

Meaning, medicine and merit

Andreas Mogensen

Global Priorities Institute | September 2019

GPI Working Paper No. 3-2019



Meaning, medicine, and merit

Andreas L. Mogensen

Global Priorities Institute / Faculty of Philosophy, University of Oxford

Given the inevitability of scarcity, should public institutions ration healthcare resources so as to prioritize those who contribute more to society? Intuitively, we may feel that this would be somehow inegalitarian. I argue that the egalitarian objection to prioritizing treatment on the basis of patients' usefulness to others is best thought of as semiotic: i.e. as having to do with what this practice would mean, convey, or express about a person's standing. I explore the implications of this conclusion when taken in conjunction with the observation that semiotic objections are generally flimsy, failing to identify anything wrong with a practice as such and having limited capacity to generalize beyond particular contexts.

1.

Given the inevitability of scarcity, how should public institutions ration healthcare resources? Who should receive priority of treatment when waiting-lists are drawn up? Which treatments should we subsidize, and for which ailments?

Two criteria are unavoidable: *need* and *cost*. Those whose medical needs are more urgent should be treated ahead of others, all else being equal. Cost is also an important factor, not because we value thrift for its own sake, but because sensitivity to cost ensures that we can help more people as opposed to fewer. Setting aside the use of tie-breaking criteria such as queue order¹ or a fair lottery,² what else, if anything, should we take into account when rationing healthcare?

One possibility would be to take into account the social utility of the potential beneficiaries: i.e. their instrumental value to others. When all else is equal, is there a case for prioritizing those who will 'contribute more to society,' so to speak? These contributions might take many forms, whether

¹ James Childress 'Putting patients first in organ allocation: an ethical analysis of the US debate', *Cambridge Quarterly of Health Ethics* 10 (2001), pp. 365-76.

² John Broome, 'Fairness', *Proceedings of the Aristotelian Society* 91 (1990), pp. 87-101.

intended or unintended. A gifted surgeon may be uniquely skilled in performing difficult, life-saving operations. Higher earners may contribute more to the state's tax revenues. Entrepreneurs may generate social benefits in the form of consumer surplus.³ Some people care for others who will suffer greatly without them. Some people are simply more altruistic or public-spirited, contributing more to the general good through volunteering, charitable giving, or activism. Giving such individuals priority could allow them to return to health sooner rather than later, ensuring that less of their valuable social contributions are lost. Obviously, no one would propose making this the unique deciding factor in allocating scarce medical resources. But might it be one factor to be weighed alongside whatever other criteria should be taken into account, capable of tipping the balance when all else is equal?

Surveys of public attitudes indicate widespread resistance to this idea. Skitka and Tetlock found that subjects on average rated prioritizing people on the basis of their contributions to the community as inappropriate, scoring this at -2.18 on a scale ranging from -8 (extremely inappropriate) to +8 (extremely appropriate).⁴ McKie and Richardson found that 87% of respondents strongly disagreed with giving people who contribute more to society higher priority for medical care via public health services.⁵

What could be a defensible objection to rationing healthcare in this way? My aim in this paper is to help answer this question. I'll argue that, insofar as there is a principled objection to prioritising people on the basis of their social utility, that objection is *semiotic*: i.e. it has to do with what this practice would mean or express. The objection derives in part from the significance we invest in healthcare provision in countries with comprehensive public health care systems, such as the United Kingdom, whose National Health Service (NHS) I treat as focal in my discussion. (I focus on the NHS partly because it is closest to home for me, but also because it is the archetype: the first

³ Thanks to an anonymous referee for suggesting this example.

⁴ Linda Skitka and Philip Tetlock, 'Allocating scarce resources: a contingency model of distributive justice', *Journal of Experimental Social Psychology* 28 (1992), pp. 491-522.

⁵ John McKie and Jeff Richardson, 'Social preferences for the inclusion of indirect benefits in the evaluation of publicly funded health services: results from an Australian survey', *Health Economics, Policy and Law* 6 (2011), pp. 449-68.

public healthcare system in the Western world providing comprehensive coverage to everyone, free at the point of use.) As well as the social significance invested in healthcare, the semiotic objection that I discuss depends on the salience of meritocratic ideals in post-industrial societies, including their salient incompatibility with the ideal of communal caring underlying public health care provision in systems like the NHS.

I think this semiotic objection is the most plausible interpretation of the widely shared intuition that prioritizing treatment on the basis of patients' usefulness to others is objectionable from the point of view of equality. Because semiotic properties are contingent and socially constructed, I also consider what my argument implies for how we should think about the ethics of prioritization considered more generally. To that end, I explore the implications of my argument for evaluating rationing decisions concerning life and health in the sphere of private philanthropy, where donors may wish to give preference to beneficiaries with greater instrumental value to others.

2.

Many people find the idea that publicly funded healthcare institutions might prioritise among patients on the basis of their social utility intuitively objectionable. It is natural to feel that doing so would be somehow inegalitarian. Public healthcare systems are typically founded on a strong commitment to equality. For example, the first principle of the NHS Constitution for England emphasizes the public health authority's 'social duty to promote equality through the services it provides'⁶. Prioritising among people on the basis of their social utility seems to offend against that ideal. However, it's surprisingly difficult to substantiate this objection. In this section, I'll run through various moral and political theories that may be understood as egalitarian in some sense or other, beginning with utilitarianism. No theory seems to confirm our intuitions.

⁶ NHS Constitution for England

<https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/480482/NHS_Constitution_WEB.pdf> (2015), p.3.

2.1

Utilitarianism is egalitarian insofar as it counts each for one and none for more than one: equivalent benefits and harms count equally no matter whose welfare is at issue. But utilitarianism provides a straightforward case for thinking that we *should* take indirect social benefits into account. It represents a challenge to our intuitions, not a potential justification.

By way of illustration, suppose that treating Arkady or Baruq would benefit each to the same extent. Treating Arkady would also generate indirect benefits for Cathy and Desmond. Treating Baruq would provide no additional benefit to anyone. There are no further benefits or harms arising from the available actions. Utilitarianism implies straightforwardly that we should prioritize Arkady.

Furthermore, if we value each individual equally, how can we justify doing otherwise? If we instead toss a coin between Arkady and Baruq, aren't we treating Cathy and Desmond as if their well-being had no moral importance?

