

ON GIVING PRIORITY  
TO  
POSSIBLE FUTURE PEOPLE

Nils Holtug

Philosophy Section

Department of Media, Cognition and Communication

nhol@hum.ku.dk

ABSTRACT: According to prioritarianism, roughly, it is better to benefit a person, the worse off she is. This seems a plausible principle as long as it is applied only to fixed populations. However, once this restriction is lifted, prioritarianism seems to imply that it is better cause a person to exist at a (positive) welfare level of  $l$  than to confer  $l$  units on a person who already exists and is at a positive welfare level. Thus, prioritarianism seems to assign too much weight to the welfare of possible future people. It is in this respect even more demanding than total utilitarianism. However, in this article, I argue that all told, prioritarianism is in fact more plausible than total utilitarianism even when it comes to population ethics.

## **1. Introduction**

It is a great pleasure for me to be able to contribute to the celebration of Wlodek on his sixtieth birthday. I have learnt a lot from Wlodek and continue to do so. In this paper I shall address two issues, both of which Wlodek has discussed in his work, namely prioritarianism and population ethics. However, as far as I know, Wlodek has not combined them, which is what I shall do here. More precisely, I shall consider some of the implications of prioritarianism for population ethics.

Let me first say a little about the framework I shall employ. I am concerned with axiology. My claims will be about the goodness of outcomes and in particular, which outcomes are better than others. Furthermore, I am concerned only with the distribution of welfare. I shall simply assume that welfare is the currency of justice. So the version of prioritarianism I shall consider is one that ranks outcomes with respect to the goodness of the welfare distributions they hold.

Unlike utilitarianism, prioritarianism is distribution-sensitive. It implies that it morally matters how a certain sum of welfare is distributed on individuals. Everything else being equal, we should favour a benefit to a worse off person to an equal benefit to a better off person. Many theorists believe, myself included, that at least when we restrict ourselves to fixed populations, such sensitivity renders prioritarianism more plausible than utilitarianism.

Furthermore, unlike (axiological) egalitarianism, prioritarianism is not vulnerable to the so-called levelling down objection.<sup>1</sup> According to this objection, egalitarianism implausibly implies that it is in one respect better to decrease the welfare of some and increase the welfare of none, so that everyone is at the level of the worst off. But how can it be in any respect better to level down, when there is no one for whom it is in any respect better? Prioritarianism, on the other hand, has no such implication and this may seem to speak in favour of this latter view.<sup>2</sup>

However, when we apply these distributive principles to population ethics (and so lift the restriction to fixed populations), it is not so clear that prioritarianism is the superior view. In fact, I once claimed that, with respect to population ethics (although not in general), total utilitarianism is more plausible than prioritarianism.<sup>3</sup> I now believe that this was a mistake. Even with respect to population ethics, prioritarianism is the more plausible view. In what follows, then, I shall consider the implications of prioritarianism for population ethics. In particular, I shall consider its implications with respect to the so-called repugnant conclusion.

## **2. Prioritarianism**

Consider:

---

<sup>1</sup> Parfit (1991), p. 17.

<sup>2</sup> I defend these claims against various objections in Holtug (2007).

<sup>3</sup> Holtug (1999), p. 36.

*Prioritarianism.* An outcome is (non-instrumentally) better, the larger a sum of weighted individual benefits it contains, where benefits are weighted such that they gain a greater value, the worse off the individual to whom they accrue.

Prioritarianism implies that moral outcome value is a strictly concave function of welfare, as illustrated in figure 1:

*Figure 1*

Moral outcome value



Welfare

Thus, benefits *gradually* decrease in moral value, the higher the level at which they fall. Also, prioritarianism is an *aggregative* principle and I shall assume that it implies that for any finite sum of benefits that fall at a lower level, it can be outweighed by a sufficiently large sum of benefits that fall at a - indeed any - higher level. In this respect, it differs from a

principle that would give absolute priority to the very worst off.<sup>4</sup>

More precisely, the sort of aggregative function applied to interpersonal (weighted) benefits is an additive function. Of course, other functions could be used, but I shall simply assume this particular one.<sup>5</sup>

Finally, I shall assume a ‘whole lives’ interpretation of prioritarianism. That is, I shall assume that an individual is worse off than another – in the relevant sense – if and only if the former individual has less welfare in her life taken as a whole. So the fact that an individual is worse off than another at a particular point in time will not automatically imply that the former individual should have priority (at that time).

