



Parfit's Repugnant Conclusion

Jesper Ryberg

The Philosophical Quarterly, Vol. 46, No. 183 (Apr., 1996), 202-213.

Stable URL:

<http://links.jstor.org/sici?sici=0031-8094%28199604%2946%3A183%3C202%3APRC%3E2.0.CO%3B2-4>

The Philosophical Quarterly is currently published by The Philosophical Quarterly.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/philquar.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

DISCUSSIONS

PARFIT'S REPUGNANT CONCLUSION

BY JESPER RYBERG

A very important question raised by Derek Parfit in the part of *Reasons and Persons* which deals with population ethics is how to compare the future outcomes of those policies which differ in the way they affect population growth.¹ Such comparisons are complicated by the fact that these outcomes may differ not only in the average levels of well-being they generate but also in the identity and number of the persons who exist.

As Parfit convincingly argues, many traditional moral principles do not provide a plausible guide in such cases. Either the principles are unable to give an answer if there is a change in the identity of people between the outcomes, or they provide answers which are implausible. However, one of the principles considered (p. 387) is

The Impersonal Total Principle: other things being equal, the best outcome is the one in which there would be the greatest quantity of whatever makes life worth living.

That the principle is *impersonal* means that it assigns value to the state of affairs which results from an act rather than to the persons affected by the act. Since the principle is, therefore, indifferent to changes in identity, it is *prima facie* qualified to give an answer in these cases; and furthermore it happens to provide quite plausible answers to many of the non-identity cases which are presented by Parfit (see pp. 358, 362).

Nevertheless, he rejects the principle because it implies that losses in the quality of well-being, or in whatever makes life worth living, can be made up for by sufficiently large gains in quantity. The principle therefore generates (p. 388)

The Repugnant Conclusion: for any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living.

Many agree with Parfit that this conclusion is repugnant. But, if that is correct, it obviously leaves a problem. On the one hand, the impersonal total principle does

¹ D. Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), Part IV.

succeed in giving plausible answers to many problems concerning non-identical future people. To abandon the principle would leave one in need of another principle of beneficence which, so Parfit's comprehensive work on the subject indicates, it would be quite hard to find. In fact, Parfit's own conclusion in *Reasons and Persons* is that he has not succeeded in developing a principle of beneficence (called 'theory X') which meets the minimal requirements he has outlined. On the other hand, one cannot just accept the repugnant conclusion. In the face of this dilemma, it is not strange that some attention has been drawn to the question whether it is really the case that the impersonal total principle implies the repugnant conclusion.

In this paper I shall comment on a theory of *discontinuity in values* which has roots in Mill's theory on qualitative pleasures, and which has, in a more recent paper, been suggested by Parfit himself and, in slightly different versions, by Glover and Crisp.² According to this theory, the repugnant conclusion is not implied by the impersonal total principle after all. Though I shall focus specifically on the theory suggested by Parfit, most of the arguments in the following sections are levelled at any theory which claims that discontinuity in values blocks the repugnant conclusion. The discussion will proceed as follows. I briefly outline in §I Parfit's suggestion, and discuss the main argument which has been presented in favour of a discontinuity axiology. In §II, I reject Parfit's explanation of why we regard the repugnant conclusion as repugnant, and go on to argue that even though the proposed axiology blocks some versions of the repugnant conclusion, it is still consistent with other versions of the conclusion. In §III, I point out the importance of clarifying exactly what kind of discontinuity the suggestion is employing; and argue that the most plausible kind of discontinuity is consistent with yet another conclusion which is just as repugnant as the repugnant conclusion itself.