2.2

Utilitarianism is not egalitarian enough for many, since it attaches no inherent moral significance to the distribution of goods across people, including the distribution of well-being. We may believe that equality in some metric of overall benefit is of intrinsic moral value. Furthermore, it's natural to worry that prioritizing individuals on the basis of their social usefulness would exacerbate existing inequalities. Is this the heart of the problem?

I don't believe so. Return to the example in section 2.1. Suppose that Arkady is better off than Baruq, and hence treating Arkady would further increase the inequality between them. Assume Arkady starts at 4, and Baruq at 2, as measured by whatever we consider to be the currency of equality. The treatment confers a benefit of 2. If given to Baruq, this would close the inequality between them. If given to Arkady, this would make the inequality even worse. Even so, it needn't follow that treating Arkady is worse overall from the standpoint of equality. That's because helping Arkady indirectly benefits Cathy and Desmond. Suppose that Cathy and Desmond are worse off than Baruq and start at 0, whereas helping Arkady will raise them up to 1.

	Arkady	Baruq	Cathy	Desmond
<i>Treat Arkady</i>	6	2	1	1
<i>Treat Baruq</i>	4	4	0	0

In this case, the resulting distribution is actually more equal, as measured by the Gini co-efficient, if we choose *Treat Arkady*.

Returning to the real world, a policy that gives priority to high earners over median earners might do so in order to use the added tax revenue to benefit the worse-off. As illustrated by our toy example, this could serve to reduce overall inequality according to some reasonable measure, such as the Gini co-efficient.

Moreover, priority for high earners is the clearest case when it comes to the worry that prioritizing patients on the basis of social usefulness would exacerbate inequality. It's not clear that this worry has much plausibility when we consider giving priority to those who contribute more to society through volunteering or activism, for example.

2.3

My discussion in section 2.2 did not attach significance to the sources of inequalities. This might seem like a mistake. Roughly speaking, *luck egalitarianism* says that we ought to be concerned about inequalities between people that are attributable to *brute luck*, but not about inequalities between people that are attributable to their own choices.⁷ This may seem to capture an important part of what concerns us. It may seem that what is objectionable about prioritizing top surgeons or others with

⁷ Richard Arneson, 'Equality and equal opportunity for welfare', *Philosophical Studies* 56 (1989), pp. 77-93, 'Luck egalitarianism interpreted and defended', *Philosophical Topics* 32 (2004), pp. 1-20; G. A. Cohen, *Rescuing justice and equality* (Cambridge, MA, 2008); Ronald Dworkin, *Sovereign virtue: the theory and practice of equality* (Cambridge, MA, 2002); Eric Rakowski, *Equal justice* (Oxford, 1993); Shlomi Segall, *Why inequality matters: luck egalitarianism, its meaning and value* (Cambridge, 2016); Kok-Chor Tan, 'A defense of luck egalitarianism', *Journal of Philosophy* 105 (2008), pp. 665-90.

similar talents is that opportunities to develop such talents are not fairly distributed and depend largely on accidents of birth.

However, the arguments in section 2.2 can be repurposed here. In the choice between helping Arkady or Baruq, we can stipulate that any inequalities here are due to brute luck. We can then argue in the same way that choosing to help Arkady because doing so will indirectly benefit Cathy and Desmond is recommended from the perspective of maximizing equality of fortune, even if this increases the inequality between Arkady and Baruq in light of circumstances beyond either's control.

In response, one might insist that this misconstrues luck egalitarianism as a form of what Parfit calls *telic egalitarianism*, a view on which equality is a good to be promoted.⁸ We may think that luck egalitarianism is better understood as a form of what Parfit calls *deontic egalitarianism*: in particular, as a form of deontic egalitarianism that posits a (defeasible) constraint against distributing goods in ways that advantage and disadvantage people in light of attributes that depend on brute luck. We may insist that when choosing between Arkady and Baruq, this constraint may not be violated for the sake of reducing the overall influence of brute luck on people's circumstances, just as the constraint on killing may not be violated just in case doing so would reduce the number of killings.⁹

A different problem confronts this view. Suppose that in the choice between helping Arkady or helping Baruq, we flip a coin. Assuming that they have equivalent medical needs, flipping a coin in this way arguably will not strike many of us as inequalitarian. Nonetheless, it involves a deliberate decision to allow the effects of brute luck to determine the outcome. Whoever insists on an egalitarian constraint against distributing healthcare in accordance with brute luck therefore needs to explain why flipping a coin in this scenario is not inequalitarian.

It could be said that the use of a coin-flip as a tie-breaker is generally permissible because when the morally significant reasons for preferring either potential beneficiary in the distribution of an indivisible good are otherwise on par, it is impossible to avoid allocating that good on the basis of

⁸ Derek Parfit, *Equality or priority? The Lindley Lecture* (Lawrence, KS, 1991).

⁹ Robert Nozick, *Anarchy, State, and Utopia* (Oxford, 1974) at pp. 30-3; F. M. Kamm, 'Harming some to save others', *Philosophical Studies* 57 (1989), pp. 227-60.

morally arbitrary criteria, and so there can be no objection to the use of a lottery just in virtue of the fact that the result is morally arbitrary. But this begs the question: namely, whether the morally significant reasons for preferring either potential beneficiary in this case should be considered on par. The proponent of prioritizing among individuals on the basis of their social utility insists that the balance of reasons favours treating Arkady because this helps Cathy and Desmond. Obviously, we may insist that Cathy and Desmond have no claim in the matter. But the problem is precisely how to justify saying that, given that they have a morally significant stake in who is treated.

2.4

According to Walzer, distributive justice does not obtain just in case the overall distribution of goods within society satisfies some general distributive condition.¹⁰ Instead, we are to think of there as being many different goods, each with its own proper distributive criterion. The criterion proper to any particular good is determined by the socially constructed meaning of that good. For example, it's by virtue of the shared meaning that we attach to the good of political office that political office cannot be bought and sold. 'Complex equality' obtains when the different goods are distributed within their sphere according to their own proprietary distributive criteria. A sufficient condition for injustice obtains if the distribution of one good determines the distribution of another when these goods belong to separate spheres governed by incommensurable distributive norms.

Viewed against the backdrop of this theory, it may seem natural to object that prioritizing treatment on the basis of patients' social utility violates the condition of 'complex equality'. It subordinates the distribution of healthcare to broader social goals, whereas the inner logic of medicine requires provision in accordance with need. With respect to socially recognized needs such as health, Walzer writes: 'The actual distribution will be limited by the available resources; but all other criteria, beyond need itself, are experienced as distortions and not as limitations of the distributive process.'¹¹

¹⁰ Michael Walzer, *Spheres of justice: a defence of pluralism and equality* (Oxford, 1983).

