

Winner of *The Philosophical Quarterly* Essay Prize 2018

ON PARFIT'S WIDE DUAL PERSON-AFFECTING PRINCIPLE

BY MICHAL MASNY 

In the posthumously published 'Future People, the Non-Identity Problem, and Person-Affecting Principles' (2017), Derek Parfit presents a novel axiological principle which he calls the Wide Dual Person-Affecting Principle and claims that it does not imply the Repugnant Conclusion. This paper shows that even the best version of Parfit's principle cannot avoid this conclusion. That said, accepting such a principle makes embracing the Repugnant Conclusion more justifiable. This paper further addresses important questions which Parfit left unanswered concerning: the relative importance of individual and collective goodness, comparisons involving unequal outcomes, how to understand individual goodness, and whether incomparability at the level of individual goodness implies incomparability at the level of overall goodness.

Keywords: Parfit, repugnant conclusion, population ethics, axiology, incomparability.

I. INTRODUCTION¹

How should we rank, in terms of their relative goodness, outcomes in which different numbers of different people with lives of different personal goodness exist? One might think that this could be done by appealing to

The Wide Person-Affecting Principle: 'One of two outcomes would be in one way worse if this outcome would be less good for people, by benefiting people less than the other outcome would have benefited people' (Parfit 2017, p. 129).

Derek Parfit first considers this Wide Principle in *Reasons and Persons* (1984), but rejects it because it implies

¹ Thanks to Tomi Francis, Jake Nebel, and Theron Pummer for excellent written comments, and to Johann Frick and Jeff McMahan for illuminating conversations and words of encouragement. My greatest debt is to Kacper Kowalczyk whose remarkable ability to explain complex philosophical ideas by drawing simple diagrams on cafe napkins was an invaluable resource throughout my time at Oxford.

The Repugnant Conclusion: For any outcome in which many people exist, all with very good lives, there is a better outcome in which many more different people exist, all with barely good lives.²

Over 30 years later, in 'Future People, the Non-Identity Problem, and Person-Affecting Principles' (2017), Parfit returns to the Wide Principle. He claims to have overlooked that there are two ways of benefiting people, not just one. This observation is accommodated in a novel version of this principle endorsed by Parfit:

The Wide Dual Person-Affecting Principle: 'One of two outcomes would be in one way better if this outcome would together benefit people more, and in another way better if this outcome would benefit each person more' (2017, p. 154).

Parfit's discussion of this principle is confined to just a few pages at the end of his article. He claims that this 'Wide Dual Principle does not imply the Repugnant Conclusion' (p. 157), but the argument has gaps and some important questions are left unanswered. This is understandable – readers should bear in mind that Parfit submitted 'Future People . . .' in a state that he regarded as unfinished and an untimely death prevented him from completing this project. Consequently, some interpretative and speculative work is necessary to uncover the most accurate and the most plausible version of the Wide Dual Principle.

This paper examines this Wide Dual Principle. Section 2 states essential assumptions and outlines the wider theoretical framework in which this principle is embedded. Section 3 offers a systematic explanation of the Wide Dual Principle and of its implications in comparisons of equal outcomes. Section 4 shows, contra Parfit, that even the best version of this principle cannot avoid the Repugnant Conclusion and argues that embracing this principle makes accepting this conclusion more justifiable. The remainder of this paper turns to issues which Parfit left unaddressed. Section 5 discusses the relative importance of the two components of the Wide Dual Principle: individual and collective goodness. Section 6 shows that the definition of individual goodness used by Parfit needs to be revised to account for our intuitive verdicts about some comparisons involving unequal outcomes. Section 7 discusses two ways of determining which of two outcomes is overall better when they are incomparable in terms of individual goodness. Section 8 delivers concluding remarks.

² This is not how Parfit formulates this conclusion in (1984, p. 388) and (2017, p. 153). The above formulation is more general and preferable for reasons that will be explained in the next section.

II. PRELIMINARIES

II.1. Assumptions

This paper makes several substantive and interpretative assumptions.

First, the Repugnant Conclusion is formulated and understood here in terms of the goodness of lives. To determine how good a life is, we need a theory of well-being and a theory of intrapersonal aggregation of well-being. This is not how Parfit states the Repugnant Conclusion in his work. In *Reasons and Persons* (1984, p. 388) and ‘Future People . . .’ (2017, p. 129), he tends to talk about the ‘quality of life’. This is either insufficiently general or potentially misleading. This term is often used to refer to a person’s average well-being per some unit of time: you might say that your quality of life was higher when you resided in Zurich than it is now that you live in Kathmandu. But, unless we already know that the average is the correct principle of intrapersonal aggregation of well-being, we should not limit our discussion in this way. For the same reason, we should not assume any particular theory of well-being. The formulation of the Repugnant Conclusion in terms of the goodness of lives is neutral in these respects.

Second, the following assumptions will be helpful. A life is *good (bad)* just in case it is good (bad) for the person living it, and *neutral* just in case it is neither good nor bad for that person. Let’s assume that one life can be better than another and that two lives can be equally good. We can understand an equivalence class of equally good lives as a *life goodness level*. Let’s further assume that there are reasonably fine-grained life goodness levels which can be represented with integers, in such a way that higher numbers correspond to better life goodness levels, the differences in numbers correspond to the differences in life goodness, positive (negative) numbers correspond to good (bad) lives, and zero corresponds to a neutral life.

Third, let’s assume that there are three basic evaluative relations: better than, equally good, and worse than. Outcomes not related in any of these ways are incomparable. As a substantive assumption, this is justified by reasons of simplicity: it can be relaxed when necessary. It is also warranted as an interpretative assumption. In *Reasons and Persons*, Parfit introduces a fourth relation – *imprecise equality* – and he discusses it further in ‘Can We Avoid the Repugnant Conclusion’ (2016).³ However, imprecise equality is not mentioned in ‘Future People . . .’ even once, so it is reasonable to assume, at least initially, that he is breaking ground on a novel approach which does not involve imprecise equality.

³ See also Ruth Chang’s discussion of *parity* in her (2002) and (2016).

II.2. *The wider theoretical framework*

It is important to understand Parfit's wider theoretical framework in which the Wide Dual Principle is embedded.

Firstly, at the foundation of this principle is the claim that to be created with a good life is an *existential benefit* for the created person and to be created with a bad life is an *existential harm*. This is a controversial move because the terms 'benefit' and 'harm' are typically used either in the *comparative temporal* sense or else in the *comparative counterfactual* sense.⁴ Parfit, following Jeff McMahan (1981; 2009; 2013), suggests that these terms can be used in the *noncomparative* sense as well. Although being created with a good life is not *better for* the created person, it is nonetheless *good for* that person.

Secondly, we can compare an outcome in which Sally has a good life with an outcome in which Sally does not exist even though in the former case Sally is benefited only in the noncomparative sense. The former outcome is *better* because the former outcome is good for Sally and the latter outcome is not good for anyone. Moreover, we can compare outcomes in which different people exist by comparing how good one outcome is for the people who exist in it with how good another outcome is for the people who exist in it.

