

DOES THE TOTAL PRINCIPLE HAVE ANY REPUGNANT IMPLICATIONS?

Douglas W. Portmore

Abstract

On the Total Principle, the best state of affairs (*ceteris paribus*) is the one with the greatest net sum of welfare value. Parfit rejects this principle, because he believes that it implies the Repugnant Conclusion, the conclusion that for any large population of people, all with lives well worth living, there will be some much larger population whose existence would be better, even though its members all have lives that are only barely worth living. Recently, however, a number of philosophers have suggested that the Total Principle does not imply the Repugnant Conclusion provided that a certain axiological view (namely, the ‘Discontinuity View’) is correct. Nevertheless, as I point out, there are three different versions of the Repugnant Conclusion, and it appears that the Total Principle will imply two of the three even if the Discontinuity View is correct. I then go on to argue that one of the two remaining versions turns out not to be repugnant after all. Second, I argue that the last remaining version is not, as it turns out, implied by the Total Principle. Thus, my arguments show that the Total Principle has no repugnant implications.

In part four of *Reasons and Persons*,¹ Derek Parfit tries to come up with a principle which can account for our moral intuitions concerning certain puzzling cases, cases where we consider one choice of action to be morally worse than another although worse for no one.² Such cases arise where our choice of action affects the identities of future people. Thus, Parfit labels the problem of accounting for our intuitions in these puzzling cases as the ‘Non-Identity Problem’. And he refers to the unknown principle which will provide such an account (while meeting some other minimal requirements) as ‘Theory X’ (p. 378).

¹ Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984). All page references in the body of this paper are to this book.

² For an example of such a case, see Derek Parfit, ‘On Doing the Best for Our Children’, in Michael Bayles (ed.), *Ethics and Population* (Cambridge, Mass.: Schenkman, 1976), pp. 100–101.

One promising candidate for Theory X is

The Total Principle. Other things being equal, the best state of affairs is the one with the greatest net sum of welfare value (p. 387).

But Parfit ultimately rejects the Total Principle, because he believes that it implies


The Repugnant Conclusion. For any large population of people, all with lives that are well worth living, there will be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members would all have lives that are only barely worth living (p. 388).

We can see how it is that the Total Principle implies the Repugnant Conclusion by considering Fig. 1 (on the following page) where a number of alternative populations are graphically represented. (The width of the blocks represents the number of people living; the height of the blocks, the net value of each life. Block Z is dashed to represent the fact that it is much wider than it is shown to be in the graph.)

On the Total Principle, the best alternative would be represented by the block with the largest area. Thus, B is better than A, C is better than B, D is better than C, and so on down the alphabet. And best of all would be Z, which is some enormous population of people, where each member has a life that is only barely worth living.

But note that there are three ways the lives in Z could be barely worth living: (1) they could be drab lives, free of pain but also devoid of all but a few simple pleasures; (2) they could be lives of extreme ups and downs, emotional roller coaster rides, where the ecstasies just barely outweigh the agonies; or (3) they could be lives which are qualitatively identical to those in A but very short-lived. I will refer to these three possibilities as Drab Z, Roller Coaster Z, and Short-lived Z, respectively.

In *Reasons and Persons*, Parfit considers only two of these three possibilities. He says, 'A life could be [barely worth living] . . . either because it has enough ecstasies to make its agonies just worth enduring, or because it is uniformly of poor quality' (p. 388). But when considering whether or not Z is a repugnant alternative to A, he asks only that we 'imagine the lives in Z to be of this second drabber kind' (p. 388). Thus, it is the conclusion that Drab Z is better than A which Parfit has in mind when he

Figure 1³

speaks of the Repugnant Conclusion. I will call this version of the Repugnant Conclusion the Drab Z Conclusion to differentiate it from the two other versions of the Repugnant Conclusion (neither of which does Parfit explicitly discuss in *Reasons and Persons*): (1) the Short-lived Z Conclusion, the conclusion that Short-lived Z is better than A; and (2) the Roller Coaster Z Conclusion, the conclusion that Roller Coaster Z is better than A.

I. Why the Total Principle Does Not Necessarily Imply the Drab Z Conclusion

A number of philosophers have suggested that we can avoid the Repugnant Conclusion (and do so even while accepting the Total Principle) so long as we hold that some of the values realised by the members of A are ‘discontinuous’ with the values realised by the members of Z.⁴ For instance, let us suppose that the members

³ I borrow this diagram from Parfit, *Reasons and Persons*, p. 388.

⁴ See Roger Crisp, ‘Utilitarianism and the Life of Virtue’, *The Philosophical Quarterly*, 42 (1992), pp. 149–52; Jonathan Glover, *Causing Death and Saving Lives* (Harmondsworth: Penguin Books, 1977), pp. 69–71; and James Griffin, *Well-being: Its Meaning, Measurement and Moral Importance* (Oxford: Clarendon Press, 1986), fn. 27, pp. 338–340.

