

Should we value population? – John Broome

The intuition of neutrality

What good does it do you to continue living? Or – an equivalent question – what harm would it do you to die now? These questions have exercised philosophers since antiquity,¹ and they are not mere philosophers' questions. An answer to them is important in practice. An answer might help you when you are wondering whether to take the risk of hang-gliding, and it would help a government that is wondering whether to increase its investment in health.

Here is an answer. The benefit to you of continuing to live – and equivalently the harm to you of dying now – is the difference between the overall goodness your life will have if you continue to live and the overall goodness it will have if you die now. This answer is hard to fault. Given one way of using 'benefit' and 'harm', this difference is simply what we *mean* by the benefit of continuing to live and the harm of dying now.

The answer leaves a lot open, because it does not specify what the overall goodness of your life will be, either if you continue to live or if you die now. One might take many views about that. For example, one might think a good life has to contain some worthwhile accomplishments, or one might think the goodness of a life is the total of the good things it contains. For my purposes here it does not matter. I only need to assume there is such a thing as the goodness of a life. Various events may affect it, and death is one. Often when a person dies, her life is less good than it would have been had she continued to live. In that case, death harms the person. Conversely, if her life is extended, that is good for her. I shall call the goodness of a person's life her 'lifetime wellbeing'.

Now think about a case where extending a person's life has the effect of adding new people to the population of the world. For example, suppose the roads are made safer, with the result that some people are saved from an early death. Many of these people will subsequently have children, who would never have existed had their parent been killed on the roads. Many of the children will have children of their own; indeed many of them will start a whole line of descendants. In a case like this, what are the benefits of saving the original people from an early death?

First there is the benefit to those people of having their lives extended; this is the amount their lives are made better by being extended. But what about the goodness of the new lives that result? Each of those lives will be good to some degree. Should we not count each of them as a good thing (or possibly a bad thing) too, to add to (or subtract from) the good of extending the existing people's lives?

Most of us have the instinct that we should not. We think it a good thing to extend the lives of existing people, but we do not think it a good thing to bring new people into existence. We value the benefit brought to the original people by extending their lives, but we think the lives of the new people are no added benefit beyond that.

You might think you have an easy explanation of this instinct. You might think it is obviously not a good thing to have extra people, because we have enough people in the world already. Each new person makes demands on the earth's resources, leaving less for the rest of us. So, on average, each new person diminishes the lifetime wellbeing of the rest of us. This harm a new person does to existing people may well outweigh the lifetime wellbeing of the new person herself. For this reason we should not be pleased to have new people in the world.

That seems plausible. But now you have to explain why you are in favour of keeping alive the people we already have. Most of us are in favour of that, and as a society we put a lot of

effort into it, in the health service and elsewhere. If the demands on resources made by new people outweigh the good enjoyed by those new people, the same must on average be true of existing people too. Each year, each person who continues to live requires resources, and the resources she uses are not available for the rest of us. We should be happy to be rid of people, rather than work so hard to keep them alive. But we do not think like that. Our instinct is to make a great difference in value between the lives of new people and existing people. We are in favour of prolonging the lives of existing people, but we are not in favour of creating new lives. Recognizing the shortage of resources does not explain why we make this difference.

I think the source of our instinct is not a shortage of natural resources. Instead, it is simply that we think the lifetime wellbeing of a person who is added to the world is in itself ethically neutral – to use a vague expression that I shall soon tighten up. It is neither a good thing nor a bad thing. We recognize that adding a new person may be good or bad for existing people; a new person may bring enjoyment to her family and to other people and she may also bring harm to other people by her demands on resources. We count these benefits and harms in favour of this person's existence or against it. But the goodness of the person's own life does not count ethically.

Our moral instinct is to care about people, and to care about making people's lives go well. This is caring for the people there are; we want their lives to go better. We have no natural interest in having more people about. A famous remark of Jan Narveson's puts it succinctly: 'We are in favour of making people happy, but neutral about making happy people'.²

I think that is our commonest intuition, and I share it. I call it 'the intuition of neutrality'.

