

Can We Avoid the Repugnant Conclusion?

by

DEREK PARFIT

University of Oxford

Abstract: According to *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. I suggest some ways in which we might be able to avoid this conclusion. I try to defend a strong form of lexical superiority.

Keywords: ethics of population, quality of life, lexical superiority

ACCORDING TO

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.

This conclusion is implied by the Utilitarian view that it would always be better if there was a greater total sum of happiness. There might be such a greater sum in the lives of many people who each had very little happiness, just as there might be some greater mass of milk in a vast heap of bottles that each contained only one drop.

To avoid this conclusion, it is not enough to reject this Utilitarian Total View. On what I shall call

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 worth living.

Someone's existence has more value, on this view, if this person's life is more worth living. But even if someone's life is barely worth living, or only slightly better than nothing, this person's existence has some intrinsic value, making the world slightly better. These claims seem to imply that, compared with a world in which everyone's quality of life would be very high, there would be other possible worlds whose enormous populations would make them better, even though, in these worlds, everyone's life would be barely worth living. That is the Repugnant Conclusion.

Several people have reluctantly accepted this conclusion. This conclusion seems repugnant, some people claim, because we cannot adequately imagine very

large numbers. When we think of lives that are barely worth living, we fail to realize how much value there could be in many billions of such lives. There are some cases of this kind. Many people underestimate what natural selection can achieve during millions of years. But that is not why most of us would find the Repugnant Conclusion hard to accept. We would believe that, compared with the existence of many people who would all have very good lives, there is no number of other people whose existence would be better, if these people's lives were barely worth living. Even if we cannot adequately imagine very large numbers, we understand the belief that there is *no* number of such people whose existence would be better.

Some other people reject the Repugnant Conclusion by denying that there is any intelligible impersonal sense in which some outcomes can be better or worse than others. But on this view, compared with a world in which everyone had very good lives, it would not be worse if everyone's life was full of suffering. That is another Repugnant Conclusion.

A third group of people claim that, for someone's existence to be in itself good, and to make the world better, it is not enough that this person's life is worth living, or above the Zero Level of well-being. This person's quality of life must be above another, higher *Critical Level*. These writers accept only a weaker form of the Repugnant Conclusion. On this view, compared with the existence of many people whose quality of life would be very high, there are some much larger numbers of people whose existence would be better, though these people's lives would all be only just above the Critical Level. This conclusion would be significantly less repugnant if these lives were, as John Broome supposes, 'reasonably good'.

This view takes two forms. When people's lives are worth living, but not above the Critical Level, they are in what we can call the *Critical Range*. On the *Neutral Critical View*, when people's lives are in this range, their existence has no intrinsic value, so it would not be worse if these people's lives were not reasonably good, but only barely worth living. We might call this the *Callous Conclusion*. On the *Negative Critical View*, when people have lives that are worth living, but below the Critical Level, their existence has disvalue. This view implies that, if there existed enough people with lives that were close to being reasonably good, that would be worse than if there existed many fewer people whose lives would be full of suffering, and worse than nothing. Gustaf Arrhenius calls this the *Sadistic Conclusion*. This name is unfair, since those who accept this view would regret the existence of such people. But I cannot see how this view could be true.

Some other writers suggest another way to avoid the Repugnant Conclusion. On what we can call

Diminishing Value Views: Though it would always be in itself better if there existed any extra person whose life would be worth living, the goodness of there being more such people would steadily diminish, and would have some upper limit.¹

Though such views are less implausible, they are open to strong objections. One is the fact that we cannot defensibly make such claims about the existence of people whose lives would be full of suffering, and worse than nothing. The existence of such wretched people would not have a badness that would diminish as the number of such people grew, so that it mattered less and less whether more such people exist. The badness of more such suffering would never decline. Though there are some asymmetries between suffering and happiness, and some of the other things that can make lives good or bad, we cannot plausibly either apply some Diminishing Value View to lives that are bad, or restrict this view to lives that are good.

There is another, better view. We can believe that, as the Simple View claims:

Anyone's existence is in itself good if this person's life 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.

We can add:

If many people exist who would all have some high quality of life, that would be better than the non-existence of *any* number of people whose lives, though worth living, would be, in certain ways, much less good.

