How We Can Avoid the Repugnant Conclusion

Short Talk

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 total 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 the Simple 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. Since the existence of all such people would have some value, it would always be better if more such people existed, and the combined goodness of these people's existence would have no upper limit. These claims seem to imply that, if there were enough people in the world in which everyone's lives would be barely worth living, that would be better than the world in which everyone's quality of life would be very high. That is the Repugnant Conclusion.

Several writers have reluctantly accepted this conclusion. This conclusion seem repugnant, some of these people suggest, 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, we should admit, some cases of this kind. Many people underestimate how great the effects of natural selection might be during a million years. But that is not why most of us would find the Repugnant Conclusion so 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 people claim that, to avoid such repugnant conclusions, we should

reject the whole idea that some possible outcomes, or worlds, could be better or worse than others. On this

> *Skeptical View*: Compared with the existence of many people who would all have wonderful lives, it would not be worse if there existed many more people whose lives would be full of suffering and much worse than nothing. This second world would not be worse because it makes no sense to claim that some possible worlds would be worse than others, or because, though such claims make sense, no such claim could be true.

If we denied that this world full of suffering would be worse, however, we would be accepting another Repugnant Conclusion.

Some other people try to avoid the Repugnant Conclusion by rejecting the Simple View and making other claims about how the goodness of anyone's existence depends on this person's quality of life, or level of well-being. According to what some writers call

> *the Critical Level View*: 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*.

This view implies that what we can call

> the Less Repugnant Conclusion: compared with the existence of many people whose lives would be very good, it would be better if there existed instead some much larger number of people whose quality of life would be just above the Critical Level.

To assess this view, we must know how good lives at the Critical Level would be. When John Broome defends this view, he suggests that such lives might be 'a lot better' than lives at the Zero Level, and be 'reasonably good'. This view might then be significantly less repugnant than the Repugnant Conclusion.

This view takes two forms. When people's lives are worth living, or above the Zero Level, but are not above the Critical Level, these lives are in what we can call the *Critical Range*. According to

> *the Neutral Critical View*: When people's lives are in this range, their existence has no intrinsic value, or does not in itself make the world either better or worse.

On this view

> when people have some level of well-being that is in the Critical Range, it does not matter what that level is. Since the existence of such people would have no value even if their lives were at the Critical Level, being in Broome's phrase reasonably good, it would not be worse if these people were all worse off by having lives that were at bottom of the Critical Range, being not worth living.

We might call this *the Callous Conclusion.*

Some Critical Level Theorists claim instead that

> since there is no value in the existence of people whose lives are at the Critical Level, and it would be worse if these people were worse off, these people's existence would then have *disvalue*.

According to this

> *Negative Critical View:* When people have lives at the Critical Level, their existence is not in itself good. When people are worse off, by having lives that are lower in the Critical Range, their existence is in itself bad.

Though this view avoids both the Repugnant and Callous Conclusions, it is hard to believe. If lives at the Critical Level would be reasonably good, that would also be true of lives that were just below this level. Such lives, Broome adds, 'may contain no suffering'. When people live such lives, why is their existence in itself bad?

This view also has implausible implications. Suppose that, in one of two worlds, there would be many people whose quality of life would be high and some other much larger number of people whose quality of life would be just below the Critical Level, and therefore reasonably good. In a second possible world, there would exist only many people whose lives all contained much suffering and were much worse than nothing. On this Negative Critical Level View, if the first of these worlds had enough people just below the critical level, this world would be worse than the second. This view implies that

> compared with the existence of both many people who would all have very good lives, and very many more other people whose lives would be reasonably good, it would be better if there existed instead many people whose lives would be full of suffering, and much worse than nothing.

Gustaf Arrhenius has called this *the Sadistic Conclusion*. That name is unfair. Though some Critical Level Theorists have accepted this conclusion, that is not because they get pleasure from imagining the existence of people whose lives are full of suffering. These writers all believe that it would be bad if there existed people with such wretched lives. They accept this conclusion only because they believe that it would be even worse if there existed many more people whose lives, though containing no suffering and being reasonably good, would be just below the Critical Level.

