

Prioritarianism and the Separateness of Persons

MICHAEL OTSUKA

University College London

For a prioritarian by contrast to a utilitarian, whether a certain quantity of utility falls within the boundary of one person's life or another's makes the following moral difference: the worse the life of a person who could receive a given benefit, the stronger moral reason we have to confer this benefit on this person. It would seem, therefore, that prioritarianism succeeds, where utilitarianism fails, to 'take seriously the distinction between persons'. Yet I show that, contrary to these appearances, prioritarianism fails, in ways strikingly parallel to those in which utilitarianism fails, to take this distinction seriously. In so doing, I draw on and develop an earlier critique of prioritarianism by disentangling and pressing two distinct separateness-of-persons objections offered there. One objection is that prioritarianism is insensitive to 'prudential justifications'. The other is that it is insensitive to the competing claims of different individuals.

I

Among the most philosophically interesting objections to utilitarianism is the charge that it 'does not take seriously the distinction between persons'.¹ Rawls, for example, says that it is a 'striking feature' of classical utilitarianism that 'it does not matter...how [the] sum of satisfactions is distributed among individuals any more than it matters...how one man distributes his satisfactions over time'.² If the only way to maximize utility is to make charitable donations to Nozick's utility monster who is already very well off rather than to those whose utility is low, then so be it, as all that ultimately matters is that the overall amount of utility had by all be as great as it possibly can be, irrespective of who experiences such utility.³

Prioritarians, by contrast to utilitarians, care about more than the sum total of utility across all sentient beings. They do not share the utilitarian's indifference to how well the lives of different individuals go. Whether a certain quantity of utility falls within the boundary

I have presented versions of this article at the British Society for Ethical Theory, UCL, the Hebrew University, Balliol College, Oxford, and the Universities of Bristol, Manchester, and Vermont. I would like to thank the members of these audiences for their comments. I would also like to thank Alex Voorhoeve for extensive discussion of the claims and arguments I advance in this article and Tyler Doggett, Peter Graham, Peter Vallentyne, Andrew Williams and Gabriel Wollner for their written comments.

¹ John Rawls, *A Theory of Justice*, rev. edn. (Cambridge, Mass., 1999), p. 27.

² Rawls, *A Theory of Justice*, p. 26.

³ See Robert Nozick, *Anarchy, State, and Utopia* (New York, 1974), p. 41. Nozick presses a separateness-of-persons objection to utilitarianism (pp. 32–3).

one person's life or another's makes the following moral difference: the worse the life of a person who could receive a given benefit, the stronger moral reason we have to confer this benefit on this person.⁴ It would seem, therefore, that prioritarianism succeeds, where utilitarianism fails, to take seriously the distinction between persons.

In this article, I shall argue that, contrary to these appearances, prioritarianism fails, in ways strikingly parallel to those in which utilitarianism fails, to take seriously the distinction between persons. In so doing, I shall draw on and develop an earlier critique of prioritarianism.⁵ In particular, I shall disentangle two distinct separateness-of-persons objections to prioritarianism that Alex Voorhoeve and I offered in that critique. As I shall explain below, one of these objections is that prioritarianism is insensitive to the presence or absence of what I now call 'prudential justifications'. I shall consider an apparently plausible and well-motivated revision the prioritarian might make in the light of this objection, and I shall show how, even so revised, prioritarianism falls to a second, distinct separateness-of-persons objection – namely, that it is insensitive to the presence or absence of competing claims of different individuals. I shall also show how each of these two objections is distinct from the objection that prioritarianism is insensitive to the unfairness of inequality between persons.

II

To illustrate Rawls's aforementioned complaint against utilitarianism, consider an example in which the government taxes some of the earnings of a young adult and invests all of the revenue in a pension fund from which this same person will be able to draw greater benefits in old age.⁶ Now contrast this with a case in which the government taxes some of the earnings of this young adult in order to channel it into a pension fund from which *someone else* will be able to draw in old age. We can suppose that this other person is identically situated, save for the fact that he himself is not taxed. For a utilitarian, if this other person would benefit just as much as the one person from this pension fund, it makes no difference to whom this money is channelled. It is, however, intuitively, much harder to justify diversion

⁴ See Derek Parfit, 'Equality or Priority?', *The Lindley Lecture* (Lawrence, Kan., 1991). Henceforth, all references to 'prioritarianism' are to the 'pure' version of this view that Parfit presents and considers in this lecture.

