

Prioritarianism and Single-Person Cases¹

(Draft as of 7 December 2016)

Per Algander and Andrew Reisner

1. Introduction

In an influential paper, Michael Otsuka and Alexander Voorhoeve (2009) argue that Prioritarianism should be rejected because it cannot account for a claimed shift in our intuitions when we move from evaluating one-person cases to evaluating many-person cases. According to Otsuka and Voorhoeve, Prioritarianism correctly implies that it may be permissible, and perhaps required, to benefit those who are worse off in many-person cases, even if doing so would not maximise overall expected utility. In single-person cases, however, they claim that is always permissible to act in such a way as to maximise expected utility. According to them, Prioritarianism implies that this is not the case and should therefore be rejected.

Their paper is an important effort to side-step one feature of the debates between prioritarrians and telic egalitarians (and between each of those and other distributive views), namely the trading of intuitions about the distributive consequences of the two views in standard, many person cases.² By comparing single-person cases to many-person cases, they introduce novel intuitions and data into the mix, i.e. those about single-person cases. If successful, this would offer an important new direction for the debate to go in.

In this paper we examine Otsuka and Voorhoeve's arguments that rely on the shift from single- to many-person cases. We observe that these arguments are only valid if they assume the following two premises:

Premise 1: Morality aligns with prudence in single-person cases.

Premise 2: The moral importance of benefits and burdens differ in single-person cases compared to many-person cases.

We dispute both these claims and argue that:

(i) *Premise 1* is question-begging against Prioritarianism. Therefore, it cannot be assumed in an argument intended to refute Prioritarianism.

(ii) *Premise 2* should be rejected. The possible views which could justify benefits and burdens differing in moral importance in single- and many-person cases are not plausible.

While the main focus of our discussion is on Otsuka and Voorhoeve, we also wish to bring out a more general point. There is no free intuitive lunch when it comes to arguing about different

¹ We would like to thank Gry Wester and Katharina Berndt Rasmussen for valuable comments on an earlier draft. We would also like to thank the participants in the Higher Seminar in Practical Philosophy at Uppsala University for their valuable feedback.

² An instance of this is Roger Crisp (2003) who offers the clever example of to whom to distribute a bottle of fine wine, when the only two possible recipients are, respectively, extremely well off and more than extremely well off. In his particular case the example is designed to separate the principle of compassion from telic egalitarianism, but it can be adjusted easily to distinguish between intuitions that would favour Prioritarianism or Telic Egalitarianism.

distributive theories. Prioritarianism and telic egalitarianism are moral views, and the relevant intuitions for and against them are moral intuitions or intuitions with moral upshots. One reason why the debate has the current shape that it does is that it is difficult to avoid relying on intuitions about the morally relevant distributive consequences of the two views. This is a difficulty that Otsuka and Voorhoeve cannot avoid, despite the apparent novelty of their argumentative approach.

2. Two cases

Otsuka and Voorhoeve ask us to consider the following case:

Case 1: A young adult in perfect health is told that she will be either slightly or very severely impaired in the near future. The slight impairment is a condition in which it is difficult to walk more than 2 km. The very severe impairment is a condition in which one is bedridden except that one can, with the assistance of others, sit up and be moved about in a wheel-chair. It is equally probable that the young adult will be very severely impaired as it is that she will be slightly impaired, but it is not possible to determine which impairment she will actually get. She is now given the opportunity to choose one of two treatments: Treatment A or Treatment B. Treatment A will improve the very severe impairment to a condition of severe impairment where one can sit up on one's own for the entire day but still requires assistance to move about. This treatment has no effect on the slight impairment. Treatment B on the other hand will completely cure the slight impairment. This treatment has no effect on the very severe impairment.³

Otsuka and Voorhoeve cite survey data which indicate that people tend to be indifferent between the two treatments. If we take a person's expressed preferences to be evidence for that person's utility,⁴ we can draw the conclusion that the two treatments provide equal increases in the person's utility. *Ex hypothesi*, the two treatments also have equal expected utility for the person. ('Utility' should here, and for the rest of this paper, be understood as referring to how well a person's life is going for that person).