None of these philosophers make a distinction between the various versions of the Repugnant Conclusion. They simply claim that the Repugnant Conclusion can be avoided by appealing to the discontinuity between certain values. But, as we will see, this is a mistake, because we can only avoid one of the three versions of the Repugnant Conclusion (namely, the Drab Z Conclusion) by making such an appeal.

of Z realise but one type of value, V_1 , and only in a quantity sufficient to make life barely worth living. And let us suppose that the members of A, on the other hand, not only realise V_1 but also a large quantity of V_2 – enough in fact to make life well worth living. Finally, let's assume that V_1 and V_2 are discontinuous, meaning that it is not the case that for any given quantity of V_2 there is always some quantity of V_1 which is equally valuable. Suppose, for instance, that a sufficient quantity of V_2 is more valuable than any amount of V_1 even where the value of an increase in the quantity of V_1 is always constantly positive.⁵

Clearly, under these suppositions, Z would be less valuable than A regardless of how populous Z is. For, given the discontinuity between V_1 and V_2 , no number of lives of the type lived by those in Z could ever be as valuable as even just one of the lives in A. Consequently, the Total Principle would not imply the Repugnant Conclusion. It would, in fact, imply that A is better than Z. Thus, we can avoid the Repugnant Conclusion provided that certain values are discontinuous.

But note that, in order to avoid the Repugnant Conclusion, we must assume that the members of Z do not realise all the same values which the members of A do. Yet this assumption is true only of Drab Z. So, although an appeal to the discontinuities between certain values may enable one who accepts the Total Principle to avoid the Drab Z Conclusion, it doesn't enable such a person to avoid the Repugnant Conclusion *tout court*. For no matter what values we take to be discontinuous, it seems that the Total Principle will still imply both the Roller Coaster Z Conclusion and the Short-lived Z Conclusion – in both cases, the members of Z realise all the same values that those in A do.

So what many philosophers have failed to realise is that an appeal to the discontinuity between certain values is only useful in avoiding but one version of the Repugnant Conclusion. And

⁵ Imagine, for instance, that V_1 is the type of value associated with eating fine food and that V_2 is the type of value associated with deep personal relationships. In this case, it is plausible to suppose that there is no quantity of V_1 which would be as valuable as a sufficient quantity of V_2 , for I doubt that anyone would trade the relationship they have with whom they are closest for any quantity of epicurean delights.

Of course, there are other ways in which a set of values could exhibit discontinuity. It could be that so long as one has enough of V_1 , no further amount of V_1 would ever be as valuable as even the slightest amount of V_2 . Also, it could be that, although enough of V_3 can compensate for the absolute deprivation of either V_1 or V_2 , no amount of V_3 can compensate for the absolute deprivation of both V_1 and V_2 . See Griffin, *Well-being*, pp. 85–89.

so, if the Total Principle is to remain in contention for being Theory X, more work has to be done: what must be shown is that either these other two versions are not repugnant or that the Total Principle does not imply them.

I believe that it is a combination of the two. I will argue that the Short-lived Z Conclusion is not repugnant or, more accurately, that it is far less repugnant than what we would have to accept were we to reject it. Secondly, I will argue that the Total Principle does not imply the Roller Coaster Z Conclusion so long as we accept a certain (plausible) axiological view. Thus, my arguments will show that so long as the Total Principle is applied to a sophisticated axiology, it will only imply one version of the Repugnant Conclusion, the Short-lived Z Conclusion, and, as I will show, this version doesn't merit the name 'Repugnant Conclusion'.

However, before proceeding with these arguments, I will first defend the claim that the Total Principle does not imply the Drab Z Conclusion. In doing so, I will have to argue that the Discontinuity View⁶ (the view that no number of lives of the type lived by those in Drab Z can ever be as valuable as even just one of the lives in A) is plausible. For all that I have argued so far is that, if the Discontinuity View is correct, the Total Principle does not imply the Drab Z Conclusion. Secondly, I will defend the Discontinuity View against some potentially devastating criticisms levelled by Parfit. Parfit argues that the implications of the Discontinuity View are just as repugnant as the Drab Z Conclusion. Thus, if sound, Parfit's arguments show that we cannot accept the Total Principle without being committed to something repugnant. (Curiously, these arguments seem to have been overlooked by those who have sought to avoid the Drab Z Conclusion by appealing to the Discontinuity View.)

II. The Plausibility of the Discontinuity View

In regards to the Drab Z Conclusion, I believe that it is plausible to suppose that some of the values realised by the members of A are discontinuous with the values realised by the members of Drab Z. This claim is supported by the fact that we do seem to prefer a certain amount of life that is well worth living to any amount of life that is so drab as to be only barely worth living.