¹¹ *Spheres of Justice*, p. 75.

I think there is an element of truth here, but I also see two significant problems with this approach.

Firstly, it seems permissible to allow other criteria besides need to determine how we distribute medical resources, because decisions will have to be made when needs are roughly equal. We may use a fair lottery to allocate goods in such cases, and this does not strike us as inequalitarian. Nonetheless, this allows the allocation of healthcare resources to be determined by something other than need. If this constitutes a violation of ‘complex equality,’ then so much the worse for Walzer’s theory. If this is in keeping with ‘complex equality,’ why shouldn’t the same be true of prioritizing treatment on the basis of patients’ social utility, if only for the sake of breaking ties?

It may be replied that although a fair lottery cannot literally divide an indivisible good between the potential beneficiaries, the provision of an equal chance is uniquely fair by virtue of dividing the good probabilistically. Giving each person a chance of attaining the good provides them with a form of ‘surrogate satisfaction,’ and a fair lottery ensures that each claim is satisfied in this way to the very same extent.¹² Therefore, the use of a fair lottery is uniquely egalitarian.

However, it is not clear in what sense having a chance of obtaining some good yields ‘surrogate satisfaction’ of your interest and/or claim in that thing.¹³ Suppose a coin had been tossed to determine whether I would be given a fair trial and I lost the toss. My interest and/or claim in being fairly tried was presumably not satisfied in some respect or to some extent. There seems to be no morally interesting sense in which the chance of a fair trial gave me partial satisfaction of my right and/or interest in procedural justice.

Furthermore, we should note that a proponent of rationing on the basis of patients’ social utility is also able to point to some aspect of their favoured procedure by virtue of which *it* is uniquely well-suited to the realization of some important aspect of egalitarian fairness. As noted in our

¹² Dan Brock, ‘Separate spheres and indirect benefits’, *Cost Effectiveness and Resource Allocation* 1 (2003); Broome, ‘Fairness’; Lewis Kornhauser and Lawrence Sager, ‘Just lotteries,’ *Social Science Information* 27 (1988), pp. 483-516.

¹³ Tim Henning, ‘From choice to chance? Saving people, fairness, and lotteries’, *The Philosophical Review* 124 (2015), pp. 169-206.

discussion of utilitarianism, it is *prima facie* mysterious how a refusal to take indirect social benefits into account does not amount to a failure to give equal weight to the interests of all individuals impacted by the decision. Even if a fair lottery is uniquely well-suited to achieving some aspect of egalitarian fairness, this does not imply that breaking ties by appeal to social utility is not.

Here is a second problem for attempting to explain the inegalitarian character of rationing healthcare on the basis of individuals' social utility via the ideal of 'complex equality.' Note that we can restrict our understanding of what it means to prioritize by the social utility of patients so that only instrumental value *with respect to public health* is a criterion of priority. A policy that prioritizes top surgeons is presumably adopted for the sake of the health benefits that these surgeons achieve through the exercise of their skills. We could prioritize high earners precisely for the sake of using the extra tax revenue to improve health services for the worst-off. In these cases, we could not be accused of subordinating the distribution of healthcare to broader social goals.

2.5

The final egalitarian theory that I'll consider as a potential justification for our intuitive unease with prioritizing treatment on the basis of patients' usefulness to others is *relational egalitarianism* of the kind endorsed by Anderson and Scheffler.¹⁴ On this view, we are not to think of equality as fundamentally concerned with the distribution of goods, but with the existence of certain kinds of social relationships between citizens. These are relationships wherein 'each accepts the obligation to justify their actions by principles acceptable to the other, and in which they take mutual consultation, reciprocation, and recognition for granted.'¹⁵

¹⁴ Elizabeth Anderson, 'What is the point of equality?', *Ethics* 109 (1999), pp. 287-337, 'The fundamental disagreement between luck egalitarians and relational egalitarians', *Canadian Journal of Philosophy* 40 (2010), pp. 1 - 23. Samuel Scheffler, Samuel, 'What is egalitarianism?', *Philosophy and Public Affairs* 31 (2003), pp. 5-39, 'The practice of equality', *Social equality: on what it means to be equals*, eds. Fourie, Schuppert, and Wallimann-Helmer (Oxford, 2015), pp. 21-44

¹⁵ Anderson, 'What is the point of equality?', p. 313.

Viewed against the background of this theory, we may object that what is really wrong with prioritizing treatment on the basis of patients' social utility is that we would thereby treat those without the relevant desirable personal qualities as second-class citizens. This violates the demands of relational egalitarianism, by undermining their equal standing within the political community.¹⁶

I think there is an element of truth in this suggestion too, but it also seems subject to a significant objection. The objection I have in mind is a particular instance of a perfectly general problem facing relational egalitarians, itself descended from the *redundancy objection* to contractualism.¹⁷

Here is the problem. If some practice involves a failure of individuals to relate to one another as equals, it seems natural to expect that this is because there is something more fundamentally wrong with that practice. Suppose that we work together running a stall at a market, without having properly formalized our business relationship. At the end of the day, you pocket all the profits. I may object that you are not treating me as an equal. Intuitively, this complaint is derivative. You fail to treat me as an equal because you take more than your fair share of the profits. What's more, the fact that you take more than your fair share of the profits is of itself sufficient to explain why you treat me unjustly. The fact that you do not treat me as an equal is not what makes your action unjust, though it may be what injustice consists in. Similarly, that some action cannot be justified to others on grounds they could not reasonably reject is understood by Scanlon as redundant in respect of explaining what makes an act wrong, as opposed to what sort of property wrongness is.¹⁸

In sum, it may well be that prioritizing treatment on the basis of patients' social utility involves a failure of individuals to relate to one another as equals, but we seem to need some deeper explanation of what is wrong with the practice in order for that to be so.

¹⁶ Compare Kasper Lippert-Rasmussen and Sigurd Lauridsen, 'Justice and the allocation of healthcare resources: should indirect, non-health effects count?', *Medicine, Health Care and Philosophy* 13 (2010), pp. 237-46.

¹⁷ Simon Blackburn, 'Am I right?', *The New York Times* 21 Feb. 1999; Philip Pettit, 'Doing unto others', *Times Literary Supplement* 25 Jun. 1999.

