

Reply to Crisp

MICHAEL OTSUKA

University College London

ALEX VOORHOEVE

London School of Economics

We are grateful for, but unconvinced by, Roger Crisp's defence of the Priority View against our critique.¹ In this reply, we show that Crisp fails to grapple with, much less defeat, the central claim of our critique. We also show that an example that Crisp offers in support of the Priority View in fact lends support to our critique of that view.

1. In sections I–II of our article, we endorse the following two claims (among others):

- (i) In a one-person case involving an *intrapersonal* trade-off in which the 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 treatment for the slight impairment or instead for the very severe impairment, it is reasonable for the stranger to provide a treatment that maximizes the single person's expected utility, at least insofar as the stranger considers this person's interests in isolation from how well others are doing (himself included).
- (ii) In a multi-person case involving *interpersonal* trade-off in which half of the people will develop the slight impairment and the other half will develop the very severe impairment, it is already known who will develop which impairment, and a morally motivated stranger can offer treatment either for the slight impairment or instead for the very severe impairment, the stranger ought to provide treatment for the very severe impairment when the two treatments offer equal increases in utility. Moreover, the stranger ought to offer the treatment for the very severe impairment even in some cases in which this

¹ Roger Crisp, 'In Defence of the Priority View: A Response to Otsuka and Voorhoeve', *Utilitas* 23.1 (2011), pp. 105–8; 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), pp. 171–99.

generates less total utility than the treatment for the slight impairment.

Crisp takes issue with (i). Although we affirm (i) and show that (i) implies the falsity of the Priority View, we make clear in section III of our article that our ‘crucial argumentative move’ against the Priority View does not depend upon the truth of (i).² Our central argument is simply that there is a shift, which the Priority View cannot accommodate, in the moral weight one ought to give to changes in utility when we move from cases of intrapersonal trade-off to cases of interpersonal trade-off where some end up less well off than others. Here is what we write:

In this section, we shall argue, contrary to the Priority View, that some such shift is justified *even if one holds, contrary to what we have supposed in Sections I and II, that it is unreasonable for a morally motivated stranger to maximize expected utility in our one-person case because he should instead give some extra weight to increases lower down the utility scale*. We shall show that, *whether or not the stranger should maximize expected utility in the one-person case*, a shift of weighting when we move to the interpersonal case can be resisted only on pain of denying the moral significance of the separateness of persons. This is because a single person has a unity that renders it permissible to balance (expected) benefits and burdens against each other that might accrue to her. A group of different people, by contrast, does not possess such unity. As a consequence, some forms of balancing benefits and burdens that are permitted when these accrue to a single person are impermissible in cases where these benefits and burdens accrue to different people.³

In section V of our article, we employ a single-child variant, involving an intrapersonal trade-off, of Thomas Nagel’s well-known two-child case in order to illustrate this shift.⁴ Crisp maintains, contrary to what he takes to be a corollary of our claim (i), that the child’s parent in our one-child variant ought to give greater weight to this child’s interests should he turn out badly off, rather than maximize that child’s expected utility. But he leaves unrefuted our distinct claim, which is central to our critique of the Priority View, that ‘one should give less weight to benefiting this child should he end up with a disability (by moving to the city) than the weight one would give to benefiting the child with

² ‘Why It Matters’, p. 181.

³ ‘Why It Matters’, p. 179, emphases added. On Crisp’s presentation of our view, our ‘separateness of persons’ critique of the Priority View appears to rest upon our claim (i) that it would be reasonable to maximize utility in the one-person case involving an intrapersonal trade-off. It is, however, clear from the passage just quoted that this critique is independent of that claim.