I. THE DISCONTINUITY VIEW

The repugnant conclusion seems to imply a certain view about values (see Griffin p. 340). It treats values on a single additive scale, where low numbers, if added to themselves often enough, must become larger than any initial larger number. Any finite number of *A*-lives (lives in the less populous outcome of the repugnant conclusion) can therefore be out-ranked by a sufficiently large number of *Z*-lives (lives in the more populous outcome), because a gain in the quantity of lower values outweighs the loss of certain higher values. But suppose this axiological assumption, that there is a single additive scale, is mistaken. As is pointed out by Griffin, values might, at least to some extent, be incommensurable in a way which suspends addition. Griffin has given (p. 85) the following account of what he calls *discontinuity*:

we [can] have a positive value that, no matter how often a certain amount is added to itself, cannot become greater than another positive value, and cannot, not because

² J. Glover, *Causing Death and Saving Lives* (Harmondsworth: Penguin, 1977), pp. 69–71; J. Griffin, *Well-Being* (Oxford: Clarendon Press, 1986), pp. 338–40; R. Crisp, 'Utilitarianism and the Life of Virtue', *The Philosophical Quarterly*, 42 (1992), pp. 149–52; D. Parfit, 'Overpopulation and the Quality of Life', in P. Singer (ed.), *Applied Ethics* (Oxford UP, 1986), pp. 145–64.

with piling up we get diminishing value or even disvalue (though there are such cases), but because they are the sort of value that, even when remaining constant, cannot add up to some other value.

If it is correct that certain values are discontinuous with others, in the way outlined by Griffin, and if the move from the populations A to Z involves the loss of this kind of values, then the inference from the impersonal total principle to the repugnant conclusion might not be sound after all. This is exactly what Parfit has suggested.

Parfit considers what he calls 'the best things in life'. These are the kinds of creative activity, aesthetic experience, relationships between people and of the other things which do most to make life worth living. Now what happens as we move down the alphabet from A to Z is, according to Parfit (Singer (ed.) p. 162), that the best things in life are lost. The first step from A to B involves the loss of Mozart's music; in the move from B to C Haydn's music is lost; in the move to D Venice is destroyed; and so on down the line. All that is left in the final move to Z is 'muzak and potatoes'. Parfit's claim is that there is discontinuity between these values. The loss of the things which do most to make life worth living cannot be made good by any gain in the quantity of muzak and potatoes. He then considers at what point the discontinuity sets in, and he objects that it is too far down the alphabet. If the move from one population to another population near the end of the alphabet involves the loss of a bad performance of Ravel's *Bolero*, then this loss could be made up for by more numerous performances of an even worse version of *Bolero*. Rather, he suggests, discontinuity sets in at the beginning; we must not lose Mozart's music in the first place. Thus, the repugnant conclusion is blocked in the first step from A to B . This is Parfit's *perfectionism*.

Before entering any discussion as to whether this suggestion actually succeeds in blocking the repugnant conclusion, it is reasonable to ask how the suggestion can be justified. Is it at all plausible to deny the view that all qualitative comparison of values must resolve itself in quantitative comparison, that is, that all values lie on the same continuous additive scale? That some values fit on to the same continuous additive scale is quite obvious. Suppose, for instance, that I like eating pears and that I like eating apples even more. Thus I shall prefer one kilo of apples to one kilo of pears. But if I am offered either one kilo of apples or two kilos of pears, then I might well prefer the latter simply because the overall benefit from eating two kilos of pears is greater. But here is another case. Suppose that I am offered either an infinite number of apples or an opportunity to listen to a performance of Wagner's *Lohengrin*. In this case I might prefer *Lohengrin*. Of course it would be natural to think that I shall lose my interest in apples the more of them I consume. But even if for the sake of argument we assume that they do not diminish marginally in value no matter how many I eat, I might still prefer *Lohengrin*. But this demonstrates a suspension of addition. No matter how often a certain kind of positive value is added to itself it cannot become greater than a finite amount of another positive value.

