

DEREK PARFIT

OVERPOPULATION AND THE QUALITY OF LIFE

How many people should there be? Can there be overpopulation: too many people living? I shall present a puzzling argument about these questions, show how this argument can be strengthened, then sketch a possible reply.¹

1. QUALITY AND QUANTITY

Consider the outcomes that might be produced, in some part of the world, by two rates of population growth. Suppose that, if there is faster growth, there would later be more people, who would all be worse off. These outcomes are shown in Fig. I.




Fig.1

The width of the blocks shows the number of people living; the height shows how well off these people are. Compared with outcome A, outcome B would have twice as many people, who would all be worse off. To avoid irrelevant complications, I assume that in each outcome there would be no inequality: no one would be worse off than anyone else. I also assume that everyone's life would be well worth living.

There are various ways in which, because there would be twice as many people in outcome B, these people might be all worse off than the people in A. There might be worse housing, overcrowded schools, more pollution, less unspoilt

countryside, fewer opportunities, and a smaller share per person of various other kinds of resources. I shall say, for short, that in B there is a lower quality of life.

Except for the absence of inequality, these two outcomes could be the real alternatives for some country, or mankind, given two rates of population growth over many years. Would one of these outcomes be worse than the other? I do not mean 'morally worse' in the sense that applies only to agents and to acts. But one of two outcomes can be worse in another sense that has moral relevance. It would be worse, in this sense, if more people suffer, or die young.

Would it be worse, in this sense, if the outcome was B rather than A? Part of the answer is clear. We would all agree that B would be, in one way, worse than A: it would be bad that everyone would be worse off.

On one view, this is all that matters, and it makes B worse than A. This view is expressed in

The Average Principle: If other things are equal, it is better if people's lives go, on average, better.

The Hedonistic version of this principle substitutes, for 'go better', 'contain more happiness'.²

On the other main view about this question, it is good if any extra life is lived, that is worth living. On this view B might be better than A. B would be in one way worse, because everyone would be worse off. But in another way B would be better, because there would be more people living, all of whose lives would be worth living. And the fact that people would be worse off might be less important than—or outweighed by—the fact that there would be more people living.

Which of these views should we accept? Could a loss in the quality of people's lives be outweighed by a sufficient increase in the quantity of worthwhile life lived? If this is so, what are the relative values of quality and quantity? These are the central questions about overpopulation.³

The Average Principle implies that only quality matters. At the other extreme is

The Hedonistic Total Principle: If other things are equal, it is better if there is a greater total sum of happiness.

This principle implies that only quantity matters. Its Non-Hedonistic version substitutes, for 'happiness', 'whatever makes life worth living'.

On the Hedonistic Total Principle, B would be better than A because each life in B would be more than half as happy as each life in A. Though the people in B would each be less happy than the people in A, they together would have more happiness—just as two bottles more than half-full hold more than a bottleful. On the

non-Hedonistic version of this principle, B would be better than A because, compared with lives in A, lives in B would be more than half as much worth living. These claims may seem implausibly precise. But lives in B would be more than half as much worth living if, though a move from the level in A to that in B would be a decline in the quality of life, it would take much more than another similarly large decline before people's lives ceased to be worth living. There are many actual cases in which such a claim would be true.⁴

2. THE REPUGNANT CONCLUSION

Consider Fig. 2. On the Total Principle, just as B would be better than A, C would be better than B, D better than C, and so on.

Best of all would be Z. This is an enormous population all of whom have lives that are not much above the level where they would cease to be worth living. A life could be like this either because its ecstasies make its agonies seem just worth enduring, or because it is painless but drab. Let us imagine lives in Z to be of this second kind. There is nothing bad in each of these lives; but there is little happiness, and little else that is good. The people in Z never suffer; but all they have is muzak and potatoes. Though there is little happiness in each life in Z, because there are so many of these lives Z is the outcome in which there would be the greatest total sum of happiness. Similarly, Z is the outcome in which there would be the greatest quantity of whatever makes life worth living. (The greatest mass of milk might be in a vast heap of bottles each containing only one drop.)

It is worth comparing Z with Nozick's imagined Utility Monster.