⁶ Actually, Parfit calls it the 'Lexical View'. See Parfit, *Reasons and Persons*, p. 414.

Parfit himself takes note of this preference. Consider what he says in his article 'Overpopulation and the Quality of Life'⁷ concerning the analogue of the choice between A and Drab Z within a life, the choice between two futures: (1) the 'Century of Ecstasy', where one lives for a hundred years, all of an extremely high quality; and, (2) the 'Drab Eternity', where one lives forever, but where each year is only barely worth living – although free of pain, these years contain only a few simple pleasures.⁸ Parfit claims that although each year of life in the Drab Eternity would be worth living and have value (and given that we are dealing with an infinite number of such years, the Drab Eternity would be of infinite value), the Century of Ecstasy would still be a better life.⁹

How can a life of finite value be better than a life of infinite value? Clearly, at least some of the values which would be realised by a person living the Century of Ecstasy must be discontinuous with the values which would be realised by a person living the Drab Eternity. In the Century of Ecstasy, but not in the Drab Eternity, a person would experience what Mill called 'higher' pleasures,¹⁰ pleasures such as those derived from falling in love, listening to Bach, engaging in philosophical thought, etc. And Parfit rightly believes that not all pleasures lie on the same scale, that is, the deprivation of certain 'higher' pleasures cannot be made good by any gain in the quantity of 'lower' (simple) pleasures, not even an infinite gain.¹¹

So, if we agree with Parfit that certain values are discontinuous with others, if, for instance, we would never trade the joy of falling in love for an eternity of simple pleasures, then we must accept that living the Drab Eternity is a repugnant alternative to living the Century of Ecstasy. And, likewise, we may suppose that Drab Z is a repugnant alternative to A for the same reason: some of the values realised by the members of A are discontinuous with those realised by the members of Drab Z.¹²

⁷ In Peter Singer (ed.), *Applied Ethics* (Oxford: Oxford University Press, 1986), pp. 145–64.

⁸ *Ibid.*, pp. 160–61.

⁹ *Ibid.*, pp. 160–61.

¹⁰ J. S. Mill, *Utilitarianism* (London, 1863), ch. 2.

¹¹ Parfit, 'Overpopulation and the Quality of Life', pp. 160–61.

¹² Here, I have gone from making a personal assessment of the value of two alternative lives to making an impersonal assessment of the value of two alternative populations. But I think I am justified in making such a leap. For I believe that the best method for making evaluative choices between lives is to imagine the analogue of that choice within

So again we see that if we take the values being compared to be discontinuous, the Total Principle does not imply the Drab Z Conclusion. In fact, it implies the opposite: that A is better than Drab Z. The fact that those in Drab Z do not experience any of the 'higher' pleasures cannot be made good by any increase in the quantity of the 'lower' pleasures they experience. Although each life in Drab Z would have value, no number of such lives would be as valuable as even just one of the lives in A.

But, although the Total Principle implies that A is better than Drab Z, it would not necessarily imply that A is better than B. The lives in B are only slightly less worth living than those in A. So, presumably, the lives in B could be as rich in terms of quality as those in A: for instance, those in B may realise all the same values that are realised by those in A but only for a shorter period of time. And if the lives in B differ from those in A only in terms of quantity and not quality, then the Total Principle would imply that B is better than A.

Of course, there will come a point (as we proceed down the alphabet) where it will not be possible to have lives which are so much less worth living than those in A be, in qualitative terms, commensurable with those in A, if only because the realisation of certain 'higher' values (such as those associated with an old friendship) requires a lengthy period of time – quality of life is to some extent dependent upon length of life. This point represents a threshold beyond which it is no longer possible for a life to be any less valuable than one of the lives in A simply in virtue of it being shorter. At this point, a life cannot be any shorter without there being the loss of one or more of the 'higher' values. I will refer to this point as the Discontinuity Threshold.

So, by applying the Total Principle to an axiology which takes account of the discontinuity between certain values, we avoid the Drab Z Conclusion. And in doing so we have the added advantage of being able to stop at some point (that is, at the Discontinuity Threshold) along the slippery slope from A to Z.

a life. Thus, the best method for choosing between two alternative populations involves first imagining leading the lives of each member of each population in serial order, and second, choosing between the two resulting imaginary lives. See C. I. Lewis, *An Analysis of Knowledge and Valuation* (La Salle, Ill.: Open Court, 1946), pp. 546–47, and Crisp, 'Utilitarianism and the Life of Virtue', pp. 150–51.

For, intuitively, I believe, we want to say that B is better than A.¹³ But once we do so, it may seem that we must also say that C is better than B and that D is better than C and so on. However, we do not want to go so far as to imply that Drab Z is better than A. But I have shown that, despite appearances, we can consistently maintain both that B is better than A and that A is better than Drab Z.