Parfit, therefore, is right in rejecting all comparison of values as quantitative comparison. The objection might be raised that, if some values fit on to the same continuous additive scale but there is discontinuity between other values, then we cannot be sure that it is the latter kind of values which are lost in the move from A to Z . Of

course the difference between A and Z might be analogous to that between apples and *Lohengrin*. But it might also be analogous to the difference between apples and pears, and if so the repugnant conclusion is not blocked.

The way Parfit rejects this possibility, like others who have considered the suggestion (e.g., Griffin, Crisp) is simply by considering an intra-personal analogue to the repugnant conclusion. Suppose that you can choose between two different future life scenarios. Either you can have 100 years, all of a very high quality (corresponding to the quality of life in A): that is, you could have all the best experiences and engage in all the best activities as often as you want. Or alternatively you can have an (almost?) infinite number of years which contain only good things, but nothing better than muzak and potatoes (corresponding to the quality of life in Z). Would you prefer the 'century of ecstasy' or the 'drab eternity'?

If the difference between the values in the two scenarios were analogous to that between apples and pears then, following the impersonal total principle, we should prefer the drab eternity. The century of ecstasy contains only a limited (although high) amount of values and must, at some point, be outranked by an infinite amount of lower values. But Parfit claims that he would prefer the century of ecstasy, and he thinks many would share this preference. However, if it is correct that the century of ecstasy would give one a better future than the drab eternity then there is no continuous additive scale. The values in the two life scenarios must be discontinuous. And this indicates that there is also discontinuity between values in the two population scenarios A and Z .

Strictly speaking, Parfit's argument involves some kind of trick. It assumes that preferences with regard to comparison of life scenarios can be transferred into comparison of population scenarios. However, when we move from the intra-personal case to the inter-personal case there is nothing analogous to the preference of a single individual. Parfit does not discuss how the result of intra-personal comparison can be extrapolated into the inter-personal sphere. However, on this point he might be assisted in different ways.

For instance, N. Holtug has suggested (in an unpublished paper) that the transference is a consequence of the impersonal total principle itself. According to the principle, the best outcome is the one in which whatever makes life worth living is maximized, independently of on whom or how many it falls. If the century of ecstasy is better than the drab eternity, then it follows that the century of ecstasy is better than *any number of drab lives*. Whether an infinite number of Z -years is distributed over one life or many lives makes no difference. All that matters is the overall amount of what makes life worth living. The intra-personal case is thereby transferred into the inter-personal case. Another possibility is to adopt C.I. Lewis' method for assessing the value of populations³: to imagine oneself leading the lives of each member of any population in serial order. According to this suggestion, there is no need for a transference because inter-personal comparison is treated as intra-personal comparison.

Both suggestions support Parfit's conclusion that the comparison of life scenarios indicates that there are discontinuities in values, and that Z therefore is not better

³ C.I. Lewis, *An Analysis of Knowledge and Valuation* (LaSalle: Open Court, 1946), p. 546-7.

than A after all. The argument seems convincing. It is apparently correct that when we consider the things which do most to make our lives worth living, then these are better than any amount of inferior values, and that the repugnant conclusion must be blocked if the move from A to Z involves the disappearance of these things.

What I fail to understand, however, is why Parfit thinks that the argument supports perfectionism. What is established by comparing the century of ecstasy and the drab eternity is that there is discontinuity between values in A and values in Z and that A is therefore preferable. But Parfit is making a much stronger claim. He says that the repugnant conclusion is blocked at the first step from A to B when Mozart's music disappears. Being a stronger claim, the suggestion is much harder to justify. In order to justify perfectionism, by the same kind of reasoning as I have just outlined, Parfit would have to argue that the century of ecstasy is preferable to an infinite number of years containing the same values as lives in B . That is, an infinite number of years in which one could listen to Haydn's music and visit Venice and do all the other things that do most to make life worth living. The only thing excepted is Mozart's music.

