Egalitarianism and the Putative Paradoxes of Population Ethics

TORBJÖRN TÄNNSJÖ
Stockholm University

The repugnant conclusion is acceptable from the point of view of total utilitarianism. Total utilitarians do not seem to be bothered with it. They feel that it is in no way repugnant. To me, a hard-nosed total utilitarian, this settles the case. However, if, sometimes, I doubt that total utilitarianism has the final say in ethics, and tend to think that there may be something to some objection to it or another, it is the objection to it brought forward from egalitarian thought that first comes to mind. But if my argument in this article is correct, then it is clear that the repugnant conclusion should be equally acceptable to egalitarians of various different bents as it is to total utilitarians.

INTRODUCTION

Do egalitarian considerations solve problems in population ethics? In particular, if some egalitarian views are correct, does it mean that we can avoid the most important paradoxes we seem to face in population ethics? No, according to Gustaf Arrhenius, in a recent article.[1] Yes, as far as I can see. However, my approach is different from the one adopted by Arrhenius. Arrhenius wants to show that, if we stick to a rejection of the repugnant conclusion (see below for an explanation), then we are stuck with the paradoxes, even if egalitarian thought is on the right track. But why reject the repugnant conclusion? Why not say that this is a misnomer? If total utilitarianism is correct, then, of course, the 'repugnant' conclusion is all right. It is in no way repugnant. The same is true, however, if egalitarian thinking is right, or so I will argue in the present article. And once we accept the repugnant conclusion, the putative paradoxes of population ethics go away.

THE NATURE OF THE PROBLEM

Problems in population ethics are sometimes cast in terms of 'betterness'. Various different possible populations are compared, and judgements are made about their relations with respect to 'goodness'. One population is better than another, a third population is at least as good as another, and so forth. This makes poor sense. There is a way of making sense of it, however. What we are interested in, in ethics in general, is what we *ought to do*. And judgements about

[1] Gustaf Arrhenius, 'Egalitarianism and Population Change', *Intergenerational Justice*, ed. A. Gosseries and L. Meyer (Oxford, forthcoming).

© 2008 Cambridge University Press
doi:10.1017/S0953820808002987
Utilitas Vol. 20, No. 2, July 2008
Printed in the United Kingdom

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

betterness can be transformed into judgements about obligations. We had better think of ourselves in the place of God, being capable of realizing various different population scenarios. When we say things such that a population of one billion very happy persons is better than a population of two billion moderately happy persons, this should be taken to mean that, if we had a choice between realizing either of these populations, we should opt for the former one. If they are of equal value it means that we are allowed to choose between them as we see fit.

In order to make sense of many of the arguments in population ethics we need to know what it is that matters in a moral assessment of populations. I will take for granted a simple hedonistic view about this: felt well-being is what is at stake. Those who do not share my preference here are free to exchange for happiness whatever *they* feel to be at stake. And I will also assume that we can measure felt well-being on a cardinal scale. It makes sense to say, then, that one person is twice as happy as another person. In some cases this requirement can be relaxed, but, throughout, just in order to simplify, I will stick to it.

EGALITARIAN THINKING

What is egalitarian thinking, then? In the present context we need to distinguish between, on the one hand, relational views to the effect that, somehow, inequality is problematic and, on the other hand, *prioritarian* views to the effect that we should pay special attention to those who are (absolutely speaking) worst off (the priority view). I will discuss one rough version of the relational approach and two different versions of the priority view (I will call them the 'lexical' and the 'aggregate' view respectively).

The number of relational views is legion, since there are many ways of measuring inequality. In the present context these differences will be of no importance. However, in order to be the least plausible, such views must avoid the levelling down objection. I will not discuss egalitarian views according to which, in order to achieve equality, we should take from those who are better off, without any gain at all to those who are worst off.

How can the egalitarian avoid this conclusion? The egalitarian must pay some respect to the sum-total of well-being as well as to equality. The best way of doing this is, it seems to me, would be to maximize some weighed sum of well-being, where the value of the well-being of an individual is a function both of his actual well-being and of his position versus the rest. Many such functions can be designed that satisfy the

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

requirement that the levelling down objection should be avoided.[2] We need not go into detail here.

