

# 6. The most important thing about climate change

John Broome

## Ethics and climate change

The title of this volume – *Public Policy: Why ethics matters* – is highly significant. Among the protagonists in the debate about public policy in response to climate change, many think ethics is irrelevant. Most of the protagonists are scientists and economists, and they think they need no contribution from moral philosophy. They are wrong.

Take as an example a criticism directed by the economist Martin Weitzman against *The Stern Review* (Stern 2007). In comparing the well-being of future generations with our own well-being, Nicholas Stern uses a lower discount rate than many economists do. This means he attaches more value to the economic consumption of future generations. A consequence is that *The Stern Review* recommends strong and immediate action to diminish climate change. Stern justifies his lower discount rate on ethical grounds (Stern 2007, ch. 2A). Weitzman (2007, p. 712) says he justifies it by:

relying mostly on a priori philosopher-king ethical judgements about the immorality of treating future generations differently from the current generation – instead of trying to back out what possibly more representative members of society ... might be revealing from their behavior is *their* implicit rate of pure time preference. An enormously important part of the ‘discipline’ of economics is supposed to be that economists understand the difference between their own personal preferences for apples over oranges and the preferences of others for apples over oranges. Inferring society’s revealed-preference value of [the discount rate] is not an easy task ... but at least a good-faith effort at such an inference might have gone some way towards convincing the public that the economists doing the studies are not drawing conclusions primarily from imposing their own value judgements on the rest of the world.

Weitzman evidently thinks the discount rate used by governments when deciding their response to climate change ought to be derived from people’s

preferences about their private savings, as revealed in the money market. This judgement of Weitzman's is an ethical one, because it is about how governments ought to behave towards future generations. An enormously important part of the discipline of economics is supposed to be that economists understand the difference between ethical judgements and judgements of empirical fact – 'normative' and 'positive' judgements in the terminology of many text-books of economics. Especially as Weitzman's judgement, on the face of it, is so implausible, he could not possibly hope to justify it without recourse to the discipline that examines how people and other agents ought to behave. That discipline is ethics. Weitzman needs ethics.

Climate change raises a wide range of ethical issues. It raises issues of justice, for instance. Our emissions of greenhouse gases directly cause climate change, which is already harming other people. People are already losing their homes, their livelihoods, and even their lives as a result of the climate change we are causing. This is an injustice we are doing to those people.

Climate change also raises many issues of value. We are worsening the lives of future people by damaging the environment in which they will live. To help us decide how much we should reduce our emissions, we need to set a value on the badness of the harm we are doing to future people, and compare it with the badness of the sacrifices we could make to reduce it. Among the bad things future people will suffer is that many will die before their time, in floods and heat waves, in droughts and famines, and through increased poverty among the poorest. We need to set a value on the badness of those early deaths. We also need to assess the value of nature, which we are impoverishing. Those are some assessments of value we need to make, and there are many others too.

In this chapter, I concentrate on just one question of value to illustrate the need for moral philosophy. My question is whether the most important thing about climate change is the harm it is likely to cause or alternatively the utter catastrophe that it may possibly – though very improbably – cause. Weitzman argues that the small chance of catastrophe is the most important thing about climate change, and this view is gaining ground among economists. I shall consider whether it is right.

## **The very most important thing about climate change**

However, since the 2009 Climate Change Conference in Copenhagen, I have come to the conclusion that the very most important thing about climate change is neither what is likely nor what is unlikely to result from it. The very most important thing is this fact: that the problem of climate change can be solved

without anyone making any sacrifice.<sup>1</sup> At Copenhagen, many nations came together and failed to reach an agreement. They were asked to make sacrifices, and they declined to do so. But no sacrifice is necessary. The nations might have been more amenable if they had understood that point.