Fig. 2

This is someone who would gain more happiness than we would lose whenever he is given any of our resources. Some Utilitarians believe that the Hedonistic Total Principle should be our only moral principle. Nozick claims that, on this Utilitarian theory, it would be best if all our resources were taken away and given to his Utility Monster, since this would produce the greatest total sum of happiness. As he writes,

'unacceptably, the theory seems to require that we all be sacrificed in the monster's maw'.⁵

How could it be true that, if all mankind's resources were given to Nozick's Monster, this would produce the greatest total sum of happiness? For this to be true, this Monster's life must, compared with other people's lives, be millions of times as much worth living. We cannot imagine, even in the dimmest way, what such a life would be like. Nozick's appeal to his Monster is therefore not a good objection to the Total Principle. We cannot test a moral principle by applying it to a case which we cannot even imagine.

Return now to the population in outcome Z. This is another Utility Monster. The difference is that the greater sum of happiness would come from a vast increase, not in the quality of one person's life, but in the number of lives lived. And this Utility Monster can be imagined. We can imagine what it would be for someone's life to be barely worth living—containing only muzak and potatoes. And we can imagine what it would be for there to be many people with such lives. In order to imagine Z, we merely have to imagine that there would be very many.

We could not in practice face a choice between A and Z. Given the limits to the world's resources, we could not in fact produce the greatest possible sum of happiness, or the greatest amount of whatever makes life worth living, by producing an enormous population whose lives were barely worth living.⁶ But this would be merely technically impossible. In order to suppose it possible, we merely need to add some assumptions about the nature and availability of resources. We can therefore test our moral principles by applying them to A and Z.⁷

The Total Principle implies that Z would be better than A. More generally, the principle implies

The Repugnant Conclusion: Compared with the existence of very many people—say, ten billion—all of whom have a very high quality of life, there must be some much larger number of people whose existence, if other things are equal, would be better, even though these people would have lives that are barely worth living.⁸

As its name suggests, most of us find this conclusion hard to accept. Most of us believe that Z would be much worse than A. To keep this belief, we must reject the Total Principle. We must also reject the broader view that any loss in the quality of life could be outweighed by a sufficient increase in the total quantity of whatever makes life worth living. Unless we reject this view, we cannot avoid the Repugnant Conclusion.

When the stakes are lower, as in the comparison between A and B, most of us believe that B would be worse. We believe that, compared with the existence of ten billion people whose lives are very well worth living, it would be worse if instead there were twice as many people who were all worse off. To keep this belief, we must again reject the Total Principle.

Suppose that we do reject this principle. Unfortunately, this is not enough. As I shall now argue, it is hard to defend the belief that B would be worse than A, and it is also hard to avoid the Repugnant Conclusion.

3. THE MERE ADDITION PARADOX

Consider the alternatives shown in Fig. 3. There is here a new outcome, A+. This differs from A only by the addition of an extra group of people, whose lives are well worth living, though they are worse off than the original group.

The inequality in A+ is natural: not the result of any kind of social injustice. Take my waves to show the Atlantic Ocean, and assume that we are considering possible outcomes in some past century, before the Atlantic had been crossed. In A+ there was one group of people living in Europe, Asia, and Africa, and another group, who were worse off, living in the Americas. A is a different possible outcome at this time, in which the Americas were uninhabited. Perhaps the Bering Straits had opened before the land was crossed.

Is A+ worse than A? Note that I am not asking whether it is better. If we do not believe that the existence of extra people is in itself good, we shall deny that the extra group in A+ makes A+ better than A. But is A+ worse than A? Would it have




Fig. 3

been better if the extra group had never existed? This is hard to believe. It may seem a bad feature that there is natural inequality in A+—that the extra group are, through no fault of theirs, worse off than the original group. But the inequality in A+ does not seem to justify the view that the extra group should never have existed. Why are they such a blot on the Universe?

You may think that you have no view about whether it would have been better if the extra group had never existed. It may help to consider another outcome:

A + Hell. In this outcome the extra group are innocent people who all have lives which are much worse than nothing. They would all kill themselves if they could, but their torturers prevent this. We would all agree that A + Hell is worse than A. It would have been better if this extra group, as they all passionately wish, had never existed. Since we believe that A + Hell is worse than A, we must be able to compare A+ and A. Unlike the extra group in Hell, the extra people in A+ have lives that are well worth living; and their existence is not bad for anyone. Most of us could not honestly claim to believe that it would have been better if these people had never existed. Most of us would therefore believe that A+ is not worse than A.

Now suppose that, as a result of changes in the environment, A+ turned into Divided B. In both these outcomes the same number of people would exist, so we are not making one of the unfamiliar comparisons which involve different numbers of people in existence. Since the numbers are the same in A+ and Divided B, our ordinary moral principles apply.