⁴ See his ‘Equality’, in *Mortal Questions* (Cambridge, 1979), pp. 106–27, at p. 124. Our single-child variant is a transformation of ‘Nagel’s two-child case into a case involving a single child who has a 50 percent chance of ending up with a disability and a 50 percent chance of ending up able bodied’ (‘Why It Matters’, p. 188).

a disability in Nagel's two-child example'.⁵ We maintain this for the following reason: 'even if the child turns out to have a disability in our one-child case', one can offer a justification to him of a decision not to move to the city 'on the grounds that one was looking after that very same child's interest in flourishing in the event that he had turned out able bodied'.⁶ One cannot, by contrast, offer such a justification of not moving to the city to the disabled child in Nagel's two-child case, since he had no prospect of benefit from such a move. Even if, as Crisp maintains, the fact that one was looking after the sole child's interest in able-bodied flourishing in the suburbs is insufficient fully to justify a decision not to move to the city, it nevertheless provides a *pro tanto* reason not to move to the city that is absent in the two-child case.

The Priority View implies *both* that (a) a morally motivated stranger should give *more weight* to the interests of someone if he turns out badly off in one-person cases involving intrapersonal trade-off such as our variant of Nagel's child case *and* that (b) such a stranger should give *no more weight than whatever weight it is appropriate to give in such one-person cases* when we shift to analogous two-person cases involving interpersonal trade-off in which some end up better off than others. We think it reasonable to reject (a) when one considers the single person's interests in isolation. We also think there is an overwhelming case for the rejection of (b); moreover, such rejection forms the core of our critique of the Priority View. Since the Priority View is defeated by the falsity of *either* (a) *or* (b), Crisp's defence of (a) is, in itself, insufficient as a defence of the Priority View against our critique.

2. Crisp's example of 500 separate decisions about orphaned children is meant to provide further support for the Priority View against our critique. Perhaps it is meant to provide the support for (b) that is otherwise absent in his response.⁷ Here we explain why his example in fact lends support to our critique of the Priority View.

Before discussing Crisp's example, it will provide a useful point of reference to describe a multi-person version of our one-person case that we present in our article. This case 'is identical to our one-person case, save for the fact that this one person and her fate have been replicated many times over to create a group of people'.⁸ Since we stipulate that the fates of the members of this group are perfectly correlated – i.e. without treatment, all will have the slight impairment, or else all will have the very severe impairment – we shall call this the *case of perfectly*

⁵ 'Why It Matters', p. 188. Crisp quotes this claim on pp. 106–7 of his response.

⁶ 'Why It Matters', p. 188.

⁷ The last sentence of Crisp's response suggests that it is meant to provide such support.

⁸ 'Why It Matters', p. 175, n. 8.

correlated risks. In this case, it is possible either to provide everyone with the treatment for the very severe impairment or everyone with the treatment for the slight impairment. We maintain that '[j]ust as it is reasonable to provide the individual in our one-person case with the treatment that maximizes her expected utility, it is reasonable to provide each member of this group with the treatment that maximizes her expected utility'.⁹

Crisp maintains that we are (also) committed to the claim that it would be reasonable to provide whichever treatment maximizes each person's expected utility in his case of 500 orphans. In this example, either treatment maximizes each person's expected utility, since Crisp stipulates that the treatment for the slight impairment (which fully cures that impairment) has the same expected utility as the treatment for the very severe impairment (which transforms the very severe impairment into a merely severe impairment). As in our case of perfectly correlated risks, there is a 50 per cent chance, in the case of each of Crisp's 500 orphaned children, that the child will end up with the very severe impairment, and a 50 per cent chance that he will end up with the slight impairment. But there is a crucial difference between Crisp's case and ours: in his case, each child's chance of being very severely impaired is *independent* of every other child's chance, rather than being perfectly correlated. As the result of this independence, maximizing each child's expected utility by giving him the treatment for the slight impairment would, as Crisp notes, be likely to result in an outcome that approximates the following: '250 children who would have suffered from the slight impairment will suffer no impairment, while 250 children who might have suffered only the severe impairment in fact suffer from the very severe impairment'.¹⁰ Crisp maintains that most would reject such an outcome and would instead prefer what the Priority View would recommend, which is that every child receives the treatment for the very severe impairment, the likely upshot of which is something similar to the following: 250 children who would otherwise have been very severely impaired end up merely severely impaired, and 250 children end up with the slight impairment rather than no impairment.