The fact no sacrifice is necessary is a consequence of elementary economics. Climate change is what economists call an *externality*. Many of our activities cause greenhouse gases to be emitted. In deciding how many of these activities to engage in, people weigh the benefits they gain from them against the costs of engaging in them. But most of the costs of emitting greenhouse gases are not borne by the people who emit them. Instead, they are distributed across the population of the world, through the damage the greenhouse gases do. When the costs of an activity are not fully borne by the person who decides to engage in it, that is an externality.

Elementary economics tells us that externalities cause *inefficiency*. When an economist says that a situation is inefficient, she means it would be possible to make someone better off without making anyone else worse off. More precisely, there is some alternative state such that someone prefers the alternative to the existing state and no one prefers the existing state to the alternative. To adopt economists' terminology, let us say this alternative is *Pareto superior* to the existing state. Moving to the Pareto superior state involves no sacrifice on anyone's part. Because climate change is an externality, there is a Pareto superior state we could move to. It will involve emitting less greenhouse gas. Moreover, there is a Pareto superior state that is *efficient*, which means no other state is Pareto superior to it. If we get to a state like that, the inefficiency caused by the externality will have been eliminated, and no one will have made any sacrifices.

The theory of externalities tells us that achieving a Pareto superior state will often require resources to be transferred from some people to others. In the case of climate change, it is obvious in broad terms what sorts of transfer are required. We the current generation benefit from emitting greenhouse gases as we do at present. Suppose we change our policies and emit less of them. That by itself would make us worse off. But the theory tells us that resources could be transferred to us from the beneficiaries of our reduction in emissions, in such a way that no one ends up worse off. In the case of greenhouse gases, most of the beneficiaries are people who will live in the future. Therefore, resources will need to be transferred from future people to present people.

How can that happen? We the current generation are set to bequeath a lot to future generations. We shall leave them artificial resources such as roads and museums. We shall also leave them natural resources, since this generation

---

<sup>1</sup> I take this point from Foley (2007). It is widely recognised. For instance, it is mentioned by Stern (2010, p. 85).

will not exhaust all the natural resources the earth possesses. To compensate ourselves for reducing our emissions of greenhouse gases, we can use more of other resources for ourselves. We shall leave less of them for the future. By itself, that would be bad for future generations. However, those future generations will suffer less from climate change because we reduce our emissions. We know from the elementary economic theory that, if we do the transfer correctly, future generations will end up better off on balance, and we shall be no worse off.

The outcome I have described is not a nice one. It is Pareto superior to the existing state, but the existing state contains all the bad consequences of emitting greenhouse gases. For instance, it contains the injustices I mentioned: we are already harming some existing people by our emissions. If we reduce our emissions, we shall stop harming those people in that way. But it compounds the injustice to expect them to compensate us for reducing our emissions. Another example is that our emissions are damaging the conditions of life of future people. I have said it would be better if we reduced our emissions and compensated ourselves for doing so. But it would be even better if we reduced our emissions and did not compensate ourselves for doing so.

Compare three alternatives: *A*, we do nothing about the externality and continue to emit greenhouses gases profligately; *B*, we reduce our emissions to eliminate the externality, and compensate ourselves for doing so; *C*, we reduce our emissions to eliminate the externality, and do not compensate ourselves for doing so. *The Stern Review* in effect compares *C* with *A* by means of a cost–benefit analysis, and finds that *C* is much better than *A*. It recommends us to choose *C*. It does not compare *B* with *C*, but it is clear that Stern considers *C* better than *B*, or he would not have recommended *C*.

The difference between *B* and *C* is nothing to do with climate change. Moving from *A* to *B* eliminates the problem of climate change. To move from *B* to *C* is simply to redistribute wealth from present people towards future people. No doubt outcome *C* would be the best. But to reach *C* the current generation has to make sacrifices, and the experience in Copenhagen shows it is unwilling to do so.

I think we should try first to develop the institutional arrangements that will make the move from *A* to *B* possible. That will allow us to eliminate the problem of climate change. Then we should try going further to *C*. But we should not encumber the process of controlling climate change with the quite different matter of transferring resources to future people.