¹⁸ T. M. Scanlon, *What we owe to each other* (Cambridge, MA, 1998) at p. 391, n. 21.

3.

In the previous section, I argued that the egalitarian moral theories with which we're most familiar don't support the intuition that there is something inegalitarian about prioritizing treatment on the basis of patients' usefulness to others. In this section, I want to show that the argument of the previous section actually involved an important oversight.

To see where the argument skipped a beat, let's return to the redundancy objection to relational egalitarianism. I said that if some practice would involve a failure of individuals to relate to one another as equals, this has to be because there is something more fundamentally wrong with that practice. But there is one important class of exceptions to this rule.

The exceptional cases are those involving patterns of communication wherein one individual represents the other as inferior *merely as such*: for example, a case in which you just tell me straightforwardly that I am beneath you. Verbal communication isn't strictly necessary, of course. You could exploit other conventional signifiers of disdain, such as the unilateral lip curl associated with the expression of contempt. In these cases, the complaint that you fail to treat me as an equal seems to have non-derivative moral significance. What's gone wrong here really has to do with how we relate to one another and cannot be traced to any deeper problem.

In my view, the egalitarian objection to rationing health care on the basis of patients' usefulness to others is best interpreted as semiotic: it has to do with what is thereby expressed about the status of the patients made subject to the rule. For this reason, it can exploit the loophole we've just noted for relational egalitarians facing up to the redundancy objection. The way we ration healthcare has some kind of meaning. The standard way of rationing this good used by public health services like the UK's NHS expresses the equality of citizens, whereas rationing on the basis of social utility would express something contrary.

Most of the remainder of this paper is devoted to substantiating this proposal. The next section will provide a fragment of a general framework for thinking about semiotic objections in legal and political philosophy. The section after that will apply the fragment to the case of healthcare rationing. My focus will be on establishing conditions under which a derogatory message may be

conveyed by an act or practice and showing that those conditions are satisfied in the case that interests us.

I won't offer anything like a fully worked-out theory of why it is normally wrong to convey derogatory messages. I just assume that it is. I do assume that the communicator of a derogatory message need not harbour the derogatory attitudes conveyed by her act in order for communication of the message to be morally criticisable, but I take this as uncontroversial. Thus, if a white man blacks up in order to appear as Malcolm X at a fancy dress party, I presume we would regard this as objectionable even if the man genuinely intended to flatter the civil rights leader by so emulating his appearance. He ought to have known better. Beyond the assumption that derogatory attitudes aren't strictly necessary for this sort of wrong to occur, I hope to remain non-committal. Some may believe that conveying a derogatory message is *pro tanto* wrong in and of itself, regardless of whether any hurt is felt by those derogated. Others may think that felt offence is key. I lean toward the view on which derogatory content and felt offense are individually necessary and jointly sufficient. However, I will not ask you to assume this view.

4.

The basic idea that acts and practices have meanings is presumably uncontroversial. The practice of kneeling in front of someone expresses reverence or respect. Refusing to look someone in the eye may express disdain. Putting up a sign that says 'Men only' means something very different if placed on a bathroom door, as opposed to the door of a courthouse.¹⁹

As noted, in the sense that interests us here, what is expressed by a given act or practice need not be what any given agent intends for it to mean in any token instance. In other words, meaning, so understood is unlike *utterer's meaning* and more like *sentence meaning*.²⁰ A person might believe that

¹⁹ Thurgood Marshall, Concurrence-dissent in *City of Cleburne v. Cleburne Living Center, Inc.*, 473 U.S. 432 (1985).

²⁰ H. P. Grice, 'Utterer's meaning, sentence-meaning, and word-meaning', *Foundations of Language* 4 (1968), pp. 225-42.

her refusal to look someone in the eyes conveys disrespect in a context where it conventionally means the opposite. In the sense that interests us here, this person's act has a meaning opposite to her intention.

There can arguably be cases in which the meaning of a practice is unobvious to *most* participants. For example, Anderson and Pildes note that it used to be commonly accepted that it was compatible with an attitude of respect that men in business settings would compliment their female colleagues and subordinates on their physical appearance.²¹ One might reasonably think that this practice was nonetheless insulting to working women. What is expressed or conveyed by some practice may therefore be thought in some instances to be opaque to the majority of its participants, requiring significant interpretative skill to discern. In this respect, the meaning at issue is more like that of a sonnet or an open-textured statute of law than that of a humdrum sentence in ordinary English.

If the meaning of a practice needn't be constrained by the beliefs of the majority, how can it be determined in any given instance? How can we adjudicate between rival interpretations? I'll now try to say something to address this question. But I will offer only a fragment of a complete theory of the meaning of social practices. The view I'll present here is so far from being a complete theory that it doesn't even attempt to explain meaning from the ground up. Instead, it attempts to specify a sufficient condition for when some practices acquire meanings from others.

Consider a thought experiment. Imagine that a political cartoonist creates a drawing, which strongly resembles Francisco Goya's painting *Saturn Devouring his Son* (ca. 1819-23). That painting depicts the Titan of Roman mythology as a frightening monster, greedily gobbling up his progeny. The cartoonist's drawing is different from Goya's painting in one key respect. The face doesn't look like the Titan's face in Goya's painting. Instead, it looks like the current Chancellor of the Exchequer. The Chancellor has drawn controversy by implementing austerity policies, requiring deep cuts to public services on which the most vulnerable in society rely.

²¹ Elizabeth Anderson and Richard Pildes, Richard, 'Expressive theories of law: a general restatement', *University of Pennsylvania Law Review* 148 (2000), pp. 1503-75.

Here are some things it seems correct to say about a drawing of this kind. Firstly, the image can be said to depict the Chancellor or the government as ‘eating its own young’, in the idiomatic sense: i.e. as desperately preying on vulnerable persons for whom they are responsible. Secondly, it’s natural to think that the image has this content by virtue of its resemblance to Goya’s painting, which depicts a powerful figure literally eating his own young.

A little reflection shows that the resemblance needn’t be intended for the drawing to have this kind of inherited depictive content. Suppose the cartoonist did not aim to imitate Goya’s painting, although she was familiar with it. Her intent was merely to depict the chancellor as a monster. Her drawing just came out that way. Given that her cartoon is created and disseminated in a context where Goya’s painting is among the icons of art history, it still seems correct to say that the cartoon has managed to depict the Chancellor or the government as ‘eating its own young.’ In fact, this could hold true even if the cartoonist didn’t intend to create any artwork at all. If she somehow inadvertently collaged together a bunch of clippings to create an image of this kind, the composite image would presumably still depict the Chancellor in just the same way when disseminated in the right context.