Thirdly, the Wide Dual Principle presupposes the view that the goodness of outcomes is *intrinsic*. That is, the goodness of an outcome depends on its intrinsic properties and not on its relation to other outcomes.⁵ On this view, 'better than' is a *transitive* relation: if X is better than Y, and Y is better than Z, then X is better than Z. Moreover, assessment of outcomes is *invariant under permutations* of identities of possible members of the domain. Put differently, it matters only what life goodness levels are occupied by members of the distribution, and not who is living any particular life or whether the compared outcomes contain some or none of the same people. An outcome in which Ann exists at 50 and Beth does not exist is as good as an outcome in which Ann does not exist and Beth exists at 50.⁶

These claims might be challenged, but they are simply assumed in the following discussion. The primary aim of this paper is to examine the implications of the Wide Dual Principle, not to question its foundations.

⁴ To be benefited in the comparative temporal sense is to be made better off than one was before. To be benefited in the counterfactual comparative sense is to be made better off than one would have been otherwise.

⁵ Earlier in the same article (especially pp. 137-144), Parfit argues at length against the rival view that the goodness of outcomes is at least sometimes *essentially comparative*. For a comprehensive discussion of this view, see Larry Temkin's *Rethinking the Good* (2012).

⁶ Parfit refers to permutation-invariance as *the No-Difference View* (Parfit 1984, p. 367).

III. THE WIDE DUAL PRINCIPLE

III.1. *An intuitive gloss*

The following example will help us get an intuitive grasp on how the Wide Dual Principle is intended to work. Suppose that you are choosing between two meals, A and B. Let us agree that two considerations determine the overall goodness of a meal: tastiness and healthiness. One of two meals is in one way better if it is tastier and in another way better if it is healthier. Although goodness in one of these respects is not reducible to goodness in the other respect, the overall evaluative relation between the two meals can be determined roughly as follows. If A is healthier and tastier than B, then A is overall better than B. If A and B are equally tasty and equally healthy, then they are overall equally good. If A is healthier than B, and A and B are equally tasty, then A is overall better than B. Most interestingly, comparisons seem possible even if A is tastier and B is healthier. For instance, if A is much tastier than B, and B is only slightly healthier than A, then A is plausibly overall better than B.⁷

The Wide Dual Principle works in an analogous way. Instead of meals, tastiness, and healthiness, we are concerned with outcomes, individual goodness (benefits to each), and collective goodness (benefits to all).

III.2. *Individual and collective goodness*

As we have seen, Parfit thinks that there are two ways of benefiting (and harming) people more: benefiting *each* more and benefiting *all* more. He thinks that there are correspondingly two basic ways in which one of two outcomes can be better: individually better and collectively better. In his discussion, Parfit sometimes uses one pair of terms and sometimes the other. For clarity, only the latter pair of terms will be used in this paper.

What are the conditions for individual and collective betterness? According to Parfit (2017, pp. 153–5), one of two outcomes is *individually better* if all people who exist in this outcome have better lives than all people who exist in the other outcome. And one of two outcomes is *collectively better* if it has the higher

⁷ To say that is to accept that there is some rate of trade-off between healthiness and tastiness of a meal and that this rate is not far from 1:1. This raises a fundamental question about the metaphysics of value: what makes it the case that two values irreducible to one another or some further common value can be traded-off at all and in any particular way? Any genuinely pluralistic axiology faces this question (Johann Frick (pc.) calls it the Deep Problem of Pluralism). This includes Parfit's Wide Dual Principle which treats individual and collective goodness as mutually irreducible. As we will see shortly, Parfit clearly thinks that these two values can sometimes be traded-off but he does not provide an answer to the aforementioned question. Unfortunately, this paper does not offer it either. Let's just assume that individual and collective goodness can be traded-off. In Section 5 we will consider the relative merits of different ways of trading-off these values.

total sum of life goodness. The relations of individual and collective betterness together determine the overall betterness relation.

III.3. *Implications of the Wide Dual Principle in equality cases*

In his article, Parfit discusses only *equality cases* in which all people who exist in any particular outcome have equally good lives. This section, too, considers only such cases.⁸ Let us see how the Wide Dual Principle handles them.

Let $C(X)$ be the level of collective goodness and $I(X)$ be the level of individual goodness of an outcome X . For any two outcomes, X and Y , there are five types of equality cases:

Type 1: $C(X) = C(Y)$ and $I(X) = I(Y)$

Type 2: $C(X) > C(Y)$ and $I(X) > I(Y)$

Type 3: $C(X) = C(Y)$ and $I(X) > I(Y)$

Type 4: $C(X) > C(Y)$ and $I(X) = I(Y)$

Type 5: $C(X) > C(Y)$ and $I(X) < I(Y)$

Type 5 cases are the most interesting because considerations of individual and collective goodness pull us in the opposite directions. Let us run through some examples of each type.

Consider a Type 1 case:

Case 1: **A** 2 people at level 60

B 2 (other) people at level 60

In Case 1, outcome A is collectively as good as B because the total sum of goodness of lives in the two outcomes is the same. A is individually as good as B because the people in A and B have equally good lives. A is thus overall as good as B.

Consider a Type 2 case:

Case 2: **A** 2 people at level 60

B 1 person at level 50

In Case 2, A is collectively better than B. A is also individually better than B because each person in A has a better life than each person in B. Since A is better in both respects, A is overall better than B.

It is worth emphasising that, in the context of individual goodness in equality cases, it does not matter how many individuals exist in each outcome. Just like the presence of an additional meal of some kind does not change the tastiness of the option in which only meals of this kind are consumed, the presence of

⁸ As we will see in Section 6, introducing inequality raises a number of important complications. Most importantly, worries of pervasive incomparability will force us to at least consider revising Parfit's simple condition for individual betterness.

an additional person with an equally good life in outcome A does not change A's individual goodness.

Consider a Type 3 case discussed by Parfit (2017, pp. 150–2):

Case 3: **A** 1 person at level 100
B 10 people at level 10

In Case 3, A is collectively as good as B, but A is individually better than B. The Wide Dual Principle implies that A is overall better than B.

Is this the right verdict? The question is whether it would be better if some amount of benefits was shared between fewer or more people. Parfit thinks the former. This is also the spirit of Larry Temkin's *Consolidate Additional Benefits View* (2012, p. 68).

Similar considerations apply to cases in which people have bad lives. Consider:

Case 4: **A** 1 person at level -100 .
B 10 people at level -10

In Case 4, A is collectively as good as B, but B is individually better than A, so B is overall better than A. Whereas it is plausible to consolidate benefits, it seems that we should *disperse harms* (compare: Temkin 2012, p. 76).

Consider a Type 4 case discussed by Parfit (2017, p. 154):

Case 5: **A** $2N$ people at level 100 exist
B N people at level 100 exist

In Case 5, A is individually as good as B, but A is collectively better than B. The Wide Dual Principle implies that A is overall better than B. As earlier, two outcomes are tied according to one criterion, and the other criterion serves as a tiebreaker. The question here is whether it would be better, other things equal, if more rather than fewer people with good lives existed. Unsurprisingly, considering the commitment to the claim that being created with a good life is a benefit, Parfit's principle implies that it would be better. It also has the plausible implication that it would be better if fewer rather than more people with bad lives existed.