## Expected utility theory and very bad possibilities

That very most important thing about climate change is a matter of economics alone; it has little to do with moral philosophy. But now I come to the real question of this chapter. It arises because the science of climate change is so uncertain. We do not know how much the world will warm, and we do not know what effects the warming will have. Fortunately, we have a well-established account of how we should take uncertainty into account in our planning. It is called expected utility theory.

Here is how expected utility theory works. Suppose you have several options to choose from. For instance, they might be alternative policies towards climate change. It is uncertain what outcome will result from each option. Each may lead to various different possible outcomes, depending on how the uncertain world develops.

Each possible *outcome* is good or bad to some degree. Each has a value, that is to say. We may also assign a value to each *option* on the basis of the values of the various outcomes it may lead to. The value of an option is the weighted average of the values of its possible outcomes, where each outcome is weighted by its probability. This weighted average is called the *expectation* of the value of the outcomes.

That is not a strictly accurate description of expected utility theory. Strictly, the value assigned to an option is the expectation of something called the 'utility' of its possible outcomes, rather than the expectation of value. The utility of an outcome is not exactly its value, but its value adjusted in a way that is designed to take account of the badness of risk. The difference between value and utility is important, but for the purposes of this chapter it is safe to ignore it.

Here is an example of the working of expected utility theory. Suppose you have a cold. Suppose a cold-cure pill is available that has a 95 per cent chance of curing your cold. In the terminology adopted by the Intergovernmental Panel on Climate Change, this means the pill is 'extremely likely' to cure your cold. But suppose that, in the extremely unlikely event – having 5 per cent probability – that the pill does not cure your cold, it will kill you. Should you take this pill?

You should not. The badness of the possible outcome of your death is so great that, even weighted by the small probability of its occurrence, it outweighs the goodness of curing your cold, even weighted by the large probability of its occurrence. The expected value of taking the pill is:

$$(\text{value of having no cold}) 95/100 + (\text{value of dying}) 5/100.$$

Your alternative option is not to take the pill, which I assume leaves you certain of continuing to have a cold. The expected value of this option is just:

(value of having a cold) 100/100.

This is far above the expected value of taking the pill.

When you make the choice I have described, the most important thing is the small chance of dying. That is a consequence of expected utility theory. What matters in making a choice is not the likelihood of outcomes but the expected value of outcomes. Consequently, the most important aspect of a choice is not necessarily what is likely to happen. It may be something that is very unlikely to happen, if it is extremely bad.

## The small chance of a catastrophe

This is an elementary point, and the scientists of the Intergovernmental Panel on Climate Change (IPCC) surely know it very well. But you would not think so to read the IPCC's Fourth Assessment Report (IPCC 2007a–c). The uncertainty in the IPCC's report is presented almost entirely in terms of likelihood. The IPCC tells us what is likely to happen, what is very unlikely to happen, and so on. It does not tell us expectations.

If we take moderately strong steps to control climate change, the most likely outcome is 2°C or 3°C of warming. I shall call this 'likely climate change'. Its effects might be manageable. However, the Fourth Assessment Report also reveals a significant probability that the warming will be much greater. It is cautious about assigning probabilities to very high temperatures, but studies referred to in the report suggest there is as much as a 5 per cent probability of warming greater than about 8°C, and perhaps a 1–2 per cent probability of warming greater than 10°C (IPCC 2007a, section 9.6).

These temperatures are far beyond the experience of human beings. The world has not been as hot as that for tens of millions of years. For comparison, during the last ice age – only tens of thousands of years ago – ice-sheets kilometres thick covered Canada and northern Europe. A 10°C warming is likely to melt Antarctica, and that will raise the seas by 70 metres. The result of so much warming might be catastrophically bad. Consequently, even multiplied by its small probability, this possibility of extreme warming may be much more important in expected utility than all the predictions about what is most likely to happen.