Now consider the following case:

Case 2: You are to assist a group of people half of which will soon develop the very severe impairment and half of which will soon develop the slight impairment. You are given a choice between distributing either Treatment A or Treatment B to this group, thereby benefiting either those who will suffer the very severe impairment (Treatment A) or those who will suffer the slight impairment (Treatment B).

Assume that the two impairments and the two treatments are the same as in Case 1. That is, the impairments and treatments are such that people tend to be indifferent between Treatment A

³ These are our own versions of Case 1 and Case 2. These versions are the same in substance as the originals but have been presented slightly different for reasons of clarity and continuity.

⁴ One need not assume a preference theory of utility for the purposes of our argument or of Otsuka and Voorhoeve's. It is sufficient if people's expressed preferences in the survey accurately track the facts about their utility.

and Treatment B when they are the only person affected. As we noted above, taking a person's preferences as evidence for her utility suggests that the two treatments yield equal increases in a person's utility. Now in contrast with Case 1, the survey data cited by Otsuka and Voorhoeve show that people are not indifferent between Treatment A and Treatment B in Case 2. When respondents consider how to allocate treatments for the entire group, they prefer to distribute Treatment A, thereby benefiting the worst-off group.

Based on the survey data, Otsuka and Voorhoeve claim that it is permissible to distribute either Treatment A or Treatment B in Case 1, but that it is not permissible to distribute Treatment B in Case 2. Remember that in Case 1, people are generally indifferent between the two treatments and that we assume this to be evidence for the two treatments having equal utility for the person (and equal expected utility). The upshot of these claims, according to Otsuka and Voorhoeve, is that in single-person cases it is always permissible to maximise expected utility. In many-person cases on the other hand, it is not always permissible to maximise expected utility. In Case 2 it is not permissible to distribute Treatment B, even though the expected total utility of Treatment B is equal to that of Treatment A.

Otsuka and Voorhoeve claim that if we take these claims about the two cases seriously, then Prioritarianism must be rejected. According to Prioritarianism, benefits matter more (morally) when the recipient of the benefit is worse off in absolute terms. More precisely, Prioritarianism is the view that the moral importance of an increase (decrease) in individual utility depends on two factors: the size of the increase (decrease) *and* the level of utility which the recipient of the increase is located at.⁵ The theory therefore implies that it is not only morally permissible, but even morally required to choose Treatment A in Case 2. Even though the two treatments provide their recipients with equal increases in *utility*, the increase in utility provided by the treating the severe impairment has greater *moral importance*, because the recipients of Treatment A are at a lower level of utility than the recipients of the Treatment B.

It is perhaps not obvious that Prioritarianism has this consequence in Case 1 since this case involves uncertainty. In Case 1 we have to assess the moral importance of *potential* benefits, not actual ones. Prioritarianism, as the theory was formulated above, gives us little guidance as to the moral importance of potential benefits. There are two main ways of adapting Prioritarianism to account for uncertainty.

One interpretation of Prioritarianism (so-called 'ex-ante' Prioritarianism) is that we should assign higher moral importance to increases in *expected utility* rather than utility. On this view, it is more important to provide an increase in expected utility to a person whose expected utility is low than it is to provide an increase in expected utility to a person whose expected utility is high. This version of Prioritarianism implies that the two treatments are equally morally important in Case 1 because they have the same expected utility. This version can therefore account for the survey data regarding Case 1.

⁵ Derek Parfit offered an influential informal definition of Prioritarianism: Benefiting people matters more the worse off those people are (Parfit 1991: 19). Note that Parfit is concerned with absolute rather than relative wellbeing, or in our parlance, 'utility' levels.

However, ex-ante Prioritarianism faces other, arguably more serious, objections, some of which are discussed by Otsuka and Voorhoeve.⁶ We will therefore set this version of Prioritarianism aside and focus on the version that is the main target of Otsuka and Voorhoeve's argument.