Crisp's 500-child case is, in fact, relevantly similar to the following case that we discuss at the end of our article:

Multi-person case with risk and inversely correlated outcomes: You are confronted with a group of people, each of whom you know will either develop the very severe or the slight impairment and each of whom has an equal chance of developing either impairment. You also know that half the people

⁹ 'Why It Matters', p. 175, n. 8.

¹⁰ 'In Defence', pp. 107–8.

will end up suffering from the very severe impairment, and half from the slight impairment. You can either supply everyone with a treatment that will surely improve a recipient's situation if and only if she turns out to suffer the very severe impairment or supply everyone with a treatment that will surely improve a recipient's situation if and only if she turns out to suffer the slight impairment. Both treatments, if effective, provide the same increase in utility over nontreatment.¹¹

In our discussion of this case, we maintained that we have decisive reason to provide treatment for the very severe impairment in these circumstances, in part because such treatment would reduce the resulting brute luck inequality between the very severely and the slightly impaired.¹² Hence we should not provide each with the treatment for the slight impairment even though such treatment maximizes each person's expected utility.¹³

It should now be clear why Crisp is mistaken to think that we are committed to the view that it would be reasonable to give each of his 500 orphans the treatment that would be effective if and only if that orphan suffers the slight impairment: the reasonableness of giving such a treatment that maximizes a person's expected utility depends upon our assumption that the fate of the individual be considered in isolation from that of others.¹⁴ Once we drop this assumption and consider the case for providing treatment in the context of how different children would fare under this policy – as Crisp does by describing how badly and how well the very severely impaired and slightly impaired orphans would end up faring – the case against giving such treatment becomes decisive.¹⁵

We conclude with the observation that the Priority View yields *just as strong a reason* to provide each child with the treatment for the very severe impairment in a case such as ours involving perfectly

¹¹ 'Why It Matters', p. 197.

¹² 'Why It Matters', pp. 197–8.

¹³ As in Crisp's orphans case, in this example, either treatment maximizes the expected utility of those treated.

¹⁴ See 'Why It Matters', pp. 173–4, where we state that the reasonableness of maximizing the expected utility of the one person in our intrapersonal trade-off case depends upon the assumption that 'you are considering her fate in isolation from any consideration of how well off or badly off anybody else is (yourself included)'.

¹⁵ Although decisive, this case is not, however, as strong as is the case against providing such treatment in our original multi-person case in which it is already known who will suffer the very severe impairment and who will suffer the slight impairment. It is less strong for the following reason: in Crisp's orphans case, one can justify provision of the treatment for the slight impairment to each on grounds that such provision maximizes that person's expected utility. One cannot provide such justification to each in our original multi-person case, since one cannot provide such justification to those of whom it's already known that they will develop the very severe impairment. This shows, contrary to the very last sentence of Crisp's response, that the justifiability of provision of a given treatment is sensitive to more than the distributive upshot of such provision.

correlated risks as it does in Crisp's case of independent risks. There would, however, surely be stronger reason to provide every child with the treatment for the very severe impairment in Crisp's case than in ours. That stronger reason in Crisp's case is the reason we have to avoid the sort of strikingly inegalitarian outcome that Crisp describes: i.e. an outcome along the lines of one in which '250 children who would have suffered from the slight impairment will suffer no impairment, while 250 children who might have suffered only the severe impairment in fact suffer from the very severe impairment'. In our case of perfectly correlated risks, by contrast, an expected-utility-maximizing outcome in which some relatively mobility-rich children get richer and other relatively very mobility-poor children stay that poor is impossible, since the fates of individuals cannot diverge.¹⁶ Crisp's example is therefore grist for our anti-Prioritarian mill: it provides confirmation that the Priority View should be rejected because it is insensitive to the moral significance of the fact that some are worse off than others.

m.otsuka@ucl.ac.uk

a.e.voorhoeve@lse.ac.uk

¹⁶ For such a critique of the Priority View as insensitive to the moral difference between cases involving perfectly versus imperfectly correlated risks, see John Broome, 'Equality versus Priority: A Useful Distinction', *'Goodness' and 'Fairness': Ethical Issues in Health Resource Allocation*, ed. Daniel Wikler and Christopher J. L. Murray, World Health Organization, forthcoming.