This is the view that is being promulgated by Weitzman (2009). Weitzman claims that the most important thing about climate change is the small chance of what

he calls a 'catastrophe'. In his argument, he adds to what I have been saying, another significant twist based on statistical theory. The IPCC's estimates of prospective temperatures are derived from observed data. But there is inevitably little data about unlikely events, because they happen only rarely. Therefore, the data we have do not allow us to set a limit on the probability of unlikely events. This gives us further reason to be concerned about them.

Is Weitzman right that the small chance of catastrophe is the most important thing about climate change? Should it dominate our calculations of expected utility? We cannot answer this question except by thinking about how bad a catastrophe it would be. To know whether it should dominate our calculations even when weighted by the very small probability that it will happen, we must have some quantitative idea of its badness. Weitzman does not try to estimate the badness of a catastrophe. He seems to assume it would obviously be so bad that an estimate would be otiose. But when we do try to make an estimate, it turns out not to be obvious that the small chance of catastrophe should dominate our calculations. An estimate is, therefore, not otiose.

In this chapter, I want to argue no more than that: the small chance of catastrophe does not obviously dominate. For this purpose, only the roughest preliminary estimate of the badness of the catastrophe is required. If my conclusion is right, the chance of catastrophe will have to be included in our expected utility calculations along with other, more likely, results of climate change. Far more precise calculations will be needed then, but here I can work with very rough figures.

## Consequences of catastrophe

I shall try to make only a very rough, preliminary judgement of badness. I shall start by identifying the various sorts of harm that would result if we suffered extreme climate change of 8°C or more. For one thing, extreme climate change would be a disaster for the natural world. Millions of species would be lost and very many ecosystems destroyed. I do not wish to belittle these great harms, but I am going to set them aside in this chapter. I shall concentrate only on harms that would befall humanity specifically. I shall even set aside the harm that humans would suffer through having to live in a barren world, where nature has been impoverished. I shall consider only harms that afflict humanity more directly.

One thing is sure: if there is extreme warming, the earth will not be able to sustain anything approaching our present population. Our population would have to shrink by billions. We cannot even be confident that humanity will

survive at all. Lots of species are already becoming extinct, and we cannot assume ours will not follow them. The possibility of extinction is one thing we shall have to consider in assessing the badness of a catastrophe.

How bad would extinction or a major population collapse be? It would be bad in various ways. If extinction were to occur, it would be the end of our species. It can be argued that this species has more value than others. For one thing, we are rational, and many philosophers think that is a property of special value. I shall not try to judge that claim.

A collapse of our population, even without extinction, would cause the loss of a great many valuable things. They include human cultures, with all their parts: their knowledge, their languages, their arts, and their ideas. Indeed, civilisation as a whole would go. These are great harms, but I am not going to concentrate on them because I do not know how to judge the badness of their loss. Instead, I am going to concentrate on two other ways in which extinction or a population collapse would be bad. These are ways I can at least begin to evaluate.

The first of these more approachable sorts of badness is that billions of people would die early. A collapse of population, and even extinction, would not necessarily involve early deaths. Women could just stop having children for a while, or forever. But the collapse or extinction that climate change may cause will be brutal and violent. It will involve killing many people.

The second more approachable badness – or at least putative badness – that extinction or a population collapse would cause is that very many people who would have existed will not exist at all. An extinction would prevent the existence of all those human beings who would otherwise have lived. Even a collapse of population that falls short of an extinction would prevent very many existences, and an extinction would mean there will be no more people ever.

We are trying to judge whether the possibility of extreme climate change would be so bad that it should dominate our thinking about climate change, even when weighted by the very small chance it will occur. To make this judgement, we need only very rough, ball-park figures. For this purpose, we need only think about the numbers of people who will be killed or whose existence will be prevented. Numbers are enough; we do not need to think about the badness of all those deaths.