III. Parfit's Argument Against the Discontinuity View

Parfit does realise that we can avoid the Drab Z Conclusion if we appeal to the discontinuity between certain values. That is, he claims that we can avoid the Repugnant Conclusion (i.e., the Drab Z Conclusion) if we appeal to

The Discontinuity View. 'There is no limit to the positive value of quantity. It is always better if an extra life is lived that is worth living'. But no number of lives below the Discontinuity Threshold could ever be as good as even just one life above the Discontinuity Threshold (p. 414).¹⁴

What Parfit fails to mention is that even someone who accepts the Total Principle can avoid the Drab Z Conclusion if s/he appeals to the Discontinuity View. For if the Discontinuity View is correct, the Total Principle does not imply the Drab Z Conclusion. Now this may seem important given that Parfit rejects the Total Principle on the assumption that it implies the Drab Z Conclusion. But, ultimately it does not matter for Parfit, because he rejects the Discontinuity View. He believes that it implies 'a weakened form of the Repugnant Conclusion' which he calls

¹³ I admit that some may not share my intuition in this case. Nevertheless, there is a series of arguments (see Griffin, *Well-being*, fn. 27, pp. 338–40) leading to the conclusion that Z is better than A, and all are sound regardless of whether or not the Total Principle is true. And one of the first conclusions in this series is that B is better than A. So, whether others share my intuition or not, I believe they will be forced to accept that B is better than A in light of these arguments.

Interestingly, Griffin believes that we can defeat these series of arguments before reaching the Repugnant Conclusion by appealing to the discontinuities between certain values – but, presumably, not before reaching the conclusion that B is better than A. However, Griffin fails to realise that an appeal to the discontinuities between certain values is only effective in avoiding one of the three versions of the Repugnant Conclusion, namely, the Drab Z Conclusion.

¹⁴ Where I talk about lives above and below the Discontinuity Threshold, Parfit speaks of 'Blissful' and 'Mediocre' lives, respectively. Nonetheless, the idea is the same.

(R): For any large population of people, all with lives that are well worth living, there will be some much larger imaginable population whose existence would be better, even though its members would all have lives that are only barely above the Discontinuity Threshold (pp. 415–16, 528).

But is (R) really repugnant? As Parfit points out, the extent to which we find (R) repugnant will depend on where we think the Discontinuity Threshold lies (p. 416).¹⁵ Now given the fact that even lives which are only barely above the Discontinuity Threshold will contain what I (following Parfit) will call ‘the best things in life’, it seems reasonable to presume that such lives will still be *well* worth living. Thus, it seems that the Discontinuity Threshold will lie somewhere close to the beginning of the alphabet on the continuum from A to Drab Z. But, in this case, it hardly seems that (R) is repugnant.

However, Parfit has yet another reason for rejecting the Discontinuity View: he believes that it implies a variant of the ‘Absurd Conclusion’ (pp. 410–11) which he calls

(A): Imagine an enormous population of people almost all of whom have lives that are just barely below the Discontinuity Threshold; these lives are as good as it is possible to be absent the best things in life. (Let us call these people the Fortunate.) However, for every fifty million such lives there is one person who, due to sheer bad luck, has a life that is not worth living. (Let us call these people the Unfortunate.) The existence of this enormous population would be worse than if no population existed instead (pp. 415–16, 528).

Parfit argues as follows. First, he rightly claims that, on the Discontinuity View, the existence of the Fortunate (no matter how numerous) would not be as good as the existence of even just one life above the Discontinuity Threshold (fn. 40, p. 528). Second, he claims that the suffering of the Unfortunate would more than cancel out the value of one life above the Discontinuity Threshold (fn. 40, p. 528) – and from this it follows that the existence of a population which included the Unfortunate plus one life above the Discontinuity Threshold would be worse than if no population existed instead. Therefore,

¹⁵ Note that Parfit’s equivalent term for what I call the Discontinuity Threshold is the ‘Valueless Level’.

Parfit concludes that, on the Discontinuity View, the existence of a population which includes the Unfortunate plus the Fortunate – that is, the population described in (A) – would be worse than if no population existed instead.

So, Parfit admits that the closer the Discontinuity Threshold lies to the beginning of the alphabet the less (R) will seem repugnant. Nevertheless, he believes that, no matter where the Discontinuity Threshold lies, we will still have reason to reject the Discontinuity View, because the further the Discontinuity Threshold lies from the end of the alphabet the more (A) will seem absurd – that is, the further the Discontinuity Threshold lies from the end of the alphabet the more absurd it is to claim that the enormous population described in (A) is worse than if no population existed instead (p. 416). Thus, Parfit argues that, no matter where the Discontinuity Threshold lies, the Discontinuity View will have unacceptable implications.

