disease that is making the other uncomfortable, I do not show equal concern if I flip a coin to decide which should have the remaining dose of a drug.” But, as I will now argue, he fails to respect this principle when different life-extensions are at stake.

A general lesson nicely demonstrated by Dworkin’s case is that moral equality doesn’t require *identical treatment*, but just that our *interests* should be considered fully and fairly—i.e., neither person’s interests discounted or neglected to the unfair benefit of the other. Simply put: We should not prevent a lesser harm (or give a lesser benefit) to one at greater cost to another. All else equal, we should prefer to provide the most help we can to people, regardless of who they are.¹

¹ Some may wish to restrict Dworkin’s principle to the prevention of harms only, rather

But now notice that a longer period of life-extension is, all else equal, a greater benefit than a shorter period would be. So we should generally prefer to grant a longer life-extension over a shorter one, if forced to choose between them. Put another way: The harm of death lies in the future goods that it deprives us of (Nagel 1970), and hence your early death is worse for you the more good life you would have had if you hadn't died at that time.² Since we should prioritize the prevention of worse harms, so we should prioritize postponing the deaths of those for whom death can be postponed the longest, because their early death thereby harms them more. (I mean to leave open

than also extending it to the bestowal of benefits. For example, you may find it plausible that equal concern for two children *could* be shown by flipping a coin to decide which of them will get (say) a rare toy that one would enjoy greatly and the other would enjoy somewhat less. That's fine: For purposes of this paper, I don't need to take a stand on that further controversy. For although I sometimes (for ease of exposition) talk of life-extension as a “benefit”, it is arguably better categorized as the prevention of a harm (namely, the harm of an early death that deprives one of such-and-such future goods that one's life would otherwise encompass—see Nagel 1970). As Dworkin's own example brings out, it is much more natural to compare curing a fatal illness to curing a non-fatal illness than it is to compare it to a happiness-boosting enhancement. So, I think we can plausibly assume that life-extension/death-postponement should be treated as a kind of harm-prevention, regardless of whether one thinks that harms and benefits should be treated symmetrically in general. Thanks to an anonymous referee for bringing this challenge to my attention.

² Does this mean that we should prioritize the life of a newborn infant over (say) that of a ten year old child? Not necessarily. On some views, such as McMahan (2002)'s *Time-Relative Interests Account*, the comparative lack of robust psychological connections between an infant and their future stages means that the infant is less harmed by death's depriving them of that future. (You might say that the future in question does not sufficiently “belong” to the infant for it to be something that the infant could be “deprived” of.) But I do not mean to take a stand on that issue here. The current paper merely argues that we should allocate medical resources so as to most help medical patients. If you doubt that we should prioritize the lives of newborns over older children, then you should presumably reject the additional assumption that newborns are more greatly harmed by death than older children, rather than rejecting my core thesis that we should allocate scarce medical resources in the way that helps patients the most. (See section 2, below, for other proposed exceptions to the general rule that more aggregate life-years = more aggregate good.) Thanks to an anonymous referee for raising this objection.

that preventing a sufficient number of smaller harms may, in the aggregate, outweigh a single large harm. But sheer numbers should not trivially win out, the way that Harris’ view would imply.) In a similar way, extending life *at full health* is plausibly a greater benefit to the recipient than is extending life for the same period of time but at a lower “quality” due to poor health. Hence the use of *quality-adjusted* life years to measure the expected benefits of various medical interventions.³

We are now in a position to present the basic argument for employing QALYs in medical allocation decisions:

(P1) In cases of medical scarcity, we should generally⁴ aim to prevent the greatest (aggregate) harms, or provide the greatest (aggregate) benefits, to patients.

(P2) QALYs are the best measure we currently have of the magnitudes of harms and benefits involved in medical interventions.

∴ (C) In cases of medical scarcity, we should generally act so as to maximize the number of QALYs gained by patients.