Finally, let's consider a Type 5 case discussed by Parfit (2017, p. 155):

Case 6: **A** 1 million people at level 1,000
Z 100 billion people at level 1

Parfit thinks that the assessment of this case is straightforward. He writes:

'In [Case 6], Z would be collectively better than A, since Z would together benefit people roughly a hundred times more. But Z would be individually worse than A, since Z would benefit each person roughly a thousand time less. The lives of the people in A would be roughly a thousand times more worth living. We can plausibly believe that this second

fact would do much more to make this outcome better. On this version of the Wide Dual Principle, Z would be worse than A' (2017, p. 156).

Parfit recognises that if Z were better than A, then this would be an instance of the Repugnant Conclusion. He comments:

'As these remarks show, the Wide Dual Principle does not imply the Repugnant Conclusion' (p. 157).

IV. THE REPUGNANT CONCLUSION

IV.1. Getting clearer on how to assess Type 5 cases

Is Parfit correct when he says that the Wide Dual Principle does not imply the Repugnant Conclusion? To answer this question, we must consider Type 5 cases in more general terms. In the cases of interest (like Case 6), X is individually better than Y, whereas Y is collectively better than X. For short: $I(X) > I(Y)$ and $C(Y) > C(X)$. Parfit does not provide a rigorous method for evaluating such cases. He merely suggests that

'We can plausibly believe that slightly greater benefits to each would be outweighed by much greater benefits to all, and slightly greater benefits to all would be outweighed by much greater benefits to each' (2017, p. 155).

Call it *Parfit's Rule*. There are at least two possible interpretations of it.

The first interpretation is in terms of *differences in value*: $I(X) - I(Y)$ and $C(Y) - C(X)$. There are two versions of this interpretation. 'Much greater' and 'slightly greater' can track either the *absolute* differences in value, or else the *relative* differences in value. On the first version, X would be individually much better and Y would be collectively slightly better just in case:

$$(I(X) - I(Y)) > k > (C(Y) - C(X)), \text{ for some elected } k > 0$$

On the second version, X would be individually much better and Y would be collectively slightly better just in case:

$$(I(X) - I(Y)) \geq n \times (C(Y) - C(X)), \text{ for some elected } n > 1$$

The 'differences in value' interpretation is unsatisfactory. In Case 6, $I(A) - I(Z) = 999$ and $C(Z) - C(A) = 99$ billion. Thus, on the first version of this interpretation, if A is much better individually, then Z is also much better collectively. And if Z is slightly better collectively, then A is only slightly better individually. It depends on the value of k . On the second version of this interpretation, either A is slightly better individually and Z is much better collectively (if $n < 99,099,099.0991$) or else A is slightly better individually and Z is

slightly better collectively (otherwise). On either version, Parfit's Rule does not apply to Case 6, and Parfit thought it does, so it not an accurate interpretation of Parfit's view.⁹

The second interpretation of Parfit's Rule is in terms of *ratios*: $|I(X)|/I(Y)$ and $|C(Y)|/C(X)$.¹⁰ Again, there are two versions. 'Much greater' and 'slightly greater' can track either the *absolute* value of ratios or else the *relative* value of ratios. On the first version, X would be much better individually and Y would be slightly better collectively just in case

$$(|I(X)|/I(Y)) > k > (|C(Y)|/C(X)), \text{ for some elected } k > 1$$

On the second version, X would be much better individually and Y would be slightly better collectively just in case:

$$(|I(X)|/I(Y)) \geq n \times (|C(Y)|/C(X)), \text{ for some elected } n > 1$$

Both versions of this interpretation are consistent with Parfit's comments about Case 6. There, $|I(A)|/I(Z) = 1,000$ and $|C(Z)|/C(A) = 100$. So, the inequalities from the previous paragraph hold for a wide range of $k > 1$ and $n > 1$.

However, the first version faces a problem. To accommodate Parfit's comments about Case 6, it would be necessary to claim that, for one outcome to be *much* better individually or collectively than the other, it would have to be more than 100 times better in that respect. After all, Z is 100 times better collectively than A and it is regarded as only slightly better collectively. But the differences in the goodness of lives of different people do not seem to be as dramatic in the real world, so the Wide Dual Principle would be left with no bite in many relevant cases.

The second version of this interpretation avoids this problem. To accommodate Parfit's comments about Case 6, it would only be necessary to claim that for one outcome to be much better individually than the other, this outcome would have to be *10 times better individually than the other is collectively*.

The second version is also more intuitive. In Type 5 cases, we are interested in how strongly considerations of individual goodness favour one outcome *relative to* how strongly considerations of collective goodness favour the other. Consequently, Parfit is best understood as implicitly committed to the following sufficient condition for overall betterness:

⁹ Parfit calls it *Case Ten* (2017, p. 155).

¹⁰ Taking the modulus (absolute value) of one of the values in each ratio is necessary to correctly account for cases in which we compare outcomes in which people have bad lives. Otherwise in Case 7 (A-: 1 million people at -1,000; Z-: 100 billion people at -1), A- would turn out to be overall better. It is not, or at least it is not according to the logic of the Wide Dual Principle. Z- is 1,000 times individually better and A- is only 100 times collectively better.

- (1) For any two outcomes X and Y, such that $I(X) > I(Y)$ and $C(X) < C(Y)$,
if $(|I(X)| / I(Y)) \geq 10 \times (|C(Y)| / C(X))$ *then* X is overall better than Y.

IV.2. *The Wide Dual Principle implies the Repugnant Conclusion*

In Case 6, one outcome is significantly better individually than the other is collectively. Parfit does not discuss cases in which the reverse is true: one outcome is significantly better collectively than the other is individually. However, as we have seen, when speaking about the Wide Dual Principle in general terms Parfit claims that ‘slightly greater benefits to each would be outweighed by much greater benefits to all’ just like ‘slightly greater benefits to all would be outweighed by much greater benefits to each’ (2017, p. 155). Considering this, it is plausible to assume, at least initially, that Parfit would also endorse another sufficient condition for overall betterness that mirrors (1):

- (2) For any two outcomes X and Y, such that $I(X) > I(Y)$ and $C(X) < C(Y)$,
if $(|C(Y)| / C(X)) \geq 10 \times (|I(X)| / I(Y))$ *then* Y is overall better than X.

That is, one of two outcomes is overall better *if* collective goodness favours it 10 times more than individual goodness favours the other outcome.

Now, consider a variant of Case 6:

Case 8: **A** 1 million people at level 1,000
Z+ 10 trillion people at level 1

This case is like Case 6, except that the number of people in Z+ is 100 times greater than it was in Z. In Case 8, $C(A) = 1$ billion, $C(Z+) = 10$ trillion, $I(A) = 1,000$, and $I(Z+) = 1$. So, Z+ is collectively better than A, and A is individually better than Z+. How much better? $|C(Z+)| / C(A) = 10,000$ and $|I(A)| / I(Z+) = 1,000$. Now, if (2) is true, then Z+ is overall better than A. But this is just an instance of the Repugnant Conclusion. So, contrary to what Parfit claims, the Wide Dual Principle implies the Repugnant Conclusion.