These claims together assert one kind of *lexical superiority*. When we say that things of kind *P* are *lexically better* than things of kind *Q*, we can mean that, though the existence of more *Qs* would always be non-diminishingly better, the existence of some sufficient number of *Ps* would be better than the existence of any number of *Qs*. There is a similar sense of *lexically worse than*. These claims state what we can call the *Imprecise Lexical View*.

This view may seem implausibly extreme, since it claims that though such less good lives would have non-diminishing value, *no* amount of this value could be as great as the value of the much better lives. Critical Level Views make no such

1 On this view, more fully stated: When people have lives that are worth living, their existence always has some value. But this value is not intrinsic, since it depends on how many other people exist at the same level of well-being. The value of each person's existence would decline as the number of such people grew. The combined value of these people's existence could never rise above some upper limit – in the way in which, if we add together a half, a quarter, an eighth, and so on, the total sum could never rise above one.

claim. But these other views are not lexical only because, rather than claiming that such less good lives have lexically inferior value, these views claim that such lives, though they are worth living, have either *no* value or *disvalue*. These claims are *more* extreme.

Lexical views are often claimed to be open to decisive objections. Before discussing these objections, I shall defend some claims about what we can call *evaluative imprecision*. That may seem a digression. But of the impressive arguments that many people have given for various repugnant conclusions, most assume precision.

There can be fairly precise truths about the relative value of some things. One of two painful ordeals, for example, might be twice as bad as the other, by involving pain of the same intensity for twice as long. But in most important cases relative value does not depend only on any such single, measurable property. When two painful ordeals differ greatly in both their length and their intensity, there are no precise truths about whether, and by how much, one of these pains would be worse. There is no scale on which we could weigh the relative importance of intensity and length. Nor could five minutes of ecstasy be precisely 7.6 times better than ten hours of amusement. If we compare different ways in which our life might go, when choosing between different careers, for example, or deciding whether to have children, there are only imprecise truths about which of these possible lives would be better or worse. And there are only imprecise truths about the relative goodness of many different acts or outcomes, such as ones that would greatly benefit a few people, or give lesser benefits to many others. Such imprecision is not the result of vagueness in our concepts, or our lack of knowledge, but is part of what we would know if we knew the full facts. When two things are qualitatively very different, these differences would often make it impossible either that one of these things is better than the other by some precise amount, or that both things are precisely equally good.

Similar claims apply to many non-evaluative facts. There may be only imprecise truths about which of two scientific theories is simpler, or explains more, or which of two rooms is more untidy, or which of two mountains is harder to climb. And when two people have different mental abilities, these people could not be precisely equally intelligent.

There is one kind of case, suggested by Ruth Chang, in which it may be easiest to see that there could not be any precise truth about some relation between two things. We may start by assuming that, given the great qualitative differences between two things, there could not be *any* such relational truth. Suppose that I ask you whether Einstein or Bach was a greater genius, or achieved more. You may first assume that this question could not have an answer, since it makes no sense to compare the genius, or achievements, of scientists and composers. But I

might then point out that Bach was clearly a greater genius than many bad scientists, and Einstein was a greater genius than many bad composers. When you realize that there *can* be truths of this kind, you would not suddenly come to believe that as geniuses, or in their achievements, Einstein and Bach might be precisely equally great. As you would see, the truth could be only that one of these people was imprecisely greater than the other, or more plausibly that they were imprecisely equally great.

Many people assume that, when there are truths about the relative goodness of different things, these truths must be precise, though we may not know what these truths are. There is one way of thinking which can make this seem the only possible view. If things of some kind can be better or worse than others, and by more or less, it may seem that the goodness of these things corresponds to their positions on some line or scale of value. On this *Linear Model*, truths about goodness must be precise because positions on a line are precise. Suppose, for example, that one of two things starts by being better than the other, but then smoothly deteriorates and ends up by being worse. If the value of these things corresponded to their position on some line, the better thing's value would start by being higher on this line, and would then move down until it becomes lower, so that, at some point during this process, the value of these things must be precisely equal. But when two things are qualitatively very different, that could not be true. So when we think about the goodness of such things, we should reject this Linear Model. Nor could the goodness of such things correspond to different real numbers, since such numbers are also precise. Nor could some of these things be better than others by some imprecise amount, or to some imprecise degree, since the concepts of an amount or a degree also imply precision. We should think only about *differences* between the value of these things, since the concept of a difference does not imply precision.