The second interpretation of Prioritarianism (sometimes called 'ex-post' Prioritarianism) is that, when evaluating acts with uncertain outcomes, we should apply the prioritarian weighting first to each possible outcome and then calculate the expected moral value of the act. On this version, what we should do when choosing under uncertainty is to perform the act which has the highest expected moral value, where the moral value of an act is the sum of the probability-weighted moral value of each possible outcome.⁷

This version of Prioritarianism implies that it is not permissible to choose Treatment B in Case 1. While Treatment B has an equal expected utility to Treatment A, Prioritarianism implies that the benefit of Treatment B matters less, morally, than Treatment A. The benefits provided by Treatment A matter more, morally, according to Prioritarianism in virtue of the fact that Treatment A would benefit the person if she were to be very severely impaired, whereas Treatment B would benefit the person if she were slightly impaired.

A further illustration of the same general point is that Prioritarianism prescribes Treatment A, even in some cases where the expected utility of Treatment B is higher than Treatment A. Suppose that the young adult in Case 1 is not indifferent between the two treatments but has a slight preference for Treatment B. Suppose further that this is evidence for the claim that Treatment B would have a higher expected utility than Treatment A. According to Prioritarianism, this increase in expected utility can, if it is small enough, be outweighed by the fact that Treatment A would benefit the person at a lower level of welfare, simply because a benefit counts for more, morally, the worse-off the recipient of the benefit is.

This concludes our presentation of Otsuka and Voorhoeve's argument. In the following section we will argue that it fails because it begs the question against Prioritarianism.

3. A moral premise

Otsuka and Voorhoeve's argument has two components. The first is the data itself about single- and many-person cases. They assume that in the many-person case, the data are clearly about moral goodness and about what one is morally required to do.

The second component arises in the single-person case. With respect to this case Otsuka and Voorhoeve claim that:

Had she [the person in Case 1] preferred the treatment for the very severe impairment, then it would have been reasonable for you to provide her with that treatment. In other words, it would be reasonable for you to provide a treatment that maximizes the expected increase in utility of the recipient. This conclusion is justified at least insofar as the individual is the only person whom you can benefit and you are considering her fate in isolation from any consideration of how well off or badly off anybody else is (Otsuka & Voorhoeve 2009: 173-4).

⁶ See Otsuka & Voorhoeve (2009: 196-7). For a more philosophical motivation for why a prioritarian should opt for ex-post rather than ex-ante Prioritarianism, see Rabinowicz (2002: 10-11).

⁷ This can be reformulated in the standard consequentialist way for ties.

According to their view it would be permissible, at least in single-person cases, for a 3rd party to act in such a way as to maximise the person's expected utility. Let us call this 'the moral premise'. By appealing to intuitions about what a third person is permitted to do, they remove an ambiguity, at least for purposes of their argument, from the data about the single person case, namely whether our intuitions in that case are moral or prudential. The case itself might elicit either kind of intuition, as it only concerns one individual's attitude towards her wellbeing and her wellbeing alone.

The moral premise helpfully provides Otsuka and Voorhoeve with a moral claim: that it is permissible to help the individual in the single-person case maximise her expected utility. This puts them in a position to run their anti-prioritarian argument. It is morally permissible to help the single individual maximise her expected utility, but it is not always morally permissible to maximise total utility in the social case. There, more weight must be given to the worse off, either relatively or absolutely, when determining the social good. Prioritarianism gets the social good case right, but not the individual good case.

This argument is problematic because it is question-begging against Prioritarianism. Put simply, the negation of the moral premise is a central component of Prioritarianism. Recall that the moral importance of an increase in utility depends on two factors according to Prioritarianism: the size of the increase and the level of utility at which the recipient is located. However, the moral premise directly contradicts this claim. If it is always permissible to maximise a person's expected utility in single-person cases then it cannot be the case that increases in utility at lower absolute levels matter more. Even if we assume for the sake of argument that the survey data accurately indicate the expected utility of the two treatments in Case 1, it is a further step – one which is the direct denial of this central feature of Prioritarianism – to assume that the two treatments matter equally from a moral perspective.