### **3. Prioritarianism and the repugnant conclusion**

While prioritarianism is usually applied only to fixed populations, it can also be applied to outcomes that differ with respect to the identities and number of people who inhabit them. Assuming a ‘whole lives’ version of prioritarianism, we can take the total life-time welfare that will accrue to an individual if she comes into existence, divide it into welfare units and

---

<sup>4</sup> Thus, on my characterisation, neither maximin nor leximin is a version of prioritarianism.

<sup>5</sup> For a much more detailed characterisation of prioritarianism, see Holtug (2007).

gradually assign less weight to these units. So the first unit has the highest moral value, the second a lower such value and so on.<sup>6</sup>

On these assumptions, there is a sense in which prioritarianism assigns priority to possible future people over people who already exist. If we can either bestow, say, ten units of welfare on an individual by bringing her into existence or bestow ten units on an individual who already exists and has a positive welfare level, prioritarianism implies that everything else being equal, it is better to do the former, that is, bring the first person into existence. This is where the ten welfare units contribute most moral value. (These claims, however, should not be confused with a claim to the effect that the ‘modal’ status of a person somehow has intrinsic significance. What does have intrinsic significance is how well or badly off a person is and her modal status is relevant only because of its impact on this.)

More generally, prioritarianism implies:

*The repugnant conclusion.* A world populated by individuals every one of whom has a life barely worth living would be better than a world populated by (say) ten billion individuals all of whom have very


---

<sup>6</sup> Note that these claims do not presuppose that a person can benefit from coming into existence. I do in fact believe that a person can so benefit (Holtug 2001) and so that there are both impersonal *and* person-affecting versions of prioritarianism that cater for increases in welfare that are due to the existence of extra people, but this is not essential here.

worthwhile lives — as long as the former population is sufficiently large.<sup>7</sup>

Thus, in figure 2 (where the columns represent groups of people; the height represents their welfare level and the length represents the number of people in the group), prioritarianism implies that B is better than A:

*Figure 2*



To see this, consider first the implications of total utilitarianism. According to total utilitarianism, B is better than A if only B contains a larger sum of welfare. Unlike total utilitarianism, prioritarianism gradually assigns less weight to benefits, the higher the level at which they fall. What this means is that it will take fewer people living lives barely worth living to counterbalance the value of A than it will on total utilitarianism. After all, on average, the benefits in B fall at a lower level than they do in A. Thus, unlike total utilitarianism, prioritarianism implies:

---

<sup>7</sup> Parfit (1984), p. 388.

*The super-repugnant conclusion.* A world populated by individuals every one of whom has a life barely worth living would be better than a world populated by (say) ten billion individuals all of whom have very worthwhile lives, even if the former population has a lower total sum of welfare — if only the former population has an appropriate size.

Therefore, while prioritarianism may well be more plausible than total utilitarianism when applied to fixed populations, it would seem to be less plausible when applied to population ethics.

#### **4. Attempts to rebut the repugnant conclusion**

Various attempts have been made to modify utilitarianism so that it no longer implies the repugnant conclusion. Here is a list of some such attempts: (1) average utilitarianism,<sup>8</sup> (2) temporal restrictions: only the welfare of present persons contributes to outcome value,<sup>9</sup> (3) modal restrictions: only the welfare of necessary (alternatively, actual) persons contributes to outcome value,<sup>10</sup> (4) variable value views: a person's contribution to outcome value depends not just on her welfare, but also on

---

<sup>8</sup> For discussion, see Parfit (1984), pp. 420-422.

<sup>9</sup> Narveson sometimes writes as if this is his view, see (1976), p. 68.

<sup>10</sup> See e.g. Heyd (1992), Bigelow and Pargetter (1988), and Warren (1978) for defences of various modal restrictions.



population size,<sup>11</sup> (5) critical level theories: there is a (positive) welfare level such that only people who are above this level contribute (positively) to outcome value.<sup>12</sup> But I do not find any of these attempts particularly promising.<sup>13</sup>

However, the point I want to make here is simply the more modest one that each of these attempts is also available to prioritarrians. Just as utilitarians may consider outcome value a function of average welfare, prioritarrians may consider it a function of average *weighted* welfare. And just as utilitarians may impose temporal or modal restrictions, so may prioritarrians. Likewise, prioritarrians may render outcome value sensitive not just to weighted welfare but also to population size, and they may impose a critical level of weighted welfare such that only people who are above this level contribute (positively) to outcome value. However, since I do not find any of these suggestions particularly plausible, I shall not consider them any further here.