When one of two things is better than the other, that is often all we need to know, since it does not matter whether this difference in value is precise. But when neither of two things is better than the other, we may need to know whether *this* relation is precise. That may make a difference to which conclusions we ought to reach. Some relation R is *transitive* when, if X is R-related to Y, and Y is R-related to Z, X must be R-related to Z. One example is *taller than*. If Tom is taller than Dick who is taller than Harry, Tom must be taller than Harry. Some relation R is *not* transitive when, even if X is R-related to Y, which is R-related to Z, X may *not* be R-related to Z. One example is the relation *friend of*. Even if Tom is a friend of Dick, who is a friend of Harry, Tom may not be a friend of Harry.

Precisely equal is a transitive relation. If X and Y are precisely equally good, and Y and Z are precisely equally good, X and Z must be precisely equally good. But if X and Y are *imprecisely* equally good, so that neither is worse than the

other, these imprecise relations are *not* transitive. Even if X is not worse than Y, which is not worse than Z, X may be worse than Z.

Such cases illustrate these imprecise relations. Two things are imprecisely equally good if it is true that, though neither thing is better than the other, there could be some third thing which was better or worse than one of these things, though *not* better or worse than the other. That is most often true when one of two things is qualitatively more similar to some third thing. Consider, for example, three ways in which your life might go. It might be true that your being a fairly successful writer would not be worse than your being a doctor, which would not be worse than your being a slightly more successful writer. But your being a fairly successful writer *would* be worse than your being a slightly more successful writer. *Not worse than* would not here be a transitive relation.

These points have wide practical implications. It is often assumed that if one of two things is in one way better, and in other ways not worse, the first thing must be better all things considered than the second thing. That is not so. When there is imprecision, the relation *not worse than* does not imply *at least as good as*, so one of two things can be in one way better and in other ways not worse without being better all things considered.

If we reject such claims, and believe that even when there is imprecision *not worse than* must be transitive, we may still be thinking in terms of the Linear Model, or assigning numbers. Like some other important truths, these truths about imprecision can be hard to understand, not because they are complicated, but because they are so simple. When some things are better than others by precise amounts, such differences are like the distances between positions on some line, and that is a simple idea. But when some things are better than others but these differences are imprecise, the truth is even simpler. Such differences in value do *not* have the further feature that they are like distances on some line. They *are not* like such distances because they are not precise.

Return now to the Lexical View which claims that

If many people exist who all have some high quality of life, that would be better than if there existed *any* number of people whose lives, though worth living, would be, in certain ways, much less good.

Many people either ignore or quickly dismiss such views. Most of these people again assume that something's goodness corresponds to its position on some scale, or to some real number. On that assumption, lexicality makes no sense. If the existence of more of these lesser goods would have non-diminishing value, but the combined value of such things could never be as great as that of some greater goods, that would imply both that

there are no limits either to how high on the scale of value this combined goodness could be, or to how great the number would be that would represent this goodness,

and that

there *are* such limits, since this position must be lower, and this number must be smaller, than those which represent the value of these greater goods.

That is a contradiction. As this shows, we cannot describe any lexical view by thinking in terms of positions on some scale, or by using real numbers. But that does not refute such views. We are trying to reach true beliefs about which outcomes would be better or worse, and about what we ought to do. We should not assume that such truths must be able to be represented by using scales or numbers. To save words, we can sometimes use scales or numbers, as I shall soon do, but we should take care not to be misled.

Some other writers claim that, though lexical views may make sense when applied to things in different categories, as in Cardinal Newman's claim that sin is infinitely worse than pain, lexical views cannot be applied to actual or possible things that are in the same category and could together form some continuum, with each thing being only slightly different from its neighbours. We cannot defensibly claim that everything above some point on such a continuum is lexically better than everything below this point, since there cannot be such a great difference in the value of the similar things that are just above or just below this point. We could not, for example, claim that pains of some intensity are lexically worse than pains that are only slightly less intense.

This objection fails, I believe, when we consider some kinds of thing which, though they can be thought of as on some continuum, are only imprecisely comparable. When that is true, lexical views can appeal, not to some single point in this continuum which separates greater and lesser goods, but to various zones in which different things would be imprecisely equally good, bounded by zones in which it would be indeterminate whether one of two things would be better, or these things would be imprecisely equally good.

In defending this *Imprecisionist Lexical View*, I shall first discuss what we can call *the Continuum Argument* for the Repugnant Conclusion. This argument's main premise claims that

(A) compared with the existence of many people who would all have lives that were equally worth living, there are some much larger numbers of people whose existence would be better, though these people would all have lives that would be *slightly* less worth living.