Terminology and interpretation

What does 'ethically neutral' mean? The term might have various interpretations. To say that doing something is ethically neutral could mean it is not the case that you ought to do it, and also not the case that you ought not to do it. This may be called a 'normative' or 'deontic' interpretation of 'ethically neutral', but it is not the interpretation I shall adopt. Instead I shall interpret the term evaluatively. When I say something is ethically neutral, I shall always mean it is neither good nor bad. The connection between the evaluative and the normative is contentious, and not a subject for this paper. In this paper I shall concentrate on goodness only. (I use the word 'value' as a synonym for 'goodness'.)

I shall identify the goodness of adding a person to the world with the goodness of the state of affairs that results from adding the person: that is to say, with the goodness of having an extra person in the world. This identification may be disputed. It is an application of a debatable view I call 'consequentialism',³ but I do not need to debate it in this paper. I am really interested in the goodness of states of affairs, and not in the goodness of actions such as adding a person. So I interpret the neutrality intuition to be the view that the presence of an extra person in the world is neither good nor bad. More precisely: a world that contains an extra person is neither better nor worse than a world that does not contain her but is the same in other respects. Nevertheless, I shall continue to speak loosely of the goodness of adding a person, just because it is a convenient expression. I assume it is the same as the goodness of the resulting state of affairs. But if I am wrong about that, it is the goodness of the state of affairs that I really mean to refer to.

These points of interpretation mean that I am treating the intuition of neutrality as an intuition about the goodness of states of affairs. That may misrepresent it. For example, our common intuition may be fundamentally normative, and not about goodness at all. If that is so, it will not damage the conclusion of this paper. I shall eventually conclude that the neutrality intuition, interpreted my way, has to be rejected. If it is not our common intuition

in the first place, no matter. However, I at least find the intuition very attractive, even when it is interpreted in my evaluative way. So I think it is worth the trouble of showing that it is ultimately mistaken.

The intuition is that adding a person to the population is ethically neutral. By this, I really mean that a world that contains an extra person is neither better nor worse than a world that does not contain her. It follows that a world that does not contain this person is neither better nor worse than a world that does. I can express this by saying that subtracting a person from the population is also ethically neutral. By ‘subtracting’ I mean not bringing into existence a person who would otherwise have existed. I do not mean removing from the population a person who already exists; that would be killing. As I mean it, to say subtracting is ethically neutral is just another way of saying that adding is ethically neutral.

The practical importance of the intuition

The intuition of neutrality is not merely an intuition we happen to have. It is deeply embedded in the way we think about value and the way we form our moral judgements. We generally simply ignore the effects of our actions on the world’s population, even when the effects are perfectly predictable. This can only be because of the intuition that they are ethically neutral. If the intuition turns out to be incorrect, it will make a huge difference to the way we should make our judgements.

Take the example of programmes for improving people’s safety. When people’s lives are saved, by making roads safer or in other ways, the wellbeing of the people who are saved is generally small in comparison to the wellbeing of all the new people, their descendants, who come into existence as a result. This is perfectly predictable. If all the descendants’ wellbeing had to be counted too, that would enormously alter the value we attach to saving people’s lives. But actually, in judging the value of safety on the roads, we routinely ignore all this wellbeing. Economists regularly make these valuations, but I do not know of a single one who has taken the wellbeing of descendants into account.⁴

If the intuition is wrong, it is not just judgements on life saving that we would have to rethink radically. A great many things we do have an influence on the future population of the world. For example, changes in the social security system or the education system influence the costs and benefits of having children, and will have some effect on people’s decisions about having them. The effect on immediate numbers may be small, but a small change is likely to be perpetuated. A few extra people now means some extra people in each generation through the future. There does not appear to be a stabilizing mechanism in human demography that, after some change, returns the population to what it would have been had the change not occurred. If our intuition is wrong, and those extra people have a value, either positive or negative, we can expect it to dominate the value of the effects on existing people, just because it continues effectively for ever. When we judge the social security system or the education system, we normally ignore its demographic effects. It is the intuition of neutrality that justifies us in doing so. If this intuition turns out incorrect, we shall have to do some radical rethinking.