## **5. Pluralist prioritarianism**

It would seem as if, relative to total utilitarianism at least, prioritarianism does well when applied to fixed populations, but less well once this

---

<sup>11</sup> Hurka (1983).

<sup>12</sup> Blackorby, Bossert and Donaldson (1997), Broome (2004).

<sup>13</sup> For a good critical account of many such attempts, see Arrhenius (2000). I, myself, critically discuss modal restrictions in Holtug (2004).

restriction is lifted.<sup>14</sup> So consider the idea that prioritarianism should be combined with a different principle to obtain an adequate comprehensive axiological theory.


Here is one version of this idea. Prioritarianism should be applied only to necessary individuals and a complementary principle – e.g. total utilitarianism (which, after all, does not imply the *super*-repugnant conclusion) - only to contingent ones. An individual is necessary relative to the comparison of two outcomes if and only if she exists in both. And an individual is contingent relative to the comparison of two outcomes if and only if she exists in one but not the other. Call this principle ‘pluralist prioritarianism’. Assuming, for simplicity, that the two populations referred to in the super-repugnant conclusion consist entirely of contingent people, pluralist prioritarianism does not imply this conclusion.

Unfortunately, pluralist prioritarianism violates the transitivity of the betterness relation. Compare the following three outcomes:

---

<sup>14</sup> This section derives from Holtug (1999).

Figure 3



Here, A holds a higher total than both B and C (the slashed line represents the average welfare in A, which is higher than in B and C, and since the three outcomes contain equal numbers of people, A also holds the highest total). Now, B is at least as good as A. This is because, since both  $p$  and  $q$  are necessary, we should assess the contributions of both groups in accordance with prioritarianism. Furthermore, while  $p$  loses more in B than  $q$  loses in A, we should give priority to benefits to  $q$ . And so, assuming an appropriate level of priority, B is at least as good as A.

Also, C is at least as good as B. This is because B and C are equally good with respect to both necessary and contingent persons. Therefore, the transitivity of the betterness relation implies that C is at least as good as A. However, according to pluralist prioritarianism, A might very well be better than C. After all, in this comparison,  $q$  and  $r$  are contingent and so benefits to them should count in accordance with total utilitarianism. Although  $r$  gains more in C than  $q$  in A,  $p$  loses more in C than the difference between  $r$  and  $q$ . Furthermore,  $p$ 's gain from C to A counts with priority. We may imagine that  $p$  is still not very well off in C, such that all of this group's

gain from C to A counts more than it would on total utilitarianism. Hence, the transitivity of the betterness relation is violated.<sup>15</sup>

Note, incidentally, that what generates the intransitivity is that welfare units have different values in different contexts, that is, in different comparisons. And this feature will of course be present even if we were to combine prioritarianism with a principle different from total utilitarianism (with which to assess the welfare contributions of contingent persons).

In the following, then, I shall assume that prioritarianism applies both to necessary and contingent persons. On this assumption, prioritarianism implies not only the repugnant, but also the super-repugnant conclusion. Nevertheless, I now want to suggest that even with respect to population ethics, prioritarianism is a more plausible principle than total utilitarianism.

## **6. The negative repugnant conclusion**

Just like total utilitarianism, prioritarianism implies what John Broome has dubbed:<sup>16</sup>

---


<sup>15</sup> Of course, not everyone agrees that betterness is a transitive relation; see e.g. Temkin (1987), (1996). I cannot go into this discussion here. For a defence of the transitivity of betterness, see Broome (2004), pp. 50-63. See also Holtug (forthcoming).

<sup>16</sup> Broome (2004), p. 213. See also Carlson (1998).

*The negative repugnant conclusion.* A world populated by individuals every one of whom has a life barely worth not living would be worse than a world populated by (say) ten billion individuals all of whom have extremely miserable lives (lives very much worth not living) — as long as the former population is sufficiently large.

To illustrate, consider outcomes A and B in figure 5 (where the line connecting the columns represents the level where life ceases to be worth living):

*Figure 5*



However, while both prioritarianism and total utilitarianism imply the negative repugnant conclusion, prioritarianism requires more people in B to render this outcome worse than A than does total utilitarianism. According to total utilitarianism, the sum of negative welfare in B need only exceed the sum of negative welfare in A to render B worse. However, according to prioritarianism, the units of negative welfare that accrue to individuals in A

on average count more than the units of negative welfare that accrue to individuals in B. This is because, on average, these units fall at lower levels in A. Therefore, in order for B to be worse, it is not sufficient that this outcome includes a higher sum of negative welfare. The negative sum in B must exceed the negative sum in A by a certain *amount*.