On this view, any slight loss in the quality of everyone's lives could be outweighed by a sufficient gain in the number of people who would exist and have lives worth living. Consider the simple worlds that are overly precisely shown in Figure 1. The width of these blocks shows the number of people who would exist, and their height shows the quality of these people's lives. On the Continuum Argument, compared with world A in which everyone's quality of life would be very high, it would be better if, in world B, there would exist instead many more people who would all be slightly worse off. Compared with world B, it would be similarly better if, in world C, there would exist instead many more people who would all be slightly worse off. This series continues down to world Z, in which there would exist some vast number of people who would have lives that were barely worth living. Since *better than* is a transitive relation, this argument implies that Z would be better than A. That is the Repugnant Conclusion.

In considering this argument, we can first distinguish between the quality of people's lives and the quantity of well-being per person. These might diverge. The best things in your life might be of a higher quality than the best things in mine, and your life might go worse than mine only because you would have many fewer of these best things. That would be most simply true if your life would be much shorter than mine. As this example shows, if we care greatly about the quality of life, being in this sense *perfectionists*, that would not make us *elitists*, who care most about the well-being of the best off people.

We can next distinguish different versions of our imagined World Z, in which some vast number of people would have lives that were only barely worth living.




Figure 1.

In what we can call *Short-Lived Z*, our imagined people would live for only as long as some flowers bloom. If these imagined beings lived for only one happy day, or one ecstatic hour, these lives could not be called barely worth living, but they would be too brief to be good. When applied to this imagined world, the Repugnant Conclusion would, I believe, be significantly less repugnant. Some of the best things in life would still be there. Though there are important differences between the quality of people's lives and the amount of well-being per person, I shall say little more about these differences here. We should suppose that, in my examples, people with lives of higher quality would also have more well-being.

Suppose next that, in *Roller-Coaster Z*, everyone would live as long as everyone in World A, and all of the good things in these people's lives would be just as good, but these people's lives would be barely worth living because their lives would also contain much that was very bad. This version of Z also raises questions that I shall not discuss here. World Z, we can assume, would take a third, simpler form.

In what I have earlier called *Drab Z*, there would be nothing in people's lives that would be bad, but there would also be very little that was good. The only good features, I suggested, might be muzak and potatoes. But that description is too simple. If the people in Drab Z would be in other ways like us, we could not plausibly assume that these people's lives would contain nothing bad, but very little that was good. Even if we lost most external goods, some of us would have inner mental resources with which we could make our lives fairly good, by composing long poems, for example, or thinking about some intellectual problems. Some political prisoners have lived fairly good lives while they were in solitary confinement. But some other people would find lives in such conditions worse than nothing. To make our questions clearer, we can suppose that lives in Drab Z would be only barely worth living, not because they would be lived by people like us who were in such deprived conditions, but because these lives would be lived by beings who would be psychologically much simpler than us. Such versions of Drab Z also help to show the full implications of the Continuum Argument.

There is one continuum between human beings and their evolutionary ancestors all the way back to some of the earliest sentient animals. When applied to this continuum, a claim like (A) would imply that

compared with the existence of many human beings who would live the best lives that humans could live, it would be better if there existed instead some much larger numbers of the earliest sentient animals who had lives that were just worth living, because these animals had enough slight pleasures like those of cows munching grass or lizards basking in the sun.

That is a repugnant conclusion.

There is another continuum between the parts our own lives. In the later parts of these lives there are some great qualitative differences, some of which involve wholly new features. But such changes are gradual. Human babies slowly develop during their childhood until, as adults, they cease to develop and change only in other ways. We can imagine a series of possible human beings who would differ from actual human beings by ceasing to develop at earlier ages. Another Continuum Argument assumes that

compared with the existence of many human beings who would live the best lives that were possible for them, it would be better if there existed instead some much larger number of human beings who would live the best lives that were possible for human beings who ceased to develop when they were slightly younger.

This argument implies that

compared with the existence of many adult human beings who would live the best lives that adults could live, it would be better if there existed instead some much larger number of human beings who lived the best lives that were possible for those who ceased to develop when they were newly born babies.

That is another repugnant conclusion. A similar argument applies to any single life. This argument assumes that

compared with living for some number of years, it would be better for each of us to live much longer, though with a slightly lower quality of life.