Here is another major example of the intuition’s practical importance. Global warming will kill huge numbers of people. It will kill them through at least three processes. The first is extreme weather and its effects. Sea levels are rising and storms seem to be becoming more violent, so devastating floods will become more frequent. Also in this category comes drought with its consequent famine. The second process is that rising temperatures will extend the range of tropical diseases. The third is the direct effect of heat, particularly in cities: heat waves kill people. Cold waves kill people too, and there will be fewer of those.

But on balance the number of deaths from heat and cold will increase.

It is easy to know there will be a lot of killing. To know how much is harder. I am not a scientist, so all I can do is look in the reports for figures, and they contain very few. The 1995 report from the Intergovernmental Panel on Climate Change (IPCC) – this is not their latest report – mentioned the figure of 215,000 deaths per year just from the third cause, starting from about the middle of this century.⁵ This is a big number. But people who die in heat waves are generally very ill in any case. A heat wave generally advances their death by only a short period. I doubt this will be the major part of the harm from killing. As a guide to the magnitude of killing by flooding and weather: on average in recent decades, 123,000 people per year have died in natural disasters.⁶ (This number does not include famines, but it does include disasters, such as earthquakes, that are not caused by weather.) Global warming will increase this number. As a guide to the magnitude of killing by tropical diseases: about two million people per year die from malaria alone. Malaria's range will increase, but not by as much as you might expect, because in most of the world outside Africa malaria is not controlled by temperature as much as by public health measures. I think we may reasonably suppose global warming will cause hundreds of thousands of extra deaths per year, starting a few decades from now.

The killing will continue into the indefinite future. When people try to predict the effects of global warming, they usually give up beyond the end of this century. But the effects will continue for much longer. The latest IPCC report shows what to expect.⁷ Even its most optimistic predictions show the global average temperature still rising in 2100, and they show global sea levels not only still rising but still accelerating at that date. Global warming will have effects for centuries. It will kill hundreds of thousands of people a year for a period that we can see no end to. It amounts to tens of millions of people killed, which is killing on the scale of a world war or a dreadful global epidemic. By any sort of accounting, this is one of the worst harms global warming is going to do.

We can reduce global warming, but only by a sacrifice. We shall have to stop using energy in our profligate way. We shall not be able to travel around so much, and people who live in hot places will not be able to have such cool houses. These are genuine reductions in our wellbeing, and their purpose will be to save people's lives in the future. The question arises: how much should we sacrifice? What degree of sacrifice is worthwhile to reduce the killing? I cannot say this is the most pressing question about global warming at present. The most pressing question is how the US, which is the main cause of global warming, can be persuaded to make any significant sacrifices at all. But even for that political purpose it might help to get some account of the amount of harm that global warming will do.

Besides killing a lot of people, global warming will affect the population of the world. Of course, killing people directly affects the world's population in a sense: at a time after they have been killed, the world contains fewer people than it would have contained had they survived. But I do not mean that. That is an effect on the world's population at a particular time: the time after the killing. I am speaking of the world's *timeless* population. By this I mean all the people who exist at some time in history. The timeless population includes Julius Caesar, me, and all the people who are yet to be born. Killing a person does not remove her from the timeless population.

But although the killing global warming will do does not directly reduce the timeless population, it will affect the timeless population less directly. For one thing, when a person is killed, all the children she would later have had, if she had lived, and all their descendants, are removed from timeless population. Also, think of this. Vast numbers of people are going to have to migrate across the world as some areas become uninhabitable. In Bangladesh, ten

million people live within one metre of sea level, vertically. When the sea level rises one metre, those people will have to go somewhere else (unless Bangladesh is protected by a sea wall like Holland's). So there will be vast movements of population, and migrations on that scale do not happen without an effect on the size of the population. Global warming will add to the population, or subtract from it. I shall not even predict in which direction the population will be affected – up or down – but I predict it will change.