On the principles which most of us accept. Divided B would be better than A+. On the Principle of Utility it is better if there is a greater net sum of benefits—a greater sum of benefits minus losses. Divided B would be better than A+ in utilitarian terms, since the benefits to the people who gain would be greater than the losses to the people who lose. On the Principle of Equality it is better if there is less inequality between different people. Divided B would be better than A+ in egalitarian terms, since the benefits would all go to the people who are worse off.

It might be objected that the Principle of Equality does not apply to people who cannot even communicate. But suppose that I know about two such people, one of whom is, through mere bad luck, worse off. Call these people Poor and Rich. I could either benefit Rich, or give a greater benefit to Poor. Most of us would believe that it would be better if I do the second. And we would believe that this would make the outcome better, not only because I would give Poor a greater benefit, but also because he is worse off than Rich. Most of us would believe this even though Poor and Rich cannot (except through me) communicate.

How could we deny that a change from A+ to Divided B would be a change for the better? We would have to claim that the loss to the best-off people in A+ matters more than the greater gain to the equally numerous worst-off people. This seems to commit us to the Elitist view that what matters most is the condition of the best-off people. This is the opposite of Rawls's famous view that what matters most is the condition of the worst-off people.⁹ Most of us would reject this Elitist View. Most of us would therefore agree that Divided B would be better than A+.

Suppose finally that the Atlantic is crossed, turning Divided B into B. These two outcomes are clearly equally good. Since Divided B would be better than A+, B must be better than A+.

Let us now combine the conclusions we have reached. Most of us believe both that A+ is not worse than A, and that B is better than A+. These beliefs together imply that B is not worse than A. B cannot be worse than A if it is better than something—A+—which is not worse than A. In the same way, you cannot be taller than me if you are shorter than someone who is not taller than me. But, as I

earlier claimed, most of us also believe that B is worse than A. We therefore have three beliefs which are inconsistent, and imply a contradiction. These beliefs imply that B both is and is not worse than A. I call this the Mere Addition Paradox.

This is not just a conflict between different moral principles. Suppose that we accept both the Principle of Equality and the Principle of Utility. There can be cases where these principles conflict—where greater equality would reduce the sum of benefits. But such a case does not reveal any inconsistency in our moral view. We would merely have to ask whether, given the details of the case, the gain in equality would be more important than the loss of benefits. We would here be trying to decide what, after considering all the details, we believe would be the better outcome.

In the Mere Addition Paradox, things are different. Most of us here believe, all things considered, that B is worse than A, though B is better than A+, which is not worse than A. If we continue to hold these three beliefs, we must conclude that B both is and is not worse than A. But we cannot possibly accept this conclusion, any more than we could accept that you both are and are not taller than me. Since we cannot possibly accept what these three beliefs imply, at least one belief must go.

Which should go? Suppose that we keep our belief that B is better than A+, because we cannot persuade ourselves that what matters most is the condition of the best-off people. Suppose that we also keep our belief that A+ is not worse than A, because we cannot persuade ourselves that it would have been better if the extra group had never existed. We must then reject our belief that B is worse than A. We must conclude that, if these were two possible futures for some society or the world, it would not be worse if what comes about is B: twice the population, who are all worse off.

The Mere Addition Paradox does not force us to this conclusion. We can avoid the conclusion if we reject one of our other two beliefs. Some people reject the belief that A+ is not worse than A, because they think that the inequality in A+ is enough to make A+ worse. These people can keep their belief that B is worse than A. Note, however, that we cannot simply claim that A+ must be worse than A, since it is worse than something—B—which is worse than A. We would here be rejecting one of our three inconsistent beliefs simply on the ground that it is not consistent with the other two. This could be said against each belief. To avoid the paradox we must believe, without considering the rest of the argument, that A+ is worse than A. We must believe that it was bad in itself that the extra people ever lived, even though these people had lives that were well worth living, and their existence was bad for no one. To the extent that we find this hard to believe, we still face a paradox.

It may be objected: 'Your argument involves a kind of trick. When you compare A and A+, you claim that the extra group's existence was bad for no one. But by the time we have moved to B the original group have become worse off. The addition of the extra group was bad for the original group.'

The argument can be restated. Suppose that A+ was the actual state of the world in some past century. A is a different state of the world which was merely possible. We can ask, 'Would A have been better? Would it have been better if the

worse-off group had never existed?' As I have said, most of us could not answer Yes. Suppose next that A+ did not in fact later change into either Divided B or B. We can ask, 'If this change had occurred, would it have been a change for the better?' It is hard to answer No. On this version of the argument, the last objection has been met. The better-off group in A+ was not an originally existing group, to which the worse-off group was added. And the existence of the worse-off group was not bad for the better-off group.

