

Efficiency and Equity in Health: Philosophical Considerations

(to appear as entry 2.8 in Anthony Culyer (ed.), *Encyclopedia of Health Economics* [Elsevier])

Paul Kelleher

University of Wisconsin-Madison

jkelleher@wisc.edu

ABSTRACT: Efficiency and equity are central concepts for the normative assessment of health policy. Drawing on the work of academic philosophers and philosophically sophisticated economists, this article identifies important philosophical questions implicated by the notions of efficiency and equity and then summarizes influential answers to them. Promising avenues for further philosophical research are also highlighted, especially in the context of health equity and its elusive ethical foundations.

Efficiency and equity are central concepts for the normative assessment of health policy.

Because there is a more extensive philosophical literature on the nature of health equity, this article devotes more attention to that debate. There are, however, important philosophical questions about the goal(s) that efficient policies are to aim at, and about the moral importance of efficiency in the context of health resource allocation. We begin with those questions.

1. The Concept of Efficiency

While the concept of efficiency is straightforward, it is difficult to define in more basic terms. It at least connotes an optimizing ratio of gains to losses. So construed, the concept cannot pick out a concrete property of distributions or policies until relevant gains and losses are specified.

Efficiency, in other words, is always efficiency toward some goal, and different contexts will implicate different goals.

Some economic analyses aspire to identify efficiencies across multiple policy domains. For example, a comprehensive cost-benefit analysis seeks that portfolio of policies which spans many domains and which maximizes overall net benefits. Such analyses convert each policy domain's benefits into monetary equivalents and assume that maximizing monetized benefits is a worthy goal. By contrast, domain-specific analyses seek locally efficient policies. For example, cost-effectiveness analyses in the domain of health care seek to identify policies that would maximize certain health-related outcomes given a fixed budget. Here there is no need to convert relevant outcomes to monetary equivalents since there is no need to express both health and non-health benefits in terms of some common unit.

This article will not discuss issues specific to cost-benefit analysis, although the discussion in section 3 about the relative moral importance of benefit-maximization bears on that enterprise as well. The discussion of efficiency that follows adopts a domain-specific perspective and investigates two questions: first, to the extent that efficiency is a worthy goal in health policy, what health-related benefits should be maximized (section 2)?; second, is a single-minded focus on efficiency morally problematic (section 3)?

2. Metrics of Maximization

Health

It is natural to think that efficiency in health policy should be construed as maximizing health itself. However, two related reasons have been adduced against that proposal. First, it can be difficult to make the assessments of overall health that it requires. Second, asking if someone is

in “good health” is often a way of asking if their health adversely affects their life. If health’s impact depends on the way it interacts with other features of one’s situation, then it may be misguided to focus on health itself rather than on the ways health, together with other factors, affects people’s lives.

Murray et al. (2002) respond to these and similar worries. With regard to health assessment, they note that “Almost everyone can agree that a person with one amputated leg is healthier than a person with two amputated legs” and that “somebody with a mild sore throat is, *ceteris paribus*, healthier than somebody with two broken arms” (736). In their view, these judgments represent plausible and useful assessments of health itself, not judgments about health states’ impact on the goodness or badness of a life. This is fortunate, since organizations like the World Health Organization are “fundamentally concerned with the improvement of health as distinct from aspects of well-being that are not intrinsic determinants of health” (737). But is it possible to build a rigorous assessment of population health around such health-focused judgments? Doing so would require a large number of health-state assessments of the sort Murray et al. describe. But many such judgments are not as clear-cut as the examples they offer. To illustrate the difficulty, Daniel Hausman (2006) presents the example of a person with a mild learning disability and someone with quadriplegia. Although the first person is presumably in better health than the second, Hausman doubts that one can meaningfully “compare units of mobility with units of cognitive functioning.” He therefore doubts there are grounds for saying that the first person is “literally healthier” than the second (251). For Hausman, this highlights the difference between saying the first person literally has more health (a descriptive judgment) and saying that it is better to be in his health state than to suffer from quadriplegia (an evaluative judgment). Hausman suggests that population health measurement should aim to evaluate

health—i.e. determine the goodness or badness of a population’s health—rather than to measure health itself.

Well-Being

If one seeks to evaluate the goodness or badness of a health state, a natural proposal is to focus on the impact the state has on individual well-being. Of course, much would turn on the nature of well-being, and philosophers have identified important problems for many accounts of it.

One central candidate is subjective well-being, i.e. the sense of satisfaction with one’s life and prospects. A central worry with this approach is that it could ignore significant health improvements that accrue to those who already enjoy high subjective well-being. Ronald Dworkin used the example of Tiny Tim to make a similar point. If the magnitudes of relevant health benefits are tied to improvements in subjective well-being, then an intervention that restores Tiny Tim’s mobility may bring very little health benefit, given Tim’s already exceedingly cheerful disposition. A similar problem concerns adaptive and even malformed preferences: intuitively significant improvements in health will be downplayed if they would go to individuals who are already subjectively satisfied with very little because of exposure to aspiration-numbing deprivation or injustice.