Now although there is no problem with Parfit's reasoning, I think we should reject Parfit's arguments against the Discontinuity View. For I do not believe that we should necessarily accept Parfit's assumption that the suffering of the Unfortunate would more than cancel out the value of one life above the Discontinuity Threshold. It is plausible to suppose that certain types of pain are discontinuous in the same way that certain types of pleasures are. That is, just as we would prefer falling in love to even an infinite quantity of some 'lower' pleasure, we might prefer an infinite quantity of some mild pain to the agony of heartbreak.

Hence, there are two significantly different ways in which we can think of the Unfortunate as suffering: we can think of them as merely suffering mild pain, or we can think of them as suffering pain like that of heartbreak where such pain (or perhaps a sufficient quantity of such pain) is taken to trump any quantity of mild pain. Now, if the Unfortunate suffer only mild pain, then we should reject Parfit's claim that the suffering of the Unfortunate more than cancels out the value of one life where the best things in life are realised. For it seems plausible to suppose that no quantity of mild pain can cancel out the value of that which is best in life. For instance, I think that most of us would be willing to live the rest of our lives with mild arthritic pain if the only cure was one which comes at the cost of losing one's capacity for love.

On the other hand, if the Unfortunate suffer pain like that which comes from heartbreak, then we should accept Parfit's

claim that the suffering of the Unfortunate more than cancels out the value of one life where the best things in life are realised. But, in this case, (A) no longer seems absurd. (A) no longer seems absurd, because it does not seem that the happiness of the Fortunate can outweigh the suffering of even one unfortunate person who suffers some of the worst pain that a person can suffer. After all, the Fortunate are not so fortunate as to have what is best in life, and it is plausible to suppose that only the joy which comes with realising what is best in life can outweigh the suffering which comes with enduring what is worst in life. Thus, in this case, it is not at all absurd to suppose that the existence of the enormous population described in (A) would be worse than if no population existed instead.

I realise that the Discontinuity View may still seem controversial. But I have tried to show why someone who takes the Discontinuity View to be correct need not think of its implications as being either absurd or repugnant. The real proof of what I claim here will come with my arguments concerning the Short-lived Z Conclusion, the case where we abstract away from complex axiological concerns by holding the quality of lives constant from A to Z (varying them only in terms of length of life). I will argue that there is good reason to conclude that Short-lived Z is better than A. And I believe that we should accept this conclusion without repugnance. I will thereby prove that the repugnance of the Drab Z Conclusion derives not from the Total Principle itself but from its application to an axiology which denies the discontinuity between certain values. I will now turn to these arguments.

IV. Why We Must Accept the Short-lived Z Conclusion

The Total Principle does imply that Short-lived Z is better than A.¹⁶ Here, there are no discontinuities to be concerned with; the lives in Short-lived Z are qualitatively identical to those in A. But, in this case, the fact that the Total Principle implies the 'Repugnant' Conclusion, that is, the Short-lived Z Conclusion,

¹⁶ Earlier I claimed that the realisation of certain 'higher' values is possible only given a rather lengthy period of time. Therefore, if the lives in Short-lived Z are going to be qualitatively identical to those in A, we must suppose that those in A never realise any of these 'higher' values. So, in this case, we must imagine that the lives in A are well worth living, not because they contain all the higher values, but because they contain such a large quantity of those values which can be realised in even the shortest period of time.

does not give us reason to reject the Total Principle, because the Short-lived Z Conclusion is something we must accept.

The conclusion that Short-lived Z is better than A is entailed by a number of compelling claims, claims which Parfit himself makes concerning some puzzling examples. And given that Parfit makes these claims, he too is committed to the conclusion that Short-lived Z is better than A.

Consider *The Two Hells*:

Hell A (Parfit's 'Hell One') is a population consisting of ten people, who each undeservedly suffer terrible agony for fifty years. Their lives are much worse than non-existence, and thus they would all kill themselves if they could. Hell B (Parfit's 'Hell Two') is a population of ten million people, who each undeservedly suffer the same agony for fifty years minus a day (p. 406).

Parfit believes that Hell B is worse than Hell A (p. 406). One way to justify this belief is to claim that a vast increase in the total sum of suffering within a population morally outweighs a very small reduction in the average suffering per life within a population (p. 406). But this claim implies that there is a Hell Z – an enormous population of people each of whom undeservedly suffer the same great agony for a little less than a day – which would be the worst of all. We are led from the position that Hell B is worse than Hell A to holding Hell Z to be the worst of all by the same parity of reasoning (the same slippery slope type argument) which lead us from the claim that B is better than A to its ultimate implication, that (Short-lived) Z is the best of all.¹⁷ (Imagine Fig. 1 inverted with the blocks labelled Hell A through Hell Z – except, in this case, you should imagine that each subsequent block is not just two, but a million, times wider than the former.)