At first, the problem of global warming seems to be a problem of weighing sacrifices made by present people against the saving of people's lives in the future. But now we can see it is more complicated. Global warming will alter the size of the population. Do we need to take this effect into account in evaluating it? Our intuition is that the size of the population is ethically neutral, because we think adding people to the population, or subtracting people from it, is neutral. But if we are wrong about that, and adding or subtracting is either good or bad, then this goodness or badness threatens to swamp the badness of the killing itself. Tens of millions will be killed. But each person is only the beginning of a whole line of descendants, so the increase or decrease in the population will almost certainly outnumber the people who are killed.

A difficulty for the intuition

It is only the neutrality intuition that saves us from having to undertake the extraordinarily difficult job of evaluating the changes in population that global warming will cause. Only this intuition allows to be confident even that global warming is a bad thing; without it, it might turn out that its effect of population outweighs the badness of all the killing it will do. But I am sorry to say I am now going to argue that the neutrality intuition cannot be sustained. It has to be abandoned, however painful that might be. At least, it has to be abandoned if it is interpreted in terms of goodness. It may have a defensible interpretation in normative terms, but it has no place in *evaluations* of global warming or anything else.

To make my argument, I am going to concentrate on a stripped-down example. It will be implausible in some respects. Philosophy often requires implausible examples. I want to concentrate on just one thing: the value of having a single extra person in the population. To do that, I want to exclude from the example all complicating factors, which might affect value in other ways. So I shall compare two alternative possible states of affairs that are the same in all respects, except that in one of them a person exists who does not exist in the other.

Figure 1 illustrates the problem. One thing we shall obviously have to take into account in assessing the value of having a new person is how well off that person is. We shall certainly be interested in her lifetime wellbeing. So I have shown her lifetime wellbeing, and other people's too, in the diagram. It is measured vertically. The diagram shows two alternatives, *A* and *B*. In *A* the extra person does not exist; in *B* she does. Everyone apart from the added person is made neither better nor worse off by this person's existence. That is implausible, of course. One would expect her parents to be affected, if no one else. But that is my implausible simplifying assumption. I shall consider different possible levels for the extra person's wellbeing.

Our intuition of neutrality is that adding the person is ethically neutral. This intuition is to some extent independent of her level of wellbeing; there is not just one level of wellbeing of which we think intuitively that adding a person at that level is ethically neutral. To be sure, we do not think this of every level. We think it is a bad thing if a person lives a poor life, full of suffering. We are not neutral about adding a person whose life is like that; we are against it. Conversely, some of us may think it is a good thing if a person lives an extremely good

life; they are not neutral about adding a person at an extremely high standard of living. But for a wide range of levels of lifetime wellbeing between a bad life and a very good life, we intuitively think that adding a person at that level is neutral. The intuition is that there is a neutral *range*: a range of levels of wellbeing such that adding a person whose wellbeing is in that range is ethically neutral. The range may extend upwards to infinity, but in figure 1 I have shown a finite neutral range.

Adding a person whose wellbeing is in the neutral range is neither better nor worse than not adding her. That is my definition of neutrality. We may naturally suppose, that if adding her is neither better nor worse than not adding her, it is equally as good as not adding her. But if there is a neutral range, that cannot be so. Figure 2 shows why not. Compare *B* with *A*. The only difference is that in *B* a person exists who does not exist in *A*. This person lives within the neutral range. So if adding a person in the neutral range is equally as good as not adding her, then *B* is equally as good as *A*. For the same reason, *A* is equally as good as *C*. It follows that *B* is equally as good as *C*, because ‘equally as good as’ is a transitive relation. But actually *B* is not equally as good as *C*. *B* and *C* are easy to compare in value, because they both contain the same population of five people. *B* is plainly better than *C*, since it is better for the fifth person and equally as good for everyone else.