Since I shall conclude that Weitzman's case is dubious, I shall be generous to it in my calculations.

## Deaths

Start with the first sort of approachable badness: the killing that extreme climate change would do. Out of generosity to Weitzman, let us assume the worst possibility, which is that everyone on earth is killed, at a time when our population is at its greatest. Our maximum population is predicted to be 9 billion in about 2050. Suppose extreme climate change kills 9 billion people. But extreme climate change is very unlikely. Suppose its probability is one in a thousand. Then the *expectation* of the number of deaths is 9 billion divided by a thousand, which is 9 million. If the probability is as high as one in a hundred, the expectation is 90 million. The expectation is, therefore, in the range of tens of millions of people killed.

These figures must be compared with the number of deaths we may expect from likely climate change of 2°C or 3°C. Likely climate change will kill in various ways. People will be killed in floods, droughts and famines, and other climate disasters. They will be killed by the increased range of tropical diseases, and they will be killed in heat waves. Also, they will be killed simply by impoverishment, since you die if you cannot afford enough to eat. Predictions about the numbers of deaths are hard to come by, but I have found a few relevant figures. The 1995 IPCC report offers a prediction that 215,000 people will be killed each year by heat waves alone, once the concentration of carbon dioxide in the atmosphere is effectively doubled (IPCC 1996, p. 198). A report of the World Health Organization attributed 150,000 deaths to climate change as long ago as 2000 (cited in IPCC 2007b, p. 407). The Food and Agriculture Organization reports that the world now contains more than 1 billion undernourished people (FAO 2009). A small increase in poverty can be expected to kill many of them. I think we might estimate that half a million people will be killed each year by likely climate change; that would be a conservative figure. As a check, about 60 million people die in the world each year. So this would amount to about a 1 per cent increase in the world's annual death rate. That increase seems plausible. No end is predicted to this killing; it will continue for decades. We must, therefore, expect likely climate change to kill tens of millions of people.

This puts the expected number of deaths that would result from likely climate change in the same ball-park – tens of millions – as the expected number that must be attributed to the small chance of the extinction of humanity. The first of two approachable badnesses I mentioned – killing – therefore, does not suggest that the small chance of catastrophe should dominate our calculations of expected value.

## Absences and the intuition of neutrality

The second approachable badness of a catastrophe is that it will prevent the existence of a great many people who would otherwise have existed. In this case, the expected numbers of people involved are much bigger. The human species could exist for perhaps hundreds of thousands of years, renewing itself every century or so. All those future people amount to at least hundreds of times the world's present population. Extinction will prevent the existence of that number of people, and a population collapse may prevent the existence of many of them. Even multiplied by a small chance of one in a hundred or one in a thousand, the expected number of 'absences' – as I shall call them – associated with extinction is vastly greater than the number of deaths to be expected from likely climate change. The same may also be true of a population collapse. These numbers suggest we have a potentially dominating consideration here.

A first response is a point that was made to me by Lukas Wallrich. It is not just extinction or a population collapse that would cause absences. Likely climate change will do the same. Likely climate change will kill tens of millions of people, and many of those people would have had children had they survived. Indeed many would have started a long line of descendants. All their potential descendants will be absent because of likely climate change. The expected numbers of these absences may well be as great as the expected number associated with extinction. That depends on the details of human demography, which I shall not go into here.

This point by itself is enough to show that the chance of catastrophe from extreme climate change is not obviously a dominating consideration. But I shall go further. I shall consider whether absences are truly a bad thing in the first place. Is it bad if people who would have existed do not exist after all? To put the same question differently: is the existence of people a good thing? That is the question I come to now.