It is worth giving another version of the argument. To ensure that there was no social injustice, we assumed that the two groups in A+ did not know of each other's existence. We could assume instead that both these groups live in the same society, and that the people in one group are worse off, not because of social injustice, but because they all have some handicap which cannot be cured. Suppose, for example, that they are deaf. If this is so, would it have been better if these people had never existed? Would this have been better even though these people's lives are worth living, their existence is not bad for anyone, and if they had never existed no one else would have existed in their place? It is hard to believe that these deaf people should never have existed. On this version of the argument, it again seems that A+ is not worse than A.

Suppose next that these deaf people could be cured, at some lesser cost to the other group. This would be like the change from A+ to B. It is again hard to deny that this change would make the outcome better. In this version of the argument, with the groups in one society, we seem again driven to conclude that, since B would be better than A+, which is not worse than A, B cannot be worse than A.

There are some other possible objections to this argument. But rather than discussing these I shall turn to another argument. This is harder to answer, and it also leads to the Repugnant Conclusion.

4. THE SECOND PARADOX

Consider the first three outcomes shown in Fig. 4. Though this argument involves many outcomes, we need to make only two comparisons.

One is between A+ and the much more populated Alpha. Suppose that Alpha will be the actual outcome at some time far in the future, after humans have colonized thousands of planets in this Galaxy. A+ is a different possible outcome at this time, in which humans have colonized only one other planet, near a distant star. As before, in neither Alpha nor A+ would the inequality between different people be the result of social injustice. Because of the difficulties of trans-Galactic travel, those who are better off could not raise the quality of life of those who are worse off.

The comparison between A+ and Alpha replaces the comparison, in the old argument, between A and A+. On one view, the natural inequality in A+ makes it worse than A. If I held this view, I would now say:

The inequality in Alpha is in one way worse than the inequality in A+, since the gap between the better-off and worse-off people is slightly greater. But in another way the inequality is less bad. This is a matter of the relative numbers of, or the ratio between, those who are better-off and those who are worse-off. Half of the people in A+ are better off than the other half. This is a worse inequality than a situation in which almost everyone is equally well off, and those who are better off are only a fraction of one per cent. And this is the difference between A+ and Alpha. Because there are so many groups at level 45 (most of them not shown in the diagram), the better-off people in Alpha are only a fraction of one per cent.

To put these claims together: The inequality in Alpha is in one way slightly worse than the inequality in A+, but in another way much better. There is a slightly greater gap between the better-off and worse-off groups, but a much better ratio between these groups. All things considered, the natural inequality in Alpha is not worse than the natural inequality in A+.¹⁰

It may be objected that Alpha is worse than A+ because the worst-off groups in Alpha are worse off than the worst-off group in A+. Many people accept Rawls's view that what matters most is the condition of the worst-off group. But there are two quite different ways in which any worst-off group might have been better off. This group might have existed, and been better off. This is the ordinary case, which Rawls discusses. Things would have been quite different if the worst-off group had never existed. This would have provided another sense in which the 'worst-off group' would have been better off, since some other group would then have been those who are worst off. Would this have made the outcome better? If we answer Yes, we must agree that it would have been even better if the second worst-off group had also never existed, and the third worst-off group, and the fourth worst-off group, and so on. It would have been best if everyone except the best-off group had never existed. Similarly, it might be best if in future only the best-off nation—such as the Norwegians—have children. Even if this would be worse for them, it might cause it to be true that, after the rest of us have died, the 'worst-off people in the world are as well off as possible. This way of raising the level of 'the worst-off group' has no moral merit. The non-existence of all but the best-off group would not, in the morally relevant sense, make the worst-off group better off.