However, Parfit believes there is another way one can justify the belief that Hell B is worse than Hell A. One can reject the claim that a vast increase in the total suffering morally outweighs a small reduction in the average suffering (and thereby avoid implying that Hell Z is worse than all its alternatives) and still hold Hell B to be worse than Hell A, but only if one appeals to

¹⁷ The slippery slope from Hell A to Hell Z and the slippery slope from A to Short-lived Z are both significantly different from the slippery slope from A to Drab Z. With the first two, we are not dealing with discontinuous values; so, there is no way to halt the slide down the slippery slope in these two cases, as there is in the case of the slippery slope from A to Drab Z.

The Limited Suffering Principle. 'It will be bad if, at any time, there is a greater sum of suffering than there might have been, unless this sum is above a certain limit' (p. 406).

The Limited Suffering Principle does enable one to stop at some point along the slippery slope from Hell A to Hell Z, but Parfit rejects this principle. He says, 'it is much more implausible than the Repugnant Conclusion' (p. 406). Thus, given that Parfit rejects the Limited Suffering Principle, it would seem that by holding Hell B to be worse than Hell A he is thereby obliged to accept the conclusion that Hell Z is the worst of all – the converse of the Short-lived Z Conclusion.

And Parfit would seem to be right in rejecting the idea that there is a limit to the badness of an increase in suffering. Consider the choice between two other possible hells:

Hell Y is a population consisting of ten billion people, each of whom undeservedly suffers terrible agony for a day. Hell Z is a population of ten quadrillion people, each of whom undeservedly suffers agony just as great for one day minus 4.73 seconds. (Here, I have made the proportions the same as those in the choice between Hell A and Hell B. Fifty years is to one day as one day is to 4.73 seconds. Ten is to ten million as ten billion is to ten quadrillion.)

Presumably, if the Limited Suffering Principle can be considered plausible, our intuitions should change at some point as we progress down the *Hell Alphabet*. However, it seems that the intuitions which cause us to think that Hell B is worse than Hell A are just as strong in our determination that Hell Z is worse than Hell Y.

Thus far, we have seen that Parfit is committed to the position that Hell Z is worse than all of its alternatives. Now, unless Parfit holds that there is a difference between the valuation of pleasure and the valuation of pain, he must accept its converse, namely, that Short-lived Z is better than all its alternatives (A being one of its alternatives). In fact, Parfit rejects the view that there is a difference between the valuation of pleasure and pain; Parfit rejects what he calls the 'Asymmetry' (the view that states of affairs having a greater quantity of pleasure is of greater value up to a point, but states of affairs having a greater quantity of pain is of limitlessly increasing disvalue) because its acceptance implies the following absurdity:

The Absurd Conclusion: For any large population of people almost all of whom have lives that are well worth living – the exceptions being one in every fifty million who through sheer bad luck have lives that are not worth living – there will be some much larger population whose existence would be a worse alternative even though there would be the same prevailing quality of life and proportionally no greater number of unfortunate people. (The limitless disvalue of the increase in the quantity of pain corresponding to the increase in the number of unfortunate people [one for every fifty million added to the population] would eventually come to outweigh the limited value of what is a proportionately greater increase in the quantity of pleasure.) (pp. 410–11)¹⁸

So, it would seem that if we assent to three claims (claims to which Parfit himself assents), namely, (1) that Hell B is worse than Hell A, (2) that there is no point (no quantitative limit) at which an increase in the amount of suffering can no longer be of added disvalue, and (3) that there is no asymmetry between the valuation of pleasure and pain, then we are obliged, on pain of inconsistency, to accept the Short-lived Z Conclusion. Either we accept that Short-lived Z is better than A or we must deny one of these three claims. If we deny (1), we would have to hold Hell A to be worse than Hell B. If we deny (2), we would have to hold Hell Y to be worse than Hell Z. And, if we deny (3), we would have to accept the Absurd Conclusion. When faced with these alternatives, the Short-lived Z Conclusion seems far less repugnant, that is, if it is repugnant at all. The Short-lived Z Conclusion is, at the very least, ‘*much* less repugnant’ than the Drab Z Conclusion; this much Parfit admits.¹⁹ But I believe that one must also admit that it is much less repugnant than the alternative, which is to reject either (1), (2), or (3).

Now it might be suggested that I have not shown that Parfit is committed to the Short-lived Z Conclusion, but only that he is committed to a paradox. And Parfit says, ‘We can’t solve such paradoxes by claiming that, if the arguments are indeed strong, we oughtn’t to find the conclusions unacceptable’.²⁰ But if what

¹⁸ This is actually a much simplified version of the Absurd Conclusion which Parfit presents, but it does nevertheless make exactly the same point.

