Valuing policies in response to climate change: some ethical issues

John Broome

A contribution to the work of the Stern Review on the Economics of Climate Change

Section 1. Introduction

Value and justice

There are various policies that the world community, or Britain individually, might adopt towards climate change. We can take little action, we can work hard to develop less polluting technologies, or we can drastically alter our energy policy. We need to set a value on each of the alternative policies that are available. That is to say, we need to assess how good or bad each of them is – I take 'value' and 'goodness' to be synonymous.

We need to value our alternative policies because their relative values contribute to determining which of them we ought to choose. Their values are not the only things that determine which we ought to choose. When we make choices, there may also be issues of justice involved. Justice is a distinct consideration from value.[1]

One principle of justice is that we generally ought not to harm other people for the sake of benefiting ourselves. This is generally so even if the benefit to ourselves is greater than the harm to them. For example, you ought not to hold a party on your neighbour's lawn without permission, even if you and your guests will enjoy a benefit that is much greater than the harm you do your neighbour by damaging her lawn. Sometimes you cannot avoid harming another person, but in that case justice generally requires you to compensate that person for the harm. In many cases, this principle of justice is enshrined in law, but it remains a principle of justice even when it is not.

It is important in the context of climate change. The people who suffer the consequences of climate change are often not the people who cause it. To speak broadly, most of the climate change that is presently happening is caused by people who live in the richer countries. Those people benefit from the energy they use, but in using it they cause climate change. However, many of the bad consequences are being suffered by people who live in poorer countries. The rich benefit, and in doing so they harm the poor. According to the principle of justice I described, they ought not to do this, and if they cannot avoid doing it at least they owe compensation to those who suffer. This is so even if, in this international context, the principle of justice is not enshrined in law.

It means that a calculation of value is not enough to determine whether a policy ought to be chosen. Suppose Britain is choosing between two policies. Suppose one is better than the other on balance: it brings extra benefits that outweigh the extra harms it causes. A cost-benefit analysis will come down on the side of this policy. If we took account of value alone, we would conclude that this policy ought to be chosen. But if the extra benefits come to Britons, whereas the harms are suffered by people in developing countries, it might be unjust of Britain to choose this policy. At least justice will require compensation for those who suffer, and if those people are not to be compensated, the policy will be unjust. It may consequently be wrong to choose it even though on balance the benefits exceed the costs.

Consequences

So we must recognize that value is not the only consideration in decision making; justice counts as well. Nevertheless, we do have to set values on alternative policies. This paper is about how to do that. So now I set questions of justice aside.

We can take it for granted that the value of a policy is given by the value of its

consequences. This is a view about value known to moral philosophers as ‘consequentialism’; in our context it seems unexceptionable. I have set aside questions of justice; once we concentrate on value, we should accept consequentialism.

Each possible policy will have consequences of various different sorts. For example, a policy might impose some costs on people who are presently living, by making them pay more for their energy. If it reduces future climate change, it will improve the quality of life of some people who live in the future. It may also do harm to other future people, because climate change may well be beneficial in some parts of the world. It will save some people’s lives in the future, and it may reduce the damage done to the world’s forests. And so on. To evaluate a policy properly, all its disparate consequences need to be taken into account.

This means we shall need to assign values to each of the various consequences, and then we shall have to put those values together – aggregate them – to arrive at an overall evaluation of the policy. When we are dealing with policies that respond to climate change, both steps are difficult. For example, it is difficult to assign a value to the destruction of natural ecologies, and it is difficult to aggregate the value of natural ecologies with the value of saving people’s lives.

As well as those difficulties, we face another. For policies that respond to climate change, it is very uncertain what their consequences will be. In section 2 I shall consider how we should deal with this great uncertainty when we evaluate policies. In section 3 I shall come back to valuing different sorts of consequence, and aggregating them together.

Section 2. Uncertainty

Expected utility theory

In principle, this question about uncertainty has a clear answer. Expected utility theory is a very well-grounded and well-established theory of value in the context of uncertainty. We should rely on it. It is not always easy to apply in practice, and it is not universally applicable, but where it is applicable, it is in principle the correct way to assess the value of policies that have uncertain results.

Expected utility theory can be applied to various different domains. It is often presented as a theory about people’s *preferences* – either a theory of the preferences people do have when facing uncertainty, or of the preferences they ought to have. But I shall treat it as a theory of *value*. It tells us about the value of alternative policies, when the results of these policies are uncertain.[2]

To understand expected utility theory, it is best to start from trivial cases. Gambling games provide good examples, because many of them involve known risks and clearly specified outcomes in the form of quantities of money. Suppose you are about to be given £10, and are now offered ‘double or quits’ as an alternative. That is to say, you can either take the £10, or alternatively accept a gamble that will give you either no money or £20, each with a 50% probability. Each of your two options has an *expectation* (a mathematical term) of the money you will receive. The expectation of a gamble is defined as the weighted average of your possible winnings or losings from the gamble, weighted by the probability of your getting them. In this case, if you play double or quits, the probability of your getting no money is equal to the probability of your getting £20, and both are 50%. So the weighted average of these two possibilities is just £10; this is your expectation from playing double or quits. If you choose not to make this gamble, you will receive £10 for sure. So your expectation from this decision is also £10. Your expectation of money from the gamble is equal to your

expectation of money from simply taking the gift.