Preference-satisfaction

Economists often equate the satisfaction of preferences with well-being, and many summary measures of population health stem from health-state valuations derived from data about

respondents' preferences. "Satisfaction" can be a misleading term here, since what is relevant is getting what one wants, not the subjective feeling that may come from getting what one wants. (This also separates the economist's notion of "utility" from the classical utilitarian's.)

There is strong reason to keep individual well-being and preference-satisfaction separate, and to avoid tying the importance of health improvements to individual preferences. For example, one may prefer a policy in part for altruistic reasons, i.e. because of its impact on third parties. In such cases, satisfying one's preferences could actually come at a cost to one's own well-being. Second, it seems possible for individuals to want and prefer things that are not in fact good for them, that do not in fact promote their well-being. This can be due both to false empirical beliefs and to misguided prudential outlooks. Prudential preferences are hardly ever brute "gut" preferences, but rather preferences stemming from judgments about what is worth valuing and pursuing (Hausman 2006, 273). As T. M. Scanlon puts it, "My preferences are not the source of reasons but reflect conclusions based on reasons of other kinds" (Scanlon 2003, 177). This opens the possibility that individuals' preferences may be insensitive to objective reasons for thinking that a given health state is better or worse.

Opportunities and Capabilities

Many take examples like the one involving Tiny Tim to justify focusing on more objective consequences of deficits in health. Regardless of its impact on his subjective welfare, Tiny Tim's impairment reduces the opportunities that are available to him in significant ways. Amartya Sen has long advocated for a metric of policy evaluation that incorporates both subjective criteria and objective criteria. Sen's metric of capabilities conveys the importance of effective abilities to

function in valuable ways, even if a particular agent is not interested (at least not here and now) in using those abilities. This framework also partially divorces the public importance of capabilities from an agent's personal preferences about them. The main questions raised by the framework are, of course, which health-related capabilities are relevant to understanding efficiency in health, and how might they be measured in a scientifically respectable way.

3. The Morality of Maximization

Much philosophical attention to the issue of efficiency in health focuses on the status of the maximization imperative itself (regardless of metric). Everyone agrees that efficiency is an important consideration, but when and why should it give way to other moral considerations? Much of the work on this question takes place within the context of cost-effectiveness analyses (CEA) that use quality-adjusted life-years (QALYs) as the metric of evaluation. This sort of CEA is also known as Cost-Utility Analysis (CUA), since the QALY is a construct reflecting respondents' preferences about health-states. QALYs are designed to integrate longevity considerations with quality of life considerations in a way that enables comparisons between health interventions targeting very different dimensions of health. CEAs using QALYs are then used to identify policies that would maximize health-related value or utility within a given budget. There are four main philosophical problems with "pure" CEAs, i.e. CEAs that identify and then recommend policies based solely upon the criterion of efficiency.

Aggregation: CEA permits small benefits to lots of people to be summed up to outweigh large benefits to a smaller number of people. For example, a government-sponsored

commission in Oregon (U.S.) in 1990 released a draft priority list of health care services that prioritized tooth capping for exposed pulp over life-saving procedures like appendectomy and surgery for ectopic pregnancy. As David Hadorn wrote at the time, “This counterintuitive preference order did not occur as a result of faulty data...but as an inevitable consequence of the application of cost-effectiveness analysis” (Hadorn 1991, 2219). Dollar for scarce dollar, providing appendectomies was not as cost-effective as those non-life-saving services.

Discrimination against the disabled: Suppose that sub-population A is disabled while sub-population B is not; each sub-population is the same size and all individuals are otherwise equally healthy. Now suppose an epidemic afflicts both populations and leaves all individuals with a life-threatening illness. Assume also that logistical limitations allow for life-saving treatment to be administered to just one sub-population; all members of the treated sub-population will be restored to their pre-illness condition and if saved each would live the same number of additional years. Pure CEA recommends against choosing the disabled population, since this generates fewer QALYs. Many find this a troubling form of discrimination against persons with disabilities.

Priority to the worse off: Pure CEA cannot explain why one should give priority to the worse off when this intuitively seems required. Suppose the individuals in Group A generate 0.3 QALYs per year and could be brought to produce 0.5. And suppose that equally numerous individuals comprising Group B generate 0.8 QALYs per year and can be brought to full health (1.0). Once again suppose that scarcity or logistics require

choosing just one group to assist. Pure CEA recommends flipping a coin, since from the standpoint of the maximizer, adding 0.2 QALYs per year to a person's life has the same importance regardless of that person's initial condition. Many find this counterintuitive and believe that there is a presumption in favor of treating the worse off.