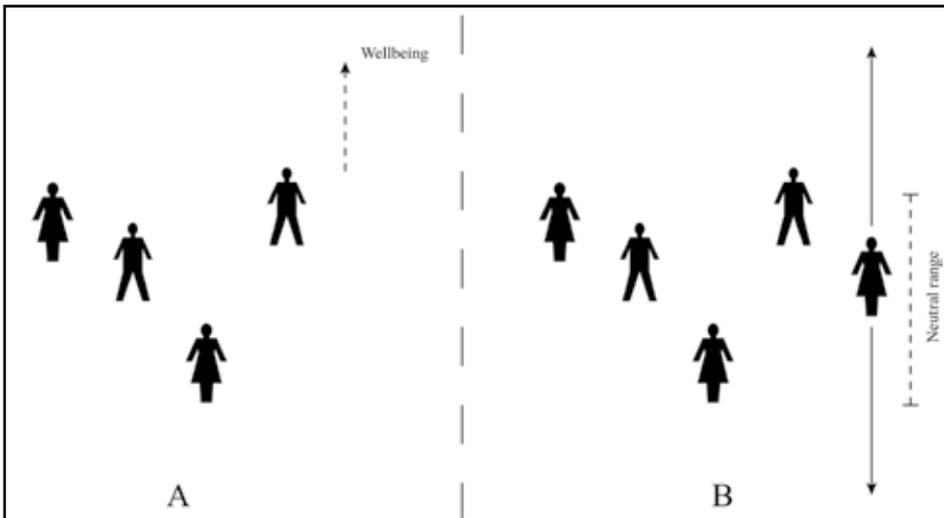
Many of us have conflicting intuitions about it. On the one hand, many of us are horrified at the thought of humanity's extinction; it strikes us as obviously a dreadfully bad thing. That seems to be Weitzman's reaction. But on the other hand, when we think of a particular person's existence, many of us think intuitively that it is neither good nor bad. We recognise that it may have good or bad effects on other people. The person may bring joy to her family, and she may detract from other people's well-being through the demands she makes on the earth's limited resources. But we do not think her existence is good or bad on its own account.

To illustrate this intuition: few of us think that the good a person would enjoy in her life constitutes a reason for creating this person. Suppose a couple who

could have a child decide not to. We may think they are doing some good things: for instance, they are refraining from creating one more mouth to feed. We may also think they are doing some bad things: for instance, their own lives might be worse in the long run as a result of being childless. But few of us think they are doing a further bad thing in preventing the existence of a person just because the person herself would have enjoyed a good life.

I call the view I am describing ‘the intuition of neutrality’. Put more exactly, it is this. Take a world *A* that has a number of people living in it, each with a particular level of lifetime well-being. Compare it with another world *B* that contains all the same people with the same levels of lifetime well-being, and also contains one more person. The intuition of neutrality is that neither *A* nor *B* is better than the other. Existence has neutral value, we might say.

This intuition is illustrated in Figure 1. The figure shows worlds *A* and *B*. Well-being is shown on the vertical dimension of the figure; better-off people are placed higher up. In this example, *A* contains four people. *B* contains those four at the same levels of well-being, and also one more person. The intuition of neutrality is that neither *A* nor *B* is better than the other.



**Figure 1: The intuition of neutrality**

Plausibly, this common intuition has limits. Many people think that the existence of a person whose life is very bad is a bad thing. So if the extra person in *B* was very badly off, we might think *B* worse than *A*. Some people may think that the existence of an extremely well-off person is a good thing, so *B* would be better than *A* if the extra person was extremely well off. But at least for a range of qualities of life, which we may call ‘the neutral range’, we think that

the existence of a person whose life is within that range is neither a good thing nor a bad thing. That is the limited version of the intuition of neutrality. It is shown in Figure 1.

I do not insist that there is a finite neutral range; I merely allow for the possibility that there is. Another possibility is that existence has neutral value, whatever the level of well-being of the extra person.

If the intuition of neutrality is correct, it tells us that extinction is neither good nor bad, provided the future people who will exist if humanity does not become extinct, live within the neutral range. So the intuition of neutrality conflicts with our other intuitive idea that extinction would be a dreadfully bad event. The intuition of neutrality threatens Weitzman's implicit view that extreme climate change would be catastrophically bad. We must, therefore, assess the truth of this intuition.