Did Parfit see it coming? It is unclear. In the last paragraph of his article, Parfit writes:

‘These remarks do not show that, if we can justifiably appeal to this Wide [Dual] Principle, we can thereby avoid the Repugnant Conclusion. If we believe that it would always be in one way better if there existed more people who would together have a greater total sum of benefits, this belief might be one premise of more complicated and forceful arguments for the Repugnant Conclusion.’ (2017, p. 157).

It is important to note that the Wide Dual Principle is not just one of the premises in a complex argument for the Repugnant Conclusion. Rather, the preceding considerations show that the Wide Dual Principle *alone* implies the Repugnant Conclusion.

IV.3. Possible amendments

Can the Wide Dual Principle be saved on Parfit's behalf? There are at least three amendments worth considering.

The first possible amendment turns on the idea that individual goodness is so much more important than collective goodness that the greater collective goodness of one outcome can *never* outweigh the greater individual goodness of another outcome. In particular:

- (3) For any two outcomes X and Y, such that $I(X) > I(Y)$ and $C(X) < C(Y)$, X is overall better than Y.

This claim, (3), avoids implying the Repugnant Conclusion but is otherwise dubious because it implies that in the following case A is overall better than B.

- Case 9: **A** 1 person at level 80.
B 1 million people at level 79.

If (3) is true, then, in equality cases, the role of collective goodness is reduced to breaking ties.¹¹

But it is difficult to make sense of this role in the context of the theoretical underpinnings of the Wide Dual Principle. This view turns on the thought that a person who is created with a good life (and we can assume that level 79 corresponds to a very good life) is genuinely benefited and that adding people with good lives to an outcome genuinely makes that outcome better. Yet, (3) implies that there is no number of people at 79 whose existence would be better than the existence of a single person at level 80, and hence a cap on how good an outcome containing people at 79 can be.

Moreover, (3) has even more implausible implications regarding comparisons of outcomes in which people have bad lives. In the following case, it implies that B- is better than A-.

- Case 10: **A-** 1 person at level -80.
B- 1 million people at level -79.

Finally, (3) is clearly not a good interpretation of Parfit's view who, as we have seen, writes that 'slightly greater benefits to each would be outweighed by much greater benefits to all' (2017, p. 155).

The second possible amendment makes room for the claim that individual goodness is more important than collective goodness in another way. The thought is that although greater collective goodness of one outcome can *sometimes* outweigh greater individual goodness of another outcome, it is *easier*

¹¹ This role would be perhaps more extensive in inequality cases. Section 7 discusses the view that when two outcomes are incomparable in terms of individual goodness, then the overall betterness relation is determined on the basis of collective goodness alone.

for a difference in individual goodness to outweigh a difference in collective goodness than conversely. In particular, the suggestion is to replace (2) with:

- (4) For any two outcomes X and Y, such that $I(X) > I(Y)$ and $C(X) < C(Y)$,
 if $(|C(Y)| / C(X)) \geq 100 \times (|I(X)| / I(Y))$ then Y is overall better than X.

This kind of amendment obviously cannot help avoid the Repugnant Conclusion, whatever value of n is chosen instead of '10'. If (4) is accepted instead of (2), then simply a greater number of people at level 1 will be required to make an outcome in which only such people exist better than another outcome in which everyone exists at level 1,000.

The third possible amendment explores a somewhat different avenue. Recall that Parfit suggests that 'slightly greater benefits to each would be outweighed by much greater benefits to all' (2017, p. 155). In Case 8, benefits to each in A are not slightly greater; they are 1,000 times greater! Perhaps such a great difference in individual goodness cannot be outweighed by *any* difference in collective goodness. Indeed, Parfit seems to be considering something like this in the penultimate sentence of his article:

'We might justifiably believe that great losses in the quality [goodness] of people's lives could not be outweighed by any increase in the sum of benefits, if these benefits came in the lives of people whose quality of life would be much lower' (2017, p. 157).

Perhaps one outcome can be at most 10 (or 5) times individually worse than another, for it to be possible that the former's greater collective goodness can outweigh the latter's greater individual goodness. Call this proposal *Local Lexical Superiority*. When treated as a constraint on something like (2), it gives us:

- (5) For any two outcomes X and Y, such that $I(X) > I(Y)$ and $C(X) < C(Y)$,
 if $(|C(Y)| / C(X)) \geq 10 \times (|I(X)| / I(Y))$
 and $(|I(X)| / I(Y)) \leq 10$
 then Y is overall better than X; otherwise X is overall better than Y.

On this view, an outcome with greater collective goodness is overall better than another outcome with greater individual goodness only if either people in both outcomes have very good lives or else people in both outcomes have barely good lives. So, the Repugnant Conclusion is false.

Despite a promising start, this amendment does not succeed. Earlier in the same article, Parfit argues at great length (against Larry Temkin) that overall goodness of outcomes must be intrinsic and that 'intrinsically better than' is a transitive relation. But one cannot accept (5) in addition to these two claims. To see that, consider:

- Case 11: **A** 1,000 people at level 100
B 1 million people at level 10
C 1 billion people at level 1

What does the Wide Dual Principle imply in this case? Compare A and B. $|C(B)|/C(A) = 100$ and $|I(A)|/I(B) = 10$, so (5) implies that B is overall better than A. Now, compare B and C. $|C(C)|/C(B) = 100$ and $|I(B)|/I(C) = 10$, so (5) implies that C is overall better than B. Now, if ‘overall better’ is a transitive relation, then C is overall better than A. However, $|I(A)|/I(C) = 100$, so (5) implies that A is overall better than C.

This is a contradiction. Proponents of the Wide Dual Principle face a choice. One of the following claims must give way: (i) goodness of outcomes is intrinsic; (ii) ‘intrinsically better than’ is a transitive relation; (iii) greater collective goodness of one outcome can sometimes outweigh greater individual goodness of another outcome; or (iv) Local Lexical Superiority. Parfit is clearly committed to the first three and his arguments for them are compelling, so Local Lexical Superiority must give way.

There is a further worry about Local Lexical Superiority. It must be asymmetric in the sense that one could not hold both that, for any two outcomes X and Y:

- (a) X can be at most n times *individually* worse than Y, for it to be possible that X’s greater collective goodness can outweigh Y’s greater individual goodness; otherwise the individually better outcome is overall better.
- (b) Y can be at most n times *collectively* worse than X, for it to be possible that Y’s greater individual goodness can outweigh X’s greater collective goodness; otherwise the collectively better outcome is overall better.

Local Lexical Superiority, as used earlier in this section, refers just to (a). To see why it is impossible to hold both (a) and (b), consider:

Case 12: **A** 10 people at level 100
B 100,000 people at level 1

In Case 12, $|C(B)|/C(A) = 100$ and $|I(A)|/I(B) = 100$. If (say) $n = 10$, then (a) implies that A is overall better than B, whereas (b) implies that B is overall better than A. Assuming that ‘better than’ is asymmetric, this is a contradiction.

The contradiction can be avoided by rejecting (a) or (b). Since (b) outright implies the Repugnant Conclusion, (b) is a more natural candidate for the axe. However, since (a) and (b) are symmetric, and individual and collective goodness are both supposed to be genuinely good-making considerations, it becomes difficult to see what considerations could justify one-sided Local Lexical Superiority. We are back to square one: if anything, we should embrace (2).