When applied to many such comparisons, this argument implies that

compared with living one of the best lives that adults could live, it would be better for each of us to live a very much longer life with only the pleasures of a never-developing baby.

We ought, I believe, to reject this Single-Life version of the Repugnant Conclusion. The best lives that adults could live are lexically better than any life, however long, with only the pleasures of a never-developing baby.

If we assume precision, it would be hard to reject these arguments. When we consider the Many-Lives version, we would have to claim that, compared with the existence of some human beings who would live the best lives that they could live, there would be no greater number of human beings whose existence could be better, if these other human beings would cease to develop when they were slightly younger. That claim would be hard to believe. But we should deny that these differences in value could be precise. We should claim that, of the worlds in this imagined continuum, many would be imprecisely equally good. We might, for example, admit that

(B) the best lives that could be lived by adults would not be lexically better than the best lives of those who ceased to develop at the age of 12, which would not be lexically better than the best lives of those who ceased to develop at the age of 11, and so on all the way down to the best lives of those who ceased to develop at the age of 3, or 2, or 1.

When there is imprecision *not better than* is not a transitive relation. That is how, even if your being a writer would not be better than your being a doctor, which would not be better than your being a slightly less successful writer, your being a writer *would* be better than your being a slightly less successful writer. So we could accept claim (B), but also claim that

(C) the best lives that could be lived by adults *would* be lexically better than the best lives that could be lived by those who ceased to develop when they were babies or young children.

Similar claims apply to other continuum arguments. These arguments assume that any slight loss in the quality of people's lives could be outweighed by a sufficient gain in the number of people who would exist and have lives that would be slightly less worth living. As we can more briefly say, any slight loss of quality could be outweighed by a sufficient gain in quantity. If we assumed precision, it would be hard to reject these arguments. We would have to claim that any slight loss of quality would outweigh any gain in quantity. As several writers claim, that would be very implausible. Compared with the existence of some number of people, it would not always be worse if instead there existed many more people who would have a slightly lower quality of life. But we should deny that such truths would be precise. We should then claim that no slight loss in quality would either be outweighed by, or outweigh, any such gain in quantity. It would not be better if there existed many more people whose quality of life would all be lower, since two such worlds would at most be imprecisely equally good. Though the larger of these worlds would not be worse, this relation is not transitive. So we could claim that it *would* be worse if, in other, larger worlds, everyone's quality of life would be *much* lower.

An argument fails if this argument has some premise whose rejection is less implausible than this argument's conclusion. These continuum arguments assume that

(A) compared with the existence of many people who would all have some quality of life, there are some much larger numbers of people whose existence would be better, though these people's quality of life would all be slightly lower.

As I have said, we could reject premise (A) by claiming

(D) such larger worlds would not be better, but would at most be imprecisely equally good.

This argument's conclusion is that

compared with the existence of many people whose quality of life would be high, there are some vast numbers of people whose existence would be better, though these people's quality of life would be much lower, since these lives would be barely worth living.

(D)'s way of rejecting premise (A) is, I believe, less implausible than this Repugnant Conclusion. If that is so, this argument fails.

We can next note that, if (A) seems more plausible than (D), that is because (A) asks us to consider only one of the many small steps that would take us all the way down to the Repugnant Conclusion. This argument is in this way like several other arguments with unacceptable conclusions. There are many acts each of whose effects, considered on their own, may seem too slight to make such acts wrong, though these acts would together make things go much worse. That is how we and others are together overheating the atmosphere, and destroying or polluting much of the natural world. We should suspect arguments that would lead us to ignore the combined effects of very many small differences.

There are, however, other arguments for the Repugnant Conclusion. Consider next the possible worlds that are too precisely shown in Figure 2. In World A there would exist N people whose quality of life we can call 100. In World Alpha, there would exist N people at the much higher level of 200, and there would also exist a million times as many other people whose lives would all be at level 1. In World Y there would exist as many people as in Alpha, who would all be at level 2.

Someone might argue:

(E) Since Alpha would contain N people whose quality of life would be much higher than that of all the N people in A, and the other people in Alpha would all have lives worth living, Alpha would be better than A.

(F) Since Y would contain the same number of people as Alpha, and these people would together have a much greater and more equally distributed sum of well-being, Y would be better than Alpha.

Therefore

(G) Y would be better than A. Compared with N people at level 100, it would be better if, instead, there existed some much larger numbers of people who would all be at level 2.




Figure 2.