But surely it is not reasonable to expect everybody to have this preference. As long as one is still very well off, I think one would gladly sacrifice a little in quality in favour of a gain in quantity. Especially for an infinite gain. What is justified by intra-personal comparison is therefore not perfectionism but rather the thesis that discontinuity sets in somewhere in between A and Z . Perhaps not between A and B , B and C , or even between C and D , but somewhere down the line. (In fact this has been suggested by Griffin and by Holtug.)

II. BLOCKING THE REPUGNANT CONCLUSION?

Does Parfit's discontinuity axiology really block the repugnant conclusion? As we have just seen, the answer is apparently in the affirmative. Suppose that somewhere down the alphabet, let us say in the move from population M to N , all higher-quality values have disappeared. According to the impersonal total principle it is now M rather than Z which is the *optimum population* and the repugnant conclusion is blocked.

However, the whole argument is built on the assumption that the move from A to Z involves the disappearance of some values which are discontinuous with others. To what extent it is possible to generate high-quality and low-quality populations within values which fit on to the same continuous additive scale has not been discussed by adherents of the discontinuity view. The only person who has offered a comment on this question is Parfit. And not even he discusses it thoroughly, though, as I shall argue, it seems to be a crucial point when considering whether the impersonal total principle avoids repugnance.

Let us consider first the possibility that the quality of life decreases when we move from A to a population Z^* , not because the best things in life disappear, but because these things are so thinly spread. While people in A can listen to Mozart's music and engage in the best activities as often as they want, people in Z^* have the same

opportunity only once in a lifetime. All the rest is muzak and potatoes. If this is what Z^* involves, which population scenario would then be the better? According to the impersonal total principle, Z^* is better than A because Z^* is the outcome which contains the greatest quantity of higher values. The conclusion that Z^* is better than A is therefore fully consistent with an axiology which admits discontinuities in values. But is it an acceptable conclusion? Parfit considers this conclusion and claims (*ibid.*) that

if this is what [Z^*] involves, it is still hard to believe that [it] would be better than a world of ten billion people, each of whose lives is very well worth living.

But Parfit regards this conclusion as less repugnant than the repugnant conclusion, and apparently he thinks it is a conclusion which it is not absurd to accept. However, what neither he nor any other adherent of the discontinuity view notices is that his axiology is consistent with an even less attractive conclusion. This conclusion follows when we start considering outcomes which not only contain positive but also negative values.

Let us consider a very large population, called Z^{**} , in which the best things in life are once again thinly spread. People in Z^{**} each, once in their lives, engage in one of the best experiences or activities. But all the rest is slightly negative values, let us say cats' concerts and rotten fish. If a life contained nothing but cats' concerts and rotten fish it would not be worth living. But if a single experience of Mozart's music outweighs the total amount of slightly negative values in a life, then the impersonal total principle implies that Z^{**} is better than A . This conclusion is once again consistent with a discontinuity axiology.

A possible objection to this conclusion is that the single experience of one of the best things cannot counterbalance the total amount of slightly negative values, and that lives in Z^{**} therefore contain a surplus of negative values over positive values. Z^{**} will therefore not be better than A after all. This objection is, however, not plausible. If there is discontinuity between an experience of Mozart's music and any amount of muzak and potatoes, and if slightly negative values such as cats' music and rotten fish have the same *numerical* value as muzak and potatoes, then it follows from the discontinuity view that no amount (not even an infinite amount) of cats' concerts and rotten fish can counterbalance an experience of Mozart's music. The only way this could be avoided would be if slightly negative values have the same numerical values as higher positive values. Then obviously a life in Z^{**} would contain a surplus of negative values. But this suggestion is quite implausible. Surely some amount of lower positive values can counterbalance a small amount of slightly negative values. A life scenario which contains only lower positive values, except for a single or a few slightly negative values, is obviously preferable to a life scenario which contains neither positive nor negative values. Lower negative values do not, therefore, have the same numerical values as higher positive values and lives in Z^{**} do, therefore, contain a surplus of positive value.