Intuitively, this may suggest your two options are equally good. Intuition suggests that the value of an option is the expectation of money you receive from it. This is indeed correct in some special circumstances. For instance, it is correct for a casino, in operating its slot machines and roulette tables. On each occasion when a gambler offers to make a bet with the casino, in principle the casino can choose either to refuse the bet or accept it. In practice, when a gambler offers to bet at a roulette wheel or a slot machine, the casino automatically accepts the bet. But it makes sure that its expectation of money from accepting it is always greater than its expectation from refusing it. Its expectation from refusing it is zero, so it accepts only bets that give it a positive expectation of money. This is the best thing for it to do. For a casino, the value of an option is the expectation of money it receives from the option. This is because a casino is in the special position of making a great many bets in a short period of time. It loses some and wins others, but so long as each has a positive expectation of money, it is virtually certain to make a profit out of all of them taken together.

But not many people have the opportunity to make so many bets in a short time. If you are not able to do this, the option with the greater expectation of money is not always the best for you. There is more than one reason for this. First, there is the fun of gambling. If you enjoy gambling, your enjoyment adds to the value of accepting a bet. In the example of double or quits, this may make it better to make the gamble rather than accept the boring certainty of taking your £10. But the fun of gambling has no relevance to the context of climate change, and we can ignore it. There is a more serious consideration.

Even ignoring the fun of gambling, it is not necessarily true that the value of receiving money or alternatively of accepting a gamble of double or quits are the same. Suppose you are offered double or quits, not over a gift of £10, but over your salary for a month. If you accept this gamble, you may lose and not have enough money to pay your rent for the month, or buy food. You may alternatively win and be affluent for the month, but that possibility may not be enough to make up for the chance of poverty. If you accept double or quits over your salary every month, you life will oscillate uncertainly between poverty and affluence, and this may be a very bad thing for you.

This example shows that, for many of us, there is benefit in having some certainty about our money income, and avoiding risk to it. For that reason, a risky gamble is not necessarily as good as a certain alternative that has the same expectation of money. This does not affect casinos much, because casinos do not actually take much risk. Because they bet so often in a short time, the risk they take on each bet balances out with other bets and comes to a small risk overall. But when we are facing significant risk, the value of a gamble may not be its expectation of money. Risk to money may be a bad thing. That is to say, of two options that lead to the same expectation of money, the less risky one may be better than the more risky one. In our example, taking your month's salary for sure may be better than playing double or quits with it, even though either has the same expectation of money.

We need to take the badness of risk into account when assessing the value of our options. Some may be riskier than others, and their riskiness is a bad feature we need to recognize.

Bernoulli's theory

The eighteenth-century mathematician Daniel Bernoulli proposed a way of taking it into account.[3] In the end his idea was not entirely successful, but it constituted a major step on the road to modern expected utility theory. To understand the modern theory it is useful to start with Bernoulli's. Bernoulli thought that the value of an action or policy is not the expectation of money it leads to, but the expectation of value or goodness. More money is better than less

money, but Bernoulli recognized that the value of money to you does not increase in proportion to the amount of money you have. To a poor person, an extra hundred pounds is extremely valuable because it can be used to buy essential food. To a rich person, an extra hundred pounds is less valuable because the only use she will have for it is to buy unnecessary luxuries. Money has 'diminishing marginal value' as economists say. Once we recognize this fact, it can be used to account for the badness of risk.

Suppose your salary is £1000 a month. If you accept it, you get the value of £1000. If you play double or quits with it, you get either the value of no money or the value of £2000. For convenience, we can take the value of no money to be zero. The expectation of value that results from playing double or quits is the weighted average of nothing and the value of £2000. The weights in this case are both one half, because this is the probability of either outcome. So the weighted average is

½ (value of no money) + ½ (value of £2000).

Since we are taking the value of no money to be zero, this is just half the value of £2000. Because money has diminishing marginal value, this is less than the value of £1000. According to Bernoulli's theory, the value of accepting the gamble is therefore less than the value of taking your salary. You do better to take your salary and not play double or quits with it.

So long as money has diminishing marginal value, risk will always be a bad thing according to Bernoulli. Of two options that have the same expectation of money, the more risky one will always have the lower expectation of value. So the more risky option will always be the less good one. This theory explains satisfactorily what is bad about risk to money. The explanation is that money has diminishing marginal value.

However, there is a point that Bernoulli's theory does not allow for. Risk to *value* may also be a bad thing. Suppose that in the course of a disaster you are faced with a terrible choice. By doing one thing you can save the lives of ten people for sure. There is an alternative, risky action you might take. If you take it, there is a 50% chance you will save twenty people, but also a 50% chance you will save no one. Let us assume that each life is equally valuable, so that saving twenty people is twice as good as saving ten. The expectation of value from taking the risky option is therefore the same as the expectation of value from taking the safe option. In either case it is the value of ten lives. Consequently, according to Bernoulli's theory, both options are equally good.