Thus, whereas prioritarianism does - and total utilitarianism does not - imply the super-repugnant conclusion, prioritarianism does better than total utilitarianism with respect to the negative repugnant conclusion.

Furthermore, those of us who are inclined towards prioritarianism may well find that we are more disturbed by the negative than by the (positive) repugnant conclusion. After all, we will find it more important to raise an individual from, say, -10 to -1 than to raise an individual from 1 to 10 (although, in each case, an individual is raised by nine units). And so there is a sense in which we will be more concerned about the difference between the worse off and the better off in the negative repugnant conclusion than about the difference between the worse off and the better off in the (positive) repugnant conclusion.

## **7. The weak asymmetry**

There is also another, related aspect with respect to which prioritarianism seems more appealing than total utilitarianism when applied to population ethics. Consider what has been called:

*The asymmetry.* While it detracts from the value of an outcome to add individuals whose lives are of overall negative value, it does not

increase the value of an outcome to add individuals whose lives are of overall positive value.<sup>17</sup>

This principle has seemed attractive to several population ethicists but, amongst other things, it implies that, everything else being equal, it is better if we put an end to the existence of humans by refusing to have children. After all, the misery of the children who would have lives worth not living cannot be counterbalanced by the happiness of the children who would have lives worth living, because the happiness of the latter children would count for nothing. This will be so even if there were only one child whose life would be worth not living, as surely there would be.<sup>18</sup>

However, prioritarianism implies a weaker, more plausible version of the asymmetry:

---

<sup>17</sup> The label is McMahan's (1981), p. 100. Strictly speaking, McMahan formulates the Asymmetry in terms of reasons rather than the value of outcomes, but for the present axiological purposes, my formulation is more appropriate. The asymmetry has been famously defended by Narveson (1967), pp. 69-71.

<sup>18</sup> It may of course be argued that the interests of presently existing people in having children would outweigh the misery of one child. But we may imagine that, for some reason, existing people are indifferent with respect to whether or not they have children. Alternatively, we may imagine that there are so many future generations that, on any plausible welfare function, the unhappiness of the miserable members of these future generations outweighs the interests of the present generation.

*The weak asymmetry.* Everything else being equal, it is better to avoid that a person comes into existence and has a life worth not living (at level  $-l$ ), than to ensure that a person comes into existence and has a life worth living (at level  $l$ ).

The weak asymmetry follows from prioritarianism because, according to this principle, benefits (and the avoidance of harms) matter more at lower levels. So each negative unit of welfare counts more than each positive unit. Furthermore, if we bring a happy individual into existence, the units of positive welfare that are thereby created gradually count less as they are piled on top of each other, whereas if we bring a miserable individual into existence, the units of negative welfare that are thereby created gradually count more.

Incidentally, this asymmetry will tend to somewhat weaken what seems to be a formidable prioritarian case for population growth. After all, presumably, amongst the people we may create some will have lives worth not living and, everything else being equal, the miserableness of such a life weighs more heavily than the happiness of a happy person we can cause to exist. In fact, since a population growth of the magnitude we are now considering would lead to a significantly lower standard of living than that with which we are familiar today, presumably *many* of the lives we would create would be below the level where life ceases to be worth living.

Furthermore, since the standard of living of the present generation is rather unevenly spread (to say the least), many presently existing individuals are likely to fall below the level where life ceases to be worth living. And on prioritarianism, of course, such lives below zero are especially bad, since we should give priority to the worse off.



Finally, a significant decrease in the average standard of living is likely to lead to various kinds of conflict, including rebellion, theft, violence and war. Especially since, as I have just pointed out, resources are very unevenly spread and so a significant reduction in the average standard of living will reduce some (many) individuals to a very low level indeed.<sup>19</sup> Presumably, this would be so even if, when reducing the average standard of living, we *aimed* for an equal distribution. Many individuals would come to live at a level comparable to that of a poor, malnourished Ethiopian. So in light of the further disastrous consequences of rebellion, violence, war etc. for welfare, the sort of population growth now under consideration may not be a very good idea, even in prioritarian terms.