The inequality in Alpha is not worse than the inequality in A+ nor is Alpha worse than A+ because the worst-off groups are worse off. Nor is there any other way in which Alpha is worse than A+.¹¹ And, in one way, Alpha is better than A+. Alpha does not differ from A+ merely by involving the existence of the very many groups at level 45. All of the people in A+ are in one of two groups, and both these

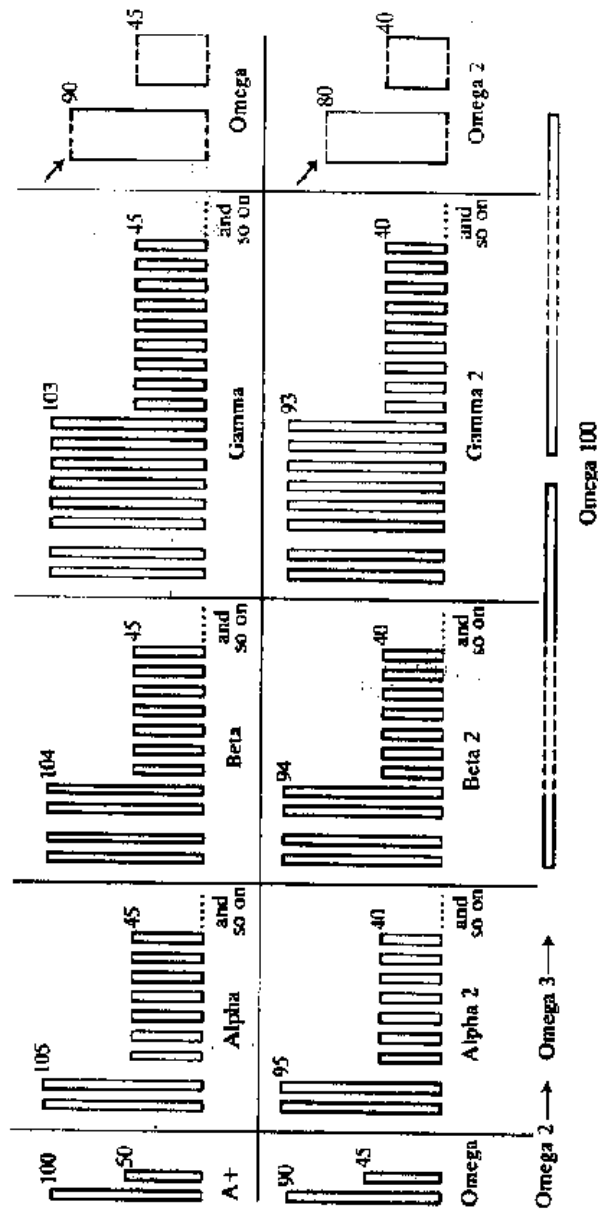


Fig. 4

groups are, in Alpha, better off. (These are the groups at level 105.) I conclude that, since Alpha is in this way better than A+, and is in no way worse, Alpha is better than A+. We are assuming that the actual outcome at some future time is Alpha. If the outcome had been A+, this would have been worse.

Could we honestly deny this conclusion? Could we honestly claim that A+ would not have been worse than Alpha? This is the claim that it would not have been worse if the worst-off people in Alpha had never existed, even though their lives are worth living, and if they had never existed, so that the outcome had been A+, the inequality would have been no better, and everyone who did exist would have been worse off. That this would not have been worse is hard to believe.¹²

Consider next whether Beta would be better than Alpha. In a change from Alpha to Beta, the best-off group in Alpha would lose a little, but an equally large worse-off group would gain very much more. If this is all we know about this change, it would need extreme Elitism to deny that it would be a change for the better.

The rest of the argument merely involves repetition. Gamma would be better than Beta in the same way in which Beta would be better than Alpha. Delta would be better than Gamma in the same way, Epsilon better than Delta, and so on down to Omega. We then run through the argument again, on the second line of the diagram, from Omega to Omega 2. (Omega is thinner on this second line only because, to make room, all widths are reduced.) Similar steps take us to Omega 3, Omega 4, and all the way to Omega 100. Every step would be a change for the better, so Omega 100 must be the best of all these outcomes.

Since this argument implies that Omega 100 would be better than A+, it leads us to the Repugnant Conclusion. A+ might be a world with ten billion people, of whom even the worse-off half have an extremely high quality of life. According to this argument it would be better if instead there were vastly many more people, all of whose lives were barely worth living.

What is wrong with this argument? To avoid its conclusion, we must either deny that A+ would have been worse than Alpha, or deny that Beta would be better than Alpha. Unless we deny one of these claims, we cannot plausibly deny the similar claims which carry us down to Omega 100. But how can we deny that A+ would have been worse than Alpha? If the outcome had been A+, everyone who existed would have been worse off. And how can we deny that Beta would be better than Alpha? In a change to Beta some people would lose a little, but as many people who are much worse off would gain much more.

