

Imprecise Lexical Superiority and the (Slightly Less) Repugnant Conclusion

James Fanciullo

Forthcoming in *Philosophical Studies*

Abstract: Recently, Derek Parfit has offered a novel solution to the “Repugnant Conclusion” that compared with the existence of many people whose quality of life would be very high, there is some much larger number of people whose existence would be better but whose lives would be barely worth living. On this solution, qualitative differences between two populations will often entail that the populations are merely “imprecisely” comparable. According to Parfit, this fact allows us to avoid the Repugnant Conclusion without violating the transitivity of *better than*. In this paper, I argue that Parfit’s view nevertheless implies two objectionable conclusions. The first is an alternative version of the Repugnant Conclusion that, Parfit suggests, may not be all that repugnant. The second is a revised version of the first that is nearly identical to the Repugnant Conclusion. I conclude that Parfit’s view offers no escape from repugnance.

Many of us think that the existence of any person whose life is worth living adds some positive value to the world. More precisely, many of us accept what Derek Parfit calls

The Simple View: Anyone’s existence is in itself good, and makes the world in one way better, if this person’s life is good to live, or is worth living. Such goodness has non-diminishing value, so if there were more such people, the combined goodness of their existence would have no upper limit.¹

On this view, if a person’s life is more worth living, or has a higher level of welfare, this person’s life contributes a greater amount of value to the world. Still, given that the addition of any person whose

¹ This combines two formulations of the view in Parfit 2016 (pp. 110, 112).

life is worth living would contribute at least some positive value to the world, it seems that a large enough number of lives that were just barely worth living would contribute a greater total amount of value to the world than any given number of lives that were very high in quality. In which case, the Simple View seems to imply

The Repugnant Conclusion: Compared with the existence of many people who would all have some very high quality of life, there is some much larger number of people whose existence would be better, even though these people would all have lives that were barely worth living. (Parfit 2016, p. 110)²

To see this, consider Fig. 1. Each block's width represents the number of lives in a population; its height the quality of each life. Though the quality of each life in B would be slightly lower than that of each life in A, in B many more people would exist. Accordingly, B would contain a much greater total amount of value. It thus seems plausible that, since B would contain a much greater total amount of value, and lives of only a slightly lower quality, B would be better than A. If we apply similar reasoning to populations B and C, and populations C and D, and so on, we reach the claim that population Z, which would contain a great number of people whose lives would be just barely worth living, would be better than Y. And since *better than* is a *transitive* relation—that is, if x is better than y , and y is better than z , x must be better than z —we reach the conclusion that Z would be better than every population preceding it—including A.³ Thus we reach the Repugnant Conclusion.

² This is introduced in Parfit 1984 (p. 388), though its formulation there is implied by the “total” Utilitarian principle according to which the best state of affairs (*ceteris paribus*) is the one containing the greatest total amount of utility. Since this conclusion also seems to be implied by the Simple View, we cannot avoid it simply by rejecting the total principle. So it presents a worry for, not just Utilitarians, but anyone that accepts the more plausible Simple View.

³ Of course, some people reject the transitivity of *better than*, and so avoid this conclusion. See, for example, Rachels 1998, and Temkin 1987 and 1996. Given these people reject such an intuitive and widely accepted relation, though, it is perhaps unsurprising that they face more pressing worries than the Repugnant Conclusion. See, for example, Nebel 2017. See also Jensen 2008 and Voorhoeve 2013.




Figure 1

Recently, Parfit has offered a novel solution to this problem. He argues, roughly, that qualitative differences between the lives in two populations will often entail that the populations are merely “imprecisely” comparable. This fact allows us to claim that A would be better than Z, on Parfit’s view, without violating the transitivity of *better than*. Parfit thus seems to have offered a solution to the Repugnant Conclusion that avoids any counterintuitive implications.

In this paper, I want to explore Parfit’s view as well as its implications. In section 1, I introduce his view, and explain how it avoids the Repugnant Conclusion. I then argue, in section 2, that the view fails to avoid two further conclusions, one of which is nearly identical to the Repugnant Conclusion. Thus, my arguments show that Parfit’s view leaves us, in effect, back where we began.