On one version of the Priority View, the lexical reading of it, we should only focus on those who are worst off. If there is a way of improving their situation, we should do so, whatever the costs to those who are better off (after the improvement has taken place). This is a very demanding view. It requires that we take all the happiness from some and give it to those who live a miserable life, if only we can improve their lot an iota.

I don't know if anyone has really defended this view. It is sometimes attributed to John Rawls (his difference principle), but, if we cast it in terms of individuals and well-being, as I have done here (because it is appropriate to do so), Rawls doesn't really adhere to it. Yet, for all that, I will make reference to it in my argument.

A more plausible version of the Priority View is the following. Total utilitarians are right when they insist that we should aggregate well-being and opt for the sum-total of well-being in our actions, but with one important proviso: we have to weigh the well-being differently. We may speak here of a diminishing marginal utility of utility. The happier a person is, the less important it is to increase the happiness of this person. The less happy he or she is, the more important it is to increase his or her happiness. In our actions we ought to see to it that we maximize the weighed sum-total of well-being. This is the aggregate version of the Priority View.

With these three rough egalitarian ideas at our disposal we may face the problems in population ethics. I start with what Derek Parfit has called the mere addition paradox. I will discuss it in Arrhenius' more general formulation, however.

THE MERE ADDITION PARADOX

The diagram below shows three populations: A, B and C. The width of each block represents the number of people, and the height represents their lifetime welfare. All the lives in the diagram have positive welfare. A is a population of people with very high welfare, B is a much larger population than A, but consists of people with very low positive welfare. C is a population of the same size as A∪B (that is, the population consisting of all the lives in both A and B). Everybody in C has very low positive welfare but they are all better off than the people in B. Moreover, there is perfect equality in C and the total and average welfare in C is higher than in A∪B.

[2] See Wlodek Rabinowicz, 'The Size of Inequality and its Badness – Some Reflections around Temkin's Inequality', *Theoria* 69 (2003).

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

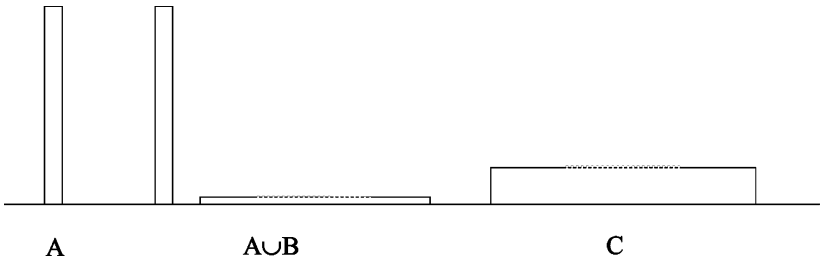


Figure 1.

How should we rank these populations? Following Arrhenius, consider first population A and A∪B. Since the B-people have lives worth living, many would agree that A∪B is at least not worse than A. As was indicated above, I take this to mean that it would not be wrong to opt for A∪B rather than for A, if it were up to me to decide which population would become real. The rationale behind this judgement is the following condition:

The Mere Addition Condition: An addition of people with positive welfare does not make a population worse, other things being equal.

What about A∪B and C? Again, following Arrhenius, since there is perfect equality in C and higher total and average utility in C as compared to A∪B, it seems reasonable to claim that C is better than A∪B. Again, in my words, if we had a choice, we should choose C rather than A∪B. The rationale behind this judgement is articulated in the following way:

The Non-Anti-Egalitarianism Condition: A population with perfect equality is better than a population with the same number of people, inequality, and lower average (and thus total) welfare, other things being equal.

Finally, how should we opt between A and C? According to Derek Parfit, to accept that C is better than A would be to accept a 'repugnant' conclusion. Arrhenius expresses this conclusion, in a slightly generalized form, as follows:

The Repugnant Conclusion: For any perfectly equal population with very high positive welfare, there is a population with very low positive welfare which is better, other things being equal.

