

Prioritarianism and the Measure of Utility*

MICHAEL OTSUKA

Philosophy, London School of Economics

HERE I present a challenge to prioritarianism, which is, in Derek Parfit's words, the view that 'we have stronger reasons to benefit people the worse off these people are'.¹ We have such reasons, according to this view, simply by virtue of the fact that a person's 'utility has diminishing marginal *moral importance*'²—i.e., that equal improvements in a person's well-being matter less, morally speaking, the better off she is in absolute terms. It follows, from this view, that one might have stronger reason to benefit someone who is less well off rather than someone who is better off, even when this benefit would amount to a lesser increase in utility than a benefit to the better off person.³

In discussions of prioritarianism, it is often left unspecified what constitutes a greater, lesser, or equal improvement in a person's utility. In his own defence of prioritarianism, for example, Parfit explicitly prescinds from 'difficult questions . . . about what it would be for some benefits to be greater than others' and 'simply assume[s] that we can distinguish between the size of different possible benefits'.⁴ Parfit just stipulates numerical benefits of different magnitudes that

*I presented an earlier version of this article as the John Passmore Lecture at the Australian National University and at Duke University, the Hebrew University, New York University, Nuffield College, Oxford, Princeton University, Rutgers University, the University of Maryland, the University of Sydney, and the annual conference of the Society for Applied Philosophy. I'm grateful to the members of these audiences for their comments. I'd also like to express my debts and thanks to Matthew Adler, Marcello Antosh, Ralf Bader, Luc Bovens, Geoffrey Brennan, David Chalmers, Thomas Dougherty, Cécile Fabre, Marc Fleurbaey, Robert Goodin, Christopher Morris, Alan Patten, Philip Pettit, Jonathan Quong, Samuel Scheffler, Nicholas Southwood, Larry Temkin, Peter Vallentyne, Alex Voorhoeve, four anonymous referees, and Princeton's University Center for Human Values, where I was a Visiting Faculty Fellow during most of the period I devoted to this piece.

¹Derek Parfit, 'Another defence of the priority view', *Utilitas*, 24 (2012), 399–440, at p. 401.

²Derek Parfit, 'Equality or priority?' *The Lindley Lecture* (Lawrence, Kansas: University of Kansas, 1991), reprinted in *The Ideal of Equality*, ed. Matthew Clayton and Andrew Williams (Basingstoke: Palgrave, 2002), pp. 81–125, at p. 105. (All subsequent page references are to the reprinted version.) On the typical version of prioritarianism, which I shall assume throughout this discussion, there are no discontinuities in the diminution in marginal moral importance of utility.

³Unless indicated otherwise, a person's 'utility' should be regarded throughout this article as a synonym for her 'well-being', where the latter, and therefore the former, is to be understood in terms of how well the life of that person is really going (or would go). I shall also assume that the measure of the size of a benefit is always the magnitude of the increase in utility (i.e., well-being) to which it gives rise.

⁴'Another defence', p. 403.

The copyright line for this article was changed on 5 August 2014 after original online publication.

comprise intervals along a whole number cardinal scale that is meant to represent the absolute levels of people's utility in linear fashion.⁵ We are supposed to assume that this scale provides an accurate representation of people's utility. But we are not offered an account of what constitutes the measure of the size of the units on this scale.

I shall argue that prioritarianism cannot be assessed in such abstraction from an account of the measure of utility. Rather, the soundness of this view crucially depends on what counts as a greater, lesser, or equal increase in a person's utility. In particular, prioritarianism cannot accommodate a normatively compelling measure of utility that is captured by the axioms of John von Neumann and Oskar Morgenstern's expected utility theory. Nor can it accommodate a plausible and elegant generalization of this theory that has been offered in response to challenges to von Neumann and Morgenstern. This is, I think, a theoretically interesting and unexpected source of difficulty for prioritarianism, which I shall explore in the remainder of this article.

I.

In presenting this challenge, it will be useful for me to begin with a case that is closely modelled on one that Thomas Nagel presents in his essay 'Equality'⁶—and with which Derek Parfit opens his *Lindley Lecture* 'Equality or Priority?'⁷ Imagine that you are the parent of thirteen-year-old twins. For the sake of interpersonal comparability, let us assume that they are identical twins with identical preferences. They differ, however, in the following crucial respect: one of them has recently been diagnosed with a condition that will soon give rise to a severe mobility-impairing disability, but the other has been given a clean bill of health. You need to change jobs and must therefore choose whether to move your family to cramped urban accommodations in an unpleasant and dangerous neighbourhood or to the blue skies and open spaces of a semi-rural suburb. Either option would be equally good insofar as your own well-being is concerned. But they would not be equally good insofar as the well-being of each of your two children is concerned. If you move to the city, your able-bodied child, who loves nature and sports, will be hemmed in and frustrated, but your (soon-to-be) disabled child will have access to special medical facilities that will somewhat, but far from wholly, alleviate the effects of his disability. If you move

⁵In his *Lindley Lecture*, whole numbers ranging from 9 to 200 are presented as representations of people's utilities without any further context. Parfit specifies that, for any whole number n , a unit increase from n to $n + 1$ will always constitute an equally large increase in utility: i.e., the cardinal scale is linear. (See 'Equality or priority?' pp. 82–3.) In 'Another defence', Parfit correlates people's absolute levels of lifetime utility with the number of years they will live, via the stipulation that 'each extra year of life would be an equal benefit' along with an implicit assumption that utility is additive. (See 'Another defence', p. 425.)

⁶Thomas Nagel, 'Equality', in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), pp. 106–27, at pp. 123–4.

⁷Op. cit.

to the suburb, your able-bodied child will flourish, but your disabled child will not receive this treatment. The marginal benefit to the able-bodied child of a move to the suburb would, I shall stipulate, be slightly greater than the marginal benefit to the disabled child of a move to the city. It follows that a *utilitarian* would opt for the suburb. Nagel would maintain that you nevertheless have overriding *egalitarian* reason to move to the city.

Parfit would agree that you have overriding reason to move to the city, but he would maintain that this reason is *prioritarian* rather than egalitarian. In other words, Parfit would maintain that you ought to move to the city simply on account of the fact that improvements in well-being matter more, the worse off someone is in absolute terms. The granting of such priority to the worse off is not, Parfit maintains, an egalitarian concern, where the latter is read as taking how well off someone is in relation to someone else to be a matter of intrinsic moral concern. All that matters, on the priority view, is how well off people are in absolute terms: the worse off they are, the more their well-being matters. It does not matter, above and beyond this, how well off they are in relation to others. In support of this priority view, Parfit maintains that ‘it would be just as urgent to benefit the handicapped child even if he had no sibling who was better off’.⁸

Now consider the following transformation of Nagel’s two-child case into a case in which you have only a single child. Moreover, this thirteen-year-old teenager has recently received a diagnosis that there is a 50 percent chance that he will soon develop the severe disability and a 50 percent chance that he will remain entirely healthy.⁹ Suppose that you must opt for the job in either the city or the suburb before it is known whether your child will develop the disability. As before, the marginal benefit to your child of a move to the suburb, if he remains in full health, would be slightly greater than the marginal benefit to him of a move to the city, if he develops the disability.