Fair chances vs. best outcomes: Suppose the members of two equally numerous groups, A and B, each currently generate 0.5 QALYs per year. Now suppose that either A can be helped or B can, but not both: members of A can be brought to generate 0.9 QALYs per year, or members of B can be brought to generate .93. Pure CEA straightaway favors assisting B, but many believe A and B should receive equal chances to be benefitted.

Each of these stylized scenarios implicates a distinctive equity consideration, but there is no consensus on how to incorporate such considerations into health-economic analysis. Consider, for example, the problem of aggregation. Employing different variations of Oregon's methodology, Peter Ubel and colleagues found that CEA can equate the successful treatment of 10 cases of appendicitis with the successful treatment of between 111 and 1,000 cases of mild hand pain. Yet when the respondents whose personal preferences were used to construct the CEAs were asked directly how many curings of mild hand pain would be equivalent to 10 curings of appendicitis, the median answer could not be computed because it was between a numerical answer and one denoting infinity (Ubel 2001, 55). This finding comports with a common response to Oregon's draft proposal: many believe that, morally speaking, no number of capped teeth could equal or outweigh saving a life with an appendectomy. But this raises a puzzle, since virtually no-one claims that it is always wrong to give priority to less serious but

more numerous needs over more serious but fewer needs. Suppose one could either restore the mobility of 10,000 people with paraplegia or one could save one person's life, but not both. It seems obvious that the relative numbers tip the ethical scales toward the 10,000. But note that there is no person among the 10,000 who, if not helped, could reasonably complain that he was left without mobility while someone else's life was saved. In that respect, this case parallels the case involving tooth capping and appendectomy: there is no person among the candidates for tooth capping who could reasonably complain that his tooth will be left uncapped if the legislature pays for appendectomies instead. But then if it can still be permissible to favor large numbers in the case involving paraplegia, why not also in the case involving tooth capping? The difficult question, therefore, is not whether aggregation can be morally permissible, but rather when and on what basis aggregation is permissible.

Partially in response to the problem of aggregation, health economists have explored ways to build respondents' direct rationing preferences into a modified CEA framework. Such preferences are elicited using so-called "personal trade-off" (PTO) exercises of the sort Ubel and colleagues used to uncover the discrepancy between pure CEA and respondents' direct rationing preferences. One notorious problem with the PTO methodology is the problem of multiplicative intransitivity. The problem is nicely described by Ubel (2000, 168-9):

Imagine a person who thinks that curing one person of condition A is equally beneficial as curing ten people of condition B, and that curing one person of condition B is equally beneficial as curing ten of condition C. To be consistent, this person ought to think that curing 1 person of condition A is equally beneficial as curing 100 people of condition C. However, when we conducted PTO measurements for three such conditions and

multiplied the PTO values of the two “nearer comparisons” (such as A vs B and B vs C), we calculated a different value for the relative importance of the “far comparisons” (such as programs A and C) than people told us when they were directly asked to compare these programs [i.e. A and C].”

Since no survey can ask respondents directly to compare every possible pair of competing health interventions, health economists seek a solution to the problem of multiplicative transitivity that could license inferences from discrete preferences about “nearer comparisons” (A vs B, B vs C, C vs D, ... , Y vs Z) to preferences about “far comparisons” (A vs Z). One problem not mentioned in the economics literature is that success in this endeavor would conflict with some of the convictions that raised the problem aggregation in the first place. Suppose that a very long chain of near comparisons begins by comparing an appendectomy that saves one person’s life with an intervention that cures some number of cases of paraplegia. Suppose the next comparison on the chain compares the curing of one case of paraplegia with the curing of some number of cases of one paralyzed arm. Now suppose the chain continues down the line until one gets to the near comparison between curing one case of mild tendonitis with curing some number of cases of individuals who suffer very mild headaches once per week. The worry now is that any solution to the problem of multiplicative transitivity would entail that there is some non-infinite number of cases of mild headaches that would be granted priority over curing a case of appendicitis. Many economists and all consequentialist philosophers will be willing to bite this bullet. Many non-consequentialist philosophers and perhaps many outside the academy will not.

4. The Concept of Health Equity

The most commonly cited definition of health equity is Margaret Whitehead's (1991, 219):

The term 'inequity' has a moral and ethical dimension. It refers to differences which are unnecessary and avoidable but, in addition, are also considered unfair and unjust.

This definition leaves open the possibility that some differences in health—health inequalities—are neither unfair nor unjust. This seems to be a virtue. It is not clear, however, that a health inequality must be avoidable before it can be counted an inequity. Here is what Whitehead says about this aspect of equity (1991, 219):

We will never be able to achieve a situation where everyone in the population has the same type and degree of illness and dies after exactly the same life span. This is not an achievable goal, nor even a desirable one. Thus, that portion of the health differential attributable to natural biological variation can be considered inevitable rather than inequitable.