So if we interpret ethical neutrality as equality of goodness, we will derive a false conclusion from the intuition of neutrality. Indeed, if neutrality is understood as equality of goodness, there can be at most one neutral level: one level of wellbeing such that adding a person at that level is neutral. (This level may depend on the context; it may depend on how big the existing population is and how well off the people are.) Adding a person with a wellbeing above that level would be a good thing, and adding a person with a wellbeing below that level would be a bad thing. But the neutrality intuition is that adding a person is neutral as a general rule, and certainly not just at a single level of wellbeing. If we are to sustain the intuition, we shall have to understand neutrality differently.

Greedy neutrality

It is easy to find a way to understand it that is at first more plausible. We are used to the idea that sometimes two things are incommensurate in value: neither is better than the other and yet they are also not equally good. The classic example is Sartre’s student, who was faced with a choice of staying in France to look after his mother, who desperately needed him, or of escaping to Britain to join the Free French Forces and fight for France.⁸ These two options would realize very different values – so different that, plausibly, they could not be precisely weighed against each other. Plausibly therefore, neither of the options was better than the other, yet they were not equally good either.

So the resource we need is available. We simply do not assume that if neither of two options is better than the other, the two are equally good. Neutrality is being neither better nor worse, but it is not necessarily being equally good. Our intuition is that adding a person is neither better nor worse than not adding her. So long as we do not suppose it is equally as good as not adding her, the problem I mentioned does not appear. Equality of goodness is a transitive relation, but the relation of being neither better nor worse is not. So we cannot derive a corresponding false conclusion. In figure 2, we can say that *B* is neither better nor worse than *A*, and that *A* is neither better nor worse than *C*, and that *B* is better than *C*. That makes good sense of this example, so for this example we can sustain the intuition that there is a neutral range.

However, a different difficulty arises. It is illustrated in Figure 3.⁹ *B* is made out of *A* by adding a person within the neutral range. We are supposing that addition within this range is

ethically neutral. That is to say, under our present interpretation, *B* is neither better nor worse than *A*. In particular it is not worse than *A*. Now compare *B* and *C*. These options contain the very same people as each other. So comparing their goodness does not pose the sort of problem we are up against in comparisons involving different populations. The difference between them is that in *C* one man is worse off and one woman is better off than in *B*. I ask you to assume *C* is better than *B*. That is to say, you are to think that in *C* I have pushed the woman's wellbeing up by more than enough to balance the man's loss. I am not asking you to be a utilitarian about this. All I ask is that you allow it to be possible for a distribution of wellbeing like *C* to be better than one like *B*, and you assume I have done enough to make this the case. I have designed the example so that, if you have egalitarian tendencies you will like it. For example, the worst off person is better off in *C* than in *B*. I hope almost everyone who is not sceptical about the idea of goodness altogether will be willing to assume *C* is better than *B*. Since *B* is not worse than *A*, it follows that *C* is not worse than *A*.

But now compare *A* and *C*. These two options differ in two respects. First, one man is worse off in *C* than *A*. For instance, we may suppose his life is shorter in *C*; he lives eighty years in *A* and sixty in *C*, say. *C* is unequivocally worse than *A* in this respect. Second, *C* contains an extra person, whose wellbeing is in the neutral range. In this respect, *C* is neither better nor worse than *A*. That is the intuition we are working on: that adding a person is neutral. So *C* is worse than *A* in one respect, and neither better nor worse in the other respect. Intuitively, therefore, *C* must on balance be worse than *A*. In going from *A* to *C* we have one bad thing and one neutral thing. A bad thing plus a neutral thing must add up to a bad thing.

However, we have actually concluded the opposite, that *C* is not worse than *A*. In going from *A* to *C*, although adding the extra woman is a neutral thing, it has managed to cancel out the bad thing that is the harm done to the man. We have found that our neutrality is *greedy*. Although neutral in itself, it is able to swallow up bad things and neutralize them. With a different example, I could equally well show how adding a person, though neutral, can cancel out good things too. Intuitively, neutrality should not behave like that; it should not be greedy.