But this is not obviously correct. It might be the case that the safe option is better just because it is safe. Risk to value – in this case the value of lives saved – may be a bad thing. At any rate, it is a substantive ethical question whether this is so. We should not accept a theory of value that rules it out, without facing up to the substantive question.

Modern expected utility theory

For this reason, modern expected utility theory is more general than Bernoulli's. It makes room for the possibility that risk to value is a bad thing. The modern theory says that the value of an option is the expectation of the *utility* of its possible outcomes. 'Utility' is a technical term, which is defined within the theory itself. Each possible outcome is assigned a number called its utility, and this is done in such a way that one option is better than another if and only if it has a higher expectation of utility.

Take the life-saving case, for example, and suppose that risk to value is indeed a bad thing. Then utilities might be assigned as follows: saving no lives has a utility of zero; saving ten lives a utility of ten; saving twenty lives a utility of eighteen. The expectation of utility from taking the risky option is then nine. It is less than the expected utility from taking the safe

option, which is ten. So the safe option is better according to expected utility theory.

In this example, we are taking it for granted that the *value* of saving twenty lives is twice the value of saving ten. But the *utility* assigned to saving twenty lives is less than twice the utility assigned to saving ten. Outcomes that have higher value will always have higher utility; utility increases with value, that is to say. But in this case utility increases less than in proportion to value. We may say that in this case value has 'diminishing marginal utility'. So long as this is a so, a safer option will always be better than a more risky one that has the same expectation of value.

This is the only significance of the diminishing marginal utility of value. In expected utility theory, utility is assigned to outcomes in just the way that is required to capture the badness of risk. So if risk to value is a bad thing, utility is assigned in such a way that value has diminishing marginal utility. But it may also be the case that risk to value is not a bad thing. Whether or not that is so is a substantive question that we have left open. If risk to value is not a bad thing, then utility is assigned to outcomes in such a way that utility increases proportionally with value.

The relation between value and utility does nothing other than represent the badness of risk to value. We can think of utility as an adjustment that is made to value, with the sole purpose of representing properly the badness of risk to value, if risk to value is indeed bad. If risk to value is bad, the effect of the adjustment is to make losses more important than gains. The possible good outcomes of a policy will count for less than its possible bad outcomes. This is the way that expected utility theory takes account of the badness of risk. Take a policy that is likely to promote economic development, but on the other hand leads to a small chance that humanity will become extinct. If value has diminishing marginal utility, then in assessing the value of this policy, the potential economic benefit is discounted but the possible very large loss is given extra weight.

Axioms

That is the way that modern expected utility theory works. It is a generalization of Bernoulli's theory, designed to allow for the possibility that risk to value is bad. I asserted that it is the correct way to evaluate policies in the context of uncertainty. But how do we know that is so? In introducing expected utility theory, I started with gambles about money. On a first look at those cases, it seems intuitively correct to judge the value of an option by the expectation of money it leads to. So with this example we intuitively recognize the importance of expectations. Our intuition generalizes easily to the idea that the value of an uncertain option is the expectation of value that it leads to. But now we have reached the further generalization that the value of an uncertain option is the expectation of utility that it leads to, and utility is defined in such a way that it may be different from value. Utility is not an intuitive idea, so by this point in the development of expected utility theory we cannot get much support from intuition. So how can the theory be justified?

Modern expected utility theory is not justified by intuitions, but instead on the basis of a number of assumptions or *axioms*. If the axioms are true, we can be sure the theory is true, because it can be proved from the axioms. So to justify the theory, we must justify the axioms. Not all the axioms are indubitable, and indeed at least one is probably false. But all can be given strong support at least within a particular range of application.

Each axiom makes a claim about the structure of values. For example, one is the axiom of *transitivity*. It says that, if one option A is better than another option B, and B is in turn better than a third option C, then A is better than C. This is certainly true. Transitivity is an extremely secure axiom because it follows from the meaning of 'better than'.

A debatable axiom is *completeness*. This says, for any two options A and B, that either A is better than B, or B is better than A, or A and B are equally good. We commonly think that certain pairs of values are *incommensurable* with each other, which means roughly that they cannot be measured on the same scale. Climate change raises important examples. For instance, the value of protecting nature may be incommensurable with the value of promoting human wellbeing. Suppose a policy A will protect nature to a high degree, but do little for human wellbeing, whereas policy B will promote human wellbeing but not protect nature. Then A and B may be *incommensurate* with each other. This means that A is not better than B, and B is not better than A, and furthermore, A and B are not equally good. If this is so, the completeness axiom is false of A and B.[4]

Incommensurable values raises severe problems for correct decision-making. Suppose a government has to make a choice between these two policies A and B. It cannot make its decision on the basis of which is the better policy, because neither is. Somehow, on some other basis, the political process will have to come up with a decision.