There are two ideas at work here. First, there is the idea of the desirability of equality: everyone being the same in some respect or respects. But, second, Whitehead also refers to the impossibility of equality, and it is this that seems to motivate the condition that an inequity in health must be an avoidable inequality.

To see the problem with this condition, suppose a subset of the population is afflicted by a health impairment that cannot be avoided or resolved medically—perhaps an unalterable

genetic defect makes amputation below both knees a necessity for this group. Suppose also that the legislature is considering whether to pay for wheelchairs for those afflicted by the disorder. On Whitehead's definition, considerations of equity might say nothing about whether the state has an obligation to provide these assistive devices. This is because wheelchairs arguably cannot eliminate the differences *in health* caused by the disorder. Whitehead's definition therefore seems flawed, since it definitionally entails that the provision of assistive devices is not a demand of equity (Wilson forthcoming).

Since the concept of health equity should not prejudge substantive issues that a theory of health equity is intended to address, it is better to start from a modified version of Whitehead's definition. Thus: health inequities are health or health-related differences that are unjust, all things considered. The "all things considered" qualification means that if a difference is an inequity, then there exists a moral requirement on the part of (certain) agents or institutions to do something about them. It follows from this definition that some view of justice is required before a health difference can be counted a health inequity, but this new definition does not rule out the possibility that unavoidable health differences raise issues of equity.

5. Unfairness and Equality

Whitehead's "necessary and avoidable" condition is therefore problematic. Recall, however, that Whitehead's definition included another condition, viz. that an inequity is an inequality that is unfair. It might seem that this unfairness condition adds nothing to the definition, since whatever is unfair is unjust. But that is too quick. Whether unfair inequalities are also unjust depends on what unfairness is, how it is related to moral obligation, and whether other considerations can

outweigh or displace fairness in the final determination of what is, all things considered, just and unjust.

How does Whitehead's definition of health equity connect up with the moral value of equality? She says: "We will never be able to achieve a situation where everyone in the population has the same type and degree of illness and dies after exactly the same life span. This is not an achievable goal, nor even a desirable one" (Whitehead 1991, 219). Setting aside the question of achievability, why would equality not be a desirable goal? Imagine that medical progress has left us with just one disease—heart disease, say—that sets in at age 100 and leaves us dead at 105. Would this not be desirable? Surely it would. Imagine a slightly different scenario in which heart disease sets in at age 100 for both men and women, but men tend to die at 105 while women die at 110. If we then had to choose between giving males an extra five years of life expectancy and giving females an extra six years, wouldn't there be something to be said in favor of closing the gap rather than widening it with the more efficient female-oriented policy? And might not the value of equality explain why it is unfair that some have inferior life prospects?

These considerations might suggest that equality is indeed intrinsically desirable, so long as its place is known. Having human beings be equal in each and every respect would surely be undesirable, and this may be all Whitehead is saying. But this does not entail that there is no respect in which it would be desirable to promote greater equality of health prospects. In some contexts, equality may be very important, and in others it may simply be less important than some other moral consideration (e.g. efficiency, individual choice, or individual responsibility).

6. Equality of Outcomes versus Process Equity

This last point is sometimes invoked in the context of sex differences in longevity. In 1994, the WHO's Global Burden of Disease team used high-income populations in low-mortality countries to peg the biologically-determined sex-based inequality in longevity at 2-3 years. It might therefore be suggested that if one is committed to equity in health, health care systems should tilt in favor of treating men, as a way to achieve equality of health. However, Amartya Sen and Angus Deaton distinguish between equality of outcomes and process equity (Sen 2002, 660-661; Deaton 2002, 24). Process equity is the idea that procedural fairness, in health care access and delivery (e.g.), is of independent importance. In Sen's and Deaton's view, process equity can sometimes be more important than equality of outcomes. This line of argument would enable one to give some value to equality of health outcomes while not letting it dictate health policies that seem intuitively unjust for other reasons.

There is, however, a response that can be made by someone skeptical of process equity. Indeed, it is a response that Deaton himself has made. He first concedes that the inequality in life expectancy between men and woman *may* justify tilting medical research toward understanding the factors that disproportionately affect men (Deaton n.d.). This is the sort of bias that seems defensible in cases where diseases disproportionately afflict racial minorities. It is therefore not clear that it should be ruled out in the context of sex-differences in longevity. But if a bias in state-funded research and development can be justified, why not a bias in health care delivery?