While we consider these outcomes in these simple terms, it is hard to answer this argument. There is little room for manoeuvre. To find an answer we must consider other features of these outcomes.¹³

5. THE QUALITY OF SINGLE LIVES

Consider first the analogue, within one life, of the Repugnant Conclusion.¹⁴ Suppose that I can choose between two futures. I could live for another 100 years, all of an extremely high quality. Call this the Century of Ecstasy. I could instead live for

ever, with a life that would always be barely worth living. Though there would be nothing bad in this life, the only good things would be muzak and potatoes. Call this the Drab Eternity.

I believe that, of these two, the Century of Ecstasy would give me a better future. And this is the future that I would prefer. Many people would have the same belief, and preference.

On one view about what makes our lives go best, we would be making a mistake. On this view, though the Century of Ecstasy would have great value for me, this value would be finite, or have an upper limit. In contrast, since each day in the Drab Eternity would have the same small value for me, there would be no limit to the total value for me of this second life. This value must, in the end, be greater than the limited value of the Century of Ecstasy.

I reject this view. I claim that, though each day of the Drab Eternity would be worth living, the Century of Ecstasy would give me a better life. This is like Mill's claim about the 'difference in quality' between human and pig-like pleasures.¹⁵ It is often said that Mill's 'higher pleasures' are merely greater pleasures: pleasures with more value. As Sidgwick wrote, 'all qualitative comparison of pleasures must really resolve itself in quantitative [comparison]'.¹⁶ This would be so if the value of all pleasures lay on the same scale. But this is what I have just denied. The Century of Ecstasy would be better for me in an essentially qualitative way. Though each day of the Drab Eternity would have some value for me, no amount of this value could be as good for me as the Century of Ecstasy.

6. PERFECTIONISM

Return to the argument about overpopulation. Should we make a similar claim, not about the value for one person of different possible futures, but about the relative goodness of different outcomes? Cardinal Newman made such a claim about pain and sin. He believed that both of these were bad, but that no amount of pain could be as bad as the least amount of sin. He therefore wrote that, if all mankind suffered "extremest agony", this would be less bad than if one venial sin was committed.¹⁷ Can we make such a claim about what is good in my outcomes A and Z?

Consider what I shall call the best things in life. These are the best kinds of creative activity and aesthetic experience, the best relationships between different people, and the other things which do most to make life worth living. Return next to A and B. Suppose that all of the best things in life are, in B, better. The people in B are all worse off than the people in A only because they each have many fewer of these things. In B, for example, people can hear good music only a few times in their lives; in A they can often hear music that is nearly as good. If this was the difference between A and B, I would cease to believe that B would be worse.

A similar claim applies to the Repugnant Conclusion. Why is it so hard to believe that my imagined world Z—or Omega 100—would be better than a world of ten billion people, all of whom have an extremely high quality of life? This is hard

to believe because in Z two things are true: people's lives are barely worth living, and most of the good things in life are lost.

Suppose that only the first of these was true. Suppose that, in Z, all of the best things in life remain. People's lives are barely worth living because these best things are so thinly spread. The people in Z do each, once in their lives, have or engage in one of the best experiences or activities. But all the rest is muzak and potatoes. If this is what Z involves, it is still hard to believe that Z would be better than a world of ten billion people, each of whose lives is very well worth living. But, if Z retains all of the best things in life, this belief is less repugnant.

Now restore the assumption that in Z, and Omega 100, most of the good things in life are lost. There is only muzak and potatoes. By appealing to the value of the best things in life, we can try to answer the argument.

The argument involves two kinds of step. One is the claim that Alpha is better than A+, Alpha 2 is better than Omega, and similar later claims. A+ contains two groups of people, all of whom are better off in Alpha. We can add the assumption that these people are better off because, in Alpha, the best things in life are even better. Appealing to the value of these best things cannot help us to reject the claim that Alpha is better than A+. And, as I argued, there seems to be no other way to reject this claim. If the actual outcome had been A+, the inequality would have been no better, and everyone who existed would have been worse off. How can we deny that this would have been worse? There seems to be little hope of answering these steps in the argument.

The other steps are all redistributive. In each step the best-off people would lose a little, but an equally large worse-off group would gain much more; Can we claim that at least one of these steps would not be a change for the better? This cannot be plausibly claimed if what we appeal to is the Elitist View. We cannot plausibly claim that it is the best-off people whose condition matters most.

What we might appeal to is not Elitism, but Perfectionism. In the movie from Alpha to Omega 100, the best things in life must have disappeared. Suppose for instance that, in the move from Alpha to Beta, Mozart's music would be lost, in the move to Gamma, Haydn's. In the move to Delta, Venice would be destroyed, in the move to Epsilon, Verona.¹⁸ We might claim that, even if some change brings a great net benefit to those who are affected, it is a change for the worse if it involves the loss of one of the best things in life.