¹⁹ This comment was made in correspondence regarding an earlier draft of this paper.

²⁰ *Ibid.*

we do have here is a paradox, it is one which is easily resolved in light of my arguments.

Of course, the ideal solution would be to dissolve the paradox altogether by convincing people that the Short-lived Z Conclusion is not repugnant. I am not sure how to go about convincing people of this, because, for me, it is just a basic intuition. Nevertheless, there is a resolution to the paradox for even those who, like Parfit, find the Short-lived Z Conclusion somewhat repugnant.

One cannot resolve the paradox in holding a set of beliefs which are all individually plausible but mutually inconsistent by arbitrarily rejecting one over the others, but such a paradox can be resolved if rejecting one of the beliefs is much less repugnant than rejecting any of the others. Such is the case here. If we refuse to accept the Short-lived Z Conclusion, then we must reject either (1), (2), or (3). And what we would have to accept were we to deny any one of these three claims is far more repugnant than the Short-lived Z Conclusion.

V. Why the Total Principle Does Not Imply the Roller Coaster Z Conclusion

Not only does the Total Principle imply the Short-lived Z Conclusion, but it also seems to imply the Roller Coaster Z Conclusion. As with the members of Short-lived Z, the members of Roller Coaster Z experience the same 'higher' pleasures that those in A experience. So, no matter what values we take to be discontinuous, the Total Principle will still imply the Roller Coaster Z Conclusion.

Now if the Total Principle does imply the Roller Coaster Z Conclusion (as it seems to), we have good reason to reject it, because, unlike the Short-lived Z Conclusion, the Roller Coaster Z Conclusion is clearly quite repugnant. For consider the analogue of the choice between A and Roller Coaster Z within a life, the choice between two futures: (1) the Century of Ecstasy, again, where one lives for a hundred years, all of an extremely high quality; and (2) the Roller Coaster Eternity, where one lives forever, but where each year of life is only barely worth living because the ecstasies just barely outweigh the agonies. Clearly, living the Roller Coaster Eternity is a repugnant alternative to living the Century of Ecstasy. Yet, given that each year of life in the Roller Coaster Eternity is worth living (although just barely)

and has value, and given that we are dealing with an infinite number of such years, it may seem that the Roller Coaster Eternity would have to be better than the Century of Ecstasy. After all, isn't a life of infinite value better than a life of finite value?

Well, as we have seen, a life of infinite value isn't always better than a life of finite value: earlier, we concluded that the Century of Ecstasy is better than the Drab Eternity in spite of the fact that the Drab Eternity is, in quantitative terms, infinitely more valuable. We did so by appealing to the discontinuity between certain values; that is, we appealed to the fact that a sufficient quantity of certain 'higher' values is better than any quantity of 'lower' values. However, we cannot avoid the conclusion that the Roller Coaster Eternity is better than the Century of Ecstasy by making such an appeal, because someone living the Roller Coaster Eternity will experience the exact same agonies and ecstasies which someone living the Century of Ecstasy will; only, in the latter case, the proportion of agonies to ecstasies is much smaller.

Nevertheless, I believe there is a way we can avoid the conclusion that the Roller Coaster Eternity is better than the Century of Ecstasy. We need only deny that well-being is additive in the sense that the welfare value of a person's life is always equal to the sum of how well off that person was during each moment of his/her life. That is, we must reject

The Summative View. The view that the welfare value of a life is necessarily equivalent to the sum of the momentary well-being²¹ contained within.

If we accept the Summative View, we must also accept that the Roller Coaster Eternity is better than the Century of Ecstasy, because clearly the Roller Coaster Eternity does contain a greater net sum of momentary well-being than the Century of Ecstasy does. But I believe that we should reject the Summative View, because there are a number of cases where our intuitions

²¹ Momentary well-being concerns how well off a person is at a particular moment. For example, were we to accept a desire-fulfillment account of well-being, a person's momentary well-being would be a measure of the extent to which his current desires, actual or hypothetical, are at that moment fulfilled. So, hypothetically, we could calculate how well off a person is at a given moment by assigning a positive number to those desires of his which are, at that moment, fulfilled and assigning a negative number to those desires of his which are, at that moment, unfulfilled and then summing them together – how great the numbers we assign will depend on the strength of the given desires. See Parfit, *Reasons and Persons*, p. 496.

concerning the value of a given life conflict with what the Summative View implies. For instance, consider the following example from J. David Velleman's article 'Well-being and Time'.