We should avoid this conclusion, then, and instead conclude that A is better than C. In the choice between A and C we ought to opt for A. And the rationale behind this judgement is as follows.

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

The Quality Condition: There is at least one perfectly equal population with very high positive welfare, which is better than any population with very low positive welfare, other things being equal.

According to Arrhenius, this means that we contradict ourselves.

If C is better than A∪B, and A is better than C, then by transitivity of 'better than', it follows that A is better than A∪B. But we said that A∪B is not worse than A, that is, A is not better than A∪B. Hence, these valuations imply a contradiction: A is better than A∪B and A is not better than A∪B.

This is not entirely correct, however, if we accept my understanding of value terms. However, if I have a choice between A, A∪B and C, it is, if not a contradiction, at least very strange, to conclude that, between A and A∪B, I may choose A∪B, while, between A∪B and C, I must choose C, and then go on to conclude that, in the choice between A and C, I must choose A. If I face all three alternatives in a situation, it is true of me that, now matter how I choose, I will have made the wrong choice. This means that the rightness and wrongness of my actions depend on how I actually choose. In a sense, there are genuine moral dilemmas, then. In the end we may find this acceptable, but if we could avoid having to conclude this, we should try to do so.

Is there a way of avoiding this conclusion? Arrhenius tries to show that egalitarian thinking cannot block the conclusion that we are allowed to choose A∪B over A, or at least that, if we look at other ways of generating the paradox, egalitarian thinking will be of no avail to us. However, it is then taken for granted that we must prefer A to C. But why should we prefer A to C? In what way is the repugnant conclusion repugnant?

According to total utilitarianism we ought to accept the repugnant conclusion and, once we do, the paradox disappears. But some people seem to find the repugnant conclusion repugnant. Is there a way of showing them that it is not? In particular, is there a way for an egalitarian to show that we ought to accept the repugnant conclusion? Indeed there is and, once we realize this, the paradox will go away; that is, not only will total utilitarians find the repugnant conclusion acceptable, but egalitarians will as well. Let us see why.

THE SIMPLE EGALITARIAN ARGUMENTS

A simple (simplistic?) way of showing that C is better than A is to have recourse to the second reading of the Priority View. According to this view, we should opt for the world with the greatest sum-total of weighed well-being, where low scores on the scale weigh heavier than high scores. If this kind of reasoning is applied to the choice between A and C it is obvious that C is better than A. The difference between C

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

and A is even greater than has been acknowledged by total utilitarians, who have neglected the diminishing utility of utility.

It must be noted, of course, that on the lexical reading of the same view, we get the opposite result. Those who are worst off in C are much worse off than those who are worse off in A. But we noted that this is probably the least plausible egalitarian view (among the three I have put forward in the present context).

Finally, on the relational view, we must conclude that there is no difference between A and C with respect to equality (both are perfectly equal), so one of the concerns held by a relational egalitarian seems to indicate that indifference is the right attitude towards A and C. However, we have also noted that relational egalitarianism, in the most plausible version, pays interest to the sum-total of well-being as well as to how it is distributed. So this means that the relational egalitarian has a reason to opt for C rather than A.

It seems, then, that two egalitarian views out of three agree that C is better than A, i.e. they imply that we may endorse the repugnant conclusion.

RESTRICTED EGALITARIANISM

I indicated, however, that these arguments might be not only simple, but simplistic as well. As a matter of fact, it seems as though many egalitarians, who have considered them, have rejected them. So what is wrong with them? The received wisdom seems to be that egalitarian considerations are not really applicable to problems in population ethics, at least not to problems where we decide between *different* populations. Egalitarian considerations are only relevant when we have a *fixed* population and when we can distribute well-being in different ways between the members of this population.