Look at the same example in reverse. Imagine moving from *C* to *A*. This involves subtracting a person from the population. In going from *C* to *A* two things are changed. First, one person is benefited, the man. This is unequivocally a good thing. Second, one person is subtracted, the woman. This is a neutral thing. A good thing plus a neutral thing should add up to a good thing. So intuitively, *A* should be better than *C*. But it is not. The neutral thing swallows up the good thing and neutralizes it. In a different example, I could show how subtracting a person, though neutral, can also neutralize a bad thing. Neutrality is in all respects greedy; it can swallow up both good and bad things. Both adding a person to the population and subtracting a person from the population can have the effect of neutralizing other good and bad changes.

Neutrality is greedy just because we have assumed there is a neutral range rather than a single neutral level. So long as there is range of wellbeings such that adding a person within that range is neutral, neutrality will be greedy. Only if there is just a single neutral level will neutrality not be greedy. We may satisfy the intuition of neutrality to the extent of having a neutral range, or at least what we may call a neutral range. But when we do that, the neutrality turns out greedy. This means it is not intuitively neutrality at all. So in the end we have failed to satisfy the intuition of neutrality.

The end of the intuition

The problem is not just that greedy neutrality is intuitively unsatisfactory. Much more serious

is that, if neutrality is greedy, it cannot do the work we need from it. Remember it was the intuition of neutrality that apparently justified us in ignoring population effects when assessing some of the practical problems I described. Take global warming. As I explained, it will kill huge numbers of people, and that is a dreadfully bad thing. For this reason alone, if no other, it seems clear we should think global warming will be terribly bad. But it will also undoubtedly alter the size of the population either upwards or downwards. The intuition is that this effect will be ethically neutral. For that reason, we think it safe to ignore the population change when we evaluate global warming. Since global warming will do clear and unequivocal harm in killing so many people, we conclude it is bad. The intuition of neutrality allowed us to maintain that conclusion.

However, we now know that, even if changing the population is ethically neutral, it has the sort of neutrality that is greedy. So it is quite possible for it to swallow up the badness of killing people. The net effect of global warming may turn out neutral, and not bad after all. It may turn out neutral so long as it changes the population, whether upwards or downwards. I am not saying the change in population will certainly cancel out the badness of killing. That depends on the figures. It depends on how many people are killed, and how that number compares with the change in the population. It depends on how much harm is done to the people by being killed, and by how good are the lives that are added to the population (or how good the lives that are subtracted would have been). It also depends on the width of the neutral range.

Think of it this way. If there is a neutral range, distributions containing different populations are not always commensurate in value. The wider the neutral range, the greater the scope for incommensurateness. If two distributions differ only slightly in their population, the scope will be small. But if they differ a lot, there will be a great deal of incommensurateness. So the more two distributions differ in their populations, the less likely they are to be commensurate in value. That is to say, the more likely are their good features and bad features to be swallowed up by the greediness of neutrality.

I said earlier that we can expect changes in population to be perpetuated; they will last for ever. For that reason, we can expect them to be large. The change in population caused by global warming will probably be large, whichever direction it goes in. Therefore, if this change is neutral, I think we have to expect its neutrality to swallow up the bad effects of global warming. We shall be forced to conclude that global warming is probably not bad, but neutral.

Likewise, we shall have to assume that all the policies and actions I mentioned earlier, which affect the population, are probably neutral in value too. For example, it is pointless to work out the costs and benefits of a change in taxes, because taxes affect the population. This effect will be perpetuated, and is therefore likely to neutralize any good or bad effect the change might have.

I think this is absurd. It spells the end of the neutrality intuition, interpreted in terms of goodness. This intuition simply cannot be fitted into a coherent account of goodness without leading to an absurd conclusion. However strong the intuition may be, we have to grit our teeth and give it up. We can no longer assume that, as a general rule, adding a person to the population is neither better nor worse than not adding her. Generally it will be either better or worse. There can at most be only one neutral level of wellbeing, not a neutral range.