IV.4. Should we despair?

The Wide Dual Principle cannot avoid implying the Repugnant Conclusion. Should we despair? Not necessarily.

Let's begin by noting that the Wide Dual Principle has some advantages over the Wide Principle which Parfit discussed in his previous work. Firstly, these two principles differ in terms of the requirements for overall betterness in Type 5 cases. Consider some outcome X in which 1 million people at level 1,000 exists (like A in Case 6). What is the minimum number of people at level 1 that must exist in another outcome Y for Y to be overall better than X? According to the Wide Principle, the minimum number is 1 billion. According to the Wide Dual Principle, assuming (2), the minimum number is 1 trillion. This is a major difference. It is partly for this reason that the instance of the Repugnant Conclusion implied by the Wide Dual Principle is *less repugnant* and perhaps easier to accept than the instance of the Repugnant Conclusion implied by the Wide Principle (and its impersonal equivalent: the Total Principle).¹²

Secondly, these two principles differ in terms of how the Repugnant Conclusion is reached. Many people find the Wide Principle (and the Total Principle) unattractive because it is not explicitly concerned with the goodness of life of each person who exists in a given outcome. For instance, the Wide Principle is incompatible with the spirit of Larry Temkin's Consolidate Additional Benefits View mentioned earlier because it regards as equally good an outcome in which fewer people with very good lives exist and another outcome in which many people with barely good lives exist, as long as the collective goodness of the two outcomes is the same. It is partly in this sense that the Wide Principle regards people merely as containers of well-being.¹³ By contrast, the Wide Dual Principle is concerned not only with collective goodness but also with individual goodness: with how good each life is. At the same time, the Wide Dual Principle recognises that, whenever one good-making consideration (collective goodness) significantly outweighs the other good-making consideration (individual goodness), a concession has to be made and the Repugnant Conclusion accepted.

V. THE RELATIVE IMPORTANCE OF INDIVIDUAL AND COLLECTIVE GOODNESS

As shown in Section 4.1, Parfit is implicitly committed to at least one sufficient condition for overall betterness, (1), in Type 5 cases in which the two components of the Wide Dual Principle conflict. This sufficient condition fails to

¹² The thought that there are more or less repugnant instances of the Repugnant Conclusion is not entirely new. For instance, Nils Holtug (2010, p. 254) argues that prioritarianism in variable-population contexts implies what he calls *the Super-Repugnant Conclusion*: for any outcome in which many people exist, all with very good lives, there is a better outcome in which many more different people exist, all with barely good lives, and in which the total sum of life goodness is lower.

¹³ Parfit illustrates this objection memorably: 'The greatest mass of milk might be found in a heap of bottles each containing only a single drop' (1984, p. 388).

determine which of the two outcomes is overall better in the following case mentioned by Parfit (2017, p. 155):

- Case 13: **A** N people at level 100.
B 2 N people at level 75.

(2) does not determine it either. In Case 13, $|C(B)|/C(A) = 1.5$ and $|I(A)|/I(B) \approx 1.33$, so one ratio is not 10 times greater than the other, and neither component is much more prominent than the other.

As Parfit aptly observes, to determine which of these two outcomes is better, ‘we must compare the relative importance of these two ways of benefiting people more’ (2017, p. 155). He adds: ‘I can only start to do that here’. Ultimately, however, Parfit does not offer a way of extending the Wide Dual Principle to such comparisons.

There are at least two moves worth discussing on Parfit’s behalf. The first one is to accept the *Incomparability View* according to which outcomes like A and B in Case 13 are incomparable in terms of overall goodness. More generally, the claim is that in Type 5 cases, whenever neither outcome is at least 10 times collectively better than the other is individually *and* neither outcome is at least 10 times individually better than the other is collectively, the two outcomes are overall incomparable.

The Incomparability View has several disadvantages. Firstly, it will make the Wide Dual Principle regard very many pairs of outcomes as overall incomparable. Very few, if any, real-life cases are like Case 6 (Parfit’s Repugnant Conclusion case). More are like Case 13 in which the differences between outcomes in terms of individual and collective goodness are not as dramatic. Secondly, we may well wonder what could possibly make it the case that the irreducible values of individual and collective goodness can be traded-off in some cases but cannot be traded-off in other cases. One would also expect a deeper metaphysical explanation of the particular claim that one ratio must be at least 10 times greater than another, which seems to be the cut-off point between comparable and incomparable outcomes on this view.¹⁴

The following *Simple View* is superior. It holds that, in Type 5 cases, overall betterness is determined by simply comparing the prominence of the two components. If one outcome is individually better to a greater extent than the other is collectively better, then the former outcome is overall better. And, if one outcome is collectively better to a greater extent than the other is individually better, then the former outcome is overall better. Finally, if one outcome is collectively better to the same extent that the other is individually better, then they are overall equally good. More precisely:

¹⁴ This brings us back to the Deep Problem of Pluralism mentioned in footnote 6.

- (6) For any two outcomes, X and Y, such that $I(X) > I(Y)$ and $C(X) < C(Y)$,
if $|I(X)|/I(Y) = |C(Y)|/C(X)$ *then* X and Y are overall equally good;
if $|I(X)|/I(Y) > |C(Y)|/C(X)$ *then* X is overall better than Y;
if $|I(X)|/I(Y) < |C(Y)|/C(X)$ *then* Y is overall better than X.

Consider the application of this Simple View to Case 13. There, $|C(B)|/C(A) = 1.5$ and $|I(A)|/I(B) \approx 1.33$. Collective goodness favours outcome B to a greater extent than individual goodness favours outcome A, so B is overall better than A.

To further illustrate this view, we can compare Case 13 with

Case 14: **A** N people at level 100
B- 2N people at level 70

The difference is that in B- people exist at 70, not at 75. This is not a major difference, but it is enough to change the verdict about overall betterness. In Case 14, $|C(B-)|/C(A) = 1.4$ and $|I(A)|/I(B) \approx 1.43$. So, A is overall better than B-.

The Simple View has a minor disadvantage relative to the Incomparability View. Although the Wide Dual Principle implies the Repugnant Conclusion when paired with either view, the Wide Dual Principle paired with the Incomparability View implies a slightly less repugnant instance of this conclusion. On the Simple View, a smaller number of people with barely good lives are required to make an outcome overall better than an outcome in which fewer people with very good lives exist. That said, that number of people with barely good lives would still be greater than what is required by the Wide Principle.

On the other hand, the Simple View has a number of theoretical virtues. Firstly, as we have seen, it avoids the worry of pervasive incomparability in equality cases. And it does so in a very intuitive way. The Wide Dual Principle paired with the Simple Proposal is premised on the idea that individual and collective goodness are equally important. Whenever these considerations pull us in the opposite directions, we compare how strongly they pull us in these directions.