³ It’s worth flagging that the move to *quality-adjusted* life years is likely to be more controversial than the basic recognition that gaining many years of life is generally a greater benefit than gaining just a few. One may, for example, worry that it violates political neutrality for the state to evaluate some lives as being of higher quality than others. On the other hand, it’s widely accepted that preventing disabilities is worthwhile, as is restoring full health and capability to someone currently beset by a disability, so there doesn’t seem any deep barrier to making and acting upon such evaluations after all. For more on the practical absurdities that result from rejecting QALYs on this basis, see [Beckstead and Ord \(2013\)](#).

⁴ I want to allow some wiggle-room here for prioritarians who give disproportionate weight to the interests of the worst-off, and hence in some circumstances allow a lesser benefit to one person to trump a greater benefit to another. We are at least in agreement that the magnitude of a potential harm or benefit *generally* matters.

Harris’ opposition to QALYs, by contrast, seems to rest entirely on his failure to appreciate that not all life-extensions are equal. Consider his core argument (p.121):

The State must, in short, treat each citizen with equal concern and respect [... and so] is not entitled to discriminate between individuals in ways that mean life or death for them on grounds which count the lives or fundamental interests of some as worth less than those of others. If for example some people were given life-saving treatment in preference to others because they had a better quality of life than those others [...] this would amount to regarding such people as more valuable than others on that account. [...]

The only way to deal between individuals in a way which treats them as equals when resources are scarce, is to allocate those resources in a way which exhibits no preference. To discriminate between people on the grounds of quality of life, or QALY, or life-expectancy, is as unwarranted as it would be to discriminate on the grounds of race or gender.

This is simply confused. In particular, Harris’ talk of “life-saving” leads him to confuse the value or worth *of a person* (as someone whose interests we should take into account, the equal of any other) with the value of *a particular life-extension for that person*. To prioritize “life-saving treatment” for people

with greater life-expectancy and quality of life is not to thereby treat *them* as more important than anyone else. It’s merely to recognize that we can *do more for* them than we can for others. One option yields a greater benefit to its recipient (or prevents a greater harm) than the alternative does, and so equal concern for both parties’ interests should lead us to prefer it—just as it should lead us to prefer to cure one child’s fatal disease rather than another’s sore throat. There’s nothing objectionably “discriminatory” about allocating resources to the people who will most benefit from them. On the contrary, it’s precisely *departures* from such equal consideration of interests (e.g., favouring a lesser benefit to one person over a greater benefit to another, on morally irrelevant grounds) that constitute unfair discrimination.

Harris is not alone in this confusion. For a more recent example, consider [Brock \(2005, 96\)](#)’s suggestion that “the disabled patient might argue that just like the nondisabled patient, she will lose everything—her life—without [life-saving treatment].” This, too, neglects the fact that not all life-extensions are equal: “Everything” is not a fixed quantity. If one person would have a better future than another, then the life—the “everything”—that early death robs her of is greater. Importantly: If QALYs are unjustly discriminatory here, it cannot be because the idea of allocation based on quality-adjusted life years is *essentially* unjust; rather, the objection can only be that some particular *implementations* of the concept exaggerate the harm involved in living with some disabilities, and so fail to accurately reflect the actual quality of life for people with those disabilities. (That may or may not be accurate—cf.

footnote 3 above. In any case, my aim here is not to defend the particular way that QALYs are currently implemented, but just the moral appropriateness of this general approach to resource allocation decisions.)

Or consider Nord (2005, 126):

I define the *worth* of a person as the value attached by society to the enhancement of the interests and opportunities of that person relative to the interests and opportunities of other persons. One important form of enhancement of interest is the protection of the person’s life. If society regards two persons as being of equal worth, it means, among other things, that it is willing to do equally much to protect their lives.

Again, this clearly confuses the concepts of *worth of* and *worth to* an individual.⁵ It’s perfectly coherent for society to attach equal value to “enhanc[ing] the interests” of all people, whilst recognizing that not all life-extensions offer equal such “enhancements”, and hence to be willing to do more to “protect” some lives than others.