Expected utility theory fails in the face of incommensurability. But that does not make it useless. Even when a problem involves incommensurability, there may be big pockets of commensurability within it, where expected utility theory has a role. Take again the comparison between A and B, where both natural values and human values are in play. Those two sorts of values may be incommensurable with each other. Nevertheless, all the human values may be mutually commensurable. If so, it will be possible to calculate an expected human value for each policy. This is an important piece of information, to be set against the natural values.

Moreover, incommensurability is rarely, if ever, total. We are supposing that the particular policies A and B are incommensurate with each other. But compare policies C and D, where C is much better for nature than D is, but D promotes human wellbeing only a little bit better than C does. Then C and D may well not be incommensurate; C may be overall a better policy than D. So expected utility theory may allow us to establish a useful partial ordering of policies.

Objective and subjective probabilities

In the examples I have used so far, each possible outcome of a gamble has an ‘objective’ probability, which I specified. Gambling devices such as roulette wheels or coin tosses are designed to be operated repeatedly, and they are designed so that their repetitions have a particular property. As the number of repetitions increases, the relative frequency of each particular outcome will converge to a particular number. This number is an objective probability. In repeated tossings of a fair coin, the relative frequency of heads converges to a half, so heads has an objective probability of a half. Gambling devices have objective probabilities because of these properties.

But hardly any of the uncertainty we face in life is associated with objective probabilities. Each day is unique, and few of the situations we face each day are repeated in a way that determines any objective probabilities. The uncertainties associated with climate change do not have objective probabilities. There is no objective probability of, say, the Gulf Stream’s shutting down.

Fortunately, the axioms of expected utility theory are not confined to situations where there are objective probabilities. The theory is justified by its axioms, and the justification of the axioms (such as it is) does not depend on the existence of objective probabilities. The theory applies to all sorts of uncertainty. Even if there are no objective probabilities, it nevertheless tells us that the value of an option is its expectation of utility. Since an

expectation is a weighted average – weighted by probabilities – this means that the value is given by probabilities, even if no objective probabilities are available. When there are no objective probabilities, we need subjective ones.

When we are trying to evaluate our options, in effect the theory tells us to assign probabilities to the various possible events as well as we can, using whatever information we have. Then we should use those probabilities consistently in making judgements about the values of the various options. In the case of climate change, we must use the same probabilities when assessing the value of each different policy. For instance, whatever policy we are assessing, we should assign the same probability to a particular value of climate sensitivity.

Among the probabilities we need to use are probabilities that are conditional on which option is adopted. For instance, there is a probability of global sea levels' rising 50cm in a century, conditional on our adopting a policy of moderate restraint. There is a different probability of sea levels' rising 50cm in a century, conditional on our adopting the policy of business as usual. When evaluating the policy of moderate restraint – calculating the expected utility of its possible outcomes – we should use probabilities that are conditional on that policy. Obviously, we should not simply use the unconditional probability of the sea level's rising by 50cm in a century.

The reason we need to assign probabilities and use them consistently, even when no objective probabilities are available, is that otherwise our judgement of value will be inconsistent with each other, in one way or another. They will not conform to the axioms of expected utility theory. For example, they may not be transitive. To avoid inconsistency, each of our judgements of value needs to be based on the same probabilities. So we need to find the best probabilities we can, and use them consistently.

When there is uncertainty without objective probabilities, other methods of evaluation have often been proposed as rivals to expected utility theory. For example, there is *maximin*. This is the view that one option is better than another if and only if the worst possible outcome the first option could lead to is less bad than the worst possible outcome the second could lead to. Among a collection of alternative policies, maximin says that the best is the one whose worst possible outcome is less bad than the worst possible outcome of any of the others. Maximin pays no attention to possible outcomes of any policy other than the worst. In a sense, it is entirely devoted to reducing risk, with no other objective at all. It has the advantage that it can be applied without making any judgements about the probability of the various possible outcomes. But maximin is inconsistent with expected utility theory; if expected utility theory is true, maximin is false. Evaluations based on maximin must violate one or other of the axioms of expected utility theory. Since those axioms are well founded, it is safe to reject maximin. We cannot escape the need to make judgements of probability.

Some versions of the *precautionary principle* imply that we should never adopt any policy that leads to any risk of serious environmental damage. This idea resembles the maximin principle in some respects. One is that it can be applied without making any judgement about the probability of serious environmental damage. However, like maximin, this version of the precautionary principle is not consistent with expected utility theory. (One exception is where the potential environmental damage is infinitely bad. But I shall explain below that this is not possible.) Since expected utility theory is correct, this version of the precautionary principle is not correct. When there is a risk of serious environmental damage, we have to make some judgement about its probability, to use in our calculations of expected utility.

Other versions of the precautionary principle are consistent with expected utility theory. I shall mention them in a moment.