1. Lexical superiority and evaluative imprecision

One popular attempt to avoid the Repugnant Conclusion appeals to the claim that the lives in A are qualitatively different from those in Z. Supporters of this type of view typically claim that A would contain an amount of A-life to which no amount of Z-life—not even an infinite amount—is to be

preferred.⁴ That is, they claim that lives of the type lived in A are *lexically better* than lives of the type lived in Z—that there is some number of A-lives to which no number of Z-lives should be preferred. Importantly, such views do not claim that Z-lives have diminishing marginal value—that is, they do not claim that the value of each additional Z-life would converge to zero. Rather, the views are typically combined with the Simple View, which denies this claim, and thus accept that each Z-life would contribute the same amount of positive value to the world.⁵ So, on these *lexical views*, while Z-lives have non-diminishing value, there is still some amount of A-life to which no amount of Z-life is to be preferred. The views then, presumably, allow us to prefer A to Z. But, while these views may seem plausible, they are susceptible to an important objection.⁶

In particular, it is not enough for proponents of these views to claim that the lives in A are lexically better than those in Z: if a qualitative difference of this sort holds between A and Z, it must hold also between two successive populations on the continuum from A to Z.⁷ That is, if the lives in A are indeed lexically better than those in Z, then there must be two adjacent populations between A and Z, one of which would contain lives that are lexically better than those in the other. This is because, if the lives in A are not lexically better than the lives in B, which are not lexically better than the lives in C ... which are not lexically better than the lives in Z, it follows, by the transitivity of *not better than*, that the lives in A are not lexically better than the lives in Z. More precisely: if the lives in A are not lexically better than the lives in B, it follows that, for any number, x , there is some

⁴ See, for example, Crisp 1992 (pp. 149-152), Glover 1997 (pp. 69-71), Griffin 1986 (pp. 338-340), and Portmore 1999 (pp. 82-84).

⁵ Notice, though, that combining these views forces a change in the Simple View: since some number of A-lives is better than any number of Z-lives, it follows that there is, in fact, *some* upper limit to the value of Z-lives—in particular, whatever value the relevant number of A-lives has. That is, since there is some value that cannot be surpassed by the value of any number of Z-lives, it follows that there is, in this sense, some upper limit to the value of the Z-lives. So, if lexical views are to be combined with the Simple View, the latter must be amended to leave out talk of the nonexistence of any upper limit. Thanks to an anonymous reviewer for pressing me to make this point clear.

⁶ Another objection of note claims that lexical views are in some cases internally inconsistent. This point is explored in Jensen 2008; and also mentioned by Parfit (2016, p. 116).

⁷ A similar point is made in more general terms by Arrhenius and Rabinowicz (2005).

number, x' , such that x A-lives are not better than x' B-lives. And if we apply similar reasoning to the lives in B and C, and so on, it will follow, by the transitivity of *not better than*, that there is some number, x'' , such that x A-lives are not better than x'' Z-lives.⁸ Thus proponents of lexical views will have to do more than claim that the lives in A are lexically better than those in Z; in particular, they will have to claim that there are two adjacent populations, say P and Q, such that the lives in P are lexically better than those in Q. Since the qualitative difference between P-lives and Q-lives is presumably slight, the claim that there is some amount of P-life to which no amount of Q-life is to be preferred will likely seem implausible. And since something similar can be said of each step on the continuum from A to Z, it seems lexical views are certain to have at least some implausible implications.

It seems, then, that lexical views are untenable. But this is only the case, according to Parfit, so long as we assume that the qualitative differences between populations on the continuum from A to Z are also *precise*. That is to say, we might avoid the above objection if we claim that this is one of many cases where the qualitative differences between two things give rise to *evaluative imprecision*. These are unlike cases in which, for instance, pain of a particular intensity is experienced first for ten seconds, and then once more for twenty seconds. Plausibly, in this case, the second experience of pain is precisely twice as bad as the first, given that it was of the same intensity and precisely twice as long. But if the intensity is altered along with the duration—if the qualitative aspects of the experience are changed—it seems that it will no longer be possible for a difference between the two experiences to be precise. Consider, for instance, five minutes of torture versus one day of pain from a paper cut. Surely the latter is not precisely twice as good as the former; nor, it seems, is one better than the other to any precise degree. In cases of this sort, Parfit claims, one thing cannot be better or worse than another by some precise number of units, nor can the two things be precisely

⁸ For more on this, see Arrhenius and Rabinowicz 2005.

equally good.⁹ Instead, when comparing two qualitatively quite different things, often the most that can be said is that one thing is better or worse than another—to some greater or lesser, yet imprecise degree—or that the two things are imprecisely equally good.¹⁰

Another example may make this notion of evaluative imprecision clearer. Consider then Parfit's question of whether Einstein or Bach was a greater genius.¹¹ If we first suppose that the only available relations are precise, the question may seem to have no true answer. After all, given the qualitative differences between scientists and composers, neither seems clearly to have been the greater genius. Nor do their geniuses seem to have been precisely equal. We might conclude, then, that this sort of qualitative comparison simply cannot be made. But of course this would be a mistake: clearly, Bach was a greater genius than many poor scientists, and Einstein was a greater genius than many poor composers. This suggests, as Parfit sees it, that comparisons of this sort can be made, though they must be imprecise. And this suggests that the geniuses of Bach and Einstein were, perhaps, imprecisely equally great.