Of course, we can all agree that public institutions should not prioritize benefits to people with high life-expectancy or quality of life when it could instead grant *equal or greater benefits* to others. That would indeed be objectionably discriminatory, and fail to treat all persons as being of equal worth.

⁵ Importantly, my claim here is not that either concept is inherently confused, but just that the two are distinct and ought not to be conflated. Thanks to an anonymous referee for inviting me to clarify this point.

But if our resource allocation decisions can *effect* different increases in life-expectancy and/or quality of life for different patients, then this constitutes a real difference in potential harms and benefits to the recipients, and one that morally ought to be taken into account. The rhetoric of “life-saving” tends to occlude the differential impact that “life-saving treatments” can have on different patients, leading to an unjust neglect of the differential impact that alternative allocation decisions would have on them. Not all life-extensions are equal, and to pretend otherwise does a grave injustice to the patients with the most to lose (or gain).

We may illustrate this point by returning to our original case of Alice, Brian, and Cody. The latter two would unquestionably benefit from the rare drug treatment: it would yield them a few extra years of (desired) life. But much *more* is at stake for Alice, who could go on to live for *decades* more if she receives the rare treatment instead. To deprive her of the treatment is to deprive her of much more than either Brian or Cody (even in aggregate) stand to gain. So, the equal consideration of interests mandates that we prioritize Alice’s treatment in this case, even though it leads to a smaller number of “lives saved”.

1.2 *Double Jeopardy?*

Harris further claims that the use of QALYs gives rise to an unfair kind of “double jeopardy” (p.120):

QALYs dictate that because an individual is unfortunate, because

she has once become a victim of disaster, we are required to visit upon her a second and perhaps graver misfortune. The first disaster leaves her with a poor quality of life and QALYs then require that in virtue of this she be ruled out as a candidate for lifesaving treatment, or at best, that she be given little or no chance of benefiting from what little amelioration her condition admits of. Her first disaster leaves her with a poor quality of life and when she presents herself for help, along come QALYs and finish her off!

Again, I think it will be informative to clear up the conceptual confusion underlying this argument. Harris makes it sound as though the worse off people are, the less inclined we (QALY advocates) will be to help them. This is of course not true. The QALY approach is an attempt to help people (generally) as much as possible. Whether this is to the advantage or the disadvantage of disabled persons (relative to an unadjusted life-years approach) is a contingent matter. To see this, note that if a class of medical treatments tended to yield a fixed prognosis (say of perfect health and capability) regardless of one’s antecedent condition, then severely disabled persons and others with a comparatively low antecedent quality of life would tend to be prioritized over those who are better off to begin with (and so whose prospects in the absence of treatment would be better). For more realistic cases, note that QALYs allow the sensible prioritization of cost-effective “quality-of-life” treatments like cataract surgery over expensive life-extending treatments for the (other-

wise able but) terminally ill. So the QALY approach is at least not *inherently* hostile to the interests of the disabled.

Indeed, on the most straightforward QALY approach, a patient’s *antecedent* condition is simply irrelevant. All that matters is the forward-looking counterfactual comparison of how many QALYs they stand to gain from treatment versus non-treatment (and how this potential benefit compares to that of other potential beneficiaries of our scarce medical resources, on a ‘per unit’ basis). Prioritarians might modify this approach by giving extra weight to those who are antecedently worse off. On no serious view do the worse off count for less.