The special uncertainty of climate change

Le me summarize what I have said about expected utility theory so far. I have been considering how to cope with the great uncertainty involved in working out the effects of climate change. My conclusion so far is this. We have to evaluate alternative policies that respond to climate change. Each policy will lead to a range of possible outcomes. We first need to assign a value to each of those possible outcomes; we shall come to how to do that later. Then we need to take account of the badness of risk, if risk is bad. We do this by adjusting values to utilities. Then we assign probabilities to each of the possible outcomes of each policy. On the basis of these probabilities, we calculate an expectation of utility for each policy. This expectation tells us the relative value of the policies: one policy is better than another if its expectation of utility is greater. All this is to follow the recommendation of expected utility theory.

Now to the particular uncertainty associated with climate change. This is often held to be unusual in at least three respects. The first is that it is very extensive; our predictions about the results of climate change are very uncertain indeed. The second is that many of the changes that climate change will cause are irreversible. The third is that climate change creates a small probability of quite terrible catastrophe: the extinction of humanity or even of all life on Earth. Expected utility theory deals correctly with all these features. They are no grounds for thinking we need a different way of dealing with the uncertainty. But they do need some special attention.

The first unusual feature means that there are no objective probabilities, and that suitable subjective probabilities are very hard to come by. Nevertheless, we cannot escape the need to make judgements of probability; we must do as well as we can. We have to make them, even when the scientific basis of them is insecure. The Rio Declaration says ‘Where there are threats of serious or irreversible environmental damage, lack of full scientific certainty shall not be used as a reason for postponing cost effective measures to prevent environmental degradation.’[5] This was intended as one expression of the precautionary principle. It is not only consistent with expected utility theory; it is exactly what expected utility theory requires.

However, we also have to recognize that reasonable people may differ widely in their judgements of probability. Expected utility theory tells us that each person needs to stick to her probabilities consistently. But it recognizes that different people might stick consistently to different probabilities. Consequently, different people might arrive at different conclusions about the values of various policies. One person might rank a particular policy above another; another rank it below.

It is pointless to insist on some particular set of probabilities unless they are widely accepted by reasonable people. I think the only adequate way of accommodating differing views is to assess the values of all the policies several times, each time using different probabilities. This is a sort of sensitivity analysis.

The second special feature of the uncertainty that surrounds climate change is irreversibility. For example, climate change is causing the extinction of species, and those species will never be seen again. It is destroying ecosystems that can never regenerate. Other events besides climate change cause irreversible effects, but climate change is a particularly potent cause of them.

What difference does irreversibility make to evaluations? One thing is that it magnifies the goodness or badness of the effects of climate change. If it is bad for the world to be without

polar bears, then this is a bad thing that the world has to suffer for ever. This adds to the importance of the decisions that have to be made, but it does not alter the structure of the problem.

A second consequence of irreversibility is that it gives us reason to be cautious in our decision making. If a decision cannot be reversed, we need to make it with special care. We should gather as much information as possible before making a decision that has irreversible consequences. Often, we should postpone the decision, to allow more information to come in. This is a further aspect of the precautionary principle. At first it may seem to cast doubt on expected utility theory, because it can seem as though expected utility theory cannot recognize the value of postponing decisions. But actually, once again, expected utility theory gives exactly the right value to it. An example from nuclear energy will show how.

Restarting our nuclear programme will help to reduce climate change. It will be beneficial in that respect. On the other hand, it may have an irreversibly bad consequence: it may permanently leave a dangerous residue of radioactive waste. We are uncertain about this consequence, because we may find a good solution to the waste problem, or we may not. We now have two alternative policies to choose between. We can restart the nuclear programme, or we can eschew nuclear energy. In evaluating these two alternatives, one of the things we need to do is form our best judgement of the probability that we will be able to find a good solution to the waste problem.

Suppose we make this judgement of probability, and suppose that, on the basis of it and our best probability judgements for all the other uncertainties in the problem, it turns out to be better to restart than to eschew. Does it follow that this is what we should do? It does not. There is a third policy we also need to consider, and that is the policy of delaying the decision for, say, ten years.

Compared with restart, the policy of delay has costs and benefits. A cost is that some of the reduction of climate change will be lost. A benefit is that a better decision will be made later on. We know now that, later on, we shall be able to form a better judgement of the probability of finding a solution to the waste problem.

How does this feed into our present calculation of value? To keep things simple, let us suppose we know now that, in ten years time, either a solution will have been found or we shall know for certain that no solution will ever be found. Furthermore, let us suppose that, on the basis of our present best judgements, we can work out that, if a solution is found by then, it will be best at that time to restart, whereas if no solution is found by then, it will be best at that time to eschew. None of this makes any difference to our present judgement of probability. We are already making that judgement as well as we can, and on our present judgements it is better to restart now than to eschew now.

You might think as follows. If in ten years it turns out better to restart, and we do so then, the result will be definitely worse than restarting now, because some of the benefit will have been lost. On the other hand, if turns out worse to restart then, and we do not do so, we shall be adopting exactly the policy of eschew. We already know, on the basis of our present probabilities, that this is a worse policy than restarting now. Either way therefore, delaying ten years can only lead to a worse result. So delay must be a worse policy than restart now.