On its most simple and straightforward application to cases involving risk, the priority view yields just as strong a reason to move to the city in this one-child case as one has in the two-child case. On this application, prioritarianism calls for a familiar sort of maximization of expected value, where the value in question is priority-weighted utility. It calls for the maximization of the sum of the priority-weighted utilities of all possible outcomes that might be the upshots of this choice, where the value (i.e., the priority-weighted utility) of each such outcome is discounted by the probability that it will obtain. The priority-weighted utility of any given outcome is the sum of everyone’s utility in that outcome, after each person’s utility has been assigned the relevant positive but decreasing marginal moral importance in prioritarian fashion. One discounts the value of each outcome by its probability simply by multiplying this value by

⁸‘Equality or priority?’ p. 108.

⁹See Michael Otsuka and Alex Voorhoeve, ‘Why it matters that some are worse off than others: an argument against the priority view’, *Philosophy and Public Affairs*, 37 (2009), 171–99, at p. 188.

the probability that this outcome will come about. On such a standard expected value calculation, the relative strength of a prioritarian parent's reason to move to the city rather than the suburb will be just as strong in the one-child case as it is in the two-child case.¹⁰

In 'Another Defence of the Priority View', Derek Parfit defends a more nuanced and complex approach to risk than the one just sketched.¹¹ On the version that Parfit defends there, a prioritarian calculates the expected value of the priority-weighted outcomes in just the manner sketched a moment ago. But in addition to valuing *outcomes* in such prioritarian fashion, Parfit proposes that the prioritarian should also attach independent value to the *prospects* of individuals, where their prospects are simply their expected utilities. Moreover, the prioritarian, on Parfit's proposal, attaches prioritarian weightings to these prospects (as well as to outcomes): the lower a person's expected utility, the greater the moral importance of a marginal increase of a given size to his expected utility.¹² In other words, not only are the actual levels of a person's utility that are the different possible outcomes of a choice morally important in a manner that diminishes at the margin, but so is a person's *expected* level of utility that is associated with a given choice in prospect. One identifies a person's expected utility now, and his actual utility later, and assigns prioritarian weight to each. Since it combines a concern with outcomes with a concern with prospects, I shall call this Parfit's *hybrid version* of prioritarianism.¹³ Parfit motivates this hybrid view by noting that, just as non-prioritarians such as egalitarians are right to care about the equality of people's chances of receiving goods as well as the equality of the goods they end up with, prioritarians can and should have priority-weighted concern for people's prospects for utility as well as for their actual level of utility *ex post*.¹⁴

Hybrid prioritarianism yields less strong reason to move to the city in the one-child case as compared with the two-child case. Nevertheless, the hybrid prioritarian will maintain that the parent ought, all things considered, to move to the city in the one-child case, and it would be wrong to move to the suburb instead.¹⁵ Hybrid prioritarianism therefore reaches the same verdict regarding

¹⁰It will be just as strong so long as we assume that relative strength is determined by the ratio of the expected value of moving to the city to the expected value of moving to the suburb. This is because the expected value of a move to a city in the one-child case is simply half the expected value of such a move in the two-child case. The same goes for a move to the suburb. Hence the ratio of the expected value of a move to the city versus the suburb remains the same across the two cases.

¹¹In his earlier *Lindley Lecture*, Parfit does not address the question of how the priority view should be applied to cases involving risk. He addresses this question only in 'Another defence', which provides a reply to Voorhoeve's and my critique of the priority view in 'Why it matters'. This article offers a rejoinder to that reply. (See n. 18 below.)

¹²'Another defence', secs. VI–VII.

¹³Wlodek Rabinowicz attributes such a hybrid interpretation of prioritarianism to John Broome. See Rabinowicz, 'Prioritarianism for prospects', *Utilitas*, (2002), 2–21, at p. 10. See also Broome, *Weighing Goods* (Oxford: Blackwell, 1991), pp. 216–17.

¹⁴'Another defence', pp. 430–36.

¹⁵For a demonstration of these claims regarding the implications of hybrid prioritarianism in the one-child case, see an unpublished note, which the author will email on request.

what one ought to do—i.e., move to the city—as the simpler version of prioritarianism that assigns prioritarian weight to outcomes only.

Contrary to both versions of prioritarianism, the countervailing position I shall defend in this article is simply this: in cases along the lines of the one-child case, involving nothing other than trade-offs between the possible lives of a single person, you are permitted simply to maximize this person's expected utility, rather than instead giving any extra prioritarian weight to benefitting him, should he turn out badly off.¹⁶ Contrary to prioritarianism in either its hybrid or its simpler form, it is therefore permissible to opt for the suburb in the one-child case. The following is a preview of my defence of this claim.

Even if one maintains, as I do, that there are objective prudential goods, one should also maintain that ideally rational self-interested preferences will provide the measure of a person's well-being.¹⁷ Moreover, such preferences, being ideally rational, will conform to sound axioms of expected utility theory. It follows that you will provide this teenager with an alternative other than that which it is in his rational self-interest to prefer if you fail to maximize his expected utility. That fact provides you with strong reason to maximize his expected utility rather than giving any prioritarian weight to bad outcomes. Moreover, this reason is not decisively outweighed by any countervailing reason that either you or he has. There are, for example, no interpersonally comparative or otherwise distributive considerations here that tell in favour of paying heed to anything other than what is in this teenager's rational self-interest. In this case, unlike Nagel's two-child case, there are no distributive considerations that might arise from the competing interests of a second child. It is also legitimate for you, the parent, to consider the child's interests in isolation from those of others to whom you do not have special ties. Therefore, you are permitted to maximize your child's expected utility in this one-child case on grounds that this is what it would be rational for him to prefer, taking all relevant normative considerations into account. That is the core of my argument against prioritarianism in either its hybrid or its simpler form.¹⁸

¹⁶Cf. 'Why it matters', pp. 173–4, where Voorhoeve and I claimed, without argument, that such maximization of expected utility is the 'reasonable' course of action.

¹⁷Such a preference-based measure of utility does not imply the dubious claim that well-being is itself constituted by—or to be identified with the satisfaction of—preferences. It is implausible to maintain, along Hamlet's lines, that there is nothing either good or bad, but wanting makes it so. There are many things that are good or bad for someone, where wanting them is not what makes them so. Rather than any radical subjectivism about the constitution or identity of utility, the claim being advanced is the more measured one that preferences provide the measure of utility.

¹⁸This article was originally conceived as a rejoinder, on Voorhoeve's and my behalf, to Parfit's 'Another defence'. (See n. 11 above.) Although, as it is now conceived, the article presents a more general and self-standing line of objection to prioritarianism, it might be useful to sketch the way in which it still serves as such a rejoinder. In a nutshell, the 'crucial argumentative move' of Voorhoeve's and my critique of prioritarianism was as follows: even if, rather than maximizing his expected utility, one ought to give priority to the child if he will be badly off in the one-child case, one ought to give *greater* priority to the badly off child in the two-child case; moreover, the priority view, on its most natural and straightforward construal, cannot account for such a shift from lesser to greater prioritarian weighting. Parfit has now offered a version of the priority view—namely, the hybrid

II.