The actual phenomenon that Harris struggles to articulate here is rather that the QALY approach is less likely to recommend treatments that offer less benefit to patients. Regardless of one’s antecedent welfare levels, if one’s *post-treatment* welfare would be very low, then this treatment is not as worthwhile as one that offers greater benefits to its recipients. In practice, this does mean that patients with low quality of life *and who would continue to have low quality of life even post-treatment* will be less eligible for life-extending treatments than patients who would instead have high post-treatment quality of life. And this could be considered “double jeopardy” in a sense. But it should not be regarded as an *objectionable* form of double jeopardy. It’s simply to prefer treatments that significantly help patients over those that do little good. If one’s condition admits of “little amelioration”, as Harris puts it, then that is certainly a tragedy—but our outrage should be directed

at the unfairness of *the world*, rather than (mis)attributing “unfairness” to the policy of prioritizing treatments that will help their patients more.

Once we appreciate that not all life-extensions are equal, the double jeopardy objection to QALYs thus collapses on similar grounds to the “discrimination” objection of §1.1. The underlying moral principle is morally innocuous, and the outcomes are good on the whole (if we count all people’s interests equally). What initially sounded so morally repugnant is instead seen to be an artifact of misleading rhetoric and selective attention.

2 Person-numbers vs. Person-years

More sophisticated critics of the QALY metric object that it wrongly neglects distributional concerns. Many would think that we should give some priority to the worst-off, rather than giving a slightly greater benefit to someone better off, for example. Reasons of fairness might similarly dictate that an extra year of life go to he who has so far lived the least. One could easily tweak the QALY approach to accommodate these ideas, if desired (cf. footnote 4). But more generally, one may worry that “giving a few life-years to many differs from giving many life-years to a few” (Persad et al. 2009, 425)—even if it isn’t always clear which is preferable. As we’ll see, I think there is some truth to this, in that not all life-years (even holding fixed one’s health status) are equally significant to the subject’s wellbeing. This means that some distributions of a fixed number of life-years may be preferable to

others, even on utilitarian grounds. Sometimes we should prefer to give a few life-years to many, and sometimes we should prefer to give many life-years to a few, depending on the precise details. But what I will dispute in this section is the idea that saving a larger *number* of people is itself intrinsically better.

Bognar and Kerstein (2010, 15) distinguish what they call ‘person-numbers’ from ‘person-years’:

To illustrate: if we save five people for two years, then on the person-years dimension we preserve 10 years, while on the person-numbers dimension we preserve five people. We believe that reflective common sense values the preservation of persons along these two dimensions.

Their terminology helps to bring out the distinction nicely, but talk of “preserv[ing] five people” still risks occluding the huge variability in value that different life-extensions have for their recipients, since “preservation” of life could be read in all-or-nothing terms, rather than essentially coming in degrees (as life-extension actually does). To avoid this risk, let us instead talk of person-numbers *aided*, alongside the total *magnitude* of the aid in person-years. This brings out more clearly the underlying fact that increases along the person-numbers dimension amount to simply *distributing life-years across a greater number of people*. Seen in this light, it is very difficult to see why this should be valued at all.

To further support this conclusion, let us consider two possibilities in turn. First suppose that all life-years are generally of equal value to a person, whether it is their 20th year or their 80th. Then, to value person-numbers independently of person-years would be to value distributing a fixed benefit across a greater number of beneficiaries. But there is no reason to value this in general. At most, one might posit “prioritarian” reasons to prefer benefits to the worst-off. But in a medical context, that would presumably mean concentrating life-extension benefits on those who have lived the least, rather than indiscriminately distributing life-years across as many patients as possible.⁶ If we assume that the first additional year is no more valuable to the recipient than the tenth, it would seem simply perverse to extend the lives of nine very elderly patients by a year each in preference to extending the life of a young adult by ten years. That would be to both choose a smaller “pie” *and* distribute it to those who have already had more. Every morally relevant consideration counts against this.⁷

⁶ Harris (1985, 93–101) actually endorses a ‘threshold’ version of this *fair innings* idea, giving all non-elderly some (equal) degree of priority over the elderly.