Here Deaton provides an answer that invokes the importance of equality of outcomes rather than process equity. He points out that longevity is not the only component of health, and that health-related quality of life matters too. He notes that although women have lower prevalence of conditions with high mortality, they have a higher prevalence of conditions with

high morbidity (Deaton n.d.). Thus, in some contexts, providing equality of access to health care could actually be one way of equalizing overall health between men and women, since the advantage in life expectancy might offset the morbidity disadvantages faced by women. Indeed, there are surely many health and non-health disadvantages faced by women that a few extra years might help offset. Note again that this argument begins from a concern about equality of outcome, and this may lead one to rethink the relative importance that Sen and Deaton originally gave to process equality over equality of health outcomes. That is, perhaps process equity seems to conflict with equality of outcomes only when one is focused on the wrong outcome. If one instead focuses on overall health (or well-being, or capability sets, or...) perhaps one would have no reason to discriminate against women at the point of care.

It was noted above that even if an inequality in health is not avoidable, equity may demand the provision of non-medical services such as assistive devices. The lesson was that whether an unavoidable inequality *in health* is equitable or not may depend upon the impact it makes on something of more fundamental importance. Likewise, a disadvantage in morbidity might be offset by the advantage of lower mortality to yield overall equality of health or well-being. In each case, a health inequality that is neither avoidable nor *remediable* may still be *compensable*, and if it can be compensated for by some non-health good, this may be required in order to equalize some third value that includes health and non-health goods as components (Hausman 2007, 50-51).

If this is correct, then Culyer and Wagstaff (1993) were mistaken to suggest that conclusions about just health distributions follow immediately from the premise that justice demands equality in well-being. They claim that if one starts with the assumption that justice requires equality of well-being (what they call “equality of flourishing”), then “Insofar as health


is a necessary condition for ‘flourishing’, it follows that a just distribution of health is an equal distribution” (452). This would not follow in cases where all individuals are healthy enough to enjoy other determinants of well-being, and where these other determinants can be allocated so as to offset the impact of residual inequalities in health.

7. Questioning the Value of Equality

So far no reason has been identified to reject a form of egalitarianism that is prominent in the philosophical literature and that nicely explains the connection between Whitehead’s reference to fairness and the close linguistic relation between equity and equality. The egalitarian philosopher Larry Temkin puts it thus (2003, 775):

The essence of the egalitarian’s view is that comparative unfairness is bad, and that if we could do something about life’s unfairness, we have some reason to. Such reasons may be outweighed by other reasons, but they are not...entirely without force.

Temkin maintains that unfairness exists when some are worse off than others through no fault of their own. Temkin identifies two objections that might be used to rebut his view that undeserved inequality is intrinsically bad. These are the so-called Raising Up and Leveling Down Objections.



Consider first the choice between scenarios A and B. There are two social groups in each of A and B. The width of the bars is related to the size of the group's population, and the height is related to the how well-off each individual within a group is. Height may here capture years of life lived, quality-adjusted years of life enjoyed, life expectancy, etc. Taking A as the status quo, one is asked to consider whether an otherwise benign policy should be implemented that would lead to scenario B. The Raising Up objection to a Temkin-style egalitarianism simply points out that, insofar as one is an egalitarian, one must condemn the move from A to B. The anti-egalitarian who makes this objection emphasizes that the move to B is a Pareto improvement: it harms no one and at least some are made better off. How, she will ask, could there be any reason not to improve the lives of some if there is no one who is made worse off by it? (We are assuming that the improvement is welcomed and not forced on anyone who does not want it.)

Temkin's response to this objection underscores a point made above, namely that if equality has value, it does so only in the context of other important values. To use an example of Joseph Raz's, it is not important that everyone be equal with respect to the number of hairs on their shirts. That sort of egalitarianism is precisely the sort that Whitehead would be right to call undesirable. So where it makes sense to talk about the value of eliminating undeserved differences, there will always be other genuine values that are also relevant. But if equality is not

the only value, it is possible that equality can be outweighed by the other values whose presence makes equality relevant. This is Temkin's response. He agrees that a move from A to B is the right choice once all values and reasons are given their due. He simply notes that one consideration, equality, counts against the move. But he then concedes that this is a scenario in which the value of equality is outweighed by the other values that would be promoted by moving from A to B.

To some, this is a fine response in defense of egalitarianism. True, it may seem strange to deflate equality's relative importance this much, but that seems necessary if one is attracted to this brand of egalitarianism.