More generally: two comparable things, x and y , are *imprecisely equally good* if, while x is neither better nor worse than y , there could be a third thing, z , which is neither better nor worse than x , but is better or worse than y .¹² As Parfit notes, this relation is most likely to apply in cases where one thing is qualitatively quite different from two other, qualitatively similar things. So, for instance, when considering one's future, one might take being a writer and being an investment banker to be imprecisely equally good careers.¹³ In this case, it might be that while being a writer would not be

⁹ In other words, these things are cardinally comparable, but not by some number of units on a relevant scale. See Chang 2016.

¹⁰ Parfit 2016 (p. 113).

¹¹ Parfit 2016 (pp. 113-114). This type of example is, as Parfit notes, attributable to Ruth Chang (1997, 2002).

¹² Parfit 2016 (p. 115). This is similar to Chang's (2002) notion of "parity." For discussion of how these notions differ, see Chang 2016.

¹³ This case is closely modeled after examples in Hare 2013 (p. 171) and Parfit 2016 (p. 115).

better than being an investment banker, which would not be better than being a slightly more successful writer, being a slightly more successful writer *would* be better than being a writer.

Importantly, then, when differences in value are imprecise, *not better than* is not a transitive relation. It is this break in transitivity that serves as the basis for Parfit's solution to the Repugnant Conclusion. As we have seen, one objection to lexical views is that they appear to imply that the lives in one population are lexically better than lives, which are very similar, in an adjacent population. These views are then defeated by the more plausible claim that the lives in A are not lexically better than the lives in B, which are not lexically better than the lives in C ... which are not lexically better than the lives in Z. So long as we assume precision, *not better than* is a transitive relation, and so we conclude that the lives in A are not lexically better than the lives in Z. Thus, lexical views are defeated. According to Parfit, however, the qualitative differences between populations along the continuum from A to Z give rise to evaluative imprecision. That is, because there is at each step of the continuum a tradeoff between the number of lives in a population and (more importantly) the quality of those lives, precise evaluations of any two populations will be impossible. It then seems plausible that, given this tradeoff between quantity and quality of lives, the adjacent populations would be at most imprecisely equally good. And since, when there is imprecision, *not better than* is not a transitive relation, we might accept the claim that A would not be better than B, which would not be better than C ... which would not be better than Z, and yet reject the claim that A would not be better than Z. Thus we might claim, for each step along the continuum, that the lives in this particular population are not lexically better than those in an adjacent population, but claim also that the lives in the first population *are* lexically better than those in the last. Since we do not assume precision, there is no contradiction in our claims. Accordingly, appealing to Parfit's *Imprecise Lexical View* allows us to claim both that there is no population along the continuum from A to Z in which the lives are lexically better than those in an adjacent

population, and also that A would be better than Z.¹⁴ So long as we accept that the qualitative differences between adjacent populations entail that these populations would be at most imprecisely equally good, then, we can retain the essence of lexical views while avoiding their implausible implications.

Unfortunately, however, Parfit's solution does imply an alternative version of the Repugnant Conclusion. And while this version may seem less repugnant than the original, it can, I think, be amended to become almost identical to the Repugnant Conclusion. Thus, in what follows, I'll argue that Parfit's Imprecise Lexical View has an implication that is nearly as implausible as the Repugnant Conclusion, and that a solution will in this case be more difficult for proponents of his view to provide. We'll turn to this argument in a moment. But before that, let me explain why Parfit's view fails to block the alternative, possibly less repugnant version of the Repugnant Conclusion.

2. Two implausible conclusions

2.1 The Short-lived Z Conclusion

Recall our example in which pain of a certain intensity is experienced first for ten seconds, and then once more for twenty seconds. Since the two experiences do not differ along some qualitative dimension—differing only with respect to length of time—it seems that our evaluation of the experiences can be precise: the second experience is precisely twice as bad as the first. We have no

¹⁴ If there is no population in which the lives are lexically better than those in an adjacent population, how exactly is this lexical superiority supposed to arise? Here is how I understand Parfit's suggestion. Along the continuum from A to Z, there are ranges of populations, each of which is such that the populations within it are imprecisely equally good, and these ranges are bounded by ranges where it is indeterminate whether one of its populations is better than another or the two populations are imprecisely equally good. Thus it might be, for instance, that the lives in P are not lexically better than the lives in Q, which are not lexically better than the lives in R, yet the lives in P are lexically better than the lives in Z. In this case, P, Q, and R are in a range of populations which are imprecisely equally good, and this range is bounded by a range—which includes P, Q, R, S ... and W—where it is indeterminate whether one of its populations is better than another or the two populations are imprecisely equally good. Still, it might be that W, X, Y, and Z are in a range of populations which are imprecisely equally good. As we can see, these ranges overlap such that no population contains lives that are lexically better than those in an adjacent population, and yet the lives in P (and, so, A) are lexically better than those in Z.

need to appeal to evaluative imprecision in cases of this sort, then, because differences that are in no sense qualitative can be represented precisely. Thus we can note at the outset that evaluative imprecision only becomes relevant in cases where two things vary along some qualitative dimension.