Though not sadistic, this view is incredible. Though the Repugnant Conclusion is hard to accept, we can understand how that conclusion might be true. We can see how, if there existed enough people whose lives would be worth living, though only minimally good, that might be better than a world in which there existed many fewer people with wonderful lives. When we are comparing things that are all good, some losses of quality can be compensated, or made up for, by sufficient gains in quantity. But it is hard to see how, if there existed enough people whose lives would be

reasonably good, and many other people whose lives would be very good, that would be *worse* than a world in which there existed many fewer people all of whose lives would be much worse than nothing. Though enough of what is minimally good might be better than smaller amounts of what is very good, how could enough of what is reasonably good be worse than smaller amounts of what is very bad?

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

> *the Diminishing Value View:* 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.

This view is also in itself implausible. We might justifiably make such claims about some of the things that can make a single person's life good. These things may have less value, doing less to make this life better, when they are repeated. But it is hard to see how the existence of more people with lives worth living might have such diminishing value.

This objection is clearer when we consider people whose lives would be full of suffering, and much 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 between some of the other things that can make lives good and bad, we cannot plausibly either apply the Diminishing Value View to lives that are bad, or restrict this view to lives that are good.

There is a better view, which I shall call the *Imprecisionist Lexical View*, or *V*. We should accept the Simple View, claiming

> 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 should add

> If there existed many people who would 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.

These claims together assert one kind of *lexical superiority*. When I say that things of kind *P* are *lexically better* than things of kind *Q*, I shall mean roughly

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.

This Lexical View may seem implausibly extreme, since it claims that though such less good lives 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 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 also claimed to be open to decisive objections.

To answer these objections, we should first defend a view that we can call *Normative Imprecisionism*. On this view, there are often no precise truths about the relative goodness of different things. There can be fairly precise truths about the relative value of some things. One of two painful ordeals, for example, might be 2.3 times as bad as the other, by involving pain of the same intensity for 2.3 times as long. In most important cases, however, relative value does not depend only on any such single, measurable property. When two pains 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 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 two hours of amusement. When we are choosing between different ways in which our life might go, by choosing between two careers, for example, or deciding whether to have children, there are only imprecise truths about which life would be better. And there are only imprecise truths about the relative goodness of most different acts or outcomes, such as acts that would benefit a few people, or give lesser benefits to many others. Such imprecision is not the result of 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 that one of these things is better than the other by some precise amount, or that both things are precisely equally as 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 is 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 couldn't 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 as great. As you would see, the truth could be only that one of these people was imprecisely greater than the other, more plausibly, that they were imprecisely equally as 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 all be precise because positions 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 couldn't be true. So when we think about the goodness of such things, we should reject this Linear Model.

Nor could the differences between the goodness of such things correspond to different numbers, since numbers are also precise. We might try to avoid precision by thinking, not of a single number, but of a range, such as 80 to 90. But if some other range is 81 to 91, that difference would be precise. Since most things can be better or worse than others by more or less, or by different amounts, it may help to use scales or numbers, as I have been doing, as rough ways to describe these differences. But we should take care not to be misled.

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 can 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 Y. 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 is precisely equally as good as Y, which is precisely equally as good as Z, X must be precisely equally as good as Z. But if X and Y are *imprecisely* equally as good, neither is better than the other, and these imprecise relations are *not* transitive. Even if X is not better than Y, which is not better than Z, X may be better than Z.

Such facts can be used to illustrate and explain imprecise equality. Two things are imprecisely equally as good if it is true that, though neither 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. If your life could go in different ways, it might be true that 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. But your being a writer *would* be better than your being a slightly less successful writer. *Not better 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 that the second, and in no way 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 thing can better in one way and not worse in other ways without being better all things considered.

If we reject such claims, and believe that even when there is imprecision *not better than* must be transitive, we may still be thinking in terms of the Linear Model, or assigning numbers. Like some of the other most 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 by amounts that 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 *aren't* like such distances because they are not precise. We should think about such differences without thinking of a line, or scale of value.