When should we make this claim? It would not be plausible when we are considering outcomes that are close to Omega 100. Suppose that, in one such outcome, the best thing left is a bad performance of Ravel's Bolero; in the next outcome, it is an even worse performance of Ravel's Bolero. We cannot claim that great benefits to those who are worst-off would not make the outcome better if they involved the loss of a bad performance of Ravel's Bolero. If such a claim is to have any plausibility, it must be made at the start. We must reject the change in which the music of Mozart is lost.

Has such a claim any plausibility? I believe that it has. It expresses one of our two main reasons for wanting to avoid the Repugnant Conclusion. When we are most concerned about overpopulation, our concern is only partly about the value that

each life will have for the person whose life it is. We are also concerned about the disappearance from the world of the kinds of experience and activity which do most to make life worth living.

Perfectionism faces many objections. One is raised by the moral importance of relieving or preventing great suffering. We should reject the Nietzschean view that the prevention of great suffering can be ranked wholly below the preservation of creation of the best things in life. What should Perfectionists claim about great suffering? But this problem is irrelevant here, since we can assume that in the various outcomes we are considering there would be no such suffering.

Another problem is raised by the fact that the good things in life do not come in quite different categories. It is because pain and sin are in such different categories that Newman believed sin to be infinitely worse. If we merely compare Mozart and muzak, these two may also seem to be in quite different categories. But there is a fairly smooth continuum between these two. Though Haydn is not as good as Mozart, he is very good. And there is other music which is not far below Haydn's, other music not far below this, and so on. Similar claims apply to the other best experiences, activities, and personal relationships, and to the other things which give most to the value of life. Most of these things are on fairly smooth continua, ranging from the best to the least good. Since this is so, it may be hard to defend the view that what is best has more value—or does more to make the outcome better—than any amount of what is nearly as good. This view conflicts with the preferences that most of us would have about our own futures. But, unless we can defend this view, any loss of quality could be outweighed by a sufficient gain in the quantity of lesser goods.

These are only two of the objections facing this view. It seems to me, at times, crazy. But at least, unlike the Elitist View, it is not morally monstrous. And without Perfectionism how can we avoid the Repugnant Conclusion?¹⁹

NOTES

¹ This essay is published by permission of the author and Oxford University Press. It was first published in P. Singer, *Applied Ethics*, Oxford University Press, 1986. The first half of the essay summarizes a longer discussion in my *Reasons and Persons* (Oxford, 1984). I have been greatly helped by J. McMahan, J. R. Richards, L. Temkin, K. Kalafski, and R. Jones.

² Of the many economists who appeal to the Average Principle, some make it true by definition. See, for example, P. A. Samuelson, (1970) p. 551. Certain writers state this principle so that it covers only the lives that are, at any time, being lived. This makes the principle imply that it would have been better if all but the best-off people had just dropped dead. My versions of the Average Principle do not imply this absurd conclusion. If anyone with a life worth living dies earlier, this causes people's lives to go, on average, worse, and to contain a smaller average sum of happiness.

³ These remarks assume that the quality of life is higher if people's lives go better, and that each life goes better if it contains a greater quantity either of happiness or of whatever else makes life worth living. 'Quality' thus means 'quantity, per life lived'. In Section 5 below I drop this assumption, thereby simplifying the contrast between quality and quantity. (If this note is puzzling, ignore it.)

⁴ In what follows I assume, for convenience, that there can be precise differences between the quality of life of different groups. I believe that there could not really be such precise differences. All that my arguments require is that some people can be worse off than others, in morally significant ways, and by more or less.

⁵ R. Nozick, (1974), p. 41.

⁶ According to some versions of the widely assumed *Law of Diminishing Marginal Utility*, we could do this. The point can be made most easily in Hedonistic terms. It is assumed that, because resources produce more happiness if they are given to people who are worse off, they would produce most happiness if they are all given to people whose lives are barely worth living. There is here an obvious oversight. Many resources are needed to make each person's life even reach a level where it begins to be worth living. Such resources do not help to produce the greatest possible quantity of happiness, since they are merely being used to raise people to the level where their happiness begins to outweigh their suffering.