⁵ Michael Otsuka and Alex Voorhoeve, 'Why It Matters That Some Are Worse Off Than Others', *Philosophy and Public Affairs* 37 (2009), pp. 171–99.

⁶ This example is inspired by some remarks of Nozick's in *Anarchy, State, and Utopia*, p. 33.

of the young adult's earnings to this other person's pension fund than it is to justify diversion of his earnings into his own pension fund.⁷

As it happens, prioritarianism is able to track this difference. This is because the person who is taxed during his youth without receiving a benefit in old age will lead a worse life in absolute terms than the otherwise identical person who is benefited in old age without being taxed during his youth. The life of a person who is both taxed and benefited will fall in between these two points. Moreover, a prioritarian will maintain that the badness of a person's welfare's falling any increment below this point in between exceeds the goodness of someone's welfare's rising the same increment above this point. Hence, it will be worse from a prioritarian point of view to divert the one person's earnings to another's pension than it would be to divert that person's earnings to his own pension.

So far so good for the prioritarian.

III

But now let us consider a different pair of cases that are along the lines of cases that feature in 'Why It Matters':⁸

- (1) *One-person case involving risk*: A single person will develop either a slight impairment or instead a very severe impairment, the impairments are equally likely, and a morally motivated stranger must now decide whether to offer him treatment for the slight impairment or instead for the very severe impairment. Suppose that the treatment for the slight impairment would, if effective, do somewhat more good than the treatment for the very severe impairment.
- (2) *Two-person case involving certainty*: One person will develop the slight impairment and the other will develop the very severe impairment, it is already known who will develop which impairment, and a morally motivated stranger can either offer the one person treatment for the slight impairment or offer the other person treatment for the very severe impairment. Again, an effective treatment for the slight impairment would do somewhat more good than an effective treatment for the very severe impairment.

⁷ This is true even though the diversion into his own pension fund might be regarded as paternalistic.

⁸ See section I and pp. 181–2 of that article.

As was true of the previous pair of cases involving pensions, these two cases should be treated differently. In the single-person case involving treatment for impairment (case 1), we can provide the following prudential justification to that person for having provided him with the treatment for the slight impairment even if he turns out to suffer from the very severe impairment: we could tell him that we were looking after *his* interest in receiving the significant benefit of a cure in the event that he turned out to have the slight impairment. We were taking a gamble on *his* behalf, on grounds that taking such a gamble maximized his expected utility.⁹ If we provide treatment for the slight impairment in our two-person case (case 2), no comparable prudential justification can be provided to the person who will develop the very severe impairment. This is because the administration of treatment for the slight impairment was never in his interest, as only a different person could ever possibly have benefited from that treatment.

Prioritarianism, however, treats these two cases as on a par.¹⁰ In weighting benefits to people whose lives go badly overall more highly than benefits to people whose lives go well (or only a slightly badly) overall, it does not draw any distinction between cases in which, as in our two-person case (case 2), it is different people whose lives will go badly or well, and in which, as in our one-person case (case 1), it is a single person whose life might go badly or might instead go well. Unlike utilitarianism, prioritarianism is sensitive to how lives go overall. Moreover, prioritarianism gives more weight to benefiting a person if his life will go badly than it does to benefiting a person if his life will go well. But prioritarianism fails to take seriously the distinction between persons in the following respect: it is insensitive to whether the life that goes well and the life that goes badly are possible lives of the same person or rather the lives of different people. When determining what we owe to others, prioritarianism is therefore insensitive to the presence or absence of a prudential justification for a given course of action. All that ultimately matters is how great a sum of priority-weighted utility this or that course of action yields. It is of no further significance whether or not those who lose out when such priority-weighted utility is maximized were compensated for this sacrifice by any prospect of gain.

⁹ See 'Why It Matters', p. 188.