In setting up this argument, I have appealed to a case involving a parent and a single child. Such a case is useful because it filters out non-prioritarian interpersonally comparative considerations that muddy the waters. The special ties of the parent to this single child provide a justification for the exclusion of non-prioritarian considerations regarding how this child will fare in relation to others. They justify a narrowing of the focus to the interests and other claims of this one child alone, considered in isolation from the claims and interests of others.¹⁹ They therefore provide us with a purer test case of prioritarianism than one in which an agent is moved by both prioritarian and non-prioritarian considerations.

There is, however, an aspect of the parent-child relation that renders it problematic for these purposes. The parent is the guardian of his child until the age of majority. Since a guardian is someone who is specially entrusted to look after the interests of another person, a parent has a special reason to promote that child's self-interest. This particular reason will fail to generalize to other cases of prioritarian agents who are deciding whether they should provide risky benefits to single persons over whom they are not guardians. Hence, even if I am able to motivate the claim that the parent may maximize the expected utility of his one child, this permission might arise solely from his guardianship and fail to support a more general critique of prioritarianism as faulty on grounds of its inability to deal more generally with one-person cases involving risk.

A case that better suits my purposes would be one that involves special ties that serve to narrow the focus to a single individual, without also giving rise to special reasons to act on behalf of the interests of a particular individual that are tied to guardianship or other roles, such as that of an attorney who is charged to promote her client's interests.

So let us assume, in the discussion to follow, that you, who must make the decision regarding the move to the city or the suburb, are the child's *aunt* rather than his parent. This teenager stands in the unique special relation to you as your only niece or nephew, and he is dear to you. You are not, however, his guardian. The parent is the guardian. As it happens, that parent is also an employee of a firm of which you are a middle manager, but someone to whom you have no

version—that nicely accounts for such a shift. In order to defeat this version of the priority view, Voorhoeve and I must establish something stronger than the claim that one may give *lesser* prioritarian weight in the one-child case than in the two-child case. We must vindicate the stronger claim that one may permissibly give *no prioritarian weight whatsoever* in the one-child case—i.e., that it is permissible to maximize expected utility in that case. This article is an attempt to vindicate this stronger claim.

¹⁹Parfit acknowledges that the existence of special ties might provide grounds for such a narrow focus. He writes: 'Egalitarians would agree that, when we must choose between acts that would affect the well-being of certain people, we can often ignore the question whether there are other people who are worse off or better off. That may be true, for example, when we must choose between acts that would benefit our close relatives or friends, or other people to whom we stand in various special relations' ('Another defence', p. 407).

special ties of kinship or friendship. (Perhaps the child's parent is a step-parent whom you do not know personally.) The branch of the firm where the parent is employed has closed down, and you must decide whether to reassign the parent to another branch in a city or in a suburb. Insofar as the interests of the parent and of the firm are concerned, grounds for each of these reassignments are equally strong. Here it is appropriate for you to focus just on the manner in which this decision bears on the fate of your one and only nephew.

In this case, I would argue that, even though the priority view mandates that you do otherwise, you are permitted to reassign the parent in a manner that maximizes the expected utility of the teenager. As I shall argue in Section IV below, you are permitted to do so on grounds that this is what it is rational for him to prefer, taking all relevant normative considerations into account.

When I argue that you are permitted to do so on grounds that this is what it is rational for your nephew to prefer, I will not also ground your permission to do so in what he *actually* prefers. Grounds having to do with actual preferences provide very different sorts of reasons from grounds having to do with idealized rational preferences. We often have autonomy-based reasons to do what someone actually prefers, even when such preferences come apart from what it would be ideally rational for him to prefer. We can, for example, have such autonomy-based reasons not to treat somebody in paternalistic fashion.²⁰

Thus far, I have made no mention of what the teenager's actual preferences are, or what preferences he would sincerely express if asked. I shall now stipulate that whatever preferences the teenager actually has regarding such a move are unknown to you, his aunt.²¹ Moreover, you cannot ask the teenager to express his preference. This is because you must make a very quick decision to reassign the teenager's parent to the city or the suburb: if you do not act now, then another middle manager will fill both of these posts with equally meritorious employees from a different department, and the parent will thereby be out of a job, to your nephew's great detriment, relative to either of the available alternatives. Under this stipulation, the aforementioned considerations of autonomy do not tell in favour of the suburb over the city, or vice versa. They do not therefore tell in favour of or against what it would be rational for the teenager to prefer, which, as I shall explain in Section IV below, is a move to the suburb.

III.

Before I turn, in Section IV, to my argument in favour of the aunt's permission to maximize her nephew's expected utility, I shall reject an argument that Parfit

²⁰See James Griffin, *Well-Being* (Oxford: Oxford University Press, 1986), pp. 10–11.

²¹Here I follow Parfit's lead. See his 'Another defence', p. 424. In 'Why it matters', Voorhoeve and I assumed that the person's actual preferences were in conformity to what it would be ideally rational for him to prefer. Parfit's assumption is more helpful in this context, since it screens out considerations of autonomy.

offers on behalf of the contrary view that she is not permitted to maximize his expected utility. Parfit would maintain that she is instead required, for the following reason, to behave in prioritarian fashion by opting for the city:

When we have to make some decision on someone else's behalf, and we don't know how this person would prefer us to act, we may believe that we ought to be cautious, or risk averse. On this plausible and widely held view, it would be wrong to take, on this person's behalf, some of the risks that this person could rationally choose to take, if he or she made this choice for purely self-interested reasons. As Prioritarians claim, we ought to give more weight to avoiding possible outcomes in which this person would be worse off.²²

In this section, I shall explain why I find Parfit's reasoning on behalf of prioritarianism in the above passage unpersuasive.

In this passage, Parfit invokes something reminiscent of the familiar view that it may be wrong for you to take risks on behalf of another that it would be rational for you to choose for yourself. This familiar view offers sound advice across a range of cases in which one's own aversion to risk is relatively low, yet not so low as to be irrationally so. But the view to which Parfit actually appeals is that it may be wrong for you to take risks on behalf of another that it would be rational for *him* to choose *for himself*. This view is a good deal less plausible than the familiar view. It is more difficult to see why it might be wrong for you to act on someone else's behalf in a manner that would be rational for *him* to choose than it is to see why it might be wrong to act on someone else's behalf in a manner that would be rational for *you* to choose.

Even if we assume that one ought to be averse to taking risks on behalf of another, such an attitude does not necessarily imply a prioritarian rather than a utilitarian weighting of the prospects of another. Sometimes it implies precisely the opposite. Suppose, for example, that you come upon a stranger in the remote wilderness after he has just suffered an accident. You have great medical expertise, so you correctly diagnose that this person has an equal chance of regaining consciousness either in full health or very severely disabled. You can either do nothing or else administer a risky treatment. This treatment would benefit him by reducing the severity of his impairment in the event that he turns out disabled. But it also has the following regrettable side-effect: it would harm him by inducing a slightly impairing stroke in the event that he wakes up fully healthy. Suppose, moreover, that the gap in utility between full health and the slight impairment is slightly greater than the gap in utility between the very severe impairment and its somewhat alleviated state.