⁷ It may help to give this illustration. Suppose that, as a *Negative Utilitarian*, I believe that all that matters morally is the relief or prevention of suffering. It is pointed out to me that, on my view, it would be best if all life on Earth was painlessly destroyed, since only this would ensure that there would be no more suffering. And suppose I agreed that this would be a very bad outcome. Could I say: 'It is true that this very bad outcome would, according to my moral view, be the best outcome. But this is no objection to my view, since we are not in fact able to bring about this outcome?' This would be no defence. On my view, I ought to *regret* our inability to bring about this outcome. Whether my view is plausible cannot depend on what is technically possible. Since this view implies that the destruction of all life on Earth would be the best outcome, if I firmly believe that this outcome would be very bad, I should reject this view.

⁸ The phrase 'if other things are equal' allows for the possibility that the existence of the larger population might, in some other way, be worse. It might, for instance, involve injustice. What the Repugnant Conclusion claims is that, though the lower quality of life would make Z in one way worse than A, this bad feature could be less important than, or be outweighed by, Z's good feature: the existence of enough extra people whose lives are—even if only barely—worth living.

⁹ J. Rawls, (1971).

¹⁰ If you believe that the inequality is worse in Alpha than it is in A+, read (when you reach it) footnote 12.

¹¹ Alpha is worse than A+ according to the Average Principle. But this is one of the cases which show that we should reject this principle. The Average Principle could also imply that it would be best if in future all except the Norwegians have no children. For further objections to this principle, see my *Reasons and Persons*, Section 143, and J.A. McMahan (1981). It may also be claimed, 'Alpha is worse than A+ because, if we had to choose in which outcome we would prefer to exist—without knowing who we would be—it would be rational to choose A+.' For objections to this claim, see my *Reasons and Persons*, Section 133.

¹² Suppose you believe that the inequality in Alpha is worse than the inequality in A+. Is this enough to justify the claim that it would *not* have been worse if the actual outcome had been A+ rather than Alpha? Which would have mattered more: (i) that the inequality would have been less bad, or (2) that everyone who did exist would have been worse off? It is hard to deny that (2) would have mattered more.

¹³ It may be objected that my argument is like what are called *Sorites Arguments*, which are known to lead to false conclusions. Suppose we assume that removing any single grain of sand cannot turn a heap of sand into something that is not a heap. It can then be argued that, even if we remove every single grain, we must still have a heap. Or suppose we assume that the loss of any single hair cannot cause someone who is not bald to be bald. There is a similar argument for the conclusion that, even if someone loses all his hair, this cannot make him bald. If my argument was like this, it could be referred to those who work on what is wrong with Sorites Arguments. But my argument is not like this. A Sorites Argument appeals to a series of steps, each of which is assumed to *make no difference*. My argument would be like this if it claimed that Alpha is *not worse* than A+, Beta is not worse than Alpha, Gamma is not worse than Beta, and so on. But the argument claims that Alpha is better than A+, Beta is better than Alpha, Gamma is better than Beta, and so on. The objections to Sorites Arguments are therefore irrelevant.

¹⁴ This section is partly based on an unpublished paper by J. McMahan.

¹⁵ J. S. Mill, (1863), Chapter II.

¹⁶ H. Sidgwick, (1907), p. 94.

¹⁷ J.H. Newman, (1885), Vol. I, p. 204.

¹⁸ If, in the move from Alpha to Beta, the best-off people lose Mozart, it may seem that their quality of life cannot, as my argument assumes, fall by only a little. But I have explained how this might be so. The loss of a few performances of Mozart could for these people be nearly outweighed by many extra performances of Haydn.

¹⁹ I would be grateful for any comments on this essay, which could be sent to me at All Souls College, Oxford.

REFERENCES

- McMahan, J. "Problems of Population Theory". *Ethics* 92 (1981).
 Mill, J. S. *Utilitarianism*. London, 1963.
 Newman, J. H. *Certain difficulties Felt by Anglicans in Catholic Teaching*. London, 1885.
 Nozick, R. *Anarchy, State and Utopia*. Oxford, 1974.
 Parfit, D. *Reasons and Persons*, Oxford: Clarendon Press, 1984.
 Rawls, J. *A Theory of Justice*. Cambridge, Mass., 1971.
 Samuelson, P. A. *Economics*. New York, 1970.
 Sidgwick, H. *The Methods of Ethics*. London, 1907.
 Sikora, R. I. & Barry, B. *Obligations to Future Generations*. Philadelphia: Temple University Press, 1978.
 Singer, P. *Applied Ethics*. Oxford: Oxford University Press, 1986