In describing these cases in the way I’ve done, I don’t mean to commit myself to any general resemblance theory of depictive content,²² nor to wholesale rejection of the view that intended resemblance is important in explaining depictive content.²³ Keep in mind that the kind of content I’m analysing is *derivative content*. It’s entirely compatible with what I’ve said that the content of Goya’s painting should not be explained in terms of resemblance, or that it should be explained in terms of intended resemblance.

My claim, then, is that something similar happens in respect of social practices: meaning can be inherited by virtue of resemblance to other practices, regardless of ‘authorial’ intentions.

²² Catharine Abell, ‘Canny resemblance’, *Philosophical Review* 118 (2009), pp. 183-223; Richard Hopkins, *Picture, image, and experience: a philosophical inquiry* (Cambridge, 1998) ; John Hyman, ‘Depiction’, *Royal Institute of Philosophy Supplement* 71 (2012), pp. 129-50; Christopher Peacocke, ‘Depiction’, *Philosophical Review* 96 (1987), pp. 383-410.

²³ As per Abell, ‘Canny resemblance’, Hopkins, *Picture, image, and experience*.

Here is an illustration of this principle in action, taken from the literature on racial profiling. Hellman argues that racial profiling of black Americans would be wrong even if it were based on accurate and well-evidenced statistical regularities and implemented without the influence of racial prejudice.²⁴ On Hellman's view, even racial profiling of this kind would be objectionable, because of its expressive content.²⁵ If the practice of profiling is grounded in accurate and well-evidenced statistical regularities, and enacted for the sake of reducing crime by persons without any prejudice against the profiled community, how can profiling express something derogatory? Hellman's answer is that it does so by virtue of its similarity to socially salient discriminatory practices characteristic of racial animus toward black Americans. As she explains:

Because its content tracks familiar cultural stereotypes about blacks and crime, the practice also expresses something more troubling such as 'blacks are naturally disposed to be criminals.' While a statistical generalization about black teens and crime itself does not entail a statement about causation, the cultural baggage of the category 'black' brings this meaning forward as well.²⁶

Before we move any further, let me quickly address two issues that'll be familiar to - and may have occurred as queries in the minds of - those acquainted with stock objections to resemblance theories of depictive content.²⁷

Firstly, we should keep in mind that resemblance is symmetric. Since inheritance of meaning is asymmetric, we need something to break the symmetry. In order to do so, we can appeal to the idea that one of the relata of the resemblance relation may be socioculturally pre-eminent with respect to the other. Meaning flows along this gradient. The cartoonist's drawing inherits its meaning from

²⁴ Deborah Hellman, 'Racial profiling and the meaning of racial categories', *Contemporary debates in applied ethics*, 2nd edn., eds. Cohen and Wellman (Malden, MA, 2014), pp. 232-43.

²⁵ Compare Benjamin Eidelson, *Discrimination and disrespect* (Oxford, 2015) at pp. 173-222.

²⁶ 'Racial profiling and the meaning of racial categories', p. 237.

²⁷ Nelson Goodman, *Languages of art: an approach to a theory of symbols* (Indianapolis, IN, 1968), at pp. 3-5.

Goya's painting and not *vice versa* because the latter is an icon of Western art history and the former isn't. With respect to racial profiling, unbiased statistical discrimination may inherit its meaning from derogatory stereotypes because the latter are more salient as a guiding framework for relations between law enforcement and blacks in American society.

Secondly, we should keep in mind that everything resembles every other in some respect or other. Which dimensions of resemblance are the ones that matter?

The view I've proposed arguably has an easier time of answering this question than standard resemblance theories of depictive content, because it tries to explain derivative content. That lets us piggyback. We can just stipulate that what counts is resemblance in respect of those properties of the pre-eminent relatum in virtue of which it has its content, whatever those might be.

For what it's worth, I don't believe that there is any mind-independent fact of the matter about whether one thing sufficiently resembles another. What matters ultimately is salient similarity in the minds of the relevant audience. Furthermore, what is salient may be capable of changing due to the very act of interpretation. An interpretation that reads a derogatory meaning into a practice may succeed in affixing that meaning to the practice for the very first time, by altering our perception of its salient resemblances to comport with those highlighted by the interpreter.

5.

I now want to get back to the issue of rationing healthcare and the question of whether there is an egalitarian objection to doing so on the basis of the social utility of patients. The aim of this section is to show that a semiotic objection fitting that description exists, drawing on the resources of the previous section.

Why should prioritizing people on the basis of their social usefulness convey a message at odds with the moral equality of citizens? I think there are two key explanatory factors to consider. The first is the significance attached to public healthcare provision in countries like the United Kingdom. The second is the culture of meritocracy that pervades post-industrial societies. I'll argue that the demeaning message conveyed by prioritizing treatment on the basis of patients' usefulness to others arises from an interaction effect involving these factors.

I'll begin with a discussion of the significance attached to public healthcare provision in countries like the United Kingdom. Some philosophers - most prominently Daniels²⁸ - believe that health by its nature is a special good, subject to more stringent egalitarian norms of distributive justice than other social goods. Whether or not this is correct, health care is certainly treated as a special good in many Western countries, including the UK. In these countries, healthcare is a socially recognized need: it is understood as something claimable by the individual against the community. It is like physical security or basic education, but unlike food or friendship - goods that may be no less essential to a flourishing life, but which are not recognized as owed to each individual by the state.

Where health is a socially recognized need of this kind, access to healthcare understandably becomes bound up with someone's status as a member of the political community.²⁹ To be a member of the political community is to be the sort of person who has a claim upon the community to help meet her health needs. Where the political community is also invested in an ideal of social and political equality, the provision of healthcare will become entangled with that ideal as well.

By way of illustration, consider the UK's National Health Service. Its chief architect was Aneurin (Nye) Bevan, Minister of Health in the Attlee government from 1945-51. Bevan placed considerable emphasis on ensuring that everyone was treated the very same by the NHS. For example, he rejected the proposal of the 1944 White Paper, under which executive responsibility for public health would rest with local governments. Instead, the National Health Services Bill of 1946 created a nationally centralised system, which was viewed as necessary to achieve a uniform standard of care.³⁰ Bevan also rejected any suggestion that those who were well-off should have to pay for NHS services. He feared that this would mean 'the creation of a two-standard health service, one below and one above the salt.'³¹

²⁸ Norman Daniels, *Just health care* (Cambridge, 1985), *Just health: meeting health needs fairly* (Cambridge, 2007).