If you are a utilitarian, you will do nothing rather than administer this treatment, as doing nothing is what maximizes the person's expected utility. If,

²²'Another defence', p. 423. In this passage, Parfit is not addressing the particular example of the aunt and the nephew that I have introduced in this article. Rather, he is addressing a structurally analogous case involving a purely intrapersonal trade-off.

however, you are a prioritarian who accords greater moral weight to changes in utility lower down the absolute scale, you will administer this treatment that risks a 50 percent chance of causing a minor stroke rather than doing any good. That is what maximizes the person's expected priority-weighted utility.²³ In this case, an individual who is averse to taking risks on behalf of another will do what *utilitarianism* rather than prioritarianism recommends. If, as utilitarianism recommends in this case, you do nothing, you will no doubt leave this stranger exposed to the peril in which you found him in the wilderness. You will not, however, be *taking* any risk on his behalf if you leave him alone rather than administer the risky treatment. More generally, we do not take risks on behalf of strangers in peril by failing to come to their aid.²⁴

A prioritarian might grant this point but insist that what he is advocating is a more general risk-aversion in one's behaviour that bears on the fate of others, where such aversion extends to one's inaction as well as one's actions. Moreover, on the economists' standard measure, risk-aversion in one's general behaviour regarding the utility of another is an implication of prioritarianism. This is because a prioritarian moral weighting of utility is captured by a concave function, where the x-axis represents utility and the y-axis represents the morally weighted value of that utility. For any good, including utility itself, whose value is captured by such a concave function, it follows from the standard (Arrow-Pratt) measure of the degree to which someone has a positive or a negative attitude towards risk that such a valuation of that good manifests risk aversion.²⁵

I would question the claim that such Arrow-Pratt risk aversion regarding the utility of others captures a 'plausible and widely held view' regarding our duty to be cautious. Rather, it is more likely that intuitions in support of caution in one's behaviour towards others are largely tracking something other than risk-aversion in the economists' technical sense: they are instead tracking convictions regarding one's special responsibility for the consequences of the exercise of one's agency. Insofar as this is the case, Parfit will not be able to recruit such intuitions in favour of a prioritarian weighting of utility. This is because, as I have just illustrated with the wilderness example, such thoughts regarding one's special responsibility will sometimes move one to do what utilitarianism rather than prioritarianism calls for. Recall that, in that case, a concern that one not be the agent of something that might turn out badly will move one to behave as the

²³It is what maximizes the person's expected priority-weighted utility on either the hybrid or the simpler prioritarian approach to risk.

²⁴Things might be different if you were this person's attending doctor, entrusted to look after his well-being during his trek in the wilderness. In that case, by doing nothing rather than intervening, you might be taking a risk on his behalf.

²⁵See Kenneth J. Arrow, *Aspects of the Theory of Risk-Bearing* (Helsinki: Yrjö Jahnssonin Säätiö, 1965), lecture 2, and John Pratt, 'Risk aversion in the small and in the large', *Econometrica*, 32 (1964), 122–36.

utilitarian would recommend, whereas prioritarianism would direct one to administer the risky treatment rather than keeping one's hands off.²⁶

IV.

As I noted at the outset of the previous section, Parfit proposes that a third party is bound by a duty of caution or risk-aversion to make a prioritarian rather than a utilitarian choice regarding another person's fate. In the remainder of that section, I cast doubt upon this suggestion. In this section, I shall press the following more fundamental difficulty with Parfit's claim that one ought to behave in prioritarian rather than utilitarian fashion in such scenarios as the aunt's choice regarding the urban or the suburban fate of her teenage nephew. The difficulty is this: given the normative soundness of expected utility theory, what prioritarianism requires in such circumstances runs contrary to what it would be rational for the nephew to prefer.

On the classic and most familiar version of expected utility theory, which traces back to von Neumann and Morgenstern,²⁷ a cardinal scale of a person's utility is constructed out of her preferences over risky prospects that conform to axioms of completeness and transitivity, plus a continuity and an independence axiom.²⁸ All but the first of these axioms have been plausibly described as requirements of rationality. Moreover, there is a valid proof that a person whose preferences satisfy these four axioms has a utility function, an essential feature of

²⁶If, moreover, the reason to give priority to the worse off even in one-person cases, where that person's interests are the only ones that are morally relevant, has to do with risk-aversion, how can one explain such prioritarian weighting when there is no uncertainty involved in the benefit one might confer on that one child?

²⁷See von Neumann and Morgenstern, *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press, 1944). For an exegesis of the relevant passages, see Peter Fishburn, 'Retrospective on the utility theory of von Neumann and Morgenstern', *Journal of Risk and Uncertainty*, 2 (1989), 127–58.

²⁸Here is one formulation of the axioms (where P , P^* , and P^{**} stand for risky prospects whose probabilities are known):

- *Axiom 1 (Completeness)*: For any prospects P and P^* , either P^* is preferred to P , P is preferred to P^* , or one is indifferent between the two.
- *Axiom 2 (Transitivity)*: If prospect P^{**} is preferred to P^* or one is indifferent between the two, and P^* is preferred to P or one is indifferent between the two, then P^{**} is preferred to P or one is indifferent between the two.
- *Axiom 3 (Continuity)*: If prospect P^{**} is preferred to P^* or one is indifferent between the two, and P^* is preferred to P or one is indifferent between the two, then there exists some number p , where $0 \leq p \leq 1$, such that one is indifferent between P^* and $p \cdot P^{**} + (1 - p) \cdot P$.
- *Axiom 4 (Independence)*: For any prospects P and P^* , prospect P^* is preferred to P or one is indifferent between the two if and only if the following holds true for all numbers p , where $0 < p \leq 1$, and for any prospect P^{**} : $p \cdot P^* + (1 - p) \cdot P^{**}$ is preferred to $p \cdot P + (1 - p) \cdot P^{**}$, or one is indifferent between the two.

In this statement of the axioms, I follow Mark Machina, 'Expected utility hypothesis', *The New Palgrave Dictionary of Economics*, 2nd edn, ed. Steven N. Durlauf and Lawrence E. Blume (New York: Palgrave Macmillan, 2008), vol. III, pp. 130–38, at p. 134.

which is that the expected utility of one prospect is greater than that of another if and only if the former prospect is preferred to the latter.²⁹

The actual preferences of an individual will often fail to satisfy all four of these axioms. Even if they satisfy them all, it would not necessarily follow that the resulting von Neumann-Morgenstern utility function would provide a measure of the *well-being* of the individual in question, where ‘well-being’ is understood to capture that to which the prioritarian attaches moral weight: namely, how well the life of an individual is really going (or would go). In order for the utility function to provide a measure of ‘well-being’, so understood, we will need to restrict the content of the preferences to those that are appropriately self-interested rather than other-regarding (i.e., altruistic) or impersonal in nature.³⁰ It will also be necessary to idealize these preferences, since a person’s actual self-interested preferences might be misinformed by factual error or marred by various forms of irrationality even if they happen to conform to the von Neumann-Morgenstern axioms. In order to provide the measure of well-being, preferences will need to be idealized along the lines of ‘self-interested preferences that the individual would have after ideal deliberation while thinking clearly with full pertinent information regarding those preferences’.³¹ Throughout this article, I shall assume that the preferences that give rise to a person’s utility function are so-restricted in content and ideally rational in form along the lines just sketched.