Consider two different lives that you might live. One life begins in the depths but takes an upward trend: a childhood of deprivation, a troubled youth, struggles and setbacks in early adulthood, followed finally by success and satisfaction in middle age and a peaceful retirement. Another life begins at the heights but slides downhill: a blissful childhood and youth, precocious triumphs and rewards in early adulthood, followed by a midlife strewn with disasters that lead to misery in old age. Surely, we can imagine two such lives as containing equal sums of momentary well-being. Your retirement is as blessed in one life as your childhood in the other; your nonage is as blighted in one life as your dotage is in the other.²²

Concerning this example, Velleman claims that most of us would agree that the life of improvement is better than the life of deterioration. Why? Well, as Velleman argues, it seems that, other things being equal, we prefer a life that gets better to a life that gets worse. Thus, the life of improvement is better than the life of deterioration, not because events which come earlier in life have a lesser impact on the value of one's life than events which come later in life, but because later events alter the meaning of earlier events. For instance, in the life of improvement, the tragedies turn out to be meaningful, because those defeats led to the triumphs. Whereas, in the life of deterioration, the tragedies are meaningless, for they have led to nothing. Thus, it is not that tragedies are more hurtful when they come later in life. Rather, what is significant about tragedies which come late in life is that they come at a time when it is too late to learn and grow from them; so, they are suffered for not.

Now if we believe that the life of improvement is better than the life of deterioration in spite of the fact that both contain equal sums of momentary well-being, then we must reject the Summative View. And so, we must instead accept the 'Global View'. On the Global View, we cannot tell whether one life is better than another simply by comparing the sums of momentary well-being contained within each. Rather, we have to look at each

²² J. David Velleman, 'Well-being and Time', *Pacific Philosophical Quarterly*, 72 (1991), p. 50.

life as a whole considering not only how much momentary well-being each life contains, but also considering the global features of each, such as, how the momentary well-being is distributed through time and the impact that that distribution has on the meaningfulness of each life.

And, importantly, if we accept the Global View, we can deny that the Roller Coaster Eternity is more valuable than the Century of Ecstasy even though the Roller Coaster Eternity is infinitely more valuable in terms of momentary well-being. We can claim that, globally speaking, the Century of Ecstasy is more valuable than the Roller Coaster Eternity. And it seems plausible to do so, for it seems that the Roller Coaster Eternity is a life without any meaningful direction, and, globally speaking, we seem to prefer that our life has some sort of meaningful or positive direction. For instance, we prefer a life of improvement to a life of deterioration, because we prefer a life with a positive direction to a life with a negative direction, other things being equal. But we also seem to prefer a life with a positive direction to a life without any direction (e.g., the Roller Coaster Eternity), other things being equal. Of course, when we compare the Roller Coaster Eternity with the Century of Ecstasy everything else is not equal: the two lives contain different sums of momentary well-being. But, it doesn't seem to matter, for we prefer a short life with meaning to even an eternity of aimlessness.

Furthermore, if we hold the Century of Ecstasy to be more valuable than the Roller Coaster Eternity, we should take A to be better than Roller Coaster Z in terms of welfare value. After all, each member of A lives a life comparable to the Century of Ecstasy, and the Roller Coaster Eternity is comparable to one person living all the lives of Roller Coaster Z in serial order. Thus, on the Global View, the Total Principle does not imply the Roller Coaster Z Conclusion.

But one may ask, what if instead of imagining the members of Roller Coaster Z leading lives where every year is a series of extreme ups and downs, we imagine that they each lead the life of improvement as described by Velleman. Wouldn't the conclusion that this population is better than A be a version of the Repugnant Conclusion which the Total Principle implies even if we accept the Global View? Well, the Total Principle would imply that such a population is better than A, but this implication is not a version of the Repugnant Conclusion. For although the life of improvement may be barely worth living in terms of momentary

well-being, it is, on the Global View, a life well worth living. Thus, this population would not be a version of Z; it would not be an enormous population where everyone has a life that is only barely worth living.

VI. Conclusion

In the concluding chapter of *Reasons and Persons*, Parfit admits to failure: he failed to find an adequate formulation of the principle of beneficence (i.e., Theory X), one which could, among other things, account for our intuition that an action can be wrong even if it doesn't cause anyone to be worse off. The Total Principle is able to provide such an account, and it seems a promising candidate but for the fact that it appears to imply the Repugnant Conclusion.

However, I have argued that, so long as we apply it to a sophisticated axiology, the Total Principle will only imply one version of the Repugnant Conclusion, the Short-lived Z Conclusion, and this version is much less repugnant than either of the other two, that is, if it is repugnant at all. And I have argued that even if the Short-lived Z Conclusion does seem somewhat repugnant, it is a conclusion we must accept in light of the alternatives.²³

Department of Philosophy & Religious Studies
College of Charleston
Charleston SC 29424-0001
USA
DPortmore@aol.com

²³ I am indebted to Matthew Hanser, Derek Parfit, and a number of anonymous referees for their many helpful comments and suggestions.