(G) is another version of the Repugnant Conclusion.

We can call this the *Up Down Argument*. This argument's premises do not assume that it would be in any way better if there existed more people who would have lives worth living. That makes this argument harder to answer. These imagined worlds differ in three other ways. If we could create these worlds, we could get from World A, through World Alpha, to World Y by first raising the quality of N people's lives, then adding many extra people whose lives would be worth living, and then redistributing to produce a greater and more equal sum of well-being. We could make the world better, this argument claims, by *raising*, *adding* and *redistributing*.

This argument is valid. If World Y would be better than Alpha, which would be better than World A, Y must be better than A.

We ought, I believe, to accept that, as premise (E) claims, Alpha would be better than A. Alpha may be in one way worse than A, by involving inequality between different people. But the badness of this inequality would be outweighed by the facts that N people in Alpha would have a much higher quality of life than anyone in A, and that everyone else in Alpha would have lives worth living. It would have been worse, we should agree, if some of these people in Alpha had been much worse off in A, and everyone else in Alpha had never existed.

It may seem that we should also accept premise (F). We may assume that it would always be better if the same number of people would together have a greater and more equally distributed sum of well-being. I believe, however, that we can defensibly deny that Y would be better than Alpha. Some people in Alpha would have a much higher quality of life than anyone in Y, since this quality would be at level 200 rather than at level 2. This higher quality of life, we should assume, would not be merely a difference in the amount of well-being per person. At level 200, the best things in life would be very good, and lives at level 2, in World Y, would not include any of these good things. There would be no art, or science, no deep loves or friendships, no other achievements, such as that of bringing up our children well, and no morally good people. World Y would be much worse than Alpha in what we can call *qualitative* or *perfectionist* terms. In one version of this case, lives at levels 1 and 2 would be like the lives of never-developing one-year-old and two-year-old children. This great qualitative loss would, I believe, make Y in itself a worse world than Alpha, even though Y would give, to the same number of people, a greater and more equally distributed sum of well-being.

This view could take two forms. Suppose that we had the power to bring about one of these two worlds. We might believe that, because Y would be a worse world than Alpha, we ought to bring about Alpha, thereby making things go better. We might instead believe that, even though Y would be worse than Alpha, we ought to bring about Y, since we would cause the same number of people to have a greater and more equally distributed sum of well-being. On this view, we ought here to do what would make things go worse.²

² Many Act Consequentialists have believed that it would be always be better if the same number of people had a greater and more equally distributed sum of well-being. These people may find it easier to reject this belief if they ceased to be Act Consequentialists, and believed instead that, if we had to choose between producing world Alpha and producing world Y, we ought to produce Y, even though Y would be worse than Alpha. This would be one way of answering this argument for the Repugnant Conclusion. Though these people would cease to be pure Act Consequentialists, their revised view would be more plausible than the Repugnant Conclusion, which they might otherwise find hard to avoid.

On both these views, Alpha would be better than A, which would be better than Y. On the first view, if we could choose between Alpha and either or both of the other two worlds, we ought to choose Alpha. On the second view, if we had to choose between Alpha and Y, we ought to choose Y, because we would thereby cause the same number of people to have a greater and more equally distributed sum of well-being. If instead we had to choose between A and Y, or we could choose between all three worlds, we ought to choose A. Compared with Y, A is a world in which there would exist many fewer people, who would all have a much higher quality of life. And if we chose A rather than Y, we would not be causing the same number of people to have a smaller and less equally distributed sum of well-being.

I shall not here discuss whether, as only the second view implies, it would be wrong to choose Alpha rather than Y. We are asking whether we ought to believe that Y would be better than A, thereby accepting one version of the Repugnant Conclusion. On both these views, Y would be worse than A, and it would be wrong to choose Y rather than A. These views are both, I believe, more plausible than the Repugnant Conclusion. By appealing to one of these views, we could answer this Up Down Argument for this conclusion.

There is, however, a third argument which combines the arguments that we have been discussing. Consider the first three worlds shown in Figure 3. In World A, as before, there would exist N people at level 100. In Raised A Plus, there would exist N people at level 101, and a million times as many other people at level 95. In World B, there would exist the same number of people as in Raised A Plus, who would all be at level 99.

Someone might argue:

(H) Since N of the people in Raised A Plus would have a higher quality of life than all of the N people in A, and everyone else in Raised A Plus would have a quality of life that would not be much lower, Raised A Plus would be better than A.