The strength of Otsuka and Voorhoeve's argument relies on that they can use survey data to give them conclusions about moral goodness. In the many-person case, the data is plausibly interpreted as being about moral goodness, since it involves inter-personal distributive judgements and because the respondents are naturally seeing the case from a 3rd person perspective. The survey data concerning the single-person case is however not clearly, and perhaps likely not at all, about moral goodness. In order for the two surveys to line up in such a way as to make their argument against Prioritarianism, the moral premise is required. But the moral premise cannot be inferred merely by looking at the surveys. It is a further normative claim, one which simply assumes what is at dispute.

In claiming that the moral premise is question-begging against Prioritarianism we are of course not assuming that it is false, or even that it is implausible. We are merely claiming that the moral premise alone implies the negation of Prioritarianism and that it cannot be assumed as an un-argued premise in an argument against the theory.

One who is sympathetic to the Otsuka and Voorhoeve line might argue that the moral premise is *prima facie* plausible and that it reveals a fatal flaw in Prioritarianism; namely that the theory implies that when acting under uncertainty it can be better to perform an act A even though there is an alternative act B which has higher expected utility for everyone. A more generous reconstruction of Otsuka and Voorhoeve's argument might therefore be that Case 1 illustrates this consequence of Prioritarianism and that any theory which has this consequence should be rejected.

However, merely pointing out that the theory has this consequence is not to point out that it is a fatal flaw, or that it is unacceptable. The question of whether it is a flaw, feature, or neither requires an independent argument in this dialectical context.⁸

In this section we have argued that Otsuka and Voorhoeve rely on a moral step when they claim that it is morally permissible to choose either Treatment A or Treatment B in Case 1. What they have to assume in order to get their argument going is that prudence and morality align in single-person cases, at least in so far as it is always permissible to do what will maximise individual utility in such cases. We claim that this assumption is not theory-neutral and begs the question against Prioritarianism. It is a central commitment of Prioritarianism that prudence and morality *does not* necessarily align and that individual utility and morality can come apart.

However, it may be surprising that Prioritarianism should imply that morality and prudence can come apart in single-person cases. In the next section we discuss Otsuka and Voorhoeve's second argument against Prioritarianism. This argument, they claim, does not rely on premise that it is always permissible to do what will maximise a person's expected utility in single-person cases. Instead, this argument aims to show that a "shift occurs in the moral importance of benefits and burdens when we move from a case involving intrapersonal to a case involving interpersonal tradeoffs" (Otsuka & Voorhoeve 2009: 181) and that Prioritarianism cannot account for such a shift.

4. *Single- and many-person cases*

Otsuka and Voorhoeve suggest a more general argument against Prioritarianism which does not rely on the moral premise. They claim that when we evaluate cases involving only one person, but where uncertainty is involved, then it is plausible to allow for certain trade-offs between the different possible outcomes. For example, they claim that it would be permissible to choose a risky option rather than a safe option, given that the potential gain in the risky option is great enough. In many-person cases however, similar trade-offs are not permissible. For example, it is not permissible to risk a significant burden to one person in order that another person should have the chance of receiving a significant benefit.

The upshot of such examples, they claim, is that there is a shift in the moral importance of benefits and burdens when we move from a single to a many-person case. Benefits and burdens have a certain moral weight in single-person cases but they need not have the same moral weight in many-person cases.

Prioritarianism, it seems, cannot explain such a shift. The moral weight of a benefit or burden according to Prioritarianism depends on the size of the benefit (burden) and on the absolute level of well-being which the recipient of the benefit (burden) is at. None of these factors can explain why benefits or harms would have a different weight in many-person cases than they do in single-person cases.