There are many kinds of difference which can make two things less precisely comparable. There is one such difference, I believe, between the possible outcomes or worlds that we have been considering. We should claim

> When two possible worlds would contain different numbers of people, this fact makes these worlds less precisely comparable.

We can call this *different-number-based imprecision*.

I have time only to sketch briefly why I believe that there is such imprecision. Some people believe that, when we consider possible future states of the world, or what I am calling possible worlds, it is enough to appeal to some familiar person-affecting principles. On one such view, it would not be worse if fewer people later existed, since the non-existence of these people wouldn't be worse for them. But this person-affecting principle fails even when applied to the easier comparisons between worlds in which the same number of future people would exist. We might know, for example, that if our society follows one of two economic policies, that would greatly lower the quality of people's lives in the further future, but that these effects would not be worse for any of these people, since if we had chosen the better policy these particular future people would never have existed. This fact could not be enough to justify our lowering the future quality of people's lives. We must therefore appeal to other, less familiar principles, and it is far from clear what these principles should be.

Suppose for example that, in world A, there would exist some larger number of people who would all have a high quality of life. In A Plus, there would also exist some extra people whose quality of life would be much lower, though these people would have lives worth living. On the Utilitarian Total

View, A Plus would be better than A, since there would be a greater total sum of well-being. But this view leads us to the Repugnant Conclusion. On the Utilitarian Average View, to which many economists have appealed, A Plus would be worse than A, because the existence of the extra people would lower the average quality of life. On this view, it would have been better if these extra people had never existed, even though their lives are worth living, and their existence is not bad for anyone else. That is also hard to believe.

Consider next a third world, Improved A Plus, in which there would exist the same number of extra people, but these people would all have a higher quality of life. Improved A Plus is clearly better than A Plus. It is much less clear, hower, how these worlds compare with A. We can plausibly believe that A Plus would not be worse than A, since it would not be better if these extra people with lives worth living had never existed. It may seem that, since Improved A Plus would be better than A Plus, which would not be worse than A, Improved A Plus must be better than A. But we are not, I believe, forced to that conclusion. We can believe that, though A Plus would not be worse than A, that does not imply that A Plus would be at least as good as A. We should instead conclude that, since these worlds contain different numbers of people, they are less precisely comparable. We can then coherently believe that, though Improved A Plus would be better than A, which would not be worse than A, Improved A plus would not be better than A. These two worlds are also, believe, only imprecisely comparable. Though Improved A Plus would be better than A Plus, both these worlds would be neither better nor worse than A. That may not seem to you obviously true. But, if these worlds are only imprecisely comparable, as I believe, we should expect that no claim about their relative value would seem obviously true.

We can next describe more fully one of the implications of different-number-based imprecision. We ought, I have claimed, to accept

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

This claim implies that, whenever extra people exist whose lives are worth living, that would make the world in one way better. This would be true, I believe, even if these people's lives would be only barely worth living. But if these people would have lives that were more worth living, their existence would do more to make the world better, giving the world what we can greater *added value.* I have also claimed

> When two possible worlds would contain different numbers of people, that makes these worlds less precisely comparable.

As well as adding value to the world, the existence of such extra people always creates what we can call an extra *margin of imprecision*.

There are now two possibilities. The amount of added value may be less than the increased margin of imprecision. When that is true, though the existence of these extra people would make the world in one way better, it would not make the world better all things considered. The resulting world would instead be imprecisely equally as good as the world in which these people never exist. Suppose next that these people's lives would not merely be worth living, but would be very good, being above what we can call the

High Level. The amount of added value would then be greater than the increased margin of imprecision. If the existence of these people with very good lives would not be bad in other ways, their existence would make the world, not only in one way better, but also better all things considered.