The second objection to Temkin-style egalitarianism seems much more damaging. Imagine that scenario C is the status quo and one is deciding whether to support a move to scenario D (which would bring everyone in C down to the level of C's worst off group). Plainly, D is superior with respect to equality. But there is also no one for whom D would be better than C. And yet the Temkin-style egalitarian is forced to say that there is something to be said in favor of moving from C to D. Here again is how Temkin navigates these waters (2003, 780):

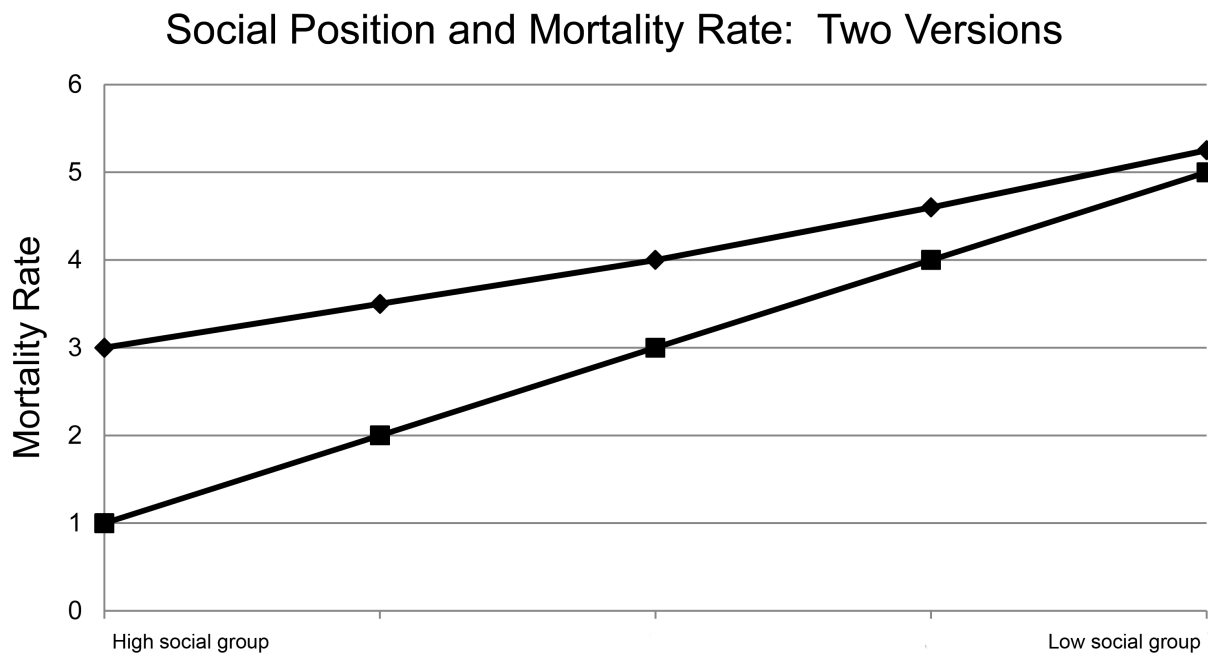
But, the anti-egalitarian will incredulously ask, do I *really* think that there is some respect in which only some being blind is worse than all being blind? Yes. Does this mean that I think it would be better if everyone else was blind? No. As noted previously, equality is not all that matters. But it does matter *some*.

Again, this rebuttal is clearly open to Temkin. But here the anti-egalitarian's reply seems even stronger. She will highlight how bizarre it is to say there could be any reason to move from

C to D, especially since no one is benefited and many people are significantly harmed. That is the Leveling Down Objection.

8. From Equality to Priority

The Raising Up and Leveling Down objections lead many to give up entirely their belief in the intrinsic value of equality. But others remain steadfast. Consider the following diagram, which replicates a diagram first drawn by Michael Marmot and discussed in his book *The Status Syndrome* (Marmot 2004, 246).



The diagram graphs the mortality effects of two different policies on four social groups arranged from left to right in descending order of social advantage. Call the two policies Square (indicated by the bottom graph) and Diamond (indicated by top graph). Square would lead to greater health

inequalities than Diamond. But Square offers Pareto improvements with respect to life expectancy, since everyone in Square has lower mortality than anyone in Diamond. Marmot drew the graph during a conversation with Deaton. Deaton wanted to know if Marmot cared more about reducing inequalities than he did about reducing sickness and death. Marmot writes:

I demurred. [Deaton] was in no doubt that all economists would choose the bottom graph because everyone is better off...[He] suspected that I went for the one with less inequality where everyone suffered more...It is my view that we should reject both alternatives and aim for a society where health for everyone has improved and inequality is less. (Marmot 2004, 245-6).

The economists Deaton refers to will likely be motivated by a commitment to Pareto improvements. In contrast, many philosophers who agree with Deaton's choice of graph will be driven by a belief in prioritarianism. There are a number of versions of prioritarianism, but its general thrust is that, morally speaking, benefiting people matters more the worse off they are. While prioritarians will agree that there is often reason to promote greater equality, they do not think equality is intrinsically important. Rather, a system that tilts in favor of helping the worse off will often end up more equal merely as a side-effect of the prioritarian focus on the more disadvantaged. But if improving the lot of worse off should require increases in inequality, prioritarians (like many economists) will not care.