⁸ That (ex-post) Prioritarianism can diverge with prudence in single-person cases is, for example, noted and discussed by Rabinowicz (2002). Rabinowicz discusses how Prioritarianism should be formulated in relation to Harsanyi's theorem, especially to a modern version of that theorem developed by John Broome (1991, 2015). He argues that a prioritarian should reject the Principle of Personal Good for Prospects; that is, the prioritarian should reject the claim that if a prospect P_1 is at least as good for everyone as a different prospect P_2 , then P_1 is better than P_2 . Otsuka and Voorhoeve mention these arguments in passing (fn. 16), but they fail to note that their central assumption just is the Principle of Personal Good for Prospects.

In what follows we will argue that the explanation of the shift offered by Otsuka and Voorhoeve do not yield the results they want when considering a further case (Case 3, described below). Consequently, the purported shift in the moral relevance of benefits and burdens in single- and many-person cases should be rejected. If there is no such shift, then it is not a fault of Prioritarianism that it cannot account for it.

Otsuka and Voorhoeve invoke Case 1 and Case 2 again to support the claim that there is a difference in moral importance of benefits and burdens in single- and many-person cases. They claim that the two cases should be evaluated differently 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 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. (Otsuka & Voorhoeve 2009: 179).

A pair of cases which is suggested as an example of when intra- and inter-personal benefits and burdens should be weighed differently is the following:

- (i) If you do A, then there is a 50-50 chance that one person receives a gain in utility or an equally large loss in utility. Otherwise, nothing happens.
- (ii) If you do A, then there is a 50-50 chance that one person receives a gain in utility and another person receives an equally large loss in utility. Otherwise, nothing happens.

With respect to (i), Otsuka and Voorhoeve claim that we should weigh the expected gain and loss against each other and, if the result comes up positive, we should opt for the gamble. In case (ii) however they think that the weighing is not so simple. In this case it is not just a matter of weighing potential benefits against burdens, it is also a significant fact that there is no particular person who is the potential recipient of both the benefit and the burden. As Otsuka and Voorhoeve put it, “there is no single person for whom the prospect of a greater gain is the desirable flip side of exposure to the risk of a lesser loss and for whom the prospect of such a gain might be worth the exposure of such a risk” (Otsuka & Voorhoeve 2009: 180).

The way Otsuka and Voorhoeve explain the shift seems to be by appealing to a difference in inter- and intra-personal aggregation. In the single-person case it is permissible to “trade off” the chance of receiving a benefit against the risk of receiving a burden because this is a form of intra-personal aggregation. In cases which involve inter-personal aggregation it is not permissible to do the same kind of trade-offs. Hence, there is a shift.

However, explaining the shift in these terms does not imply that there is necessarily a difference between single- and many-person cases. Consider the following case:

Case 3. You are to assist a group of people *all* of which will have either the very severe impairment or the slight impairment with an equal probability of .5. You are given a choice of distributing Treatment A which will slightly benefit a person if and only if she has the very severe impairment or distributing Treatment B which will slightly benefit a person if and only if she has the slight impairment.

Roger Crisp (2011) presents a similar case in his reply to Otsuka and Voorhoeve. An important difference between this case and the one presented by Crisp is that Crisp's version of the case allows for mixed outcomes. That is, the way Crisp presents the case it is possible that some members of the group will get the very severe impairment and that the other members will get the slight impairment. In Case 3 this is not possible. A way to conceive of a situation like the one we are presenting in Case 3 is that a group of people have the same symptoms, but it is uncertain what these symptoms are symptoms for.

In their reply to Crisp, Otsuka and Voorhoeve (2011) point out that there is a relevant difference between Case 1 and an "iterated" version of Case 1. While it is permissible on their view to choose either Treatment A or Treatment B in Case 1, it would not be permissible to consistently choose Treatment B on several occasions. In the iterated version of Case 1 each person's risk of being severely impaired is independent of any other person's risk of having the same impairment. This means that, over many iterations, the end result would approximate an equal distribution of the very severe and the slight impairment. The choice between the two treatments in the iterated version of Case 1 is therefore similar to a case where we know that half of the group will get the very severe impairment and the other half the slight impairment, but we do not know of any particular person which impairment she will get.