## The intuition of neutrality is false

And actually, it is false. Unfortunately, the full demonstration of its falsity is too long for me to present in this chapter. It appears in my book *Weighing Lives* (Broome, 2004, chapters 10–12, summarised in Broome, 2005). Here, I shall set out only the first step, which demonstrates that one particular, strong version of the intuition is false.

The strong version is this. Compare the same two worlds *A* and *B* as before. The strong intuition is that they are equally good, at least if the extra person's well-being is within the neutral range. The weak version of the intuition is only that neither *A* nor *B* is better than the other. It leaves open the possibility that *A* and *B* are incommensurate in value. The strong version excludes this possibility.

The strong version is false for a reason that is illustrated in Figure 2. The worlds *A* and *B* in Figure 2 are the same as *A* and *B* in Figure 1. World *C* contains the same five people as *B* contains. The first four of them are equally as well off in *C* as in *B*, but the last person in *C* is worse off than she is in *B*. Obviously, *B* is better than *C*, since it is equally as good as *C* for four of the people and better for the fifth person. However, according to the strong version of the intuition of neutrality, *B* is equally as good as *A*, and *A* is equally as good as *C*. It follows, obviously, that *B* is equally as good as *C*. Yet we know already that this is not so: *B* is better than *C*. So the strong version of the intuition of neutrality entails a falsehood. Therefore, it is false itself.



**Figure 2: The strong version of the intuition of neutrality**

This is a knock-down argument against the strong version. I think the weak version is also false, but the argument for that conclusion, presented in my *Weighing Lives*, is more open to debate (Rabinowicz, 2009). Still, in this chapter I shall take it for granted that the weak version is also false.

What follows? It follows that there is no such thing as a neutral range of well-beings of the sort shown in Figure 1. Compare the worlds *A* and *B* again. We have learnt that there can be at most one level of lifetime well-being such that, if *B*'s extra person lives at that level, neither *A* nor *B* is better than the other. If the extra person lives above this level, *B* is better than *A*; if below, *B* is worse than *A*. I assume there is indeed such a level of well-being, and I call it the 'neutral level'. The only other possibilities are that *B* is worse than *A* whatever the extra person's well-being, or that *B* is better than *A* whatever the extra person's well-being. Neither of those possibilities is plausible.

So the neutral range collapses to a single neutral level. What does this tell us about the value of extinction? It tells us that extinction would almost certainly be either a good thing or a bad thing, and not a neutral thing. But it does not tell us whether it would be good or bad. That depends on how well off future people will be if humanity survives. If their well-being will on balance be above the neutral level, extinction would be a bad thing. If their well-being will on balance be below the neutral level, extinction would be a good thing.

To know whether extinction is a bad thing, we need to know what is the neutral level of well-being, and we also need to know what will be the levels of well-being of future people. The latter question is partly a matter for empirical prediction and partly a matter for moral philosophy. It is not entirely empirical, because to answer it we need first to know what a person's well-being consists

in. That is one of the major topics of moral philosophy. On the other hand, the former question – what is the neutral level – is entirely a matter for moral philosophy.

Some existing philosophical views imply a conclusion about what the neutral level is. One is a version of the theory known as ‘hedonism’. Hedonism is the theory that the only good thing is pleasure and the only bad thing pain. It can be interpreted as a theory about people’s well-being. Under that interpretation, it tells us about the relative goodness for a person of different lives she might lead. It says that one life is better than another if and only if it contains a greater preponderance of pleasure over pain. I call this ‘personal hedonism’. It does not tell us about the goodness or badness of a person’s existence.