Once prioritarianism is introduced, an intrinsic concern with equality can seem like an aesthetic preference rather than a moral conviction. Where the egalitarian claims that things have gotten more unfair even though everyone is doing much better and even if the worst off are as

well off as possible, the prioritarian demands to know who (other than the egalitarian!) is complaining about unfairness. It cannot be the best off, since they are doing better than anyone and so have no right to complain. And it is unlikely to be the worst-off, since they surely would not demand to be worse off than they already are.

9. The Value of Equality Revisited

Without concluding that Temkin-style egalitarianism is false, consider further the alternative of well-being prioritarianism, i.e. the view that social policy should tilt in favor of promoting the well-being of those worse off (at least to the extent consistent with respecting individual choice and personal responsibility). Such a view sees nothing intrinsically valuable in distributive equality.

Consider now an objection to well-being prioritarianism. When one looks outside the narrow sphere of personal well-being, one encounters other spheres of life within which equality seems to have intrinsic importance. Consider the spheres of political liberty and social mobility. Many believe there is a presumption in favor of equality of access to political influence and equality of opportunity (whereby no child is systematically disadvantaged in their life prospects due to their parent's socioeconomic status). If it seems correct to stress the intrinsic importance of equality in these non-welfare domains, should this be interpreted as support by analogy for the intrinsic importance of equality within the narrower domain of personal well-being?

The first thing to note is that the spheres of political influence and social mobility are zero-sum. So even if one is a consistent prioritarian across the three domains of personal well-being, political influence, and social mobility, equality will be the only distribution available for

the last two domains. It is simply impossible to boost one social group's share of political influence or social mobility without making another worse off. This does not of course prove that equality in these realms has no intrinsic value. But it might explain why one would remain attracted to distributive equality in some spheres even if one's most fundamental ethical framework was prioritarian.

Further, the recurring theme of trade-offs between different types of goods is relevant here as well. If inequalities in life prospects led to unequal political influence, unequal social mobility, *and* optimal improvements in the absolute well-being of those in lower- and middle-income groups, then the trade-off might be worth it on prioritarian grounds. This tolerance of inequality is usually embraced more readily in the context of unequal social mobility stemming from efficiency-enhancing inequalities, but it might be warranted with respect to political influence as well.

Of these two considerations—(1) that prioritarian inequalities are not possible in the local spheres of political liberty and social mobility; and (2) that there may be high-order prioritarian reasons to tolerate inequalities in more local spheres—neither proves conclusively that equality in these spheres is without intrinsic importance. Indeed, there is one more way of conceiving of equality and its importance that differs from Temkin's approach and that raises the possibility that egalitarian and prioritarian concerns are both valid and in fact closely related.

10. The Possibility of an Egalitarian Prioritarianism

It was suggested near the end of section 8 (above) that once prioritarianism is introduced, a commitment to distributive equality can begin to resemble an aesthetic preference for uniformity

rather than a commitment to the real needs of individuals. However, many who hold egalitarian views about equal political influence and equal social mobility are not primarily motivated by a general desire to eliminate undeserved disadvantages between individuals. Rather, they are often moved by the independent values of non-domination, reciprocity, and equal social status.

According to an increasing number of philosophers, these are the values that should ground egalitarian political convictions, as they are specially relevant for societies that care about treating all persons as moral equals. To say that each person is the moral equal of all is not yet to say that goods should be distributed in a particular way. So the ideal of moral equality is not, at bottom, a distributive ideal, although many claim that it has implications for distribution. For example, distributive implications may flow from considerations about the demands of reciprocity and benevolent concern that are warranted when moral equals occupy distinctive social and political relationships.

Some philosophers suggest that when prioritarian distributions are demanded by justice, this demand is ultimately grounded these non-distributional premises about moral equality and ethically mandated concern (Miller 2010). A famous example of Thomas Nagel's provides a useful illustration. In an essay that in many ways sparked the contemporary philosophical debate between distributive egalitarians and prioritarians, Nagel describes a fictional scenario in which he has one healthy child and one suffering from a painful disability. He imagines that he must make a choice between moving to a city where the second child could receive medical treatment but which would be unpleasant for the first child, or moving to a semi-rural suburb where the first child alone would benefit. He stipulates that "the gain to the first child of moving to the suburb is substantially greater than the gain to the second child of moving to the city." Nagel then claims that, "If one chose to move to the city, it would be an egalitarian decision. It is more

urgent to benefit the second child, even though the benefit we can give him is less than the benefit we can give the first child” (Nagel 1991, 124).

In response, Derek Parfit (1997, 215) writes:

Since this distinction [between equality and priority] has been overlooked, some writers have made claims that are not really about equality, and would be better stated as claims about priority...[Nagel’s] remarks suggest that, to the question ‘Why is it more urgent to benefit [the second] child?’, Nagel would answer, ‘Because this would reduce the inequality between these two children’. But I doubt that this is really Nagel’s view.