¹⁰ At least, it does so on a standard interpretation, on which the actual rather than the expected lifetime utility of individuals is that which receives prioritarian weight. In section VIII, I consider a non-standard interpretation of the view, which attaches prioritarian weight to expected rather than actual utility.

IV

In the above two-person case involving certainty (case 2), there are distinct persons, one of whose lives goes less well than the other's through no choice or fault of his. Luck egalitarians would condemn this inequality between different persons as bad because unfair. There is no such badness in the above one-person case involving risk (case 1), as there is only one person, who is not worse off or better off than anyone else. The complaint that prioritarianism is insensitive to the unfairness of inequality between persons is distinct from our charge that prioritarianism ignores the separateness of persons on account of its insensitivity to the presence or absence of a prudential justification. To illustrate the distinctness of these two complaints, consider the following case that illustrates prioritarianism's insensitivity to prudential justification without also illustrating any insensitivity to luck egalitarian unfairness:

- (3) *The case of one actual but two possible persons:* In this case, as in our above one-person case involving risk, there is only one impaired person who will actually exist. There's a 50 per cent chance that the impaired person who exists will be slightly impaired, and a 50 per cent chance that the impaired person who exists will be very severely impaired. But *the slightly impaired person and the very severely impaired person would be different people.*

The prioritarian case for administering the treatment for the very severe as opposed to the slight impairment is just as strong, in this case, as it is in the above case involving a single person who might turn out either very severely or slightly impaired (case 1). Nevertheless, contrary to prioritarianism, one has stronger moral reason to administer the treatment for the very severe impairment in this case than one does in that earlier case.

To see why there is stronger reason, let us suppose that our morally motivated stranger administers the treatment for the slight impairment on grounds that this will yield more expected utility than the alternative course of action. Now suppose that a person with a very severe impairment comes into existence. He would have a greater complaint than an analogous very severely impaired person in our one-person case (case 1), given that one can't justify the choice of the treatment for the slight impairment to him on the prudential grounds that one was looking after *his* interests by maximizing *his* expected utility *ex ante*. Rather, the expected utility that one maximized was instead that of the collective consisting of him and the other person who might have benefited from one or another of the treatments. One

can, by contrast, offer a prudential justification that one was looking after the interests of the very person who loses out when, in our one-person case, it is the same person who would be either very severely impaired or slightly impaired.

Though it is harder to justify administration of the treatment for the slight impairment in the case of one actual but two possible persons (case 3) than it is to justify this in our one-person case (case 1), this cannot be on luck egalitarian grounds of unfairness. We haven't got luck egalitarian reason to administer the treatment for the very severe impairment in the case of one actual but two possible persons, since only one impaired person will come into existence, and he won't be better or worse off than any other existing impaired person. Luck egalitarian considerations don't apply across actual and possible people. The luck egalitarian complaint of unfairness is that, through no fault or choice of an individual, the benefit he enjoys is less than that of another (or others). If, however, this is a matter of his level of enjoyment's falling short of that of a merely possible person, there is no compelling complaint of comparative unfairness, given that this level is lower than that of someone who isn't around to enjoy this higher level of benefits. Assuming that the individual's life is worth living, it would not be rational for him to envy this other person's non-existence at this higher level of benefit. In comparing ourselves with the gods atop Mt Olympus, we might complain that it is unfair that, through no fault or choice of ours, they enjoy powers that far exceed our own. But when we climb this mountain and discover that these gods are beings who do not inhabit the actual world but instead occupy some distant and merely possible world, this particular complaint is silenced.¹¹

V

In addition to its insensitivity to the presence or absence of a prudential justification, there is a further and distinct respect in which prioritarianism fails to take seriously the distinction between persons: it is insensitive to the presence or absence of competing claims of different individuals to benefit.¹² To illustrate this presence versus

¹¹ Note that we're considering the interests of the actually impaired in isolation from the interests of other actual people.