Secondly, the Simple View has 'less' metaphysical explaining to do. Although one still needs to explain what makes it the case that irreducible values can be traded-off at all and in this particular way, there is no need to explain why in some cases they can be traded-off and in others, they cannot. And, as long as we accept that individual and collective goodness can always be traded-off by way of comparing ratios, not much hinges on the assumption that these two values are equally important. For instance, we would not avoid the Repugnant Conclusion by giving greater weight to individual goodness. Assuming that they are equally important, however, will simplify some of the following discussion and has some intuitive pull to it, so we might as well do that now.

We can call the conjunction of the Wide Dual Principle with the Simple View *the Revised Wide Dual Principle*. This principle is ‘revised’ because accepting the Simple View is a possible departure from Parfit’s view and in particular from Parfit’s Rule. Although Parfit’s Rule is a pair of merely sufficient conditions for overall betterness, the claim that one good-making consideration needs to be *much* more prominent than the other in order to outweigh it seems to occupy a central role in his thinking about the Wide Dual Principle. The Simple View rids us of the ‘much’ qualification.

Let’s summarise the implications of the Revised Wide Dual Principle in equality cases:

- Type 1: *if* $C(X) = C(Y)$ and $I(X) = I(Y)$, *then* X is overall as good as Y;
- Type 2: *if* $C(X) > C(Y)$ and $I(X) > I(Y)$, *then* X is overall better than Y;
- Type 3: *if* $C(X) = C(Y)$ and $I(X) > I(Y)$, *then* X is overall better than Y;
- Type 4: *if* $C(X) > C(Y)$ and $I(X) = I(Y)$, *then* X overall better than Y;
- Type 5: *if* $C(X) > C(Y)$ and $I(X) < I(Y)$, *and* . . .
- . . . $|C(X)|/C(Y) > |I(Y)|/I(X)$ *then* X is overall better than Y.
- . . . $|C(X)|/C(Y) < |I(Y)|/I(X)$ *then* X is overall worse than Y.
- . . . $|C(X)|/C(Y) = |I(Y)|/I(X)$ *then* X is overall as good as Y.

Despite its complex appearance, in equality cases, the Revised Wide Dual Principle is extensionally equivalent to:

X is overall at least as good as Y *iff* $|I(X)| \times C(X) \geq |I(Y)| \times C(Y)$.¹⁵

It is worth emphasising that, although these principles are extensionally equivalent, the latter is not *how* the Revised Wide Dual Principle arrives at the aforementioned implications. Simply multiplying individual and collective goodness does not seem to have any normative significance whereas the previously discussed comparisons in terms of individual and collective goodness do.

VI. UNEQUAL OUTCOMES AND INDIVIDUAL BETTERNESS

VI.1. *The problem*

The preceding discussion of the Wide Dual Principle was limited to equality cases in which all people in any particular outcome have equally good lives. How can we apply this principle to inequality cases in which comparisons involve at least one unequal outcome? Parfit does not provide an answer. He merely notes that ‘inequalities raise complications that are best considered elsewhere’ (2017, p. 152). These complications do not arise at the level of

¹⁵ Once again, taking the modulus (absolute value) of one of the values is necessary to correctly account for cases in which we compare outcomes in which people have bad lives.

collective betterness because comparing outcomes in this respect is just a matter of comparing the sum totals of life goodness. Rather, they arise at the level of individual betterness.

VI.2. *The Simple Definition*

Parfit's simple definition of individual betterness can be paraphrased as follows:

The Simple Definition: For any two outcomes, X and Y, X is individually better than Y just in case all people in X have better lives than all people in Y; X and Y are individually equally good just in case all people in these outcomes have equally good lives; otherwise X and Y are individually incomparable.¹⁶

This definition fared well in equality cases, but it runs into trouble in inequality cases. Firstly, very many pairs of unequal outcomes will be deemed incomparable in terms of individual goodness. Secondly, often only a minute change to a life goodness level of one of many people will be sufficient to force a shift from comparability to incomparability. To see that, consider:

- Case 15: **A** 10 people at level 100
B 1 person at level 101 and 9 people at level 100

Thirdly, on the initially plausible assumption that when two outcomes are individually incomparable they are also overall incomparable, the Simple Definition implies that A and B are incomparable in terms of overall goodness. But, intuitively, they are not: B is overall better than A. Finally, on the same assumption, and if the same people exist in A and B, then the Simple Definition violates a highly plausible version of the Weak Pareto Principle which holds that if one of two outcomes is better for someone and worse for no one, and the same people exist in both outcomes, then it is overall better. These considerations suggest that the Simple Definition should be rejected in favour of some more relaxed definition of individual betterness.

VI.3. *Two dominance definitions*

In Case 15, some people in B have better lives than all people in A and no people in B have worse lives than any people in A. Plausibly, B is individually better than A. This reasoning applies not only to same-number but also to different-number cases. It underlies the following candidate definition of individual betterness:

The Simple Dominance Definition: For any two outcomes, X and Y, X is individually better than Y just in case some people in X have better lives than all people in Y and no people

¹⁶ Compare Parfit (2017, pp. 153-5).

in X have worse lives than any people in Y; X and Y are individually equally good just in case all people in these outcomes have equally good lives; otherwise X and Y are individually incomparable.

This definition would be an improvement on the Simple Definition, but it also faces several problems. Firstly, consider:

Case 16: **A** 5 people at level 100 and 5 people at level 50
B 5 people at level 100 and 5 people at level 50

The Simple Dominance Definition implies that A and B in Case 16 are individually incomparable. But, intuitively, A and B are individually equally good even if different people exist in these two outcomes. The individual betterness relation should be invariant under permutations of identities of possible members of the domain.

Secondly, consider:

Case 17: **A** 5 people at level 100 and 5 people at level 50
B+ 5 people at level 100 and 5 people at level 60

The Simple Dominance Definition implies that A and B+ in Case 17 are individually incomparable. Intuitively, however, B+ is individually better than A. If the 5 people at level 100 were the same people, and the remaining 5 people were the same people, B+ would be a Pareto Improvement on A: it would be better for some and worse for no one. Plausibly, B+ is better than A even if different people exist in these outcomes. To capture this verdict, we need a concept akin to that of Pareto Improvement which is insensitive to people's identities. We can call it

Permutation-invariant Pareto Improvement: For any two outcomes which contain the same number of people, X and Y, X is a Permutation-Invariant Pareto Improvement on Y just in case there exists a bijection between X and Y such that every person in X is at least as well off as their counterpart in Y, and there are some people in X who have better lives than their counterpart in Y.

Note that all Pareto Improvements are Permutation-invariant Pareto Improvements but the converse is not true.

These two considerations suggest another definition of individual betterness which we might call

The Sophisticated Dominance Definition: For any two outcomes, X and Y, X is individually better than Y just in case (i) some people in X have better lives than all people in Y and no people in X have worse lives than any people in Y or (ii) X is a Permutation-Invariant Pareto Improvement on Y. X and Y are individually equally good just in case (iii) all people in in these outcomes have equally good lives or (iv) X and Y are permutations of each other. Otherwise X and Y are individually incomparable.

This Sophisticated Dominance Definition of individual betterness accounts for the intuitive verdicts about all cases considered earlier in this section as well as those about equality cases considered earlier. Conditions (ii) and (iv) cover only same-number cases, whereas (i) and (iii) also extend to different-number cases. (i) and (ii) do not entail one another; neither do (iii) and (iv).