It may look as though this is how expected utility theory tells us to calculate. So it may look as though expected utility theory cannot recognize the value of delay. But this is not actually how the theory tells us to calculate. We need to consider these five possible outcomes:

(a) Restart now, and a solution to the waste problem is found.

(b) Restart later, and a solution to the waste problem is found.

(c) Eschew.
(d) Restart now, and no solution to the waste problem is found.
(e) Restart later, and no solution to the waste problem is found.

These are arranged in order from best to worst. Now consider the possible outcomes of the three policies. Eschew leads to (c) for certain. Restarting now leads to either (a) or (d). At present we assign a probability to (a), which is our best judgement of the probability that a solution to the waste problem is found. Call this probability *p*. We assign a probability of (1–*p*) to (d). Finally, there is delay. Let us assume that, if we delay, in ten years' time we shall make the decision that in ten years' time turns out to be best. Then delay leads to either (b) or (c). It leads to (b) if a solution to the waste problem is found, and to (c) if no solution is found. So, from our present perspective, the probability of (b) is *p* and the probability of (c) is (1–*p*), if delay is chosen.

So, restart now gives (a) with probability *p* and (d) with probability (1–*p*), whereas delay gives (b) with probability *p* and (c) with probability (1–*p*). (b) is worse than (a), but (c) is very much better than (d). So a calculation of expected utility may well show that delay is the better policy. Expected utility theory works well in this case. However, there is one important thing to learn from the example: in calculating the expectation of utility, we sometimes need to take into account our own expected future decisions.

The possibility of catastrophe

The final special feature of climate change is that it may possibly lead to a quite terrible catastrophe. One might think such an extreme negative value calls for a special way of thinking, but once again expected utility theory is up to the job.

Indeed, the first stirring of expected utility theory in history involved an even more extreme value than this. Blaise Pascal used the expectation of value to argue it is better to believe in God than not to believe in him.[6] His argument is hard to follow, but the common interpretation of it is this. There are two possibilities: God exists or he does not. You have two alternatives, to believe in God or not. If you believe in God, and he exists, then you will gain an infinitely long and happy after-life. If you believe in God and he does not exist, then you neither gain nor loose anything. Therefore, the expectation of value from believing in God is the product of infinity times whatever is the probability that God exists. So long as this probability is not zero, the product is infinite. So believing in God leads to an infinite expectation of value. On the other hand, not believing in God leads to only a finite expectation. It is therefore better to believe in God. Given the premises, this argument seems valid. So expected utility theory seems able to cope with even infinite values.

Climate change does not raise the possibility of an infinite loss. The extinction of life on Earth would be dreadfully bad, but not infinitely bad. Life on Earth has only a finite time to survive anyway, and Earth has only finite space, so there isn't really the opportunity for it to contain infinite value. Moreover, the extinction of life on Earth is not the worst thing that could happen. For example, a worse possibility is that the Earth continues for a very long time to be occupied by as many people and animals as it has at present, but all of them live all their lives in horrible torment. So we are dealing with the possibility of a very large but finite loss. If expected utility theory is appropriate for infinite values, it is appropriate for finite ones, however large.

However, although expected utility theory is in principle appropriate for the problem, the very fact that the value at stake is finite makes the theory harder to apply in practice. Pascal was able to show that the value of believing in God is infinite, because the product of infinity and any positive number is infinite. So he could calculate the expectation of value from

believing in God, even without knowing the probability that God exists. On the other hand, climate change leads to an extremely small probability of an extremely large loss. Both the probability and the value (and hence utility) of the loss are extremely hard to estimate. The expectation of utility from the loss is the product of the probability and the utility. The product of an extremely small but incalculable number and an extremely large but incalculable number could be just about anything. On some judgements of probability and value, the possibility of extinction could be by far the worst consequence of climate change; on other judgements it could be unimportant.

The major difficulty here is not in the estimation of probability but in the estimation of utility. Assigning a probability to the extinction of life on Earth is a task within science. It is a matter of estimating the probability of runaway global warming caused by positive feedback effects. It is within the class of questions that climate scientists are already working on. But the task of estimating the badness, and the negative utility, of extinction is a task within the theory of value. No one has seriously approached it yet, and it is afflicted by at least one extremely difficult problem that I shall come to.

Section 3. Values and aggregation

Types of value at stake

This problem I have just mentioned is at the extreme edge of the range of extremely difficult problems of value that climate change raises. I now turn to those problems. Expected utility theory tells us how we should account for the uncertainty that afflicts all work on climate change. When we set out to evaluate a policy, expected utility theory requires us to start by assigning a value to each of the policy's possible outcomes. In doing this we should abstract from uncertainty. We should imagine history unfolding in one definite way, and set a value on that way. We should do this for each of the ways in which history might unfold. Then we should put all those separate values together in the way expected utility theory specifies, in order to arrive at a judgement of the policy's value. Having discussed uncertainty, now I come to evaluating possible outcomes.