¹² Cf. Thomas Nagel's objection to utilitarianism that it 'depends on an application to interpersonal conflicts of the same principles which are used to settle conflicts between reasons arising from the interests of a single person. The conditions of choice corresponding to this principle are that the chooser should treat the competing claims arising from distinct individuals as though they all arose from the interests of a single individual, himself. He is to choose on the assumption that all these lives are to be amalgamated into one life, his own. . . . But this . . . completely distorts the nature of the

absence, consider our original pair of cases (1 and 2), one of which involves a single individual who might turn out either very severely or slightly impaired, and the other of which involves two individuals, one of whom we already know will be very severely impaired, and the other of whom we already know will be slightly impaired. It is, as we already noted, harder to justify treatment for the slight impairment in this two-person case than in this one-person case because one cannot offer a prudential justification of such treatment to the person who turns out very severely impaired in this two-person case, whereas one can offer such justification to the one person even if he turns out very severely impaired in this one-person case. But, in addition to the *absence* of such a prudential justification for administering treatment for the slight impairment in this two-person case, there is the following: the *presence* of what we call a competing-claims-based justification for administering the treatment for the very severe impairment in this same case, where such justification is lacking in this one-person case.

As we explained in 'Why It Matters':

one must justify any claim on resources in light of the comparative strength of the claims of others. Those who are relatively worse off have stronger claims to a given increment of improvement simply by virtue of the fact that it is, other things equal, harder to justify improving the situation of someone who is better off rather than someone who is worse off. How, one might ask rhetorically, 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? It is telling that an analogous complaint cannot be formulated against providing the treatment for the slight impairment in the one-person case.¹³

As is implicit in the passage above, the competing claims approach has purchase only when there is a choice between alternatives which will either benefit one person or benefit another person.¹⁴ It is also implicit in this passage, and I shall now make it explicit, that the strength of an individual's claim is a function of how much he might benefit in welfare terms, and from what baseline level of well-being, relative to another who might benefit if goods are distributed to him instead. It is morally relevant that there are distinct persons with competing claims

competing claims, for it ignores the distinction between persons' (Nagel, *The Possibility of Altruism* (Princeton, 1970), p. 138). In his ensuing discussion on pp. 138–40, Nagel appears, however, to blend elements of each of the two separateness-of-persons objections that I have just distinguished.

¹³ 'Why It Matters', pp. 183–4. I have slightly modified this passage in order to make it more directly relevant to the cases under discussion in this article.

¹⁴ More precisely: either benefit one or more persons or benefit another or more other persons.

to receive benefits. Such competing claims ground moral complaints on the part of those who would be worse off, relative to others, and the case for giving benefits to people with such complaints is stronger than it otherwise would be in analogous cases in which the prioritarian value of distributing goods in one way rather than another is equally great, yet such complaints are lacking.

VI

The competing claims complaint is distinct from the complaint of luck egalitarian unfairness.¹⁵ This can be shown as follows: there is a competing claims complaint against treatment for the slight impairment in the above case of one actual but two possible persons (case 3), whereas, as shown earlier, there is no complaint of inequality between persons in that case. Here is the competing claims complaint against administering treatment for the slight impairment in this case: how, one can ask, can one justify providing a treatment that might only improve the situation of someone who would be slightly impaired, when one could instead provide a treatment that, *ceteris paribus*, might improve the situation of someone else who would be much more severely impaired, and who would not even reach the level of slight impairment if benefited? This rhetorical question shows that the person who would be very severely impaired has a stronger claim to the treatment for the very severe impairment than the person who would be slightly impaired has to the treatment for the slight impairment.¹⁶

Here I am appealing to a complaint on behalf of a merely possible person who will never be actual. One might wonder whether complaints on behalf of such merely possible people carry any moral weight. I think they do when there is some chance that these merely possible people might have been actual, but not otherwise. The moral significance of complaints on their behalf is directly proportional to the likelihood of their existence.¹⁷

¹⁵ Not only are these two complaints distinct, but the complaint that prioritarianism is insensitive to luck egalitarian unfairness should be regarded as a further way, distinct from the two that I have already identified in this article, in which prioritarianism fails to take the distinction between persons seriously. See Alex Voorhoeve and Marc Fleurbaey, 'Egalitarianism and the Separateness of Persons' (this issue), for a discussion of the precise way in which luck egalitarian considerations of unfairness are sensitive to the distinction between persons.