Importantly, and as Parfit notes, there are a number of ways in which we can understand the lives in Z to be just barely worth living. The most common and, it seems, most natural way is to represent these lives as much lower in quality than the lives in A. After all, the lives in Z would contain just a small fraction of the amount of welfare contained in A-lives, and so it seems plausible to suppose that the lives in Z would simply be much drabber than the lives in A. In this case, the best features of Z-lives might be muzak and potatoes, whereas the lives in A represent some of the best lives that human beings could have.¹⁵ We can call the world in which people's lives would be in this way just barely worth living *Drab Z*.

When most people think about the Repugnant Conclusion, this is, I think, the possible world that they imagine. Thus to set terms, we can say that the Repugnant Conclusion is the conclusion that *Drab Z* would be better than A. But there is another way that the lives in Z could be just barely worth living. In what we can call *Short-lived Z*, people would have lives that were qualitatively identical to the lives in A, but these lives would be so short that they were barely worth living.¹⁶ Obviously enough, in order for a life to be barely worth living due just to its duration, this life would have to be quite short. It would indeed have to be, I think, even shorter than the duration Parfit suggests, which is one hour. Perhaps, then, the lives in *Short-lived Z* would last for only one minute. Intuitions may vary, but, as I see it, a life of ecstasy that would last for only this long would indeed be short enough to be barely worth living. If there is disagreement here, at any rate, the lives can be shortened to whatever length one's intuition suggests is appropriate.

¹⁵ See Parfit 2004 (pp. 17-20).

¹⁶ This alternative version of Z is introduced in Portmore 1999 (pp. 81-82).

Perhaps because he realizes that evaluative imprecision arises only in cases where the things in question differ along some qualitative dimension, Parfit suggests that the conclusion that Short-lived Z would be better than A—or, what we can call the *Short-lived Z Conclusion*—is considerably less repugnant than the Repugnant Conclusion. Regardless of his motivation for making this suggestion, though, it seems clear that Parfit's Imprecise Lexical View will indeed imply the Short-lived Z Conclusion. For, as we've noted, evaluative imprecision does not seem to arise in cases where the differences between things are non-qualitative and can thus be represented precisely. And, clearly, the lives in populations on the continuum from A to Short-lived Z, differing merely with respect to how long they would last, represent a case of just this sort. That is, each population along this continuum would be slightly larger than the population preceding it, and would contain lives that were qualitatively identical to, though slightly shorter than, the lives in the preceding population. And since the lives in each population would be qualitatively identical to all others, any difference between these populations must be non-qualitative and must therefore, it seems, be precise. In short: just as pain of a certain intensity experienced for twenty seconds is twice as bad as pain of that same intensity experienced for ten seconds, so each population on the continuum from A to Short-lived Z would be better than the population preceding it to the precise degree to which it would contain a greater total amount of A-life.

Thus, without any qualitative differences between the lives on the continuum from A to Short-lived Z, it seems there is no imprecision in our evaluations of these populations. This means that, if we accept the plausible claim that the lives in A are not lexically better than the lives in Short-lived B, which are not lexically better than the lives in Short-lived C ... which are not lexically better than the lives in Short-lived Z, we cannot also, in this case, reject the claim that the lives in A are not lexically better than the lives in Short-lived Z. Since the differences in values here are not imprecise, *not better than* remains a transitive relation, and so it follows that the lives in A are not lexically better

than the lives in Short-lived Z. Thus, if we accept Parfit's Imprecise Lexical View, we still must accept that Short-lived Z would be better than A—that is, we must accept the Short-lived Z Conclusion.

Admittedly, my argument here may seem just a bit too quick. More specifically: it might be objected that evaluative imprecision may plausibly seem to arise, not just when there is a difference between the quality of the lives in two populations, but also when those populations themselves differ in certain respects. It might be, for example, that one population would have a higher average level of welfare than another, yet the other would have a higher total level of welfare than the one. Still, the quality of the lives in these populations might be the same. In this case, the two populations might seem merely imprecisely comparable, despite the fact that each population would contain lives of the same quality. The suggestion, then, is that tradeoffs between values such as average and total welfare might give rise to evaluative imprecision, even if the quality of the lives in the relevant populations is held fixed. If this is right, it seems that we will have grounds to reject my claim that, because there is no difference in the quality of the lives on the continuum from A to Short-lived Z, the proponent of the Imprecise Lexical View must accept the Short-lived Z Conclusion.¹⁷