A hard conclusion

The intuition has to go. Abandoning it will save us from the absurdity I mentioned, but it will not make things easy for us. It takes us out of the frying pan into the fire. It was only the

neutrality intuition that allowed us to ignore all the effects of global warming on population, and all the effects on population of other things such as the social security system or education. If the neutrality intuition goes, we cannot ignore those effects. We shall have to take them into account when we assess the goodness of an action or policy, and we shall have to remember they can be expected to be large. This will call for extraordinary adjustments to our beliefs about goodness. We shall have to make a serious effort to predict their effect on population, and to assign a value to those effects.

In many cases, the factor we typically ignore – population – is likely to turn out the most important of all. And at present we are not in a position to make sound judgements about its value, since our theory of the value of population is in an underdeveloped state. We are not even clear whether adding to the population should be counted as good or bad. Many of our present views are seriously at risk of being quite wrong. For example, we have no right to assume global warming will be bad just because it will kill so many millions of people. Perhaps its effects on population will be so beneficial as to cancel out this badness.

This – the fire – is almost as uncomfortable as the frying pan. Must our moral judgements be overturned in this radical way?

Is there any alternative left? So far as judgements of goodness are concerned, I see only one more possibility: to discount future wellbeing. We might suppose wellbeing that comes later in time is less valuable than wellbeing that comes earlier. We might assume its value decreases exponentially at some modest, constant rate. Under that assumption, wellbeing that comes a few centuries from now will count for virtually nothing compared with present wellbeing. Since the effects of our actions on population will mostly develop over a very long time, discounting will very much diminish their importance in our judgements of goodness.

However, it is very hard to believe discounting is the right solution to our difficulty. For one thing, it is an opportunistic solution; it is not essentially connected to the difficulty's source. The source is changes in population; the solution is supposed to be the passage of time. These two are only contingently connected. If the world's population was to change unnaturally rapidly, discounting would offer no solution at all, but we would still have the same unease about the way population enters our judgements of goodness.

Secondly, it is any case incredible that later wellbeing is less valuable than earlier wellbeing. Suppose it is so, and suppose the rate of discount is enough to serve our purpose of diminishing significantly the value of a change in population. Then the killing of thousands during Caesar's Gallic Wars was a hugely worse event than the killing of millions during the two World War, the later killings are so much discounted compared with the earlier ones. That is incredible.

I think most people who favour discounting would refuse to accept this implication of it, because they discount later times compared with earlier ones only in the future, not in the past. This bias gives them a relativist theory of goodness: it means that the value of one event compared with another depends on the time when the valuation is made. From the perspective of a time before Caesar, Caesar's wars were indeed worse than the World Wars, but from the perspective of the present they are not worse.

Relativism brings its own problems. Suppose one of the acts that is available to you will bring about better consequences than any other, and suppose on that account you ought to do this act. Suppose you do. When your act is evaluated from the perspective of a later time, it may turn out to have had much worse consequences than the alternative acts you could have done. So from the later perspective it might turn out you ought not to have done it. This might happen, not because something unexpected occurs or some new information becomes

available, but simply because values differ when assessed from the later perspective. It may turn out you should undo your previous act if you can, even though you originally acted exactly as you ought. Moreover, at the time you acted, you may have known all this. You may have calculated the value your act would have from the future perspective, and seen that from the future perspective it would have much worse consequences than the alternatives. Nevertheless according to relativism, from your present perspective, you still ought to do it, because from that perspective it has the best consequences. That is most implausible.

In sum, discounting is incredible unless it is combined with relativism, and relativism is itself most implausible. Discounting is not a plausible way out of our difficulty. I can see no other way out, so far as judgements of goodness are concerned. We shall have to include the value of population in all our judgements, and be prepared to revise our judgements radically.