¹⁶ In order to show that the competing claims complaint is distinct from the complaint of luck egalitarian unfairness, one can also construct cases in which there is a complaint of unfairness but no competing claims complaint against treatment for the slight impairment. One such case is the *two-actual-person case in which either of two possible persons might be very severely impaired* (case 5) that I introduce in section VIII.

¹⁷ Among such cases in which they have no significance are those Parfitian non-identity cases in which it is because of our choices that these merely possible people have no chance

To motivate this answer, first consider the case of one actual but two possible persons (case 3). Suppose we choose the treatment for the slight impairment and, as luck would have it, the person with the slight impairment comes into existence. The person with the very severe impairment never becomes actual in this case. Nevertheless, one can cogently complain on his behalf in the manner I sketched two paragraphs ago. This complaint has moral significance because there was a real chance – in this case, a 50 per cent chance – of his being actual.

Contrast this with a variant of the case of one actual but two possible persons in which we can choose whether a person with a very severe impairment or a different person with a slight impairment comes into existence. Here we ought to choose to bring the person with the slight impairment into existence and treat that impairment. It would not make sense to press a complaint against this on behalf of the merely possible person who would be very severely impaired because our choice ensures that there is no chance of his coming into existence.¹⁸

VII

I have claimed above that there are two respects in which prioritarianism fails to take seriously the distinction between persons: (i) it is insensitive to the presence or absence of prudential justifications and (ii) it is insensitive to the presence or absence of competing claims. I would now like to demonstrate that these two respects are, in fact, distinct. I can show this by, once again, constructing a case in which the one factor is present but not the other:

- (4) *Two-person case with risk and inversely correlated outcomes:* There are two people, each of whom you know will develop either the very severe or the slight impairment and each of whom has an equal chance of developing either impairment. You also know that their risks are inversely correlated: i.e., whenever one of

of being actual. For a discussion of non-identity cases, see Parfit, *Reasons and Persons* (Oxford, 1986), ch. 16.

¹⁸ But what, one might ask, if we choose to bring the person with the very severe impairment into existence and treat his impairment? Couldn't someone press a valid complaint on behalf of the merely possible person with the slight impairment even though our choice ensures that he never come into existence? The answer is no, since a possible person doesn't have any claim to be brought into existence, either on grounds that such existence would be very good in absolute terms or on grounds that his existence would be better than the existence of someone else we instead create. Only those who either do exist or stand a chance of existing have morally significant claims, where these claims are tied to their interests in the event, as chance or certainty would have it, that they exist, and the weight of these claims is a function, among other things, of the likelihood that they will exist.

them would suffer the very severe impairment, then the other would suffer the slight impairment. You can either supply both with a treatment that will surely improve a recipient's situation if and only if he turns out to suffer the very severe impairment or supply both with a treatment that will surely improve a recipient's situation if and only if he turns out to suffer the slight impairment. An effective treatment for the slight impairment would provide a somewhat greater increase in utility than would an effective treatment for the very severe impairment.¹⁹

In this case, unlike our original two-person case involving certainty (case 2), one can offer a prudential justification for providing the treatment for the slight impairment that appeals to the fact that this is in each person's ex ante self-interest. In this respect, this case is no different from our original one-person case involving risk (case 1). Nevertheless, this case is like our two-person case involving certainty (case 2), and unlike our one-person case involving risk (case 1), in the following respect that gives rise to a competing-claims-based complaint against providing treatment for the slight impairment. If we were to come to know what the outcome will be in this two-person case involving inversely-correlated risks (case 4), then this case would be transformed into our two-person case involving certainty (case 2). There would, when the case is so transformed, be exactly the same competing-claims-based complaint against treating the slight impairment as in that two-person case involving certainty (case 2). Though we do not actually know who will be slightly impaired and who will be very severely impaired in this two-person case (case 4), we do know the following: however things turn out, the distribution of burdens will be the same as in our two-person case involving certainty (case 2). Hence, however things turn out, there will be a competing claims complaint against providing treatment for the slight impairment.²⁰

¹⁹ This is a slightly simplified version of a case that we present on p. 197 of 'Why It Matters'. In his PhD thesis, Gabriel Wollner constructs an analogous two-person case involving inversely correlated risks and notes that, in such a case, the option that increases ex post inequality between the worse off and the better off can nevertheless be justified to each as in his ex ante self-interest. Wollner shows that, while Voorhoeve and I were aware of this feature of this sort of case, we failed to see how it implies that the 'separateness of persons' objection involving insensitivity to prudential justifications that we advanced in 'Why It Matters' is separable and distinct from our 'separateness of persons' complaint that prioritarianism is insensitive to relational and comparative considerations. (See Gabriel Wollner, 'Egalitarianism with a Human Face', UCL PhD thesis (2010), pp. 58–61.) Here and elsewhere in this article, I attempt to remedy this earlier failure.