²⁹ Walzer, *Spheres of Justice*, pp. 64-83.

³⁰ Aneurin Bevan, 'Proposals for a national health service', Cabinet Memorandum 13 Dec. 1945 <<http://filestore.nationalarchives.gov.uk/pdfs/small/cab-129-5-cp-45-339-39.pdf>>

³¹ Aneurin Bevan, *In place of fear* (London, 1952) at p. 76.

Similar ideals were affirmed by the Brown government when it drew up the NHS Constitution in early 2009 to consolidate the values and commitments of the UK's public health system. In its most current edition, the first principle of the constitution states that the NHS 'provides a comprehensive service, available to all irrespective of gender, race, disability, age, sexual orientation, religion, belief, gender reassignment, pregnancy and maternity or marital or civil partnership status ... designed to improve, prevent, diagnose and treat both physical and mental health problems with equal regard.'³² The egalitarian character of the NHS is given precedence over the other guiding principles set out in the constitution, foregrounding and reaffirming the status of the NHS as a vehicle for social equality.

As noted earlier, this is only one half of the puzzle. The other is belief in meritocracy. Although it has few adherents in political philosophy,³³ there is a widespread belief that those who are talented and productive deserve to be better off, and that an economic system that assigns greater rewards to the highly skilled is not merely efficient but just. Moreover, there is a widespread belief that current social conditions are conducive to meritocratic justice, allowing those who are talented and hardworking to rise to the top, regardless of background.³⁴

Exactly what role does the cultural prominence of meritocratic ideas play in grounding the semiotic objection to prioritizing treatment on the basis of patients' usefulness to others? To answer this question, we need to appeal to the resemblance theory of derived content that I set out in the previous section. I argued then that meaning can be inherited by one practice in virtue of its resemblance to another, regardless of the intentions of those who enact the practice, provided the resemblance is salient and the resembled item is pre-eminent. We'll now apply this principle to the present case.

A salient fact about the practice of prioritizing treatment on the basis of patients' usefulness to others is that it would reward those with high earning potential and/or special talents by assigning

³² NHS Constitution for England, p. 3.

³³ A notable exception being David Miller, *Principles of Social Justice* (Cambridge, MA, 1999).

³⁴ Louise Bamfield and Tim Horton, *Understanding attitudes to tackling economic inequality* (York, 2009).

priority to their health needs. It would therefore mimic a meritocratic principle of distribution and thereby convey judgments of merit, representing health care as something that some people deserve to a greater extent than others. This would obtain even if the underlying rationale for the practice were entirely forward-looking and wholly divorced from any thought that some individuals deserve favourable treatment. Even if the policy is justified in the minds of its architects by utilitarian principles that contradict any conception of desert, the practice would express a principle of merit under which people differ in what level of medical care they deserve.

The similarity to which I'm pointing obviously wouldn't be perfect. For example, prioritizing treatment on the basis of patients' usefulness to others would arguably give priority to parents of young children, whereas becoming the parent of a young child isn't ordinarily viewed as indicative of hard work or unique skill. But this need not pose a significant objection to the present account.

Firstly, the results obtained by McKie and Richardson indicate that there is much less opposition to prioritizing parents of young children.³⁵ Only 39.2% of respondents in their sample strongly disagreed with that proposal, as compared with the 87% who strongly disagreed with the general proposal to prioritize those who contribute more to society. This is as we should expect if my view is correct, since prioritization of parents with young children scores poorly in respect to resemblance to a meritocratic principle of distribution.

Secondly, the existence of some exceptions falls far short of establishing that the overall similarity between a principle prioritizing treatment on the basis of patients' social utility and a principle of meritocratic distribution is insufficient for a derogatory meaning to become attached to the former. Similarity is never perfect and exact. A political cartoon that sufficiently resembles Goya's *Saturn* to depict someone as sacrificing the most vulnerable in society needn't achieve perfect similarity to have the meaning it does, even setting aside the fact that it must differ from the original somehow or other if it's not to count as a simple reproduction. The cartoon could be in black and white, whereas the painting is in colour. The cartoon could be much smaller. By analogy, the practice of prioritizing treatment on the basis of patients' social usefulness needn't perfectly mimic a

³⁵ McKie and Richardson, 'Social preferences for the inclusion of indirect benefits'.

meritocratic principle to carry forward the same meaning when implemented in a social context wherein meritocratic principles are culturally ascendant.

Assume this is correct. Then we have everything we need to complete the explanation. We just have to consider how the meritocratic meaning of the practice of prioritizing treatment on the basis of patients' social usefulness would interact with the pre-existing significance attached to medical services in societies like the United Kingdom.

Here is the idea. If we were to prioritize treatment on the basis of patients' usefulness to others, we would thereby represent some individuals as more deserving than others of a good the claim to which is bound up with each person's standing as an equal member of the political community. Doing so would therefore convey a message at odds with the ideal of social and political equality hitherto expressed through the public provision of healthcare. As before, this need not be the intended meaning of the practice. It is nonetheless what the practice would convey. This, I believe, is why prioritizing treatment on the basis of patients' usefulness to others is out of keeping with an ideal of equal regard and respect.

6.

I want to use this section to reflect on the significance of the conclusions I've drawn so far. In particular, I want to consider the plausible but paradoxical observation that semiotic objections are both powerful and flimsy, and to think through the implications of this observation for the issue of rationing the means to health.

Why do I say that semiotic objections are both powerful and flimsy? Let's start with the 'powerful' part. As I see it, semiotic objections that appeal to the derogatory meanings of certain acts and practices are powerful because human beings are extremely sensitive to status and hierarchy.³⁶ Where you rank relative to others matters greatly to your sense of self-respect. It's also a key

³⁶ Christopher Boehm, *Hierarchy in the forest: the evolution of egalitarian behavior* (Cambridge, MA, 1999).

determinant of subjective well-being,³⁷ and has a surprisingly strong effect on health.³⁸ Status is something we encode and express symbolically. It's something gained and lost in virtue of how we are publicly presented and represented. Because status is so important to us and depends on what is conveyed about us by the social practices in which we take part, I think semiotic objections in legal and political philosophy should be taken very seriously. Although they might feel most at home within relational egalitarian theories, they should weigh heavily with moral and political theorists of any persuasion.