But is it an acceptable conclusion that Z^{**} is better than A ? Since Parfit himself finds it hard to believe that Z^* is better than A he must, surely, find it even harder to accept that Z^{**} is better than A . If the repugnant conclusion is repugnant to a

degree that is sufficient to reject the impersonal total principle, it seems strange that we should simply accept the conclusions that Z^* and Z^{**} are better than A . Nevertheless, I think that this is the answer which Parfit and other adherents of the discontinuity view would give. If lives have been deprived to the extent in which they contain nothing of what does most to make life worth living – that is, they contain none of the best experiences, the best activities or relationships between people, but only inferior values – then it is absurd that a sufficiently large number of such lives should be better than a smaller number of lives which contain all the best things. But if some of the best things are in fact retained in these lives, then it is no longer absurd that a sufficient number of them outweighs a smaller number of better-off people. After all, the answer might go, people in Z^* and Z^{**} have better lives than people in Z , and it is therefore not strange that the conclusions that Z^* and Z^{**} are better than A are more tolerable than the repugnant conclusion.

The whole idea in this line of arguing could also be put much more simply: we find the repugnant conclusion repugnant because the move from A to Z involves the disappearance of higher values. In fact this is exactly what Parfit suggests. He claims (*ibid.*) that the repugnance of the conclusion is due to a conjunction of two things (which expresses our general concern about overpopulation). It is hard to believe that Z is better than A because in Z

people's lives are barely worth living, and most of the good things in life are lost.

He clearly underlines the point that the first condition, that people's lives are barely worth living, is not sufficient. In fact he seems to regard the second condition as especially important. He could therefore maintain that, though the conclusions that Z^* is better than A and Z^{**} is better than A are not attractive, they are nevertheless not repugnant. But should we accept this answer? Is it at all plausible to claim that the fact that we find the repugnant conclusion repugnant is so closely related to our concern about the disappearance from the world of the best things? If we find repugnant the conclusion that Z^{**} is better than A , then surely the disappearance of higher values is not a necessary condition for generating repugnance. However, if one is still in doubt whether or not to accept Parfit's answer, I think that the following example, which involves negative values, does establish that the suggestion is implausible.

Let us consider a large population Z^{***} in which each person's life contains lower values, such as muzak and potatoes, but also one of or a few of the best experiences and activities. But lives in Z^{***} also, for each of the best things, contain one of the worst things in life. By the 'worst things in life' I mean the kind of experiences and activities which do most to make life worth not living. Thus the life of a person in Z^{***} might, for instance, contain an experience of Mozart's music but also the experience of undergoing a serious operation without anaesthetic. Now, if the negative values are higher negative values, in the sense that they have the same numerical value as higher positive values (surely it is reasonable to think that there are such higher negative values), then the higher positive values will be counterbalanced by the higher negative values. But, since a life in Z^{***} *ex hypothesi* contains the same amount of higher positive values and higher negative values, there will not be a

surplus of higher values (either positive or negative), but only a surplus of lower positive values (since a Z^{***} life also contains muzak and potatoes). Now, would it be repugnant if the impersonal total principle implied that Z^{***} is better than A ? (In fact, if the discontinuity axiology is correct the principle would not lead to this conclusion because there is no surplus of higher positive values in Z^{***} ; however the fact that a conclusion is blocked is surely not a criterion of whether or not we find it repugnant. And it is the question about the diagnosis I am now attempting to clarify.)

If the conclusion that Z is better than A is repugnant, I see no reason why the conclusion that Z^{***} is better than A should not be repugnant as well. And now adherents of the discontinuity view can no longer give the same answer as with regard to the conclusions that Z^* and Z^{**} are better than A , namely that, after all, people in Z^* and Z^{**} are better off than people in Z because their lives do contain some (though small) amount of higher values. Because in Z^{***} the higher positive values are *ex hypothesi* counterbalanced by higher negative values, which means that the net result of a life in Z^{***} is only an amount of lower positive values. People in Z^{***} are therefore not better off than people in Z and the conclusions, that each population is better than A , are therefore equally repugnant.