The outcomes of a policy will have many different features that are relevant to its value. I shall at first concentrate on these three: the condition of the natural world, people's quality of life, and the lengths of people's lives. Each of these is a value that a policy can promote. I shall comment on each in turn, and then come to the question of aggregating them together. Finally, I shall come to a fourth sort of value.

Nature

Nature first. Climate change is set to do great damage to the natural world. Plants and animals are already becoming extinct at an unprecedented rate, and ecosystems are being impoverished. Policies that reduce climate change will reduce the damage. How are we to evaluate this benefit?

Nature has two sorts of value. First, it brings benefits to human beings in many different ways. The beauty of nature refreshes and inspires us. Wetlands process our waste products. Wild plants provide us with powerful new medicines. And so on. These benefits contribute to people's quality of life, and I shall place them under that heading. Here I shall consider the other sort of value. Nature is valuable for its own sake, quite apart from the benefits it brings to human beings. This is the value that environmental economists have come to call 'existence' value. I shall call it the 'intrinsic' value of nature.

It must be taken into account. Indeed it may be a central consideration; damage to nature may well be one of the most harmful consequences of climate change. However, there is a great deal of disagreement about the intrinsic value of nature. Some people think nature has little or no intrinsic value; others that is very valuable indeed.

One extreme view can perhaps be safely rejected. Some people believe nature has no intrinsic value because they believe nothing has any value except in so far as it contributes to the benefit of human beings. This radical view is hard to sustain. If an animal suffers pain, that seems clearly a bad thing, even if no human being is affected. If it is indeed a bad thing, it is an example of a negative value that has nothing to do with human beings. Because the pain of animal seems clearly bad, it seems safe to reject the radical view. But a less radical view is more defensible, and it attributes little intrinsic value to nature. This is the view that only the experiences of sentient creatures have value. It rules out much of the putative value of nature, since sentient creatures are only a small part of nature.

So there is a defensible view that nature has little intrinsic value. On the other hand, views can also be defended that value nature very highly. So the disagreement is very extensive. We cannot expect to resolve it soon. I think this means it is best not to try and measure the value of nature in any formal way. Doing so would be far too difficult and the results would be too controversial.

In particular, I suggest we should ignore the measure of value that has become established in environment economics. This measure extends to nature a standard method of valuing public goods: by people's willingness to pay for the good. Suppose we wish to know the value of a city park. We find out what each person is willing to pay for the existence of the park. We then add up all the people's respective willingnesses to pay, and the total is a measure of the value of the park in money terms. We can do the same for a piece of nature such as a natural forest, and thereby measure the forest's value. Some people may directly benefit from the existence of the forest; they may walk in it or enjoy watching birds in it. They may be willing to pay for the forest because of these benefits. Alternatively, they may hope to benefit from the forest in the future, and be willing to pay to keep it available. But these benefits are not part of the intrinsic or existence value of the forest. To measure this value, environmental economists exclude what people are willing to pay because of actual or potential benefits to themselves. They find out what people are willing to pay for the existence of the forest *apart from* any benefit they expect to derive from it. Then they add up all the people's willingnesses to pay for existence, and this gives them a measure of the existence value of the forest. In principle, this method might be extended to all of the effects that climate change will have on nature.

This way of assessing the intrinsic value of nature is mistaken. We may assume a city park is valuable only because it brings benefits to the people. Its value is therefore the total of the benefit it brings to everybody. A person's willingness to pay for the existence of the park is a measure of the benefit to her. The total of people's willingnesses to pay is a measure of the total benefit to everybody, and that is the value of the park. This is why it is appropriate to measure the value by adding up people's willingnesses to pay. Similarly, if we were trying to value the benefits a forest brings to people, it would be appropriate to do so by adding up people's willingnesses to pay for the existence of the forest. But we are actually trying to find the intrinsic value of the forest, not the benefit it brings to people. We are specifically excluding its benefit to people. There are therefore no grounds at all for adding up people's willingnesses to pay.

People are indeed often willing to pay for the existence of a forest even if it brings them no benefit. Of course, in that case, their willingness to pay does not measure the benefit the

forest brings to them. Instead, it reflects their judgement of the intrinsic value of the forest. If you think a forest has intrinsic value, you may feel some responsibility to help preserve it, and for that reason be willing to pay to preserve it. Your willingness to pay does not measure the forest's value to you; it reflects your judgement of the forest's intrinsic value, together with your judgement about your own responsibility towards preserving that value. The total of people's willingnesses to pay is consequently not a measure of anything. It is a meaningless quantity.

Some economists assume that no one is willing to pay for anything unless it brings a benefit to themselves. This leads them to assume that a person's willingness to pay for the existence of a forest must measure a benefit that the forest gives the person. This is confused, since it is the intrinsic value of the forest that is at issue, and this specifically excludes the benefit to people. Actually, people are willing to make payments for all sorts of altruistic purposes that bring no benefit at all to them. Since their willingnesses to pay for these reasons do not measure the benefit to them, the total of these willingnesses to pay measures nothing.