To explain this restriction, let us refer to people who live in all alternative outcomes as 'necessary' people (relative to a choice). We can then see that there are two possible versions of the restriction. On a narrow understanding of it, egalitarian considerations are irrelevant to a choice as soon as it involves some not necessary but merely 'contingent' persons. On this understanding, egalitarian considerations are altogether irrelevant to problems of population ethics. On a less narrow understanding, egalitarian considerations are relevant as soon as a choice concerns some necessary persons. Then egalitarian considerations apply to *them*. This means that, when contingent persons are involved, the egalitarian judgement about the situation is not final. It must somehow be supplanted by other considerations. This is as it should be. And if there are egalitarian considerations indicating that we should opt for A rather than B, and if there are

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

no countervailing considerations to do with merely contingent persons, then we ought to opt for A, according to egalitarian thought. We could also restrict this view to same number cases (irrespective of whether the people involved are necessary or contingent). I see no merit in thus restricting the view, however; it seems to me to be too wide (in allowing that we apply the view to contingent people) and too narrow (in that it has nothing at all to say when the numbers in different outcomes vary) and I will not discuss it.

For the sake of the argument, let us accept the less narrow restriction (which, by the way, strikes me as more plausible than the narrow one), and turn to more complex egalitarian arguments in defence of the repugnant conclusion.

THE COMPLEX ARGUMENTS

Let us now leave the mere addition paradox and consider a different possible route to the repugnant conclusion. Assume that we have a situation (I) with one population, A, consisting of one billion very happy people. However, some members of A are happier than others. As a matter of fact, there is much variation (at a high level) and also a considerable gap between those who are best off and those who are not as well off. Now, compare this scenario with another one, (II), where another billion people have been added to the A-population. We may speak of them as the B-people. However, the addition of the B-people is not a matter of mere addition. The addition of the B-people has an impact on the situation of the A-people. They are necessary people, in relation to this choice, and they get their hedonistic status improved, when the B-people are added. However, the hedonistic status of the A-people is not improved in an equal manner. As a matter of fact, when the B-people are added, the A-people come to live on an *equally* high level of well-being, *higher* than *anyone* of them did in scenario I. The added contingent (with respect to this option) B-people will enjoy lives worth experiencing, and equally worth experiencing, but at a level considerably lower than the level at which the A-people are living. Ought we to make the move from I to II?

From a total utilitarian point of view, then answer is simple: yes. The sum-total of well-being in II is larger than the sum-total of well-being in I. But what should egalitarians say about this?

On the most plausible understanding of the various different egalitarian views, the egalitarians should reason in the following manner.

The move from I to II means, from an egalitarian point of view, improvement. Those necessary people who live in both scenarios, to which egalitarian thought is applicable, the A-people, get their situation improved, and they get it

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

improved in a manner that, among them, establishes perfect equality. The weighed sum-total of well-being among them is increased. Furthermore, the lot of those who are worst off among the A-people get their situation improved most. So there are strong egalitarian reasons of all sorts to prefer II to I.

But what about the B-people? Well, on the chosen interpretation of egalitarian thought, egalitarian thought has nothing to say about them coming into existence. We could still query, of course, whether something wrong is being done to them, when they are added. Certainly, nothing wrong is being done to the A-people, but what about the B-people? Have we wronged *them* by bringing them into existence?

How can we settle this question? I suggest that we ask them. Suppose they answer that they are happy to be around since, as they correctly report, they live lives worth living.

It might be objected that the world, with which we started, is different from Parfit's A world, in that the A-people live on *different* levels of well-being. This is true, but we may start with a world where everyone lives on the same very high level, and improve the lot of each and every one in a manner that takes us to my starting point, from which we can run the argument. I cannot see that anyone could object to such a Pareto-improvement, neither from an egalitarian point of view nor from any other point of view. A move in the opposite direction would amount to a case of levelling down, and, as we have noted, no plausible egalitarian view can accept that equality is obtained through levelling down.

I conclude that there are strong egalitarian reasons to move from I to II, then, and no non-egalitarian reasons not to do so. Now, consider the following option.