²⁰ For a similar argument against the maximization of expected utility in cases such as this one, see Marc Fleurbaey and Alex Voorhoeve, 'Decide as You Would with Full Information! An Argument against Ex Ante Pareto', *Health Inequality: Ethics*

VIII

How might a prioritarian respond to the claim that his view fails to take seriously the distinction between persons in each of the two respects that I have identified? In this section, I shall explore the following reply. The prioritarian might concede that his view would be indefensible if the presence or absence of a prudential justification failed to make a difference to the justifiability of a course of action. But he might propose a clarification of his view whereby such presence or absence makes such a difference, thereby immunizing it from the first of the two accusations that it does not take the distinction between persons seriously. Having so clarified his view, he might then argue that such sensitivity to the presence or absence of a prudential justification is all that is necessary in order to take due account of the distinction between persons. There is no further need to be sensitive to the competing claims of different individuals.

On this clarification, the prioritarian would maintain that, when we are uncertain how well or badly a person's life will go, prioritarian weight should track the *expected* value of a person's life: the lower the expected value in absolute terms, the more moral weight we assign to that level of expected utility. This weighting should track the expected value of a life irrespective of how well or badly the person's life actually turns out.²¹ In so doing, the prioritarian can cite a precedent in Rawls's specification of his difference principle, which gives priority to the worst off members of society. On Rawls's specification, the worse off are identified as those whose representative member has the lowest *expectation* of goods over a lifetime.²²

A prudential justification is a justification that appeals to the higher expected value for a person of a given course of action. By weighting a person's *ex ante* expected well-being rather than the well-being he actually turns out to enjoy, an *ex ante* prioritarian is therefore attuned to the presence of a prudential justification. For example, an *ex ante* prioritarian acknowledges decisive reason to opt for the more beneficial treatment for the slight impairment over a less beneficial treatment for the very severe impairment in our one-person case involving risk (case 1), given that the former is of greater expected utility than the latter for the only person whose interests are at stake. In our two-person case involving certainty (case 2), by contrast, the *ex ante*

and Measurement, ed. Nir Eyal, Samia Hurst, Ole Norheim and Dan Wikler (Oxford, forthcoming).

²¹ In a section of 'Why It Matters' entitled 'Have We Correctly Described How the Priority View Deals with Choice under Risk?' (pp. 195–8), we label this the '*ex ante* Priority View'. This section builds upon and in some respect revises that section.

²² See Rawls, *A Theory of Justice*, p. 56.

prioritarian will weigh the two alternatives in just the same manner as the standard prioritarian, since expected and actual utility do not diverge here. Hence, he will follow the standard prioritarian in opting for the treatment for the very severe impairment in that case.²³ An *ex ante* prioritarian is therefore sensitive to the presence versus the absence, in these pair of cases, of a prudential justification for treatment for the slight impairment.

An *ex ante* prioritarian will not, however, be moved by the mere presence versus absence of a competing-claims-based complaint against a given option. This insensitivity can be illustrated by noting that, for an *ex ante* prioritarian, the case for providing treatment for the slight impairment is just as strong in our two-person case with inversely correlated risk (case 4) as it is in our one-person case involving risk (case 1). This is because the profile of expected benefits to persons is identical across these two cases: however much expected utility a given option will provide our one person who features in our one-person case involving risk (case 1), that option will provide just as much expected utility to each of our two persons who feature in our two-person case involving inversely correlated risks (case 4). Hence, for an *ex ante* prioritarian, the case for administering a more beneficial treatment for the slight impairment will be no less decisive in this two-person case (case 4) than in our one-person case (case 1). Yet, as I explained above, there is a competing-claims-based complaint against administering treatment for the slight impairment in this two-person case (case 4), which is absent in the one-person case (case 1).