I therefore depart from the assumption of many economists that the preferences that serve as the inputs of expected utility theory are unrestricted in content and revealed by what people would actually choose. Here I follow John Broome, who maintains that the ‘doctrinal texts’ and ‘best textbooks’ of economists define ‘utility’ as ‘*that which represents preferences*’, while also noting that “[u]tility” in this sense need not be confined to a representation of a person’s *actual* preferences. A function can also be called a utility function if it represents the preferences a person would have if she were rational and self-interested.³² When preferences are rational and self-interested in this respect, it is not by virtue of the economists’ arguably merely stipulative definition of ‘utility’ that the function representing these preferences provides a measure of the utility to which prioritarians attach weight. Rather, we can offer the following substantive grounds in support of the claim that the utility function provides a measure of the person’s utility that is of relevance to prioritarianism, which is to

²⁹See Machina, ‘Expected utility hypothesis’, and Broome, *Weighing Goods*, chs. 5 and 6.

³⁰For an excellent survey of the literature on the need to so-restrict preferences, see Matthew Adler’s discussion of the problem of ‘remoteness’ in *Well-Being and Fair Distribution* (New York: Oxford University Press, 2012), pp. 174–81. See also Parfit’s classic discussion of this problem in *Reasons and Persons* (Oxford: Oxford University Press, 1984), appendix I.

³¹The quoted words are from Richard Arneson, ‘Primary goods reconsidered’, *Noûs*, 24 (1990), 429–54, at p. 448.

³²Broome, ‘Utility’, *Economics and Philosophy*, 7 (1991), 1–12, at p. 10.

say her well-being or personal good: ‘But if she were rational and self-interested, she would prefer, of two alternatives, the one that is better for her. So a function that represents the preferences she would have if she were rational and self-interested also represents . . . her good.’³³

If we now also suppose, for the sake of argument, that the von Neumann-Morgenstern axioms are among the requirements of rationality, then prioritarianism immediately runs up against the following difficulty. Parfit concedes that prioritarianism would be in trouble if its dictates clashed with the nephew’s rational preferences, since it ‘would be implausible to claim that we ought to treat people in ways in which they could not rationally choose to be treated’. Parfit maintains, however, that the nephew is not rationally required to choose that the aunt maximize his expected utility by moving him to the suburb. Rather, according to Parfit, the nephew could instead rationally choose that his aunt act in a prioritarian manner by moving him to the city, where his grounds for such a choice are that ‘there is a one in two chance that [this move] would actually be better for [him], and this act would be better for [him] if he is very badly off’.³⁴ Note that here Parfit appeals to grounds that invoke nothing other than what is at stake for the nephew, considered in isolation from how well he would fare in comparison with others.³⁵ In other words, he appeals to the rational self-interest or prudence of such a choice.³⁶ But, given that the move to the suburb is of higher expected utility than a move to the city, we can infer from expected utility theory that it is in the nephew’s rational self-interest to choose a move to the suburb over the city. What it is in the nephew’s rational self-interest to choose and what would uniquely maximize his expected utility cannot come apart in the manner that Parfit suggests. This is because expected utility theory derives the magnitudes of a person’s expected utilities from her axiom-conforming self-interested preferences over risky prospects. If, therefore, it would be in the nephew’s rational self-interest to choose a move to the city over the suburb, on Parfit’s grounds that ‘there is a one in two chance that [this move] would actually be better for [him], and this act would be better for [him] if he is very badly off’, then it follows that such a move to the city would maximize his expected utility.³⁷

³³Ibid., pp. 10–11.

³⁴‘Another defence’, pp. 424–5. Here I apply what Parfit says about a similar case to the case of the nephew and his aunt.

³⁵If, moreover, Parfit were to ground the rationality of the nephew’s choice of the city in considerations to do with how well the nephew would fare in comparison with others, this would undermine his claim to be providing non-egalitarian and otherwise non-comparative grounds for giving priority to the worse off.

³⁶He does not ground the rationality of the nephew’s choice of the city in an appeal to the *impersonal goodness* of such a move. Moreover, for reasons I offer in Section VI below, such grounds are unpromising.

³⁷Parfit might maintain that it would be in the nephew’s rational self-interest to choose *either* the city or the suburb. But if the nephew’s rational preferences conform to the axioms of expected utility theory, it would follow that he is indifferent between these two options, and hence either choice would maximize his expected utility. This would run contrary to the stipulation of the example that the move to the suburb would uniquely maximize the nephew’s expected utility.

Let us now add, to the assumptions of the previous paragraph, the further supposition that the aunt's own utility function is identical to her nephew's. This supposition is legitimate and provides a good test of the priority view, for the following reason. The prioritarian case for being more cautious than the person on whose behalf one is choosing cannot depend on any differences in risk aversion between the chooser and the potential beneficiary of the choice, nor can it depend on any other differences in their preferences that determine their utility functions. This is because prioritarianism is the view that, whatever one's own utility function as compared with the utility functions of others, one ought to accord more moral weight to the utility of others the less well off they are in absolute terms. Prioritarianism therefore obtains even on the supposition that everyone's utility functions are identical.

Given this supposition, not only would it be in the nephew's rational self-interest to prefer a move to the suburb over the city for himself, but it would be equally rational for his aunt to prefer a move to the suburb over the city for herself if she were in his predicament. Moreover, it would be rational for the nephew to prefer that his aunt make that choice on his behalf, and it would be equally rational for his aunt to prefer that her nephew make that choice on her behalf if she were in his predicament and he in hers. Since all of these preferences are idealized, rational preference over risky prospects, they encode a rational degree of risk aversion regarding the benefits of the city versus the suburb in each case. In light of all this, it is now very difficult to see how it could also be wrong, as a prioritarian would maintain, rather than permissible, for the aunt to choose the suburb over the city for her nephew.

The difficulties for prioritarianism to which I have been pointing have rested upon the assumption that the von Neumann-Morgenstern axioms are requirements of rationality. It is, I grant, a matter of live controversy whether or not all of the axioms are such requirements. As I suggested earlier, few would maintain that the completeness axiom is such a requirement. In other words, for any given pair of options, it does not appear to be rationally mandatory that an individual either prefer one to the other or be indifferent between the two. The phenomenon of plural values that are not precisely commensurable is the most common explanation for why a rational person's preferences need not be complete.³⁸ The axioms of transitivity and continuity have stronger claims than the axiom of completeness to be requirements of rationality. Even in the case of these two axioms, however, challenging counterexamples have been raised against such claims.³⁹ Nevertheless, for the purposes of my critique of prioritarianism, I am entitled, for reasons I shall spell out in the next two paragraphs, to assume the soundness of the axioms of completeness, transitivity, and continuity.

³⁸See, for example, Broome, *Weighing Goods*, pp. 92–3.

³⁹See, for example, Larry Temkin, *Rethinking the Good* (New York: Oxford University Press, 2012), chs. 6–8.

The diminishing marginal moral value of utility that prioritarrians affirm is standardly formally represented as a continuous, strictly increasing, concave transformation of the utilities of individuals. These utilities are standardly formally represented as real numbers on a continuous, linear, cardinal scale. Such a representation is, no doubt, a simplifying idealization that supplies more precision than is actually warranted by the subject matter.⁴⁰ A prioritarian would, however, be hard pressed to maintain that the priority view would be inapplicable under these idealized assumptions that are invoked in order to illustrate the view. It would, for example, be odd to maintain that prioritarianism would not apply if there were such things as Benthamite hedonometers that issued accurate and precise cardinal readings of our levels of utility on a linear scale. The prioritarian therefore lacks grounds to deny the legitimacy of the assumption, made for the sake of assessing the soundness of his doctrine, that utility can accurately be represented in such precise fashion.