These claims may resolve the conflict between two widely held views. Suppose that, if we had some extra child, this child's life would be worth living, and his or her existence would not be on balance worse for us or other people. If we ask whether such an extra person's existence would make the world better, some people take the answer to be obviously Yes, but others take it to be obviously No. Both sides may be partly right. Those who answer Yes may believe truly that, if such an extra person exists, that would make the world in one way better. Those who answer No may believe truly that this person's existence would not make the world better all things considered. These people may disagree because they ignore what I have called different-number-based imprecision. They may assume that, if such an extra person's existence would make the world in one way better, and would have no bad effects, this person's existence must also make the world better all things considered. That, I have claimed, is not true.

If these claims are true, the Simple View does not imply the Repugnant Conclusion. Even if the existence of extra people with lives worth living would always make the world in one way better, if these people's lives were below the High Level, their existence would not make the world better all things considered. We can therefore deny that any loss in the quality of people's lives could be outweighed by the existence of enough extra people, either with lives that are barely worth living, or even with lives that were much better than that, being above the Critical Level and more than reasonably good. These truths would partly solve our problem.

There are, however, some other arguments for the Repugnant Conclusion.

On this Imprecisionist Lexical View.

> If there existed many people who would 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 to live.

Most writers on population ethics either ignore or quickly dismiss such lexical views. Of those who dismiss such views, many believe that all truths about relative goodness are precise, so that something's goodness corresponds to its position on some scale, or to some number. On that assumption, lexicality makes no sense. If certain lower goods had such non-diminishing value, but the combined goodness of these things could never be as great as that of certain higher goods, that would imply both that

> there are no limits either to how high on the scale this combined goodness could be, or to how great the number might 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 higher goods.

That is a contradiction. As this shows, we cannot represent the claims of any lexical view by thinking in terms of positions on some scale, or by using numbers. But that does not refute such views. We are trying to reach true beliefs about what is better or worse, and about what we have most reason to want, or do, and what we ought to do. We should not assume, in advance, that such truths must be able to be represented by using scales or numbers.

Many 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, no such view could be defensibly applied applied to things which could together form one continuous series, each differing only slightly from its neighbours. Lexical superiority, these people claim, requires some deep and sharp discontinuity.

That is not, I believe, true. It may be enough if there is imprecision. To defend this belief, I shall now discuss what we can call *the Continuum Argument* for the Repugnant Conclusion. This argument's main premise claims

> Compared with the existence of many people who would all have lives that were equally worth living, there is some much larger number of people whose existence would be better, though these people would all have lives that would be slightly less worth living.

Consider a sequence of such possible worlds. Compare with the existence of population A, there must be some much larger population B whose existence would be better, thought these people would all be slightly worse off than the people in A. Compared with B, there must be some much larger population C, whose existence would be better, thought these people would all be slightly worse off than the people in A. And so on down to Z, an enormous population whose lives are barely worth living. Since 'better than' is a transitive relation, if B is better than A, C is better than B, and so on down, there must be some enormous population whose existence would be better than A, though these people would all have lives that were only barely worth living

When we describe a series of possible worlds like A, B, C, and so on down to Z, there are different ways in which people might be worse off, by having lives that were less worth living. The good things in these people's lives might be less good, or it might instead be true that, without any such loss of quality these people would be worse off because they would each have many fewer of these good things. That would be most simply true if, in later worlds in our imagined series, as the numbers of people grew, the length of these people's lives would shrink. In what we can call *Short-Lived Z*, our imagined people might live for only as long as some flowers bloom. If these imagined beings lived for only one happy day, or one ecstatic hour, these brief lives could not be very good, but we could not plausibly describe these lives as barely worth living. Though this version of World Z raises interesting questions, I shall not discuss these today.

In what we can call *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. These people's agonies would be only just outweighed by their ecstasies. This version of Z would also raise questions that I shall not discuss today. I shall assume that World Z would take a third, simpler form.

In what I have earlier called *Drab Z*, there would be nothing that would be bad in these many people's lives, 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 not enough. If the people in Drab Z would be in other ways like us, we could not plausibly assume that all 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 some long poem, for example, or making progress in thinking about some mathematical, philosophical, or practical problem. Some political prisoners have lived fairly good lives while they were in deprived 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 in such deprived confinement, but because the beings who would live these lives would be psychologically much simpler than us. This version of Drab Z also helps to show the full implications of the Continuum Argument.