⁷ Nothing I’ve said here rules out that it may be permissible to end up helping the nine *as a result of conducting a weighted lottery* or the like. I’m merely pointing out that we clearly shouldn’t straight-out *choose*—or *prefer*—to help the nine over the one in this case, merely because they’re much greater in number (while being close in terms of total years gained), the way that intrinsically valuing the person-numbers dimension would seem to suggest we should. Finally, while Bognar and Kerstein can allow additional moral principles to affect what we should do all-things-considered, and hence aren’t strictly committed to what I identify as the morally worse choice in the above described scenario, I take the scenario to nonetheless have significant dialectical force insofar as it casts doubt on their claim that there are (necessarily) *any* significant pro tanto reasons to prefer increases along the person-number dimension, *per se*.

Of course, it's plausible that not all life-years are equal. Our time is a means to realizing what we value in life. So the value of an additional year depends on what we would achieve with it, or how it would contribute to the shape of our life as a whole (Kappel and Sandøe 1992). It's ultimately an empirical question how much each year of life typically contributes to the realization of our life plans and core values, but it seems plausible that the years between late adolescence and early retirement are typically the most “central” (normatively speaking, and not just chronologically) to our lives. Extending the life of a thirty year old by two decades is *transformative*, allowing him to (e.g.) raise a family and pursue excellence in his chosen career. Two more healthy decades for an eighty year old, while undoubtedly valuable (e.g., to enjoy a leisurely retirement, and see one's lineage extend to include great-grandchildren), are unlikely to be nearly so transformative or essential for the realization of the person's core values and life plans. If that's correct, it suggests that extending these “middle-aged” years would be a greater benefit than would extending an equal number of either very young or very old years. (But again, I stress, the precise details are open to empirical revision. Perhaps our working years are typically an unrelenting grind, which only the promise of a pleasant retirement can redeem.)

A related point is that many important life goods require time investments of several years to attain (either in full or to a significant degree). This can help explain our intuition that granting five more years to each of six people is generally preferable to granting six years each to only five people. We may be

implicitly assuming that the difference between five or six extra years would not be transformative for the individuals in question, whereas the jump for the sixth person from zero to five years would be transformative (and hence its loss would constitute a disproportionate harm: more than five times worse than any of the others stand to lose by having one less year).

Given those assumptions, even the straightforward consequentialist principle of maximizing net wellbeing would have us favour the former over the latter option. But this is not to say that we should uniformly prefer to distribute a fixed number of years over more people. At some point, the time granted to each would be too short for us to expect significant gains to be achieved from it. For example, offering just one year more to each of thirty patients might well be thought to offer less total benefit than would offering five years to each of six patients. (Whether this expectation is accurate will of course depend on the particular situations of the individuals in question, such as how much time they would need to make significant progress on important life projects.) As an extreme case, few of us would think that offering a single extra *second* to each of a billion people would be more worthwhile than the above-mentioned options (despite adding up to a bit over 30 years in total). I take this to show that what our intuitions are tracking here is not *person-numbers* as such, but rather expected opportunities for *significant*, life-changing benefits—opportunities that we generally expect to require significant amounts of time.⁸

⁸ Of course, [Bognar and Kerstein](#) can accommodate this extreme case by dismissing a mere second as an “irrelevant utility” (p.18, drawing on [Kamm 1993](#), 101-102). But the

So we see that—as in the case of Prioritarianism—these departures from treating all life-years equally do nothing to justify an indiscriminate broadening of the distribution across person-numbers, but instead call for a *more discriminating* distribution of life-years across patients, depending on their circumstances (e.g., their age).

There is one consideration that may count in favour of a broader distribution, in practice if not in principle. This is the thought that there are certain values—relating to “closure” or wrapping up any “loose ends” in one’s life—that may depend upon receiving *any* (non-trivial) amount of life extension. That is, if someone suddenly finds that their life is unexpectedly in danger, having a few additional months to “put their affairs in order” before dying might be time that is of disproportionate value to them. So, allowing many people such extra time would plausibly be better than simply giving one of the people in this position a great number of extra life-years, and having all the others die instantly.