If, however, one grants this assumption that the utilities of individuals are accurately represented in such fashion as particular numbers on a continuous cardinal scale, then grounds for denying the rationality of the completeness, transitivity, and continuity of preferences regarding one's self-interest would be lacking. Such a cardinal representation is, for example, incompatible with the phenomenon of imprecise commensurability to which people appeal in order to justify the incompleteness of even rationally idealized preferences regarding one's self-interest. If the prudential value of any option can be represented as a number on a cardinal scale, then any one option will be precisely commensurable with any other. Presumably, one's ideally rationally self-interested preferences will also be such that one will rationally prefer one option to another if and only if the former option is at a point higher on this cardinal scale than the latter. If, however, this is the case, then such preferences could not fail to be transitive. The continuous nature of the cardinal scale of utility would also provide grounds for affirming the continuity axiom.

The soundness of the independence axiom cannot, however, be assumed on similar grounds. This is because reasons to reject this axiom are not undermined by the assumption that utility can be represented in precise cardinal fashion. Moreover, given its bearing on the issue of the rationality of attitudes towards risk, the independence axiom is more directly relevant than the other axioms to the adjudication between prioritarianism and utilitarianism in one-person cases involving risks. A prioritarian might, therefore, maintain that any argument in favour of the maximization of expected utility in such one-person cases that assumes this axiom begs the question against his view. A prioritarian might also offer a more direct challenge to the claim that the independence axiom is a requirement of rationality.

⁴⁰See Parfit, 'Equality or priority?' p. 83.

The independence axiom implies that preferences are ‘linear in the probabilities’. In other words, for any increase (or decrease) in a given magnitude of the probability of receiving a good, one will have the same preference for such an increase (decrease) however high or low the baseline probability. For example, an increase of 1 percent in one’s probability of receiving a good will be preferred to the same extent whether the increase is from a baseline of no chance to a 1 percent chance of receiving that good, from a baseline of a 50 percent chance to a 51 percent chance, or from a baseline of a 99 percent chance to a certainty of receiving that good.

The rational necessity of preferences that are linear in the probabilities has been called into question, most notably by Maurice Allais, who famously elicited apparently independence-axiom-violating preferences over a pair of gambles from the decision theorist Leonard Savage. Allais defended Savage’s preferences as rational. Savage’s response was to maintain that Allais had induced an irrational, because independence-axiom-violating, pair of preferences which ought to be disavowed upon reflection.⁴¹ Others since Savage, such as Broome, have offered powerful defences, against Allais, of the independence axiom as a requirement of rationality.⁴² I share their view that this axiom is a requirement of rationality.

Even if, however, one instead sides with Allais over the likes of Savage and Broome regarding the rationality of preferences that fail to be linear in the probabilities, it does not follow that one must reject expected utility theory altogether. There is the option of a more modest response, which rejects only the particular version of expected utility theory that corresponds to the von Neumann-Morgenstern axioms. There have, in fact, been fruitful generalizations or revisions of expected utility theory that relax or replace their independence axiom in a manner that accommodates the preferences over gambles that Allais elicited from Savage. Some of these departures from this axiom have been defended as consistent with rationality. One well-regarded departure involves a generalization of von Neumann and Morgenstern’s theory that has come to be known as rank-dependent expected utility theory.⁴³ Although this theory relaxes the assumption of linearity in the probabilities, it shares, with the classical expected utility theory of von Neumann and Morgenstern, the following crucial

⁴¹See Maurice Allais, ‘Le comportement de l’homme rationnel devant le risque’, *Econometrica*, 21 (1953), 503–46, and Leonard Savage, ‘The Allais paradox’, *Decision, Probability and Utility*, ed. Peter Gärdenfors and Nils-Eric Sahlin (Cambridge: Cambridge University Press, 1988), pp. 163–5.

⁴²Broome argues that attempts to show that the preferences that Allais elicited from Savage are, in fact, rational involve the drawing of a distinction (e.g., regarding the regret that would be rational) between losing outcomes across the two gambles that cannot, in fact, be so-distinguished if the preferences in question are to serve as counterexamples to the independence axiom. Broome concludes: ‘Plainly, therefore, the case against the [independence axiom] is absurd. It depends on making a distinction on the one hand and denying it on the other’ (*Weighing Goods*, p. 107).

⁴³See John Quiggin, ‘A theory of anticipated utility’, *Journal of Economic Behavior and Organization*, 3 (1982), 323–43.

feature: the magnitude of utility is a construct of preferences over risky prospects. It is precisely this feature that has been the linchpin of my argument in this section for the permissibility of departing from what prioritarianism requires in the one-person risk-involving case under discussion. An analogous version of my argument against prioritarianism will go through if one replaces the axioms of von Neumann and Morgenstern with the axioms of the Allais-preference-accommodating rank-dependent generalization of their theory. One can construct a one-person case, along parallel lines to my aunt/nephew case, in which the risky prospect of a move to the suburb—that it would be in the rational self-interest for someone who conforms to these rank-dependent axioms to prefer—will be at odds with the requirements of prioritarianism. It will be at odds for the same reason as before: namely, that prioritarianism calls for a concave transformation of the moral weights of utilities whose magnitudes have been fixed by rational trade-offs among risky prospects. In some cases, such a transformation will yield a prioritarian duty for a third party to act in a manner that is contrary to what it would be in the rational self-interest for the individual to prefer. I have argued, however, that it is permissible for a third party to choose in accord with such preference.

As I noted at the outset of this article, Parfit prescind, in his defence of prioritarianism, from ‘difficult questions . . . about what it would be for some benefits to be greater than others’.⁴⁴ I hope to have now shown that such questions cannot be ignored. Rather, the soundness of the priority view turns on their answer. It is a striking feature of this view that, for the reasons I have just sketched in this section, it cannot accommodate one of the most fruitful and powerful constructs of decision and rational choice theory: expected utility theory.⁴⁵ In order to justify the application of prioritarian weightings to cases involving trade-offs between the possible lives of a single person, of the sort that has been the focus of this article, Parfit must reject a normatively compelling von Neumann-Morgenstern measure of utility along with the alternative rank-dependent expected utility theory that has been offered as an antidote to doubts regarding the normative soundness of this measure. He must supply a different measure of utility altogether, which is not a construct of preferences over risky prospects that conform to axioms of rational choice.

It is plausible to maintain that any such alternative measure will require trade-offs of some sort in order to generate the cardinal scale of utility that

⁴⁴‘Another defence’, p. 403.