However, the intuition of neutrality may be able to survive in a different guise, and give us some practical relief from the difficulty. In so far as it is an intuition about goodness, it has to go. But at least part of it may be normative rather than evaluative. It might be partly an intuition about our moral responsibilities, rather than about goodness.

Think about a couple who might have a child. Our intuition is that their having a child is neither better nor worse than their not having one. But we now know this intuition is mistaken except in the special case where the child happens to live at exactly the single neutral level. So if the couple have a child, that will generally be either better or worse than their not having a one. Suppose it is better. Then the couple are in a position to make the world better by having a child. But even so, we might think they have no moral responsibility to do so. We might think they are doing nothing wrong if they choose not to. This normative conclusion about rightness and wrongness may be part of what the neutrality intuition is pointing to.

Possibly the intuition might be given a coherent interpretation in these normative terms. And possibly it may apply to grand issues such as global warming as well as to a couple's decision about a child. Global warming will be very good or very bad because of its effect on population. But possibly we may have no moral responsibility towards population, and we may be entitled to ignore the goodness or badness of this effect.

Notes

I wrote this paper while I was a Visiting Fellow at the Swedish Collegium for Advanced Study in the Social Sciences. I am very grateful to SCASSS for its generous support. The paper grew from lectures I gave at the Vancouver Institute in 2002 and at the Harvard Center for Ethics and the Professions in 2003. Much of its argument is set out in more detail in my *Weighing Lives*.

1. It particularly exercised Epicurus. See his 'Letter to Menoeceus'.
2. 'Moral problems of population'.
3. My own use of this term is perhaps not the commonest. For mine, see *Weighing Lives*, pp. 41–2.
4. For economists' work on safety, see Jones-Lee, *The Economics of Safety*. Arthur, 'The economics of risks to life', is one exception to the rule that population is ignored, but this is a theoretical piece, not a guide to valuation.
5. See the survey in Pearce et al, 'The social costs of climate change'.
6. McMichael and Githeko, 'Human health'.
7. See IPCC Working Group 1, *Summary for Policymakers*.
8. Sartre, 'The humanism of existentialism'.
9. This example is a version of Derek Parfit's 'mere addition paradox'; see his *Reasons and*

Persons, chapter 19.

References

- Arthur, W. B., 'The economics of risks to life', *American Economic Review*, 71 (1981), pp. 54–64.
- Broome, John, *Weighing Lives*, Oxford University Press, 2004.
- Epicurus, 'Letter to Menoeceus', in *Epicurus: The Extant Remains*, translated and edited by Cyril Bailey, Oxford University Press, 1926, pp. 83–93.
- Intergovernmental Panel on Climate Change, Working Group 1, *Summary for Policymakers*, 2001.
- Jones-Lee, M. W., *The Economics of Safety and Physical Risk*, Blackwell, 1989.
- Narveson, Jan, 'Moral problems of population', *The Monist*, 57 (1973), pp. 62–6.
- McMichael, Anthony, and Andrew Githeko, 'Human health', in *Climate Change 2001: Impacts, Adaptation and Vulnerability*, by the Intergovernmental Panel on Climate Change Working Group II, Cambridge University Press, 2001, pp. 451–85.
- Parfit, Derek, *Reasons and Persons*, Oxford University Press, 1984.
- Pearce, D. W., W. R. Cline, A. N. Achanta, S. Fankhauser, R. K. Pachauri, R. S. J. Tol and P. Vellinga, 'The social costs of climate change: greenhouse damage and the benefits of control', in *Climate Change 1995: Economic and Social Dimensions of Climate Change*, edited by James P. Bruce, Hoesung Lee and Erik F. Haites, Cambridge University Press for the Intergovernmental Panel on Climate Change, 1996.
- Sartre, Jean-Paul, 'The humanism of existentialism', in his *Essays on Existentialism*, Citadel Press, 1986, pp. 31–62.

Figure 1




Figure 2




Figure 3