But even this is not a reason to value person-numbers as such. It is merely another respect in which some life-years might be more important (for the person living them) than others. In particular, it is to suggest that we (at least, some of us) highly value having *some time in which to prepare for death*,

less-extreme variants remain—a year, for example, is certainly not *irrelevant*—and I think the alternative account I offer has the theoretical virtue of being more unified and less ad hoc. It is also more theoretically neutral, in the sense that it proposes an explanation for our moral intuitions which leaves open whether or not they are ultimately vindicated. Perhaps we are often mistaken in our intuitive assumptions about the marginal impact of short life-extensions on a person’s central life projects. I remain neutral on this.

rather than being taken unawares. But this doesn't *in general* depend upon receiving life-extending medical care. If one has a chronic condition that was previously diagnosed, or is otherwise aware that their time is running out (e.g., due to very old age), then additional time is merely desirable for the standard reasons. There is no special reason relating to unprecedented opportunities for “closure” in this case. So there would be no reason of that sort for favouring the distribution of a fixed number of life-years across more person-numbers in such a case. It is only in the case of *unexpected* and *immediate* lethal threats that life-extension (for even short periods) has this distinctive extra value.

Conclusion

Beneficence calls on us to help people as much as possible, given our limited resources. This will not be the same thing as helping as many people as possible, if we are able to help some much more (and more efficiently) than others. This is obvious, and widely appreciated, when comparing (say) a lethal disease vs. a sore throat. Everyone agrees that curing the lethal disease is more important—offers a greater benefit, or prevents a greater harm, to the patient in question—than curing a sore throat. What seems less widely appreciated is that the same holds true when comparing two “life-saving” treatments for patients with very different prognoses (in terms of expected quantity and quality of life) from the treatment. Not all life-extensions are

equal, and this fact must be recognized if we, as a society, are to allocate our scarce medical resources as best we can.

References

- Beckstead, Nick and Ord, Toby. 2013. “Rationing and Rationality: The Cost of Avoiding Discrimination.” In Nir Eyal, Samia Hurst, Ole Norheim, and Dan Wikler (eds.), *Inequalities in Health: Concepts, Measures, and Ethics*. Oxford University Press.
- Bognar, Greg and Kerstein, Samuel J. 2010. “Saving Lives and Respecting Persons.” *Journal of Ethics & Social Philosophy* 5:1–20.
- Brock, Dan W. 2005. “Preventing Genetically Transmitted Diseases while Respecting Persons with Disabilities.” In Wasserman et al. (2005).
- Broome, John. 1984. “Selecting People Randomly.” *Ethics* 95:38–55.
- Dworkin, Ronald. 1977. *Taking Rights Seriously*. London: Duckworth.
- Harris, John. 1985. *The Value of Life*. Routledge & Kegan Paul.
- . 1987. “QALYfying the Value of Life.” *Journal of Medical Ethics* 13:117–123.
- Kamm, F. M. 1993. *Morality, Mortality: Volume I: Death and Whom to Save From It*. New York: Oxford University Press.
- Kappel, Klemens and Sandøe, Peter. 1992. “QALYs, Age and Fairness.” *Bioethics* 6:297–316.
- McMahan, Jeff. 2002. *The Ethics of Killing: Problems at the Margin of Life*. New York: OUP.
- Nagel, Thomas. 1970. “Death.” *Noûs* 4:73–80.
- Nord, Erik. 2005. “Values for Health States in QALYs and DALYs: Desirability versus Well-Being and Worth.” In Wasserman et al. (2005), 125–141.
- Persad, Govind, Wertheimer, Alan, and Emanuel, Ezekiel J. 2009. “Principles for allocation of scarce medical interventions.” *The Lancet* 373:423–431.
- Wasserman, David, Bickenbach, Jerome, and Wachbroit, Robert (eds.). 2005. *Quality of Life and Human Difference*. Cambridge University Press.