A different version of hedonism does. I call it ‘general hedonism’. General hedonism is a theory about the goodness of worlds, rather than people’s well-being. It says that one world is better than another if and only if it contains a greater preponderance of pleasure over pain. Think of a person who lives a life that contains no pleasure and no pain. According to general hedonism, this person’s existence leaves the world equally as good as it would have been had she not existed. So her existence is neutral. According to general hedonism, the neutral level of well-being is therefore the level of a life that contains no pleasure and no pain.

I do not mention general hedonism to commend it, but to show that the neutral level is a matter for debate and discussion within ethics. There is work for moral philosophy to do in trying to figure out what the neutral level is. This work is a contribution towards assessing the badness of the catastrophe that might result from climate change.

## Conclusion

What have we learnt? We have learnt, first, that we cannot just assume the small chance of catastrophe is the most important thing about climate change. It may or may not be. To know which, we must work out just how bad the catastrophe would be.

We have learnt, second, that working this out will be difficult. Extreme climate change will certainly cause a collapse of the human population. It may cause the extinction of humanity. Naively, we think of these as terrible disasters. Perhaps they are. But to know whether they are terrible, and if they are, how terrible, we must investigate how good or bad it is for a person to exist. This is a difficult task for moral philosophy.

One further lesson follows. Moral philosophers have worked on the value of existence for some decades, largely stimulated by the pioneering work of Derek Parfit (1984). The subject is extremely controversial, and we cannot expect any consensus to emerge soon. Yet the problem of climate change is very pressing; we must act soon. So we shall have to act while we remain very uncertain about what is the correct moral theory of the value of existence. We therefore need to know how we should act when our moral theory about how we should act is uncertain. This is in itself a problem for moral philosophy.

As yet, it has not been much explored by moral philosophers. But some interesting work has recently been appearing on this subject (for instance, Sepielli 2010). The problem of climate change makes it urgent.

## References

- Broome, J. 2004. *Weighing Lives*. Oxford: Oxford University Press.
- Broome, J. 2005. 'Should we value population?' *Journal of Political Philosophy* 13: 399–413.
- FAO. 2009. *The State of Food Insecurity in the World 2009*. Rome: Food and Agriculture Organization of the United Nations. [www.fao.org/publications/sofi/en](http://www.fao.org/publications/sofi/en).
- Foley, D. 2007. *The Economic Fundamentals of Global Warming*. Working paper 07-12-044. Santa Fe, New Mexico: Santa Fe Institute. [www.santafe.edu/media/workingpapers/07-12-044.pdf](http://www.santafe.edu/media/workingpapers/07-12-044.pdf).
- IPCC (Intergovernmental Panel on Climate Change). 1996. *Climate Change 1995: Volume III: Economic and social dimensions of climate change*. Cambridge, UK: Cambridge University Press.
- IPCC (Intergovernmental Panel on Climate Change). 2007a. *Climate Change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK: Cambridge University Press.
- IPCC (Intergovernmental Panel on Climate Change). 2007b. *Climate Change 2007: Impacts, adaptation and vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK: Cambridge University Press.
- IPCC (Intergovernmental Panel on Climate Change). 2007c. *Climate Change 2007: Mitigation of climate change. Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK: Cambridge University Press.

- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Rabinowicz, W. 2009. 'Broome and the intuition of neutrality.' *Philosophical Issues* 19: 389–411.
- Sepielli, A. 2010. 'Normative uncertainty and intertheoretic comparisons.' *PhilPapers: Online research in philosophy*. <http://philpapers.org/rec/SEPNUA>.
- Stern, N. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge, UK: Cambridge University Press.
- Stern, N. 2010. *A Blueprint for a Safer Planet*. Vintage Books.
- Weitzman, M. L. 2007. 'A review of *The Stern Review on the Economics of Climate Change*.' *Journal of Economic Literature* 55: 703–24.
- Weitzman, M. L. 2009. 'On modeling and interpreting the economics of catastrophic climate change.' *Review of Economics and Statistics* 91: 1–19.