But I do not think this is right. If the quality of the lives in the two populations is held fixed, it does not seem to me that a tradeoff between average and total welfare in the populations could give rise to evaluative imprecision. This is because if the quality of the lives in each population would be the same, such a tradeoff could only be the result of one population containing some lives that would last longer than those in the other, and the other containing a greater total number of lives than the one. That is, in this case, a higher average level of welfare would require some lives that lasted longer—since this would be the only way to increase the welfare of each life—and a higher total level of welfare would require a greater total number of lives—since the only alternative,

¹⁷ Thanks to an anonymous reviewer for pressing this objection.

increasing the welfare of each life, would also increase the average. That being so, it seems that the differences between the populations—the total number of lives they contain, and length of each of these lives—can be represented precisely. And so, if there is a tradeoff between total and average welfare in two populations whose members would have lives of the same quality, the populations will be precisely comparable: the population containing more life of that quality would be better to the precise degree to which it would contain more such life. Accordingly, such a tradeoff cannot give rise to evaluative imprecision, and thus cannot plausibly be appealed to in attempt to avoid the Short-lived Z Conclusion.

Let me try to put this point in a slightly different way. As we have seen, if I were to experience pain of a certain intensity first for ten seconds, and then once more for twenty seconds, it seems plausible that the second experience of pain would be precisely twice as bad as the first. Now consider a variation of this case, where there would now exist two more people who would instead experience the second instance of pain, such that each of us would experience the same ten seconds of pain. Would this outcome be better than the one where I alone experience the two instances of pain? Setting aside considerations of desert, fairness, and the like—looking at the subjects as mere vessels for pain—it seems to me that these outcomes would instead be precisely equally good. So long as the amount and type of pain experienced would be the same in either case, I don't see a plausible way of separating the outcomes with respect to their goodness. Consider finally a third case where, while one of our other two subjects would still experience the same pain for ten seconds, the other would not exist, and I would now experience this pain for fifteen seconds. Would this outcome be better than either of those we've just considered? Again setting aside considerations of desert, fairness, and the like, it seems to me that this outcome would indeed be better than either of those just considered—and, in fact, that it would be precisely 1.2 times better, given the amount of pain in each. Keeping all this in mind, notice how we have gotten from the second case to the third:

while holding fixed the qualitative aspect of the experiences, we've both raised the total level of welfare and lowered the average level of welfare in the outcomes. And this tradeoff has resulted in precise evaluative differences between the outcomes. Contrary to the objector's suggestion, then, such tradeoffs do not seem to give rise to evaluative imprecision. Moreover, and importantly, notice that the continuum from A to Short-lived Z develops in just the same way as these last two cases: the quality of the lives is held fixed, the total level of welfare is raised, and the average level of welfare is lowered.¹⁸ Thus, it seems that the evaluative differences between these populations, too, will be precise. And so we remain without reason to reject my claim that the Imprecise Lexical View implies the Short-lived Z Conclusion.

Of course, proponents of the Imprecise Lexical View may have no problem with this conclusion. So long as a world would contain a greater total amount of A-life than the amount that would be contained in A, they might claim, the former would be for that reason in one way better than the latter, and should thus, if it would also be in other ways not worse, be preferred to the latter. That is to say, so long as the members of a population would experience a greater amount of A-life than the amount that would be experienced in A, it does not matter how many people would experience this total amount of A-life, or for how long these people would experience it. What matters, according to this person, is that the quantity and quality of life that would be experienced in A would be experienced at all.

¹⁸ Of course, this analogy might seem to break down when we consider other sorts of cases. It might be, for example, that pain of a certain intensity is experienced first for ten seconds, and then once more for twenty minutes. In this case, we might think differently as to whether it would be better if, rather than one person having both of these experiences, one person had the first experience while 120 others had equal parts of the second. This is because, it might be thought, after, say, ten minutes of this type of pain, the experience becomes worse in a way that is not attributable to the cumulative number of seconds involved. Notice, however, that this suggestion requires a qualitative difference between the types of pain involved in the two experiences. That is, this suggestion might find evaluative imprecision between the imagined outcomes, but only by assuming that there is also a qualitative difference between them—in particular, a difference in the type of pain experienced. Since the argument that I've presented concerns only those cases that hold fixed the qualitative features of the outcomes, this suggestion presents no threat to the argument. More generally, I suspect that any ostensible evaluative imprecision that arises between populations whose quality of life is held fixed must be attributable to an implicitly assumed difference between the quality of those lives. Otherwise, as I see it, such evaluative imprecision cannot arise. And so it cannot arise, in particular, on the continuum from A to Short-lived Z.

Regardless of whether we find this response plausible, it is, as I see it, the only response available to proponents of the Imprecise Lexical View. After all, the view implies the Short-lived Z Conclusion. And so it seems proponents of the view can either argue that this conclusion is not all that repugnant, or else give up the view because of its repugnant implications. I've just offered one suggestion of what taking the former option might look like. Let me now explain why, in the end, the latter option is the only one left to take.