However, despite being an improvement on the Simple Definition and the Simple Dominance Definition, even this Sophisticated Dominance Definition faces the worry of pervasive incomparability. Consider:

- Case 18:* **A** 10 people at level 100
B 1 person at level 101 and 9 people at 1

The Sophisticated Dominance Definition implies that A and B in Case 18 are incomparable in terms of individual goodness, whereas many people have the intuition that A is individually better than B. Perhaps we can do better.

There is a further worry. There is no apparent way of measuring *how good* individually an outcome is and *how much better* individually one outcome is than another.¹⁷ This is important because in Type 5 cases considerations of individual goodness and collective goodness pull us in the opposite directions. To determine which of the two outcomes is overall better, we must compare how much better one outcome is individually relative to how much better the other outcome is collectively. If we cannot find a way of quantifying the difference in individual goodness in inequality cases, we may be unable to determine the overall betterness relations between many such outcomes.

VI.4. *The Average Definition*

One definition, which avoids the problem with quantifying individual goodness and which initially seems to avoid the worry of incomparability is

The Average Definition: For any two outcomes, X and Y, X is individually better than Y iff the average life goodness level in X is higher than in Y; X and Y are individually equally good just in case the average life goodness in these outcomes is the same; otherwise X and Y are individually incomparable.

The level of individual goodness of an outcome is according to this definition simply the average life goodness level. This definition accounts for the intuitive verdicts about the cases considered earlier in this section but it faces other significant worries.

Firstly, the Revised Wide Dual Principle paired with this definition violates

¹⁷ The same worry arises in inequality cases for the Simple Definition and the Simple Dominance Definition.

The Mere Addition Principle: Adding people with good lives to an outcome does not make that outcome overall worse, other things equal (Parfit 1984, p. 420).

To see that, consider:

- Case 19: **A** 10 people at level 100
B 10 people at level 100 and 10 people at level 10

Suppose that the people at level 100 in these outcomes are the same people. Now, $|C(B)|/C(A) = 1.1$ and $|I(A)|/I(B) \approx 1.82$. Since $1.82 > 1.1$, A is overall better than B. This is a surprising result because B is just like A except that some further people with good lives exists in B.

Secondly, and much more worryingly, the Revised Wide Dual Principle paired with this definition implies

The Sadistic Conclusion: It would sometimes be overall better to add people with bad lives rather than people with good lives to an outcome, other things equal (Arrhenius 2000, p. 251)

To see that, consider the following alternatives:

- Case 20: **A** 10 people at level 100
B 10 people at level 100 and 10 people at level 10
C 10 people at level 100 and 1 person at level -100

Again, suppose that the people at level 100 in these outcomes are the same people. Now, $|C(B)|/C(C) \approx 1.22$ and $|I(C)|/I(B) \approx 1.49$. Since $1.49 > 1.22$, C is overall better than B. This is highly implausible because both B and C are outcomes which result from adding extra people to A. In C we are adding a person with a very bad life whereas in B we are adding some people with good lives.

These two problems for the Revised Wide Dual Principle paired with the Average Definition arise *only* in inequality cases. However, this provides little consolation. Most cases we will ever encounter in the real world are inequality cases. One of the key aims of investigating population axiology is that the results of doing so might at some point inform our decisions about how many people will ever exist, and at what life goodness levels.

VII. THE SILENT COMPONENT APPROACH

The Simple Definition, the Simple Dominance Definition, and the Sophisticated Dominance Definition all face the worry of pervasive incomparability in terms of individual goodness.

How significant is this worry? A natural thought is that when two outcomes are incomparable in terms of one good-making consideration they are also

incomparable in terms of overall goodness. We can call this *the Contaminating Component Approach* because incomparability at the level of one component 'contaminates' the overall betterness relation. If this approach is correct, then accepting any of the aforementioned definitions of individual betterness comes at the price of pervasive incomparability at the level of overall betterness.

Note, however, that outcomes are only incomparable in terms of individual goodness. They are never incomparable in terms of collective goodness because all we have to do is to calculate the total sum of life goodness for each outcome. And we are ultimately interested in comparisons in terms of overall goodness.

These observations provide a basis for the following suggestion: whenever two outcomes are incomparable according to one criterion, then this criterion is 'silent' and the other criterion determines the overall betterness relation. Call this *the Silent Component Approach*.

Why accept this approach? Firstly, it allows one to adopt the Sophisticated Dominance Definition of individual betterness without worrying about incomparability at the level of overall goodness. The resulting view accommodates our intuitions about cases considered throughout this essay, does not violate the aforementioned version of the Weak Pareto Principle or the Mere Addition Principle, and avoids the Sadistic Conclusion.

Secondly, an attentive reader might have noticed that the worry of incomparability was looming even before we turned to comparisons involving unequal outcomes. To see that, consider:

- Case 21:* **A** No one exists
B 1 person at level 100

Recall, from Section 2.2, that Parfit claims that we can compare an outcome in which no one exists with an outcome in which someone with a good life exists. The latter outcome is supposed to be better because it is good for someone whereas the former outcome is good for no one.

But how are we supposed to capture this thought with the Wide Dual Principle or its Revised version? Comparing collective goodness is straightforward: the total sum of life goodness in an empty outcome is 0, so **B** is collectively better than **A**. But what about individual goodness? Is it also 0 or is it undefined? Though the former answer would simplify applying the Wide Dual Principle to such cases, the latter answer is more plausible. This is because individual goodness is supposed to track how much each person is benefited or how good each life is. Sensitivity to this issue is what allows the proponent of the Wide Dual Principle to avoid the objection that persons are merely containers of well-being. But there is no answer to the question 'how good are the lives of people?' in an empty world. Individual goodness of an empty outcome is undefined and hence empty and non-empty outcomes are individually incomparable. If the Contaminating Component Approach were true,

empty and non-empty outcomes would be incomparable in terms of overall goodness. Thus, all of the previously considered definitions of individual betterness, including the Average Definition, need to be paired with the Silent Component Approach to deliver the verdict that B is overall better than A in Case 21.

That said, the Silent Component Approach is not without challenges. Neither is decisive but it will be instructive to discuss them anyway.

Firstly, the Silent Component Approach has its limitations. In particular, this approach paired with a principle that recognises at least three independent good-making considerations that admit of incomparability gives rise to cycles at the level of overall betterness.

To see that, consider the following case. Let C_1 , C_2 , and C_3 be three components used to determine the overall goodness of some outcome. The following table displays how each of the three outcomes (X, Y, and Z) ‘scores’ in these respects. ‘—’ denotes an undefined value.

<i>Case 22:</i>	C₁	C₂	C₃
	X	2	1
	Y	—	2
	Z	1	—

How do these outcomes compare? Firstly, X and Y are incomparable in terms of C_1 and C_3 , and Y is better than X in terms of C_2 . If the Silent Component is true, then *Y is overall better than X*. Second, Y and Z are incomparable in terms of C_1 and C_2 , and Z is better than Y in terms of C_3 . So, *Z is overall better than Y*. Third, X and Z are incomparable in terms of C_2 and C_3 , and X is better than Z in terms of C_1 . So, *X is overall better than Z*. Thus, there is a cycle at the level of overall betterness. If ‘better than’ is transitive and asymmetric, then this is a contradiction.