Of course, it is one thing to insist that such objections should weigh heavily, and another to say that they should weigh more heavily than the sum of competing considerations. Thinking that there is a weighty semiotic objection against prioritizing treatment on the basis of patients' usefulness to others is in principle compatible with thinking that doing so is all-things-considered permissible in light of the benefits that might be achieved by such a policy. The question of how to weigh the competing considerations at issue is not one that I feel can be satisfactorily answered in the remainder of this paper, and will therefore be left as a matter of judgment for the reader.

Why, then, do I say that semiotic objections are flimsy? This is not intended as a verdict about their absolute or relative weight. Instead, what I have in mind is this. Assuming that meaning is socially constructed, semiotic objections fail to identify anything wrong with a practice *as such*. If the only thing wrong with a practice is that it has an objectionable meaning, and there are otherwise good reasons to adopt it, the natural conclusion is that we should rid the practice of this meaning, as opposed to ridding society of the practice.³⁹ Furthermore, the objection won't travel. It is tied to a particular context wherein a good has a particular meaning. Outside of that context, the objection most likely has no purchase.

³⁷ Richard Ball and Kateryna Chernova, 'Absolute income, relative income, and happiness,' *Social Indicators Research* 88 (2008), pp. 497-529.

³⁸ Michael Marmot, 'Health in an unequal world', *The Lancet* 368 (2006), pp. 2081-94.

³⁹ Jason Brennan and Peter Martin Jaworski, 'Markets without symbolic limits', *Ethics* 125 (2015), pp. 1053-77.

I believe the semiotic objection to rationing healthcare on the basis of social usefulness is flimsy in each of the respects noted in the previous paragraph, though less so in the first than in the second. The remainder of this section will elaborate on this idea.

Let's begin with the idea that the proper response to the conclusions drawn in this paper is to change what it would mean to distribute healthcare in accordance with the social utility of patients. I think there is much to agree with in this proposal, although it's ultimately too strong.

For the reasons outlined in section 2, I think it's hard to come up with other plausible principled objections to the practice of rationing healthcare in accordance with social utility, besides the semiotic objection I've outlined. On the other hand, there are strong grounds for adopting a practice of this kind insofar as doing so would increase aggregate utility and promote distributive equality. These are exactly the conditions under which it makes sense to focus on changing the meaning of a practice so as to reap its moral benefits without incurring semiotic costs.

There are admittedly some objections to rationing healthcare on the basis of patients' social utility that I didn't consider in section 2 - objections that aren't naturally interpreted as egalitarian in character. For example, Kamm⁴⁰ and Brock⁴¹ suggest that prioritizing the more useful may violate the Kantian injunction never to treat people as mere means. If some patient is dispreferred as a recipient of treatment because she is not as socially useful as another whom we prioritize instead, the former could be said to be wronged by virtue of being treated solely as a means and not as an end-in-itself. But I think it is clear that this would be confused. Prioritizing among patients on the basis of their social utility involves treating people *as a means*, but there is nothing wrong with that. I use other philosophers as a means for testing out my ideas, and that seems fine. What is plausibly viewed as objectionable is treating someone as a *mere means*: behaving toward them in something approaching

⁴⁰ F. M. Kamm, *Morality mortality, vol. 1: death and whom to save from it* (Oxford, 1993).

⁴¹ Brock, 'Separate spheres and indirect benefits'.

the manner appropriate to an object without sense or reason.⁴² Rationing healthcare in the way we're considering obviously requires no such thing.

Let's assume, then, that there is a good case for trying to change what it would mean if we were to ration healthcare in this way. Based on what I've said in the previous section, if we wanted to change the meaning of this practice, we would have two options. We could alter the social significance of public healthcare provision. I don't think we should do this. The idea of health care as something to which every citizen has a claim against the state is one of the most important achievements of the social democratic movements of the past century. Alternatively, we could decrease the salience of meritocratic principles and ideals. I think there are already compelling reasons for working toward that aim. For example, the rhetoric of meritocracy represents social life as a competition and thereby tends to undercut a sense of community and solidarity. Decreasing the salience of meritocratic principles and ideals is something I think we have strong reason to work towards, quite apart from the any ambition to ration healthcare in a more socially beneficial way without simultaneously representing some individuals as inferior in status. But, it needn't follow that this is what we ought to do, at least not insofar as that means we would be criticisable for failing to follow through. After all, we may have stronger reasons to work toward other goals that compete for our time and resources.

Let me now turn to the idea that semiotic objections don't travel. The thought here is that objections of this kind are tied to a particular context and the meaning of a good within that context. These objections shouldn't be expected to generalize to other goods and/or other contexts. This matters because the public provision of healthcare isn't the only context in which we might be led to wonder whether it's appropriate to prioritize among individuals on the basis of their indirect utility.

Consider the case of private philanthropy. Suppose that you have been convinced by proponents of effective altruism to set aside some significant proportion of your income, which you

⁴² G. A. Cohen, *Self-ownership, freedom, and equality* (Cambridge, 1995) at pp. 239-40; Derek Parfit, *On what matters, vol. 1* (Oxford, 2011) at pp. 212-3.

hope to use to improve the lives of others to the greatest extent possible.⁴³ There is good reason to believe that health interventions in developing countries are highly cost-effective.⁴⁴ But standard cost-effectiveness calculations only take account of immediate health benefits. Even charity evaluators like GiveWell, who seek to rank organizations in terms of their impact per dollar, give little if any consideration to the indirect social utility of potential beneficiaries.

Should you? For example, should you prioritise saving lives and preventing disability in countries where growth rates are higher? Additional wealth in these countries compounds more quickly and therefore generates considerably more wealth in the long run. For this reason, you might prefer programs that treat neglected tropical diseases, like soil-transmitted helminth infections, in India, as opposed to sub-Saharan Africa. You might prefer these programs not because the immediate health benefits are more significant in India, but because workers there are more productive and their productivity is socially valuable. You might even wonder whether it is better to donate to Deworm the World Initiative rather than Schistosomiasis Control Initiative because Deworm the World runs programs that prevent the spread of NTDs in India as well as sub-Saharan Africa, whereas Schistosomiasis Control Initiative currently operates exclusively in Africa.