2.2 The Slightly Less Repugnant Conclusion

Reconsider the lives in Short-lived Z. Each would consist of a mere minute of ecstasy. To more easily envision this possible world, we might think of each member of Short-lived Z as experiencing a one-minute time-slice of the lives that would be experienced in A—including even the memories associated with the relevant life in A.¹⁹ Some of these time-slices would need to be repeated, of course, given that there would be much more A-life in Short-lived Z than in A. But details aside, thinking of the lives in Short-lived Z in this way makes it particularly clear how it could be that these lives would be qualitatively identical to the lives in A. The members of each population would collectively have the same exact experiences. And so, as we have seen, since Short-lived Z would contain a greater total amount of A-life, it would be, on the Imprecise Lexical View, better than A.

Keeping this understanding of the lives in Short-lived Z in mind, suppose now that we add one hundred drab years to each of these lives. And imagine that the one minute of ecstasy they experience is now sprinkled throughout these one hundred drab years.²⁰ Call the resultant population

¹⁹ Here is another way of looking at it. Given the lives in these populations would be qualitatively identical, there would be no way of determining which of the two types of life one was living. Thus, just as our own lives may have started within the last minute—all memories from before this having been implanted—so the lives in Short-lived Z may (and, indeed, must) have started within the last minute. And, thus, just as this does not restrict the quality of our lives, so it does not restrict the quality of the lives in Short-lived Z.

²⁰ Alternatively: add one drab day to each of the lives in Short-lived Z, and repeat until one hundred drab years have been added.

Short-lived Z+. The members of this population would, to be sure, have lives with a much lower average level of welfare throughout. But these people would also, importantly, have lives with a higher total level of welfare than they would in *Short-lived Z*—they would contribute more value to the world. Most importantly, *Short-lived Z+* would not only contain a greater total amount of welfare than *Short-lived Z*, but would also contain the same amount of A-life as *Short-lived Z*. This means that, on the Imprecise Lexical View, *Short-lived Z+* would be better than *Short-lived Z*, which would be better than A. And so it follows, by the transitivity of *better than*, that the Imprecise Lexical View implies

The Slightly Less Repugnant Conclusion: Compared with the existence of many people who would all have the best lives that human beings could have, there is some much larger number of people whose existence would be better, even though these people would have lives—consisting of one minute of ecstasy sprinkled throughout one hundred drab years—that were barely worth living.²¹

This is not much better than the Repugnant Conclusion. Given that the Imprecise Lexical View has this implication, we find ourselves, in effect, back where we began. After all, we have gained just one minute of ecstasy that would be sprinkled throughout each of the lives in *Drab Z*. And so, it seems that the Imprecise Lexical View hardly improves on the Simple View at all.

Of course, the way we have reached this conclusion leads also to a natural objection. It may be asked: so what? Why is it a problem for proponents of the Imprecise Lexical View that their view has this implication? After all, these are the same people who, it seems, must endorse the *Short-lived*

²¹ When Parfit considers a similar conclusion, he claims that this conclusion is “hard to believe” but still “less repugnant” than the Repugnant Conclusion (2004, p. 19). Accordingly, it seems Parfit’s view has an implication that he himself wishes to avoid (though perhaps one he would ultimately be willing to tolerate).

Z Conclusion. And so, if all we have done to reach the Slightly Less Repugnant Conclusion is add to the lives in Short-lived Z a bit of life that, while drab, is nevertheless worth living, it might seem that we have done little to inspire any feelings of repugnance in proponents of this view. To put the challenge slightly differently: if all we have done here is make the lives in Short-lived Z better, how could *that* make for a conclusion that is any more repugnant than the Short-lived Z Conclusion?²²

Well, for one thing, I suspect that proponents of the Imprecise Lexical View will find the Slightly Less Repugnant Conclusion repugnant, not because of the similarities between Short-lived Z+ and Short-lived Z, but rather because of the similarities between Short-lived Z+ and Drab Z. So while it must be admitted that these people might, much like the current objector, not immediately have the intuition that I suspect they will, the important question is whether they would have the same intuition on reflection. And so, more specifically, we should ask: after reconsidering the fully and slightly less repugnant conclusions, and looking closely at the difference between them, could these people plausibly embrace one while abhorring the other? The answer, I think, is no. This is because, for those who find the Repugnant Conclusion truly repugnant, sprinkling one minute of ecstasy throughout each of the lives in Drab Z is unlikely to make the conclusion seem much better. If this is right, the objection that proponents of the Imprecise Lexical View may not find the Slightly Less Repugnant Conclusion repugnant loses much of its force, since, after a closer look and a longer thought, these people will find themselves far less comfortable with the idea of embracing this implication of their view.