(I) Since World B would contain as many people as Raised A Plus, and these people would together have, in B, a greater and more equally distributed sum of well-being, B would be better than Raised A Plus.

Therefore

(J) B would be better than A. Compared with the existence of N people at level 100, it would be better if there existed instead a million times as many people all at level 99.

(J) is one instance of the wider claim that




Figure 3.

(A) compared with the existence of many people who would all have some quality of life, there are some much larger numbers of people whose existence would be better, though these people's quality of life would all be slightly lower.

If we accept (A), we can then apply this claim to a long sequence of other possible worlds. Compared with the existence of many people at level 99, it would be similarly better if instead there existed many more people at level 98, and it would be similarly better if instead there existed many more at level 97, and so on all the way down to another version of World Z. Since *better than* is transitive, this argument implies the Repugnant Conclusion. Z would be the best of all these worlds, (E) implies, even though everyone in Z would have lives that were barely worth living.

The Continuum Argument, I claimed, fails because the rejection of premise (A) is less implausible than the Repugnant Conclusion. We cannot answer this third argument in this way. Rather than merely assuming (A), this argument defends (J) and the wider (A) by appealing to premises (H) and (I). That makes this argument harder to answer.³ To answer this argument, we must claim that we could defensibly reject either (H) or (I).

³ That is true in part because, while the Continuum Argument merely assumes a claim like (R), (P) and (Q) are more plausible than (R) and support (R).

We could not defensibly reject (H). Since N of the people in Raised A Plus would have a higher quality of life than all of the N people in A, and everyone else in Raised A Plus would have a quality of life that would be not much lower, we could not defensibly deny that Raised A Plus would be better than A.

We might, however, reject premise (I). Though World B would be better than Raised A Plus in utilitarian and egalitarian terms, B would be worse in qualitative terms, since the best things in people's lives would be worse in B. We might claim that

(K) given the conflict between these values, Worlds B and Raised A Plus are only imprecisely comparable, and would be imprecisely equally good.

This claim would not be in itself plausible. There is a contrast here with my suggested answer to the Up Down Argument. In denying that Y would be better than Alpha, we could appeal to the fact that Y would be qualitatively much worse, since in a change from Alpha to Y the best lives would fall from level 200 all the way down to level 2. There would be no such great qualitative loss in a change from Raised A Plus to B, since the best lives would fall only from level 101 to level 99. Since this qualitative loss would be so slight, we may doubt that it could offset the greater rise in the quality of very many other people's lives, thereby making World B, not better than Raised A Plus, but only imprecisely equally good.

The important question, however, is not whether claim (K) is in itself plausible. In deciding whether we could defensibly make this claim, we should consider how (K) is related to some other claims. Suppose we believe that, as premise (H) claims, B would be better than Raised A Plus. Suppose next that we also believe that, as premise (I) claims, Raised A Plus would be better than A. We would then have to admit that B would be better than A. It would be better if there existed many more people who would all be slightly worse off. This argument could then be reapplied to other larger worlds with lower qualities of life, carrying us all the way down to the Repugnant Conclusion.

An argument fails, as I have said, if this argument has some premise whose rejection is less implausible than this argument's conclusion. We could reject this argument's premise (I) by claiming (K). Though (I) is not in itself plausible, our main question should be whether this way of rejecting premise (I) is less implausible than the Repugnant Conclusion.

The answer, I believe, is Yes. (K) seems implausible because, in a change from Raised A Plus to B, there would be only a slight qualitative loss. The best lives would fall only from level 101 to 99. It may be hard to believe that this slight qualitative loss could make Raised A Plus not better than B, but only imprecisely equally good. But it would be much harder to believe that, compared with the

existence of many people whose quality of life would be very high, it would be better if there existed instead some vast number of people whose lives were barely worth living. Since (K)'s way of rejecting premise (I) is less implausible than the Repugnant Conclusion, this argument fails. By appealing to this Imprecisionist Lexical View, we can justifiably reject this conclusion.

We can justifiably make other, wider claims. We can believe that, compared with the existence of many people whose quality of life would be high, it would be worse if there existed instead many more people whose lives would be well worth living, but whose quality would be much less good.

Acknowledgements

This paper is the text of a talk. When I publish a longer and revised version of this material, I shall thank those whose comments have helped me. But I shall acknowledge here the great help that I have been given by Jake Nebel and Jimmy Goodrich.