We now start with the A- and B-people. Assume that they live and fare as was described in scenario II. We are now invited to contemplate scenario III. In scenario III two billion people, the C-people, have been added to the existing population of A- and B-people. Now, in relation to this choice, the A- and B-people are necessary people while the added C-people are contingent. Again, the addition is no matter of 'mere' addition. This time, the addition has the following effect on the necessary A- and B-people. The A people get their hedonic status slightly worsened. However, the B-people get their hedonistic status much more improved. The net result is that both the A-people and the B-people live at the same level of well-being. The sum-total of well-being among the A-people and the B-people has been increased. The contingent C-people all live lives worth experiencing, but at a lower level of well-being than the necessary A- and B-people.

Again the total utilitarian will not hesitate to state that we ought to move from scenario II to scenario III. The sum-total of well-being is (much) larger in III than in II. But what should the egalitarian say

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

about this? Well, once again there are strong egalitarian reasons to make the move.

First of all, considering the individuals to whom egalitarian principles are applicable, i.e. to the necessary A- and B-people, we can notice, when we move from II to III, that perfect equality is established. It is true that the A-people lose some well-being, but their loss is compensated for by the increase of the well-being among the B-people; it is compensated for if we just sum gains and losses, and it is even more compensated for, if we weigh their gains and losses of well-being in an egalitarian manner; the same is true according to the aggregate version of the Priority View, when we sum weighed gains and losses. Finally, those who are worst off in III (all the A- and B-people) are better off than those who are worst off in II (the B-people). So the move is sanctioned as well by the lexical version of the Priority View.

Once again we may query whether it was all right to create the additional contingent people (in this case the C-people). And once again, the best way to find out whether they have been wronged when they were brought into existence would be to ask them. We assume again that they say they are happy to be around, having noted correctly that they live lives worth experiencing.

In a similar vein, we argue until we have come close to the repugnant conclusion.

Now it might be tempting to argue that, unless the added contingent people live at a *sufficiently high level of well-being*, their introduction should be avoided. And this is true even if they live lives that are worth living. It is true in particular if they live lives that are barely worth living.

But this argument is a non-starter. It leads us to the sadistic conclusion that it may be better to add a few people who are tormented, than very many people living lives worth living. This is pure sadism and has nothing going for it whatever.[3]

THE ORIGINAL COMPARISON

We have seen that, on total utilitarianism, the repugnant conclusion is acceptable. In the original example, the mere addition paradox, the total utilitarian can accept that the C-world is better than (to be preferred to) the A-world. And this is true even in a situation where the members of the C-world live lives that are barely worth living, provided there are enough of them around. This world, when compared to a world with at least ten billion people who are extremely happy (the A-world), has been called, for obvious reasons, the Z-world.

[3] See Arrhenius, 'Egalitarianism and Population Change'.

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

We have now seen that there exist a series of pairwise choices leading from the A world, down the moral alphabet, to the Z-world, which would be endorsed by egalitarian thought. This means that the egalitarian must reject the quality condition. The egalitarian should accept the repugnant conclusion. This does not mean that egalitarianism has collapsed into utilitarianism, however. In some cases, as we will see below, they are extensionally equivalent, in others they are not.

What are the egalitarians to say about a choice between the A-world and the Z-world, when there are no necessary persons involved? We have seen that, if we skip the restriction to necessary persons, those who adhere to the priority view in the aggregate version, inviting us to maximize a weighed sum-total of well-being, will opt for the Z-world. It is even more superior, on their understanding, than the total utilitarian has recognized. Those who adhere rather to relational views will make the same judgement: Z should be preferred to A. In both the A world and the Z world there is perfect equality but the sum-total of well-being (which is part of their concern) is larger in the Z-world, so they will chose it over the A-world. As a matter of fact, this version of egalitarianism collapses into total utilitarianism when all the outcomes are equal. The lexical prioritarian view, on the other hand, would recommend the A-world, of course. This view in its unrestricted form pays no interest to the sum-total of well-being. But I can't help finding this view extremely implausible. It seems to transform the question: 'Which world is the best one?' into the question 'In which world would I like most to live?' But this is not what population ethics is all about.