We will not consider the potential merits or flaws of this line of reply since it is obvious that it is beside the point when it comes to Case 3. In Case 3, there is no possibility of a mixed outcome: either everyone will be severely or slightly impaired. Otsuka and Voorhoeve's reply to Crisp's version of this case is therefore not relevant to Case 3.

Case 3 is similar to Case 1 except that it involves several people instead of just one. It is therefore a case where you essentially have to make a gamble. If you choose Treatment A and everyone turns out to have the very severe impairment, then everyone benefits. If everyone turns out to have the slight impairment then your choice has been completely ineffective (and vice versa for Treatment B). Case 3 does therefore not involve inter-personal aggregation of benefits and burdens. Rather, it is a case with only intra-personal aggregation where we have to assess the moral relevance of potential benefits and burdens to a group of people.

This shows that any purported shift between many- and single-person cases cannot be fully accounted for in terms of the moral relevance of inter- and intra-personal aggregation. If there is such a shift then we should expect Case 1 and Case 3 to differ in some significant way. However, Case 1 and Case 3 do not differ with respect to inter- and intra-personal aggregation.

Otsuka and Voorhoeve might concede that the distinction between single- and many-person cases cuts across the distinction between inter- and intra-personal aggregation. They might maintain, however, that the latter distinction marks a morally relevant difference. That is, cases involving only inter-personal aggregation should be evaluated differently than cases involving only intra-personal aggregation, regardless of whether they are single- or many-person cases.⁹ This would lead them to accepting the claim that it is permissible to distribute Treatment B in Case 3.

This reply invites the question what the moral relevance of benefits and burdens do depend on according to Otsuka and Voorhoeve. A possible view, which is suggested by their reply to

⁹ This is suggested by Otsuka and Voorhoeve in a footnote (Otsuka & Voorhoeve 2009: 175, fn.8).

Crisp's version of the case, is that the moral importance of a benefit or burden depends on whether it increases or decreases inequality. Benefiting the better off is less morally important than benefiting the worse-off, because benefiting the better off increases inequality while benefiting the worse off lessens inequality. This would be congenial with their claim that the difference between inter- and intra-personal aggregation is morally relevant, since questions of decreasing or increasing equality simply does not arise in intra-personal aggregation. It would also support the claim that many, though not all, many-person cases should be evaluated differently than single-person cases.

However, this view of the moral relevance of burdens and benefits comes with some significantly counter-intuitive consequences. To borrow an example from Wlodek Rabinowicz (2009: 17), it is very hard to believe that the moral significance of the benefits and burdens which Robinson Crusoe (a single-person case) might enjoy change as soon as Friday enters the picture, even if these benefits and burdens have no effect on Friday. This is not to say that *other* moral considerations – considerations that might have nothing to do with Crusoe's well-being – might enter the picture as soon as Friday does. What is hard to believe is rather that Friday's *mere* presence can affect the moral importance of Crusoe's benefits and burdens.

A more plausible suggestion is hinted at by Otsuka and Voorhoeve on several occasions. At one point, when discussing a case similar to Crisp's version of Case 3, they claim that

“[W]e do know that whoever does end up with this condition [i.e, the severe impairment] *has a stronger claim* on the treatment for the very severe impairment than whoever will end up with the slight impairment. This provides us with a decisive reason to provide everyone with the treatment for the very severe impairment, since we know that we will thereby be providing treatment to those who will turn out to be the people who had the strongest claim on it” (Otsuka & Voorhoeve 2009: 197-8. Emphasis added).

Otsuka and Voorhoeve could argue that Treatment A is more morally significant than Treatment B because the potential beneficiaries of Treatment A have a stronger claim on it than the potential beneficiaries of Treatment B. However, this cannot be in virtue of the fact that the potential beneficiaries of Treatment A would be worse off than anyone else. As we saw, Case 3 involves perfect equality whatever we choose. Rather, the most straightforward explanation of these claims by Otsuka and Voorhoeve seems to be that what makes the potential beneficiaries of Treatment A have a stronger claim is that they would be worse off in *absolute* terms than the potential beneficiaries of Treatment B. But, of course, that is just Prioritarianism!