For another thing, even if they're comfortable with embracing it, the Slightly Less Repugnant Conclusion still may present a problem for proponents of the Imprecise Lexical View. For if it is conceded that this conclusion does not improve much on the Repugnant Conclusion, yet it is also claimed that this conclusion is not all that repugnant, it follows that the Repugnant Conclusion is

²² I thank an anonymous reviewer for pressing me on this.

also not all that repugnant. And, in that case, it seems that the Imprecise Lexical View does not improve much on the Simple View, since the latter does not have any repugnant implications in the first place. Consequently, regardless of whether proponents of the Imprecise Lexical View are willing to embrace the Slightly Less Repugnant Conclusion, it seems that this conclusion may undermine their view.

In light of this, of course, these proponents might instead try to resist my argument that their view implies the Slightly Less Repugnant Conclusion. And there are at least two ways they might do this. They might object: it is far from clear that the lives in Short-lived Z+ would be better than those in Short-lived Z. Indeed, many are likely to have the opposite preference, which may suggest that we have reason to reconsider which lives would contribute more value to the world. Perhaps, then, we can avoid the claim that Short-lived Z+ would be better than Short-lived Z simply by denying that the lives in the former population would be better than those in the latter.

But on what grounds could we claim that the lives in Short-lived Z+ would not be better than those in Short-lived Z? We could not appeal to the preferences of those living in each population, for it seems that everyone who would exist in either world would prefer to continue living at a low level of welfare rather than to cease living after a single minute, and would thus prefer Short-lived Z+ to Short-lived Z.²³ Alternatively, we could of course claim that, in spite of their preferences, what matters is in fact their average level of welfare; and since this level would be lower in Short-lived Z+, we should not prefer it to Short-lived Z. But this suggestion is implausible for much the same reason that average utility approaches to the Repugnant Conclusion are implausible: they have sadistic implications.²⁴ Such a view might imply, for instance, that adding a small amount of torture to a life in Short-lived Z+ would be better than adding a great deal of drabness to it,

²³ In fact, if we apply a plausible Pareto principle here, we have an argument in favor of preferring Short-lived Z+ to Short-lived Z. See Huemer 2008 (p. 903).

²⁴ For more on this, see Arrhenius 2000 and Huemer 2008.

because the former would lower the life's average welfare to a lesser extent than the latter. If there is another appeal that can be made here, it is unclear to me what it might be. Thus it seems we are without any plausible grounds for the claim that the lives in Short-lived Z+ would not be better than those in Short-lived Z.²⁵

There is a second way of objecting to my claim that the Imprecise Lexical View implies the Slightly Less Repugnant Conclusion. It may be said: while the Imprecise Lexical View does imply that Short-lived Z would be better than A, it does not imply that Short-lived Z+ would be better than Short-lived Z, and so does not imply that Short-lived Z+ would be better than A. This is because there is a qualitative difference between the lives in Short-lived Z+ and Short-lived Z, and it follows from this that any evaluation of these two populations will be imprecise. It then seems plausible that, given this qualitative difference, Short-lived Z+ and Short-lived Z would be at most imprecisely equally good. And since, when there is imprecision, *not better than* is not a transitive relation, it might be the case that Short-lived Z would not be better than Short-lived Z+, which would not be better than A, yet Short-lived Z would still be better than A. Put differently, since Short-lived Z+ would be imprecisely equal to—rather than better than—Short-lived Z, the fact that Short-lived Z would be better than A does not entail that Short-lived Z+ would be better than A. Thus the Imprecise Lexical View does not imply that Short-lived Z+ would be better than Short-lived Z, or that Short-lived Z+ would be better than A.

Here is where this objection goes wrong. It hinges on the plausibility of the claim that Short-lived Z+ and Short-lived Z would be at most imprecisely equally good. Yet this claim is clearly false. For, the claim that two qualitatively different things are imprecisely equally good only seems

²⁵ It is perhaps worth mentioning that we are prone to a particular bias that might explain this intuition. Specifically, we tend to undervalue the duration of lives when assessing their desirability by focusing on their average level, rather than total sum, of happiness. This suggests that we are likely to judge some lives, such as those in Short-lived Z+, to be less desirable than other lives, such as those in Short-lived Z, despite the fact that the former just are the latter plus an additional duration of happy life. See Diener et al. 2001; and for discussion, see Kahneman 2011 (pp. 386-388).