However, we have seen that most adherents of egalitarian thought want to restrict their views to necessary people. How would this affect their choice between A and Z, when there are no necessary people involved in the choice? They would have to say that there are no egalitarian considerations pointing in any direction. Where does this answer take them? Should they conclude that the A world and the Z world are not comparable? Or, perhaps more plausibly, should they just say that, from an egalitarian point of view, there is nothing to say about this choice, and then have it decided on other grounds? Which other grounds, in that case?

Well, to egalitarians of the relational variety, and to aggregate prioritarians who are interested in maximizing a weighed sum-total of well-being, the most natural candidate would be total utilitarianism. They would conclude that Z is the better option since it contains more well-being and since there are no egalitarian reasons not to opt for it.

Consider finally a situation where we have a choice between the A-world and the Z-world and where all the A-people are necessary people. The rest of the people in the Z-world are contingent. Now, what would the egalitarian have to say about the choice between them?

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

In the circumstances, there would be strong egalitarian considerations counting against a move from the A-world to the Z-world. They mostly focus on the necessary people, and the sum-total of *their* (weighed) well-being will be small in Z as compared to what it would be in A; on the relational view, we could calculate it directly from their actual well-being, since there is no inequality. On the Priority View in the aggregate version where it invites us to maximize the weighed sum-total of well-being, we calculate it differently, and the difference would be less between the two worlds, but, certainly, necessary people in the Z world would live lives that are worse, also from this perspective, than people living in the A-world. Finally, from the point of view of the lexical prioritarian view, there is no denying that those necessary people who are worst off would be much worse off in the Z-world than they would be in the A-world. So I suppose that, in a choice such as this one, most egalitarians would rather opt for the A-world than the Z-world.

Actually, the fact that they would opt for the A-world rather than the Z-world, where there is a choice between these two worlds, and where the A-people are all necessary people, may help to explain why, falsely, many seem to believe that we have to reject the repugnant conclusion. Those who find it repugnant may hold egalitarian views and, mistakenly, they may be thinking that the option between the A-world and the Z-world must involve necessary people. Given that assumption, they are correct, then, when they claim that, according to their views, the Z-world is worse than the A-world. However, this is consistent with accepting the repugnant conclusion. For there are other instances, where there are no necessary people involved, where the choice should rather be for Z (given the egalitarian considerations combined with some interest in the sum-total of well-being as well). Finally, there is a possible series of pairwise choices leading up from the A-world to the Z-world, where egalitarian considerations will take one from the A-world to the Z-world and where, in each step, there are no sound countervailing reasons to the egalitarian ones.

CONCLUSION

It has often been claimed that population ethics brings with it serious paradoxes. The reaction to this putative fact has been varied. Perhaps we should put the paradoxes to one side and rely on our ordinary thinking in ordinary ethical cases (involving merely necessary people)? Perhaps we should become moral nihilists? Perhaps we should give up the idea of finding a theoretical answer to moral questions? These reactions are not only fanciful – they are premature. The reason why some people have thought that population ethics brings paradox with

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press

it has been the fact that these people have been reluctant to accept the repugnant conclusion. However, this conclusion is in no way repugnant.

We know that the repugnant conclusion is acceptable from the point of view of total utilitarianism. Total utilitarians do not seem to be bothered with it. To me, a hard-nosed total utilitarian, this settles the case. However, if, sometimes, I doubt that total utilitarianism has the final say in ethics, and tend to think that there may be something to some objection to it or another, it is the objection to it brought forward from egalitarian thought that first comes to mind. But if my argument in this article is correct, then it is clear that the repugnant conclusion is equally acceptable to egalitarians of various different bents as it is to total utilitarians.

But then we are allowed to conclude that population ethics doesn't, after all, draw us into paradoxical thinking.

torbjorn.tannsjo@philosophy.su.se

https://doi.org/10.1017/S0953820808002987 Published online by Cambridge University Press