The only way this conclusion could be avoided would be if one claimed that a life containing both lower positive values and an amount of higher positive and negative values counterbalancing each other is better than a life which contains only lower positive values. But I fail to see why this should be the case. Surely it is not implied by a discontinuity axiology. In my view it is reasonable to regard that conclusion, that Z^{***} is better than A , as equally repugnant as the repugnant conclusion itself.

However, if this consideration is correct, then Parfit's way of explaining why we regard the repugnant conclusion as repugnant falls apart. The claim that a move from a high-quality population to a low-quality population is repugnant because it involves the disappearance from the world of the best things cannot be correct if the conclusion, that Z^{***} is better than A , is repugnant. In Z^{***} the best things have not disappeared. To claim that the conclusions, that Z is better than A and that Z^{***} is better than A , are repugnant for quite different reasons is hardly plausible. Rather, it is reasonable to conclude that a plausible diagnosis of the repugnant conclusion does not involve reference to the disappearance of values which are discontinuous with others. But this conclusion is extremely important, because it means that the impersonal total principle might have implications which are repugnant and which are consistent with a discontinuity axiology.

Suppose, for instance, that what is objectionable about the repugnant conclusion is the comparative evaluation of populations: that is, the claim that the more populous outcome is superior to the outcome in which people are all better off. If this diagnosis is correct, then surely a discontinuity axiology would be consistent with repugnant conclusions. The conclusions that Z^* is better than A and that Z^{**} is better than A would both be repugnant, and neither conclusion would be blocked by discontinuity.

Personally, I do not think that this diagnosis would be sufficient, because, in my view, it is not repugnant to compensate for a decrease in quality by gain in quantity

if people in the more populous outcome are still very well off. An even stronger diagnosis, therefore, would be that we find it repugnant to compensate for a loss in quality if people in the more populous outcome have low-quality lives. According to this diagnosis, it is only the first condition in Parfit's diagnosis which should be retained. We are simply concerned about the fact that the people in \mathcal{Z} have lives only barely worth living. This diagnosis would, contrary to Parfit's suggestion, explain why we regard the conclusion that \mathcal{Z}^{***} is better than A as repugnant. People in \mathcal{Z}^{***} have low-quality lives exactly as do people in \mathcal{Z} . However, if this suggestion is correct, once again there are versions of the repugnant conclusion which will not be blocked by discontinuity. Let us consider a population A' in which each person experiences only good things but nothing better than muzak and potatoes. However, persons in A' have access to muzak and potatoes as often as they want. In another larger population, called \mathcal{Z} , each person experiences nothing tolerable but muzak and potatoes but those only once or a few times in life. People in \mathcal{Z} are therefore worse off than people in A' . But if \mathcal{Z} is a sufficiently large population then, following the impersonal total principle, it is preferable to A' , even though each person in \mathcal{Z} has a life that is only barely worth living. Unlike people in A , those in A' do not have an extremely high quality of life because their lives do not contain any higher values; however, if the suggested diagnosis is correct, what matters, when a loss in quality is compensated for by gain in quantity, is not the quality of life in the less populous outcome but rather the quality of life in the more populous outcome. Since people in \mathcal{Z} have lives barely worth living, the conclusion that \mathcal{Z} is better than A' is repugnant and it is obviously not blocked by an axiology which admits discontinuity in values. In my view, this example establishes that Parfit's discontinuity view does not provide a plausible way to avoid the repugnant conclusion.