Recall that I maintained above that the competing-claims-based complaint, as it applies to this two-person case (case 4), parallels the complaint that is present in our two-person case involving certainty (case 2). However things turn out in this two-person case involving inversely correlated risks (case 4), we are certain that one person will be worse off than the other. We can then ask, rhetorically, how can one justify providing a benefit of a given size to someone who will already be better off in order to make him better off still, when one could instead provide nearly as large a benefit to someone else who will be worse off, and who would not even reach the (unimproved) level of the better off person if she (the worse off person) is benefited? One might secure the assent of George W. Bush's favourite political philosopher in maintaining that for he that hath, to him more shall be given, at the expense of he that hath not. But this does not appear to be a sound moral principle.

²³ See 'Why It Matters', p. 196.

The ex ante prioritarian might say that he has the following ready answer to the rhetorical question I have just posed:

Here's how we can justify this. In this case, treatment for the slight impairment was in each person's ex ante self-interest. Our morally motivated stranger can and should take such gambles on their behalves, in order to promote their rational self-interests. The competing-claims-based complaint against administering treatment for the slight impairment is overridden by such an ex ante justification for taking such rational gambles on each person's behalf.

In response, I would point out that, in order to vindicate my charge that the prioritarian ignores the moral significance of the separateness of persons because he is insensitive to the presence or absence of competing claims, I don't need to reject this claim that ex ante prudential justifications override competing ex post claims when the two come into conflict. It would be sufficient, for my purposes, to show that the presence of competing claims provides a reason of some force, even if not necessarily an overriding reason, to provide treatment for the very severe impairment in the case of inversely correlated risks. The ex ante prioritarian, by contrast, must show that the presence or absence of competing claims is of no moral significance whatsoever. He must show this because his view attributes no moral significance, rather than a significance that is outweighed, to competing claims. Establishing this stronger claim is a tall order. If competing claims have any positive moral weight, even if these claims are outweighed by an ex ante prudential justification whenever the two come into conflict, then it follows that there is a respect in which prioritarianism ignores the moral significance of the separateness of persons.

In order to show that ex ante prioritarianism is vulnerable to the separateness-of-persons objection under discussion, I don't even need to establish that competing claims complaints carry any weight *at all* in cases in which there is an ex ante prudential justification for a conflicting course of action. All I ultimately need to show is that there are *some* pairs of cases – perhaps cases in which there are no such prudential justifications – in which the presence or absence of competing claims makes a moral difference that the prioritarian fails to register. I shall introduce the following case in order to show this:

- (5) *Two-actual-person case in which either of two possible persons might be very severely impaired.* One person will suffer the slight impairment and another person will suffer the very severe impairment. The person who will suffer the very severe impairment, is, however, either of two possible persons. This is because one's choice of treatment makes a difference to the identity of the person with the very severe impairment. If you provide the treatment for the slight rather than the very

severe impairment, a very severely impaired person will come into existence whose position could not have been improved by treatment for the very severe impairment. It could not have been improved for the following reason: if you provide treatment for the very severe rather than the slight impairment, then a *different* very severely impaired person will come into existence, and his position will be moved from very severe impairment to severe impairment by such treatment.²⁴ The identity of the slightly impaired person, however, is unaffected by one's choice of treatment. There is, moreover, no uncertainty regarding the identity of those who would exist and would benefit under each alternative.

In this case there would be no competing claims complaint against the provision of treatment for the slight impairment. This is because, if one provides treatment for the slight impairment, then the person who is thereby brought into existence, suffering the very severe impairment, could not have benefited from the alternative treatment for the very severe impairment. He would therefore have no claim to treatment for the very severe impairment. For reasons that I have offered in section VI, the merely possible person who could have benefited from the treatment for the very severe impairment would not have a competing claims complaint against provision of treatment for the slight impairment either, since he has no claim to be brought into existence.