⁴⁵As Temkin writes: ‘Expected Utility Theory is an enormously attractive theory that underlies game theory, decision theory, and much of modern economics. Although Expected Utility Theory has been the subject of much scrutiny and criticism, the power and successes of the theories relying on it give us good reason to believe that Expected Utility Theory is *essentially* correct, even if it requires some “tinkering with” to handle certain problems to which it gives rise’ (*Rethinking the Good*, p. 232).

prioritarianism presupposes.⁴⁶ A method of ‘time trade-offs’ has been adopted by a number of health economists as an alternative to expected utility theory’s construction of a cardinal scale of utility via trade-offs among risky prospects. On the time trade-off method, a cardinal scale of utility is constructed from people’s trade-offs between numbers of years of additional life in full health versus larger numbers of years that are constantly blighted by particular illnesses or disabilities. Since here the measure of utility is derived from trade-offs among outcomes of fixed lengths that are certain to come about rather than among risky prospects, the time trade-off method will not necessarily yield a conflict between prudential rationality and the mandates of prioritarianism in one-person cases involving risk. In its standard form, however, the time trade-off method rests on the controversial assumption that, for any given health state, extra years of life in that condition are of constant marginal utility. This assumption is both descriptively inaccurate and normatively suspect. It is more plausible to assume that the marginal utility of extra life years will eventually diminish. Moreover, the standard method for determining the rate at which such utility diminishes involves an application of expected utility theory via the eliciting of preferences over risky prospects in which one is not always certain how many more years one will live.⁴⁷ Hence, the time trade-off method, in its most plausible form, does not provide a clean alternative to expected utility theory after all.

V.

I have argued in the previous section that von Neumann-Morgenstern expected utility theory cannot provide the measure of well-being to which Parfit attaches prioritarian moral weights. A defender of Parfit might grant the soundness of this argument. He might then reply, on Parfit’s behalf, that it doesn’t follow, from this argument, that Parfit must reject the very axioms of that theory. Rather, it is open to Parfit to embrace the view that the axioms are an accurate statement of conditions of the rationality of preferences, while at the same time denying that rational preferences provide the measure of well-being, even when the preferences in question are restricted to those that are self-interested. The so-called ‘utility function’ of expected utility theory might simply provide a cardinal measure of the strength of the individual’s rational self-interested preferences rather than a measure of her well-being. Parfit can, therefore, escape my critique. He can

⁴⁶Cf. Broome: ‘I cannot see what use we can have for the notion of quantities of good except when we weight differences in good in comparing alternatives. So it is in weighing up differences that we can expect the notion to get its meaning’ (*Weighing Goods*, p. 215).

⁴⁷See Sylvie M. C. van Osch, et al., ‘Correcting biases in standard gamble and time tradeoff utilities’, *Medical Decision Making*, 24 (2004), 511–17; A. M. Stiggelbout, et al., ‘Utility assessment in cancer patients: adjustment of time tradeoff scores for the utility of life years and comparison with standard gamble scores’, *Medical Decision Making*, 14 (1994), 82–90; and J. M. Miyamoto and S. A. Eraker, ‘Parameter estimates for a QALY utility model’, *Medical Decision Making*, 5 (1985), 191–213.

simply reject the von Neumann-Morgenstern measure of well-being at no great cost, since he needn't incur the cost of denying that the axioms are axioms of rational preference.⁴⁸

Such a decoupling of a person's well-being from what it would be in her rational self-interest to prefer is not, however, available to Parfit, since this would be at odds with the account of reasons, rationality, and the good that he offers and defends in *On What Matters*. Central to this account is the 'reason-implying sense' of the good on which he builds the main arguments of his book. In applying this notion of the good to a person's well-being, Parfit maintains that when 'we call some possible life "best for someone" in the *reason-implying* sense, we mean that this is the life that this person would have the strongest self-interested reasons to want to live'.⁴⁹ Moreover, Parfit draws a tight connection between rational preferences and reasons. In a section entitled 'Irrational Preferences' he writes: 'Our desires are rational . . . when we want events whose features give us reasons to want them. Our desires are not rational . . . when we want some event that we have reasons *not* to want, and no reasons, or only weaker reasons, to want.'⁵⁰ This combination of commitments closes off the escape route of the previous paragraph.

Even if we set Parfit's own commitments to one side, there remain good independent reasons to affirm a tight connection between rational self-interested preference and the measure of a person's well-being. On behalf of a such a connection, James Griffin maintains that a person's rational and informed preferences regarding different lives that she might lead will typically be 'quantitatively basic: that is, they are judgments that do not depend upon other judgments about the amount of some quantity each option has'. Rather, these preferences form the 'raw materials' out of which 'the construction of a scale of measurement of well-being begins'.⁵¹ This is so, Griffin contends, because there is no pre-existing, preference-independent common measure of the plurality of very different objective prudential goods that contribute to a person's well-being.⁵²

While I agree with Griffin that the fact of value pluralism supports the claim that idealized preferences provide the measure of well-being, I do not think an assumption of value pluralism is essential to the case for a preference-based measure. Even if the contribution to one's well-being of only a single type of prudentially good or bad thing were at issue, the cardinal measure of its goodness or badness would, I think, nevertheless still be provided by idealized preferences.

⁴⁸I am indebted to an anonymous referee for suggesting this line of reply on Parfit's behalf.

⁴⁹Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), vol. 1, p. 102.

⁵⁰*Ibid.*, p. 56.

⁵¹*Well-Being*, p. 103.

⁵²See *Well-Being*, pp. 30–31. It is worth noting that Griffin combines an objectivism regarding the nature of the prudential goods that constitute well-being with a preference-based measure of the magnitude of well-being.

Suppose, for example, that the contribution to one's well-being of a single type of physical pain were the only thing at issue. There is still the further question of how to determine the cardinal measure of the disutility of this pain. It would be a mistake to insist on a cardinal measure that tracks something quantifiable in purely physical terms, such as the number of C fibres that are firing. The firing of twice as many C fibres might not give rise to suffering that is experienced as twice as great. Moreover, even in the case of token pains of the exact same fine-grained qualitative type, e.g., equally painful token headaches or toothaches of the same type insofar as they involve the same felt qualities, intensities, and durations of pain—these token pains might be of greater or lesser disutility depending on how badly or well one's life is going. This phenomenon would be explained by the fact that it might be rational, in attending to one's own self-interest, to prefer relief of such pain if one's life is going badly to its relief if it is going well, even if we assume that these pains are equally bad in their further causal effects.⁵³ It would not, therefore, follow from the assumption that pain is the only prudential bad at issue, that equal amounts of pain whose further effects are equalized must be equally detrimental to one's well-being. Rather, fixed quantities of pain might have diminishing marginal disutility, relative to the increasing absolute level of well-being of the person who experiences it. We would need, here, to draw a distinction between quantities of pain and the magnitude of their disutility, in order to avoid lapsing into the incoherence of claiming that the disutility of pain has diminishing marginal disutility.

These remarks regarding the preference-based measure of the disutility of pain are pertinent to the following consideration that Parfit advances on behalf of prioritarianism. Parfit maintains that prioritarianism gains support from the fact that we have greater reason to alleviate someone's suffering if that person is otherwise badly off than we have to alleviate that person's comparable suffering if she is otherwise well-off. As an instance of the latter, he asks us to imagine a case of someone whose 'great pain, caused by the freezing wind on the summit of the mountain that she has just climbed[,] . . . is outweighed by her sense of achievement, and by her seeing the sublime view'. Parfit would maintain that a moral agent would have more reason to alleviate that person's equally great pain if her life were otherwise going badly rather than well.⁵⁴

The rational self-interested preferences of an individual that provide the measure of her utility will, however, already have taken into account the importance of receiving relief from a given magnitude of pain if one is well off or badly off. This is something that registers from the first person perspective of one's self-interest as well as the third person perspective of a benevolent,

⁵³Even if, in other words, we equalize by bracketing various further effects, such as self-pitying thoughts of 'woe is me' that might well up if someone whose life is going very badly develops a headache or a toothache on top of everything else that is going wrong.