plausible in cases where there is some tradeoff between the values making this qualitative difference up. This is why, for instance, A and B can only be imprecisely equally good: because in moving from one population to the other, there is a tradeoff between the number of lives in a population and the quality of those lives. It is also why being an investment banker and being a writer might be imprecisely equally good careers. But of course the difference between Short-lived Z+ and Short-lived Z is not like these differences. After all, in order to reach Short-lived Z+, we simply *added* to the lives in Short-lived Z a type of value that is qualitatively different from the type that was already present in those lives. We did not *replace* any of the value that was already there. Thus a more appropriate analogy to the difference between these two populations might be the difference between being an investment banker and being an investment banker who would also have some mildly rewarding writing duties on the side. There is of course a qualitative difference between the two careers—just as there is between Short-lived Z+ and Short-lived Z—but not one that matters. More generally: even if x and y are qualitatively quite different things, this is irrelevant if $x+y$ is better than y . Since there is no tradeoff between the two qualitatively different types of thing, and instead one thing merely subsumes and adds to the other, the qualitative difference between the two things is irrelevant: one thing is simply better than another. Thus the qualitative difference between Short-lived Z+ and Short-lived Z is irrelevant, and the two would not be imprecisely equally good. Accordingly, the appeal to evaluative imprecision provides no escape from the conclusion that Short-lived Z+ would be better than Short-lived Z, which would be better than A.

3. Conclusion

In sum: while Parfit's Imprecise Lexical View can indeed plausibly avoid the Repugnant Conclusion, the view also has a pair of objectionable implications of its own. As I have shown, there is an alternative version of the Repugnant Conclusion, the Short-lived Z Conclusion, which this view

cannot avoid—that is, the view implies that Short-lived Z would be better than A. And if we add one hundred drab years to each of the lives in Short-lived Z, it seems that the resultant population, Short-lived Z+, would be better than Short-lived Z. This means that, on the Imprecise Lexical View, Short-lived Z+ would be better than Short-lived Z, which would be better than A. And this means, of course, that the view implies the Slightly Less Repugnant Conclusion—that is, it implies that Short-lived Z+ would be better than A. As the title suggests, I suspect that proponents of the Imprecise Lexical View will take little solace in exchanging the Repugnant Conclusion for this marginally better alternative. Thus it seems we are, in effect, back to square one.²⁶

²⁶ Many thanks to Doug Portmore and several anonymous reviewers for most helpful comments.

References

- Arrhenius, G. (2000). "An Impossibility Theorem for Welfarist Axiologies" *Economics and Philosophy* 16: 247-266.
- Arrhenius, G., and W. Rabinowicz (2005). "Millian Superiorities." *Utilitas* 17: 127-146.
- Chang, R. (1997). "Introduction." In R. Chang (ed.), *Incommensurability, Incomparability, and Practical Reason*. Cambridge: Harvard UP. 1-34.
- — —. (2002). "The Possibility of Parity." *Ethics* 112: 659-688.
- — —. (2016). "Parity, Imprecise Comparability and the Repugnant Conclusion." *Theoria* 82: 182-214.
- Crisp, R. (1992). "Utilitarianism and the Life of Virtue." *The Philosophical Quarterly* 42: 139-160.
- Diener, E., D. Wirtz, and S. Oishi (2001). "End Effects of Rated Life Quality: The James Dean Effect." *Psychological Science* 12: 124-128.
- Glover, J. (1977). *Causing Death and Saving Lives*. Harmondsworth: Penguin.
- Griffin, J. (1986). *Well-being: Its Meaning, Measurement, and Moral Importance*. Oxford: Clarendon.
- Hare, C. (2013). *The Limits of Kindness*. Oxford: Oxford UP.
- Huemer, M. (2008). "In Defence of Repugnance." *Mind* 117: 899-933.
- Jensen, K. K. (2008). "Millian Superiorities and the Repugnant Conclusion." *Utilitas* 20: 279-300.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Nebel, J. M. (2017). "The Good, the Bad, and the Transitivity of *Better Than*." *Noûs* 51: 1-26.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon.
- — —. (2004). "Overpopulation and the Quality of Life." In J. Ryberg and T. Tännsjö (eds.), *The Repugnant Conclusion: Essays on Population Ethics*. Dordrecht: Kluwer Academic. 7-22.
- — —. (2016). "Can We Avoid the Repugnant Conclusion?" *Theoria* 82: 110-127.
- Portmore, D. W. (1999). "Does the Total Principle Have Any Repugnant Implications?" *Ratio* 12:

80-98.

Rachels, S. (1998). "Counterexamples to the Transitivity of Better Than." *Australasian Journal of Philosophy* 76: 71-83.

Temkin, L. S. (1987). "Intransitivity and the Mere Addition Paradox." *Philosophy & Public Affairs* 16: 138-187.

— — —. (1996). "A Continuum Argument for Intransitivity." *Philosophy & Public Affairs* 25: 175-210.

Voorhoeve, A. (2013). "Vaulting Intuition: Temkin's Critique of Transitivity." *Economics and Philosophy* 29: 409-425.