It is important to stress that such cycles arise only if the overall betterness relation is determined on the basis of at least three components that admit of incomparability. The Revised Wide Dual Principle has only two components, only one of which admits of incomparability, and therefore avoids this worry. Perhaps this is serendipitous. But that the Silent Component Approach has troubling implications when paired with some other view is not enough to force us to reject this approach as the supplement to the Revised Wide Dual Principle.

Secondly, one might worry that, when paired with the Silent Component Approach, the Revised Wide Dual Principle loses its advantages over the Wide Principle (and the Total Principle). After all, whenever two outcomes are incomparable in terms of individual goodness, the Revised Wide Dual Principle issues the same verdicts as the Wide Principle. This is perhaps most troubling in cases which are akin to the Repugnant Conclusion ones, such as

Case 23: **A-** 1 million people at level 1,000 and 1 person at level 1
Z+ 100 billion people at level 2

Admittedly, the resulting view loses one of its advantages advertised earlier. The instances of what we might call *the Inequality-involving Repugnant Conclusion* implied by the Revised Wide Dual Principle paired with the Silent Component Approach are not 'less repugnant' than the instances implied by the Wide Principle. The total sum of life goodness needed to make an outcome in which only people at level 2 exist better than A- is the same for both views.

That said, the Revised Wide Dual Principle paired with the Silent Component Approach retains the other, more important advantage. Unlike the Wide Principle, this view remains explicitly concerned with the goodness of life of each person in a given outcome. It is only when two outcomes are incomparable in terms of individual goodness that sensitivity to this matter does not translate into reasons to favour one outcome over the other.

VIII. CONCLUDING REMARKS

The Wide Dual Person-Affecting Principle is an ingenious proposal. It is unfortunate that Parfit's project of articulating this principle was terminated by his untimely death. This paper may be seen as an attempt to complete this project. It identified various theoretical choices that need to be made by a proponent of this principle, and discussed the relative merits of each alternative. Let us summarise some of these results first and then see where they leave us.

Section 4 argued that the Wide Dual Principle cannot avoid the Repugnant Conclusion unless one of the following claim is rejected: (i) goodness of outcomes is intrinsic; (ii) 'intrinsically better than' is transitive; or (iii) the greater collective goodness of one outcome can sometimes outweigh the greater individual goodness of another outcome. Parfit is clearly committed to and offers compelling arguments for the first two claims in 'Future People, ...' and his earlier work. The third claim seems to be at the core of Wide Dual Principle and rejecting it has highly implausible consequences.

This section further argued that the instances of the Repugnant Conclusion implied by this principle are 'less repugnant' than the instances implied by the Wide Principle (and the Total Principle) and that the Wide Dual Principle does not regard people merely as containers of well-being. Thus, accepting the Repugnant Conclusion may be more justifiable than it is commonly advertised.

Section 5 discussed the relative importance of individual and collective goodness. The proponents of the Wide Dual Principle owe us an explanation of what, if anything, makes it the case that the irreducible values of individual

and collective goodness can be traded-off at all. Any genuinely pluralistic axiology faces this problem, so perhaps it can be temporarily set aside. On the assumption that these two good-making considerations can sometimes be traded-off, it is important to determine whether they can always be traded-off. According to the Incomparability View, the answer is *no*: whenever neither outcome is much better individually nor collectively than the other, the two outcomes are overall incomparable. According to the Simple View, the answer is *yes*: whenever one outcome is individually better to a greater extent than the other is collectively better, the former is overall better (and *vice versa*). The latter view is more extensionally adequate and has 'less' metaphysical explaining to do. The Simple View does not presuppose that individual and collective goodness are equally important, but it is helpful to assume, at least initially, that they are. Because accepting the Simple View could be a departure from Parfit's outlook, it was dubbed the Revised Wide Dual Principle.

Section 6 discussed comparisons involving unequal outcomes. In such cases, the Simple Definition of individual betterness delivers a number of counterintuitive verdicts, regards very many pairs of outcomes as individually incomparable, and there seems to be no way of measuring how much better individually one outcome is than another. The Sophisticated Dominance Definition is an improvement in the first respect and partially in the second respect, but it does not avoid the third worry. The Average Definition avoids the latter two problems but comes at the price of rejecting the Mere Addition Principle and accepting the Sadistic Conclusion.

Section 7 considered whether incomparability at the level of individual goodness implies incomparability at the level of overall goodness. According to the Contaminating Component Approach, it does. According to the Silent Component Approach, it does not: whenever two outcomes are incomparable in terms of individual goodness, the comparison in terms of collective goodness alone determines which of two outcomes is overall better. Taking this latter approach allows one to accept something like the Sophisticated Dominance Definition and avoid pervasive incomparability at the level of overall goodness.

Where does this leave us? It is unclear. The combination of the Revised Wide Dual Principle with the Sophisticated Dominance Definition and the Silent Component Approach might be the best version of the Wide Dual Principle. But it remains to be seen whether our best metaphysics of value makes room for this view. And there is the Repugnant Conclusion, too. A large part of Parfit's intellectual life was spent trying to find ways to avoid it. The Wide Dual Principle is an addition to the long list of theories that share this implication. Notably, this implication is not avoided despite treating individual goodness as one of two values of the highest importance. Considering this, perhaps we should embrace the Repugnant Conclusion after all. Principles and intuitions that underlie arguments in favour of this conclusion are part of a complex moral landscape. The intuition that the Repugnant Conclusion

is false, however strong it may be, is only a solitary rock in the middle of an ocean, repeatedly crushed by waves.

REFERENCES

- Arrhenius, G. (2000) 'An Impossibility Theorem for Welfarist Axiologies', *Economics and Philosophy*, 16: 247–66.
- Chang, R. (2002) 'The Possibility of Parity', *Ethics*, 112: 659–88.
- (2016) 'Parity, Imprecise Comparability and the Repugnant Conclusion', *Theoria*, 82: 182–214.
- Holtug, N. (2010) *Persons, Interests, and Justice*. Oxford University Press. doi:10.1093/acprof:oso/9780199580170.001.0001.
- McMahan, J. (1981) 'Problems of Population Theory', *Ethics*, 92: 96–127.
- (2009) 'Asymmetries in the Morality of Causing People to Exist', in M. A. Roberts and D. Wasserman (eds), *Harming Future Persons: Ethics, Genetics, and the Nonidentity Problem*, 49–68. London, New York: Springer.
- (2013) 'Causing People to Exist and Saving People's Lives', *Journal of Ethics*, 17: 5–35.
- Parfit, D. (1984) *Reasons and Persons*. Oxford University Press. doi:10.1093/019824908X.001.0001.
- (2016) 'Can We Avoid the Repugnant Conclusion?', *Theoria*, 82: 110–27.
- (2017) 'Future People, the Non-identity Problem, and Person-Affecting Principles', *Philosophy & Public Affairs*, 45: 118–57.
- Temkin, L. (2012) *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press. doi:10.1093/acprof:oso/9780199759446.001.0001.

Princeton University, US