III. A FURTHER PROBLEM

The final argument I shall present is not levelled specifically at Parfit's axiology but points to the need for further discussion and justification of the suggestion. Parfit does not thoroughly discuss how the discontinuity axiology is justified but, as outlined in §I above, refers to the preferences we have when we compare different life scenarios. The question, however, is whether this sort of comparison justifies that kind of discontinuity Parfit suggests. Parfit indicates the kind of discontinuity he has in mind by referring (Singer (ed.) p. 161) to Cardinal Newman's view on pain and sin:

[Newman] 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'. Can we make such a claim about what is good in my outcomes A and \mathcal{Z} ?

Parfit's answer is in the affirmative, which indicates that the kind of discontinuity he is suggesting is what Griffin calls *trumping*: any (even the smallest) amount of μ

outranks any (even an infinite) amount of π . The values μ and π are measured on different scales, where one lexically dominates the other. But can trumping be justified? In order to answer this it is not sufficient to compare the century of ecstasy and the drab eternity. Rather we should ask whether we would prefer even the smallest interval of ecstasy to the drab eternity. Would we, for instance, prefer a month of ecstasy to the drab eternity? Or perhaps a day of ecstasy to the drab eternity? Or, to put it differently, does even the smallest amount of Mozart's music outrank an infinite amount of muzak and potatoes? I do not think so. In my view the drab eternity is preferable to the day of ecstasy.

This does not mean that I definitely deny that some values trump others, but I do not find the suggestion very likely. In his discussion of different forms of incommensurability, Griffin denies (p. 83) that some values trump others: 'people who would call certain values "trumps" or give them "lexical priority" probably do not mean these terms entirely seriously'. Suppose he is right. The kind of incommensurability which adherents of the discontinuity view may then have in mind, where the judgements that we prefer the century of ecstasy to the drab eternity and at the same time prefer the drab eternity to the day of ecstasy would be made consistent, may have the following form: *enough* of μ outranks any amount of π . But this kind of discontinuity happens to be consistent with yet another version of the repugnant conclusion.

Let us consider a large population Z^{\wedge} in which each person has some of the best experiences and engages in some of the best activities but only to an extent which is insufficient to generate discontinuity. Besides this amount of the best things, each life also contains an amount of lower negative values which exactly counterbalances the amount of the best things (since there are not enough of the higher values to generate discontinuity, these values can be counterbalanced by lower negative values). Finally, each life contains an amount of lower positive values as well. Now, if Z^{\wedge} is sufficiently large, would it then be better than A ? In order to answer this, let us once again consider an intra-personal analogue to the scenarios. This time one could either have the century of ecstasy or an infinite life in which, once every century, one experienced a small amount (*per se* insufficient to generate discontinuity) of the best things; but besides that the life would contain both lower positive and lower negative values. Since the latter life scenario is an eternal life it will at some point contain a sufficient amount of higher values to generate discontinuity and this amount will obviously exceed the finite amount of higher values in the century of ecstasy. Suppose, therefore, that we prefer the eternal life scenario to the century of ecstasy. This preference can now either be transferred into the inter-personal case, by the same procedure as when our preference for the century of ecstasy rather than the drab eternity is transferred, that is by applying Holtug's method outlined above. Or we may instead apply Lewis' suggestion that we should consider ourselves living each life in the populations in serial order. In both cases the conclusion will be that Z^{\wedge} is better than A . But is this an acceptable conclusion? I do not think so. Since each life in Z^{\wedge} *per se* only contains an insufficient amount of higher values which is counterbalanced by an amount of lower negative values, all that is left is the amount of lower positive values. A life in Z^{\wedge} is therefore (unlike lives in Z^* and Z^{**}) no

better than a life in Z . Therefore, if the conclusion that Z is better than A is repugnant, the latter conclusion that Z^{\wedge} is better than A must be repugnant too. But the latter conclusion is consistent with a discontinuity axiology.