Given that no one would have a competing claim to provision of treatment for the very severe impairment in the event that treatment for the slight impairment is provided, it is, intuitively, much less difficult to justify provision of the treatment for the slight impairment in this case (5) than it is in our original two-person case involving certainty (case 2). In that case (2), by contrast to this case (5), there is a strong competing claim to provision of treatment for the very severe impairment, but other moral considerations are equal.²⁵ The *ex ante* prioritarian, however, registers no difference in the strength of reasons to provide treatment for the slight impairment versus the

²⁴ Let us assume that the lives of those who are very severely impaired or severely impaired would be worth living and that such people would not regret their existence.

²⁵ Unlike our two-person case involving certainty (2), this case shares the following feature with our one-person case involving risk (1), in which it is reasonable to provide treatment for the slight impairment: one person has a claim to the treatment for the slight impairment, and that claim doesn't compete with anyone else's claim for treatment for the very severe impairment.

severe impairment in these two cases.²⁶ He thereby ignores the moral significance of the separateness of persons.

I have argued in this section that *ex ante* prioritarianism should be rejected on grounds that it ignores the separateness of persons by failing to register the moral significance of competing claims. I shall conclude by raising the following further difficulty with this view. The *ex ante* prioritarian owes us an explanation, beyond the fact that it renders his view invulnerable to the first of our two separateness-of-persons objections, for why significance should be accorded to expected rather than actual utility, such that the former, rather than the latter, is the bearer of prioritarian weightings. This is a difficulty for the prioritarian because it would seem, on the face of it, that how well or badly a person's life actually ends up going is what matters, rather than how well or badly it is expected to go. The latter seems especially insignificant compared to the former when our expectation is simply down to ignorance rather than any objective risks with which the world confronts us.²⁷

Suppose that who will end up very severely impaired and who will end up slightly impaired is already fated because genetically predetermined in our case of inversely correlated risks (case 4), but that we lack the information as to which person has which fate. Here the 50 per cent chance that each has of developing the very severe impairment is merely an epistemic rather than an objective chance: it is simply on account of their ignorance – that neither knows which fate is already in store for him – that it is in each person's *ex ante* self-interest to take this gamble.²⁸ Shouldn't the prioritarian maintain that what really matters is how well each person's life is already fated to go rather than how well we can anticipate that their lives will go, where our expectations diverge from their actual fates simply on account of our ignorance? Although the morally motivated stranger doesn't know the identity of these people, he knows, with certainty, that one person will end up very severely impaired and another will end up slightly impaired. He can assign prioritarian weight to these anonymously specified outcomes rather than to the expected utility of

²⁶ He registers no difference for the following reason. Since these cases involve no uncertainty regarding who would benefit, and to what degree, under any given course of action, expected utility is identical to actual utility. Moreover, treatment for the slight impairment yields just as much actual utility in each case, as does treatment for the very severe impairment.

²⁷ In Michael Otsuka, 'Risking Life and Limb' (unpublished manuscript), I offer an elaboration of this point regarding the significance of epistemic versus objective risks.

²⁸ I am assuming that one's *ex ante* self-interest is calculated on the basis of epistemic chances. I understand 'epistemic chances' to be those that are constituted by one's justified belief, given the evidence, in what the objective chances are. Such justified belief need not be true.

each and then maximize the priority-weighted value of these outcomes. Such actual outcomes seem a firmer, less illusory foundation on which to base decisions than expectations of utility based on ignorance.²⁹

m.otsuka@ucl.ac.uk

²⁹ The ex ante prioritarian might offer the following *ad hominem* rejoinder: 'In so downgrading the significance of expected utility, you undermine the first of your two separateness-of-persons objections to prioritarianism, since, in making this objection, you rely on the claim that the fact that a given policy maximizes the expected utility of an individual provides a justification to that person. How then, can you consistently turn around and belittle the significance of mere expected utility in questioning the prioritarian's move from an ex post to an ex ante view?' My reply is that if expected utility is all that we have to go on, then it's justifiable to be guided by it. But if, as in the case of inversely correlated risks, we can overcome the ignorance on which the expected utility calculations are based, we should do so. We should do so for reasons that Fleurbaey and Voorhoeve spell out in 'Decide as You Would with Full Information!'.