Of course, the question of whether to take indirect social benefits into account isn't unique to donors who elect to focus their giving on global health, but I expect that the question is likely to strike us as most troubling when life and limb are at stake. Consider instead a donor browsing the website of Kiva Microfunds, a non-profit that facilitates microloans to low-income entrepreneurs and students. Their website presents potential donors with a wide selection of borrowers seeking support for

⁴³ Joe Horton, 'The all or nothing problem', *Journal of Philosophy* 114 (2017), pp. 94-104; William MacAskill *Doing good better: a radical new way to make a difference* (London, 2015); Theron Pummer, 'Whether and where to give', *Philosophy and Public Affairs* 44 (2016), pp. 77-95; Peter Singer, *The most good you can do: how effective altruism is changing ideas about living ethically* (New Haven, CT, 2015).

⁴⁴ Dean Jamison, Hellen Gelband, Susan Horton, Prabhat Jha, Ramanan Laxminarayan, Charles Mock, Rachel Nugent, eds., *Disease control priorities: improving health and reducing poverty* (Washington D.C., 2018); Toby Ord, 'Considering cost-effectiveness: the moral perspective', *Priority setting in health: building institutions for smarter public spending*, eds. Glassman and Chalkidou (Washington D.C., 2012), pp. 15-19.

projects including the purchase of livestock, installation of solar panels for heating water, or university tuition fees. I think we are unlikely to be troubled by the idea of a donor preferring to support a person hoping to attain a degree in medicine as opposed another aiming to study English, based on the belief that doctors do more good. By contrast, if the question concerned saving a person's life, a preference to help a prospective doctor over a prospective English teacher because doctors are more useful might well give us pause.

If the argument of this paper is correct, prioritizing among potential beneficiaries on the basis of the social utility of their health should arguably be thought *unobjectionable* when done by private individuals deciding where to donate. That's because semiotic objections don't travel. The social meaning that attaches to the distribution of healthcare resources in democratic states wherein health is a socially recognized need can't be exported to this new context. It involves individuals in different countries with different cultures. There is no shared social context encoding failure to meet health needs as imbued with symbolic content bearing on the status of potential beneficiaries. As a result, the semiotic objection that applies if public institutions choose to prioritize among potential beneficiaries on the basis of the social utility of health fails to carry over to the case of private donors.

By contrast, note that if we viewed rationing healthcare on the basis of recipients' social utility as objectionable because it violates the Kantian injunction, then there would presumably be no significant difference between the case of a government distributing healthcare resources within a nationalized service and private individuals choosing where to donate. Each would presumably violate the same injunction to the same extent, because each would instrumentalize human persons in the way forbidden by the Formula of Humanity.

Once this is pointed out, some may feel that their intuitions are actually better aligned with the latter view. As I've already hinted, we may experience a similar sense of unease when considering a private donor prioritizing among people whose life and health are at stake on the basis of their social utility as when we consider governments doing the same. But I don't think this should surprise us very much, even if my account is correct.

I say this because we tend to project meanings beyond their proper bounds and reify conventions. For example, beginning with Piaget, developmental psychologists have documented a

widespread tendency in children toward *nominal realism*.⁴⁵ This is the belief that the names we give to things could not be different from what they actually are: in other words, that the names we use are uniquely appropriate to the things they name. Even we, who know better, find it very hard to uncouple a foreign word from the meaning it has in our mother tongue when encountering a word in another language that's a duplicate of one in our own. More generally, extensive psychological research indicates that we are predisposed to construe the way things are by convention as inherent in their nature, treating what something means for us as what it must mean as such.⁴⁶

For this reason, I don't think we should attach much, if any, weight to lingering feelings of unease that might accompany the idea of a private donor prioritizing among health needs on the basis of their social utility. That lingering unease could simply be a symptom of the over-projection of meaning, similar to the misplaced unease Western travellers may feel when finding swastikas adorning Buddhist temples in Korea.

7.

Here is a recap of what I've argued in this paper.

The central question on which I've focused is that of why it would be wrong for publicly funded healthcare programs to take into account the social utility of patients in setting priorities and rationing resources. Clearly, this should not be the unique deciding factor. The question is whether it might be one factor to be weighed alongside those others that we think ought to be considered. That list might be no longer than *need* and *cost*, though some may believe that other factors, such as personal responsibility or past contributions, ought to count as well.

Although it's natural to think that there is some kind of egalitarian objection to prioritising among patients on the basis of their social utility, I argued in section 2 that it is much harder to give

⁴⁵ Jean Piaget, *The child's conception of the world*, transl. Tomlinson and Tomlinson. (London, 1929).

⁴⁶ Andrei Cimpian and Erika Salomon, 'The inherence heuristic: an intuitive means of making sense of the world, and a potential precursor of psychological essentialism', *Behavioral and Brain Sciences* 37 (2014), pp. 461-80; Howard Gabbenesch, 'The perception of social conventionality by children and adults', *Child Development* 61 (1990), pp. 2047-59.

substance to that thought than it might first seem. In section 3, I pointed to a loophole left open by the arguments in section 2, and I noted that cashing out the objection as semiotic in character could allow us to exploit that loophole. Section 4 set out a resemblance theory of derived content and section 5 made use of that theory in explaining why prioritizing treatment on the basis of patients' usefulness to others is inegalitarian. There I argued that the significance attached to medical care in societies where health is a socially recognized need interacts with the prevalence of meritocratic ideals. The result is that prioritizing treatment in this way represents some people as more deserving of a good that is bound up with each person's standing as a citizen, thereby expressing something contrary to the ideal of equal citizenship.

Finally, in section 6, I explored the implications of the observation that semiotic objections are paradoxically both powerful and flimsy. I stressed the importance for individual human flourishing of how we are publicly represented, insisting that semiotic objections should therefore weigh heavily with moral and political theorists of any stripe. But I also stressed the point that semiotic objections don't travel. I explored the implications of this for the case of private donors who may consider prioritizing among potential beneficiaries by taking account of what the potential beneficiaries can do to improve the lives of others. I argued that because semiotic objections don't travel, this should be unobjectionable, even if there is a significant egalitarian objection to carrying out a similar system of triage in publicly funded healthcare systems.⁴⁷

andreas.mogensen@philosophy.ox.ac.uk

⁴⁷ For valuable comments on previous drafts of this paper, I'm grateful to Hilary Greaves, William MacAskill, Tom Sinclair, Christian Tarsney, and Teru Thomas, as well as the audience for my presentation at University College Dublin on the 25th of October 2018. I'd also like to thank two anonymous referees at *Utilitas* for their helpful comments and suggestions.