It is often claimed, I have said, that for some things to be lexically better than others there must be some deep discontinuity between these two kinds of thing. In the development of each our own lives, from birth to adulthood, there is no such discontinuity. Nor is there such discontinuity in the evolutionary sequence between human beings and their many ancestors. In both these developments there are striking qualitative changes, some of which involve wholly new features. But such changes are gradual. These sequences are like some rope of which one end is made of fibre and the other made of gold, with overlapping strands in between of fibre, copper, bronze, silver and gold.

We can first consider the sequence between us and our evolutionary predecessors, such as pre-humans, earlier primates, and many other kinds of animal all the way back to some of the earliest sentient animals. According to the main premise of what we can call the *Evolutionary Continuum Argument*:

> Compared with the existence of many animals of some kind who would all have lives worth living, there is some larger number of animals of the preceding kind whose existence would be better, if these animals would also have lives worth living.

This argument implies

> *the Third Repugnant Conclusion*: 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 the earliest sentient animals who would all have 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.

Here is another similar argument. Human babies slowly develop throughout 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. According to the main premise of what we can call the *Age Continuum Argument*:

> 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 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 one day younger.

If we apply this premise to a long series of possible human beings who would cease to develop at slightly earlier and earlier ages, this argument would imply

> *the Fourth Repugnant Conclusion*: Compared with the existence of many adults who would live their best possible lives, 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 newly born babies.

This argument's main premise may seem undeniable. It may seem absurd to claim that, compared with the existence of many human beings who would live their best possible lives, it would *not* be better if there existed any much larger number of human beings who would also live their best possible lives, even though these other human beings would cease to develop when they were younger by only a *single day*.

This argument's main premise is one version of the claim that

> compared with the existence of many people who would all have lives that were equally worth living, there is some larger number of people whose existence would be better, though these people would all have lives that were *slightly* less worth living.

This premise may seem undeniable because it claims only that there is *some* larger number of people, all slightly worse off, whose existence would be better. It may seem that, since the existence of each these extra people would be in itself good, there would be no limit to the combined goodness of these people's existence, so that, with enough such extra people, this goodness must be able to outweigh the fact that these people would all be slightly worse off.

If this reasoning seems compelling, that is because we assume that these possible worlds would be precisely comparable. That, I have claimed, is not so. The existence of these extra people would, I believe, both make the world the resulting world in one way better, and make this world less precisely comparable with the world in which these extra people did not exist.

Though the existence of these people would add value to the world, if the quality of their lives would not be above what I call the *High Level*, as this continuum argument needs to suppose, the existence of these extra people would not make the world better all things considered. We can therefore appeal to the claim that there is what I call different-number-based imprecision. We can say:

> Each of these larger worlds would be in one way worse than the neighbouring world with fewer people, since everyone would be slightly worse off. These worlds would be in another way better, since there would exist many more people with lives that were only slightly less worth living. But though the existence of more such people would always make the world in one way better, this added value would always be less than the increased margin of imprecision. The existence of these people would never make the outcome better all things considered. So it would not be true that if there were *enough* such people, the goodness of their existence would outweigh the fact that everyone would be worse off.

If this Different Number Based Imprecisionist view is true, as I am inclined to believe, we can reject the main premise of the Age Continuum Argument. We can similarly reject what I called the Evolutionary Continuum argument, and any other similar argument.

There are some other arguments for the Repugnant Conclusion which are, in some ways, harder to answer. I have no time to discuss these today. I shall end by saying only that these other arguments can all, I believe, be answered in similar ways. There would always be at least one way of rejecting one premise of such arguments which, even if it seems implausible, would be less implausible than the Repugnant Conclusion.

We have been comparing worlds in which

> there would exist many people who would all have some very high quality of life,

or

> there would exist instead many more people whose lives would be barely worth living, or many more human beings who ceased to develop when they were newly born babies, or many more sentient animals who had only slight pleasures like those of munching grass and basking in the sun.

We can justifiably believe that the existence of these other worlds would not be better, but much less good.

6, 797