The argument can be summarized thus: if it is correct that trumping is too strong a form of incommensurability, but that some values are discontinuous in the sense that *enough* of one value outranks any amount of another value, then the argument which is supposed to block the repugnant conclusion faces a problem. The whole idea that our preferences in intra-personal cases justify a conclusion in inter-personal cases, either by transference or by Lewis' method, does not account for the possibility that higher values might be distributed in such way that each member of a population only experiences an amount of higher value insufficient to generate discontinuity. The preference in the intra-personal analogue (in which the pieces of higher value are added up and consequently generate discontinuity) therefore leads to an unacceptable conclusion in the inter-personal case.

This argument may strike some as highly speculative. It might be objected that what counts as enough of higher values to generate discontinuity cannot be reached by adding up smaller separated pieces of value. But even if this is true for some values I do not see why it should necessarily be true for all higher values. Would it necessarily be incoherent to claim that, if I could only visit Venice for a few days, then this experience would be outweighed by a sufficient amount of less attractive experiences, but if I could visit Venice for a few days every year in my (finite) life, then this would be better than any amount of less attractive experiences? I do not think so. But if that is correct, the example could be used to generate exactly the kind of repugnant conclusion I have just outlined. The conclusion clearly demonstrates the importance of further clarification of the discontinuity view. It is necessary to clarify what kind of discontinuity one has in mind and how the suggested axiology can be justified. As I have argued, the most plausible axiological suggestion is consistent with versions of the repugnant conclusion.

IV. CONCLUSION

In the previous sections I have questioned Parfit's suggestion that an axiology which admits discontinuities in values succeeds in blocking the repugnant conclusion. I briefly outlined in §I the line of argument which is supposed to justify the suggestion. In my view the question of how exactly our preferences with regard to intra-personal versions of the repugnant conclusion should be interpreted needs much further elaboration. So does the question of whether a transference from intra-personal comparison into the inter-personal sphere is legitimate. However, besides noticing that the procedure hardly justifies Parfit's perfectionism, what I have concentrated on is whether Parfit's discontinuity axiology, if assumed to be plausible, actually succeeds in blocking the repugnant conclusion. As I have argued, it is not plausible to claim that the repugnant conclusion is repugnant because the move from A to Z involves the disappearance of some values which are discontinuous with others. It is possible to formulate versions of the repugnant conclusion

in which higher values are retained in the more populous outcome. Thus a plausible diagnosis does not involve reference to the disappearance of higher values. The most plausible diagnosis seems to me that the repugnant conclusion is repugnant because it is repugnant to compensate for a loss in quality by an increase in quantity of low-quality lives. But if this diagnosis is correct, then Parfit's suggestion cuts no ice. The impersonal total principle will have implications which are repugnant and which are fully consistent with a discontinuity axiology. Therefore if it is correct that the repugnant conclusion is a serious challenge facing utilitarianism and some other consequentialist moral principles, then this challenge is not removed by the axiology suggested by Parfit.⁴

University of Copenhagen

⁴ I would like to thank Roger Crisp for commenting on an earlier draft of this paper.

HOW NOT TO OPPUGN CONSEQUENTIALISM

BY MATTHEW H. KRAMER

Much of moral philosophy has involved debates between consequentialist and deontological approaches to the subject. On the former side are philosophers who evaluate each action by reference to the consequences or states of affairs which the action brings about (including the action itself, however we choose to define it); on the latter side are philosophers who evaluate each action by reference to its accord or discord with entitlements and obligations that obtain irrespective of the consequences that will flow from our heeding of them. In the course of these debates, each side has not only sought to commend the merits of its own broad position but has also sought to expose the shortcomings on the other side. Often, the attempts to discredit either position have proceeded via the construction of scenarios in which the particular position leads to ghastly results. Alternatively the discrediting proceeds through revelations (or putative revelations) of analytical shortcomings.

The present paper criticizes an important argument that has tried to disclose an analytical shortcoming in consequentialist reasoning. That anti-consequentialist argument has been propounded by John Finnis (following the lead of Anselm Müller), who has maintained that the consequentialist version of one of our basic moral principles is self-contradictory. I seek to expose the fallacies in the critique put