Would it be just as urgent to benefit the handicapped child, even if he had no sibling who was better off? I suspect that, on Nagel’s view, it would.

In Parfit’s view, Nagel has misdescribed Nagel’s own moral commitments. Nagel says that the duty to attend to the disabled child’s needs is an egalitarian duty. Parfit insists that Nagel is not concerned with distributive equality between the two children at all, and that Nagel instead appears motivated by prioritarian concern for the worse off child.

One might respond on Nagel’s behalf by claiming that Parfit works with a false dichotomy. In insisting that Nagel must be a prioritarian, Parfit ignores the brand of egalitarianism that stresses the moral demands of distinctive interpersonal relationships, including relationships that call for the display of equal and robust concern for those to whom one is specially related. When multiple individuals compete for that concern (as they may when they are our fellow citizens or our children) it is reasonable to conclude that treating all of them as equals requires prioritarianism in practice. Although the parent-child relationship has a

different moral basis than the co-citizen relationship, Nagel's illustrative example at least suggests that one can display equal and robust concern for members of a group while also putting the needs of the worse off at the front of the line. This, Nagel can argue, is why an egalitarianism of the right sort is consistent with the prioritarian concern to address the health needs of the disabled child. Properly treating a group of persons as equals, and equitably showing them morally required concern, may in practice demand a prioritarian response to their separate needs.

This last brand of egalitarianism may hold great promise to unify and explain many intuitions about the demands of equity across multiple policy domains, including the domain of health policy. If, for example, it can be shown that compatriots or indeed "global citizens" owe robust duties of equal concern for one another, then understanding the source, scope, and degree of that concern should aid in the just allocation of scarce time, energy, and resources.

ACKNOWLEDGEMENTS

I thank Daniel Hausman and the Encyclopedia's editors for extremely helpful comments on a prior draft; and David Wilcher for research assistance. I am grateful to Sir Michael Marmot for permission to reproduce a variation of a diagram presented on p. 246 of his book *The Status Syndrome* (2004). I am also grateful to Oxford University Press for permission to reproduce a variation of a diagram adapted from p. 247 of Larry Temkin's book *Inequality* (1993).

REFERENCES and FURTHER READING

- Culyer, A.J., Wagstaff, A. (1993). Equity and equality in health and health care. *Journal of Health Economics* 12, 431-57.
- Deaton, A. (2002). Policy implications of the gradient of health and wealth. *Health Affairs* 21, 13-30.

- Deaton, A. (n.d.). What does the empirical evidence tell us about the injustice of health inequalities? (Unpublished manuscript.)
- Hadorn, D.C. (1991). Setting health care priorities in Oregon: cost-effectiveness meets the rule of rescue. *JAMA* 265, 2218-25.
- Hausman, D. (2006). Valuing health. *Philosophy & Public Affairs* 34, 246-274.
- Hausman, D. (2007). What's wrong with health inequalities? *Journal of Political Philosophy* 15, 46-66.
- Marmot, M. (2004). *The Status Syndrome: How Social Standing Affects Our Health and Longevity*. New York: Henry Holt.
- Miller, R. W. (2010). *Globalizing Justice: The Ethics of Poverty and Power*. Oxford: Oxford University Press.
- Murray, C.J.L., Salomon, J.A., Mathers, C.D., Lopez, A.D. (2002). Summary measures of population health: conclusions and recommendations. In Murray, C.J.L., Salomon, J.A., Mathers, C.D., Lopez, A.D. (eds.). *Summary Measures of Population Health*, pp. 731-756. Geneva: World Health Organization.
- Nagel, T. (1991). Equality. In Nagel, T. (1991). *Mortal Questions*, 106-127. Cambridge: Cambridge University Press.
- Parfit, D. (1997). Equality and priority. *Ratio* 10, 202.
- Scanlon, T.M. (2003). Value, desire, and the quality of life. In Scanlon T.M. (2003). *The Difficulty of Tolerance: Essays in Political Philosophy*, pp. 169-186. Cambridge: Cambridge University Press.
- Sen, A. (2002). Why health equity? *Health Economics* 11, 659-666.
- Temkin, L. (2003). Egalitarianism defended. *Ethics* 113, 764-782.
- Ubel, P. (2000). *Pricing Life*. Cambridge, MA: The MIT Press.
- Whitehead, M. (1991). The concepts and principles of equity and health. *Health Promotion International* 6, 217-228.
- Wilson, J. (Forthcoming). Health inequities. In Dawson, A. (ed.). (Forthcoming). *Public Health Ethics: Key Concepts in Policy and Practice*. Cambridge: Cambridge University Press.