## **8. Higher level repugnant conclusions**

Let me now turn to a final attractive feature of prioritarianism. Consider again the repugnant conclusion, more specifically figure 2. Suppose we increase the welfare level in A and B while keeping the difference in average welfare constant. Suppose, for instance, that while everyone in B has an *excellent* life, everyone in A has an *ecstatic* life. Here, the conclusion that there must be some number of individuals such that, if they

---

<sup>19</sup> For a similar argument against vast increases in population size, see Hare (1993), p. 79.

populated B, B would be better than A does not seem very disturbing (and in any case less disturbing than the ‘original’ repugnant conclusion).

In a sense, total utilitarianism accommodates this claim. It would take less individuals in B at this higher level to counterbalance the value of A than it would at a lower level (because, when we raise the average by an equal amount in A and B, we raise this level for more individuals in B).

However, this may not be the whole explanation of why ‘repugnant’ conclusions are less repugnant at higher levels. It seems to me that another reason is that the worse off are better off in higher than in lower level repugnant conclusions. Prioritarianism captures this reason by giving priority to the worse off. When we raise the average in both A and B, the increase in welfare that accrues to every individual in B falls at a lower level than the increase that accrues to every individual in A and so has greater moral value.

Thus, perhaps prioritarianism provides a more adequate explanation of why repugnant conclusions are less repugnant at higher levels.

## **Conclusion**

Unlike total utilitarianism, prioritarianism implies the super-repugnant conclusion. However, with respect to the negative repugnant conclusion, prioritarianism is the more plausible view. Furthermore, unlike total utilitarianism, prioritarianism implies the plausible weak asymmetry. Finally, prioritarianism seems to provide a better explanation of why repugnant conclusions are less repugnant at higher levels. Thus, all told,

prioritarianism may well be the more plausible view, even when applied to population ethics.<sup>20</sup>

## Literature

Arrhenius, Gustav, (2000), *Future Generations: A Challenge for Moral Theory*, dissertation, Uppsala.

Bigelow, J. and R. Pargetter, (1988), “Morality, Potential Persons and Abortion”, *American Philosophical Quarterly*, vol. 25.

Blackorby, C., W. Bossert and D. Donaldson, (1997), “Critical Level Utilitarianism and the Population-ethics Dilemma”, *Economics and Philosophy*, vol. 13.

Broome, John, (2004), *Weighing Lives*, Oxford: Oxford University Press.

Carlson, Erik, (1998), “Mere Addition and Two Trilemmas of Population Ethics”, *Economics and Philosophy*, vol. 14.

Hare, Richard, (1993), “Possible People”, in his *Essays on Bioethics*, Oxford: Clarendon Press.

---

<sup>20</sup> This paper was presented as part of my talk on “Prioritarianism, Animals and Future Generations” at the Oxford University Moral Philosophy Seminar 2006. I would like to thank the audience for helpful discussion.

Heyd, David, (1992), *Genethics. Moral Issues in the Creation of People*, Berkeley: University of California Press.

Holtug, Nils, (1999), "Utility, Priority and Possible People", *Utilitas*, vol. 11.

Holtug, Nils, (2001), "On the Value of Coming into Existence", *Journal of Ethics*, vol. 5.

Holtug, Nils, (2004), "Person-affecting Moralities", in Jesper Ryberg and Torbjörn Tännsjö (eds.), *The Repugnant Conclusion. Essays on Population Ethics*, Dordrecht: Kluwer.

Holtug, Nils, (2007), "Prioritarianism", in Nils Holtug and Kasper Lippert-Rasmussen (eds.), *Egalitarianism. New Essays on the Nature and Value of Equality*, Oxford: Clarendon Press.

Holtug, Nils, (forthcoming), "A Note on Conditional Egalitarianism", *Economics and Philosophy*.

Hurka, Thomas, (1983), "Value and Population Size", *Ethics*, vol. 93.

Narveson, Jan, (1976), "Moral Problems of Population", in M.D. Bayles (ed.), *Ethics and Population*, Cambridge, Mass.; Schenkman Publishing Company Inc.

Parfit, Derek, (1984), *Reasons and Persons*, Oxford: Clarendon Press.

Parfit, Derek, (1991), *Equality or Priority?*, The Lindley Lecture, University of Kansas.

Temkin, Larry S. (1987), "Intransitivity and the Mere Addition Paradox", *Philosophy & Public Affairs*, vol. 16.

Temkin, Larry S. (1996), "A Continuum Argument for Intransitivity", *Philosophy & Public Affairs*, vol. 25.

Warren, Mary Anne, (1978), "Do Potential People Have Moral Rights?", in R.I. Sikora & Brian Barry (eds.), *Obligations to Future Generations*, Philadelphia: Temple University Press.