We should therefore reject the second premise in Otsuka and Voorhoeve's argument against Prioritarianism, namely that there is a shift in the moral importance of benefits and burdens when we move from single-person to many-person cases. The difference between inter- and intra-personal aggregation does not map onto the difference between single- and many-person cases, the moral importance of decreasing inequality makes the moral relevance of benefits and burdens depend on the mere presence of other people, and appealing to the strength of different people's claims is compatible with Prioritarianism.

So far we have intentionally avoided what may be the most important question: what should we do in Case 3? Taking the difference which Otsuka and Voorhoeve draw between inter- and intra-personal aggregation seriously would lead to the conclusion that it is permissible to

distribute Treatment B. This is because in Case 3 there *are* people for whom the chance of slight impairment with treatment is the “desirable flip-side” of the risk of being very severely impaired. Prioritarianism, on the other hand, implies that we should distribute Treatment A. Our considered moral intuitions side with Prioritarianism on this issue, though we do not take these intuitions to settle the matter. Progress on this issue will not be had by merely considering intuitions about particular cases, we need to dig deeper and examine what each theory has to say about the relation between morality and prudence.

Conclusion

We have argued that the two arguments presented by Otsuka and Voorhoeve against Prioritarianism are unsuccessful. The first argument is question begging in the dialectical context of the argument. It assumes that the very point at dispute between Prioritarians and either Utilitarians or Egalitarians is settled in favour of the latter two kinds of theories.

The second argument is unsuccessful because Otsuka and Voorhoeve fail to give any support for their central claim: that the moral importance of benefits and burdens differ in single- and many-person cases. Otsuka and Voorhoeve’s suggestion that the difference concerns inter- and intra-personal aggregation comes with some counter-intuitive consequences and does not map onto the difference they are faulting the prioritarian for failing to maintain.

Perhaps more importantly, the failure of the second argument also illustrates what the fundamental difference between the prioritarian on the one hand and telic egalitarians on the other, actually is. The prioritarian is committed to the view that an individual’s utility makes the same contribution to the total moral value of a distribution, irrespective of whether there are other people present. Telic Egalitarianism, the rival view that Otsuka and Voorhoeve appear to favour, denies this claim. That is, the fundamental difference between Prioritarianism and Telic Egalitarianism is whether the contribution of an individual’s utility to the moral value of an outcome depends on whether the outcome includes other people as well.

We contend that Otsuka and Voorhoeve’s arguments rely on the assumption that the relevant core commitments of Prioritarianism are false. It is unsurprising that Prioritarianism comes out false once those assumptions have been made, as those are the very points of dispute between Prioritarianism and its main rival views. They cannot be assumed to be false when trying to settle which view is right from a dialectically neutral standpoint. The real philosophical work to be done is in that very task itself.

References

Broome, John. (1991). *Weighing Lives*. Oxford: Wiley-Blackwell.

Broome, John. (2015). “General and Personal Good: Harsanyi’s Contribution to the Theory of Value”. In *Oxford Handbook of Value Theory*, eds. Iwao Hirose & Jonas Olson. Oxford: Oxford University Press.

Crisp, Roger. (2003). “Equality, Priority, and Compassion”. *Ethics* 113(4): 745-763.

Crisp, Roger. (2011). “In Defence of the Priority View: A Response to Otsuka and Voorhoeve”. *Utilitas* 23(1): 105-108.

- Otsuka, Mike & Voorhoeve, Alex. (2009). "Why It Matters That Some Are Worse off than Others: An Argument against the Priority View". *Philosophy & Public Affairs* 37(2): 171-199.
- Otsuka, Mike & Voorhoeve, Alex. (2011). "Reply to Crisp". *Utilitas* 23(1): 109-114.
- Parfit, Derek. (1991). "Equality or Priority". *Lindley Lecture*.
- Rabinowicz, Wlodek. (2002). "Prioritarianism for Prospects". *Utilitas* 14(1): 2-21.