⁵⁴'Another defence', pp. 420–21.

morally-motivated agent. I noted in my discussion of tokens of the same type of pain that it might be rational for a person to prefer relief from a fixed quantity of physical pain if his life is going badly, over relief from that same quantity of pain if his life is going well. The self-interested preferences it would be rational for the nephew to have, therefore, will already have taken into account the importance of receiving goods or bads if he is badly off. So there would appear to be no call for the aunt to weight relief of the nephew's suffering in the event that he is badly off even further, in prioritarian fashion. That would be double counting. The superfluousness of such double-counting provides further support for my claim that it would be permissible for the aunt to maximize her nephew's expected utility by opting for the suburb rather than choosing the city in accordance with the mandates of prioritarianism.⁵⁵

VI.

In arguing that the aunt is permitted to maximize her nephew's expected utility rather than act in prioritarian fashion, I have been assuming that the rationality of her nephew's preference for this course of action is grounded in nothing other than his self-regarding preferences, were he to deliberate about his stake in the matter in ideal conditions. Among other things, I have set to one side considerations regarding how well the nephew would fare in comparison with others. One further possibility I have ignored is that the prioritarian case for moving the nephew to the city over the suburb might be grounded in something other than the interests or claims of the nephew or of anyone else—that it might instead be grounded in purely impersonal values. Prioritarianism might be construed as a form of consequentialism according to which the maximization of priority-weighted utility is justified via appeal to the claim that this is what is best from an impersonal point of view rather than from the standpoint of the claims or interests of any particular moral agents. Such a consequentialist would not be dissuaded by the fact that the view mandates a choice that is irrational from the point of view of the self-regarding preferences of the only person—the nephew—who has a stake in the matter. Rather, he would maintain that such a choice is justified on the impersonal grounds of the expectation that the world

⁵⁵Even if the rational self-interested preferences of individuals already take into account the importance of receiving goods in the event that one is badly off or well off, there remains reason to act in prioritarian rather than utilitarian fashion in Nagel's two-child case with which I opened this article. This contrast, between the one-child and the two-child case, is explained by the fact that it matters whether or not there are others with competing claims to benefit. In the two-child case, one can ask the following rhetorical question: 'How can one justify providing a benefit of a given size to someone who is already better off in order to make him better off still, when one could instead provide nearly as large a benefit to someone else who is worse off, and who would not even reach the (unimproved) level of the better off person if she (the worse off person) is benefited?' An analogous complaint cannot be formulated against the maximization of expected utility in the one-child case. (See 'Why it matters', p. 184.) I shall return to this theme in the concluding section of this article.

would end up a better place if the nephew were moved to the city, in spite of the fact it would be rational, from the perspective of his own interests and claims, for him to reject such a move.

I have ignored the possibility of such a purely impersonal case for prioritarianism for the following three reasons. First, unlike cases involving such things as complex ecosystems devoid of persons or unobservable stretches of wilderness of outstanding natural beauty, it is hard to identify any traces of purely impersonal value in the case of the nephew and his aunt—much less an impersonal value that is sufficient to override the rational self-regarding preferences of the only person whose interests are at stake. Second, as I noted in Section IV above, Parfit's own defence of the prioritarian choice of the city appeals to nothing other than what is at stake for the nephew: i.e., that 'there is a one in two chance that [this move] would actually be better for [him], and this act would be better for [him] if he is very badly off'. Third, and finally, prioritarianism gains its persuasive force because and insofar as it is grounded in a concern for what is at stake in the lives of persons. In advancing the case for prioritarianism over egalitarianism in his *Lindley Lecture*, for example, Parfit approvingly quotes the following passage from Joseph Raz:

what makes us care about various inequalities is not the inequality but the concern identified by the underlying principle. It is the hunger of the hungry, the need of the needy, the suffering of the ill, and so on. The fact that they are worse off in the relevant respect than their neighbours is relevant. But it is relevant not as an independent evil of inequality. Its relevance is in showing that their hunger is greater, their need more pressing, their suffering more hurtful, and therefore our concern for the hungry, the needy, the suffering, and not our concern for equality, makes us give them priority.⁵⁶

It is when the focus is on a concern for persons and the pressing nature of their needs and other interests that the case for prioritarianism is strongest. The doctrine loses its appeal when it takes an impersonal form that abstracts from such concerns.

VII.

A principal aim of Parfit's elaboration and defence of prioritarianism is to show that various sound intuitions about cases that lead many to infer egalitarianism can, in fact, equally well be captured by a prioritarianism that is fundamentally non-comparative in its moral concerns. His hybrid view retains this distinctive feature of prioritarianism to which Parfit drew attention in his *Lindley Lecture*: one does not need to know how well off *others* will or might be in order to determine the strength of one's reason to benefit a given individual, because it

⁵⁶Parfit, 'Equality or priority?' pp. 123–4, n. 35, quoting from Raz's *The Morality of Freedom* (Oxford: Oxford University Press, 1986), p. 240.

does not matter, fundamentally, how well off individuals are in relation to others.⁵⁷ The strength of one's moral reason to provide someone a benefit can be inferred simply from facts about his absolute level of well-being—in prospect as well as in fact, in the case of the hybrid view. Hence, in its hybrid form, prioritarianism remains an alternative to egalitarianism and other comparative approaches to distributive ethics. My charge in this article against prioritarianism, either in its simple or its hybrid form, has been that it does not provide a sound alternative to approaches that register the moral significance of interpersonally comparative considerations. This is because both versions stretch credulity too far in their approach to one-person cases, of the sort that has been the focus of this article, involving purely intrapersonal trade-offs.

Contrary to Parfit, in such a one-person case, one may be utilitarian rather than giving priority to that person in the event that he turns out badly off. One should, by contrast, give varying degrees of priority to a person if he turns out worse off in other cases involving the competing claims of different individuals, such as the two-child case with which I began. This contrast provides excellent grounds for the claim that it is not simply on account of their absolute level—either expected or actual, or a combination of the two—that we weight benefits to the badly off more heavily than benefits to the well off. Rather, it is only when and because some are, or will be, or might be, worse off than *others* that we must accord more weight to benefits to them. We should give priority to the worse off because it matters that some are worse off than others, and not simply because of their absolute level. Therefore, prioritarianism should not take priority over interpersonally comparative approaches.⁵⁸

⁵⁷See 'Equality or priority?' pp. 103–5.

⁵⁸There is a version of prioritarianism that is insensitive to interpersonally comparative considerations and also gets the right answer in one-person cases: an 'ex ante' version that applies prioritarian weighting to nothing other than prospects (i.e., expected utility). Unlike both the simple and the hybrid version of prioritarianism, the ex ante version accords no weight to outcomes as distinct from prospects. On this version, one should always maximize expected utility in the sort of one-person cases involving risk under discussion in this article. The problem with this ex ante version is that it fails to arrive at the right answer in certain multi-person cases involving competing claims. See 'Why it matters', sec. X. See also Michael Otsuka, 'Prioritarianism and the separateness of persons', *Utilitas*, 24 (2012), 365–80, at pp. 375–80.