So we can set aside this way of measuring the value of nature. We have no alternative formal measure available for the value of nature. Some partial measures may be useful. For example the number of species that become extinct may be a useful (negative) partial measure. It may even be a useful general index of the state of nature as a whole. But the number of species is not the only thing that is plausibly valuable in nature.

Quality of life

The second value that a policy can promote is the quality of people's lives. We typically assume that the higher a person's income is, the higher will be the quality of her life. So we can use a person's income as rough measure of the quality of her life. We need a measure of the general quality of life in a nation or the world, and for this we generally add together, or average, the incomes of different people. An assumption is implicit in doing so: the assumption that a particular amount of income has the same value to each person. This is unfortunately false. A particular amount of money is worth more to a poor person than it is to a rich person. This means that the effect of adding up or averaging people's incomes is to give too much weight to the quality of life of richer people.

Quantity of life

The third value is lengthening people's lives. We may call this alternatively the quantity of life as opposed to the quality of life. Climate change will kill many people, so policies that reduce climate change will save many people. This may well be their principal benefit. How can we measure it?

As a very rough approximation, we could simply count the number of lives saved. Implicitly, this measure assumes that the benefit to one person of having her life saved is the same as the benefit to another of having her life saved. It is sometimes said to be correct: that each life has the same value to the person whose it is. But actually it is not correct. Saving a particular person's life may only give her a few more days or weeks of life, whereas saving another's may give her many decades. The benefit received by the first person is less than the benefit received by the second.

Saving a person's life is prolonging her life, and the benefit she gains is the extra time she spends alive. So a more accurate measure of a policy's benefit, in respect of life-saving, is the total number of years of life it adds to people's lives. This is a less rough approximation, but it is still only an approximation. Its implicit assumption is that a year of life to one person is

equally as valuable as a year to any another person. But that is not true either. Some ways of extending a life are better than others. Think of a person who is suffering from a crippling but curable disease. It is better for her to be cured before the disease has given her a permanent disability, rather than afterwards. Even if the cure leaves her with the same length of future life in each case, life with a disability is generally less good than life without. To generalize, the value of a year of life depends on the quality of life.

An even more accurate measure of the benefit of life-saving therefore takes account of the quality of life. The World Health Organization uses a measure that is more accurate in this respect. Its measure of the benefit of life-saving is the 'disability-adjusted life-year' or *daly*.[7] This is still not entirely accurate, because the WHO takes account only of medical disabilities in its measure. A more accurate measure would also take account of other aspects of the quality of life.[8]

So we have potentially a range of measures for quantity of life: number of lives, number of life-years, number of disability-adjusted life-years, and number of life-years more accurately adjusted. Each of these measures is more accurate than the previous one, but also more difficult to predict. The number of lives a policy will save is easier to predict than the number of disability-adjusted life-year, for example. We need to judge what degree of approximation we can tolerate. When we are comparing policies, this depends on where the differences between the different policies lie.

Suppose, for example, that one possible response to climate change tends to save the lives of people who do not have long to live anyway, whereas another response tends to save people who have longer to live. Then it will be unsatisfactory to measure the benefits only by the number of lives saved. There may indeed be this sort of difference between alternative policies. One predicted effect of climate change is to increase the number and severity of heat waves in temperate climates. Each heat wave kills many people who suffer from respiratory difficulties; this happened recently in France. But these people are often old or sick, and do not have long to live in any case. Compare a policy that is targeted at saving people's lives in heatwaves, with one that tends to save the lives of younger people – malaria prevention, say. If these policies save an equal number of lives, the second will be more valuable than the first. This suggests we should be cautious about measuring the benefit of policies just by the number of lives saved.

Similarly, if some policies systematically tend to save lives of better quality than others do, it will be unsatisfactory to measure the benefits of policies by unadjusted life-years.

Aggregating different values

Those are three values that a policy can achieve. There are other values too. Each policy will promote each value to a greater or lesser degree, and perhaps some to a negative degree. Each value is a component in the overall value of a policy. To judge the overall value, the components need to be aggregated together. How successfully can this be done?

There is first of all the difficulty that we have no measure of some values. The value of nature is one. If we cannot measure a value at all, we clearly have no means of aggregating it with other values. All we can do is describe a policy's effect on nature.

Even with values that we can measure, there may be a fundamental difficult over aggregating them together. I have mentioned it already: they may be incommensurable with each other. It is very plausible that some pairs of values are incommensurable. For example, the number of species on Earth may be valuable, and the quality of people's lives is certainly valuable, but it is not plausible that these two values can be put together on one scale. When policies promote different values that are incommensurable with each other, they cannot all

be arranged in one complete ranking from best to worst. At best we can hope for a partial order.

Two of the values I have considered are the quality of people's lives and the quantity of their lives. It is sometimes thought that these two values are incommensurable with each other. The preservation of life is thought to have a sort of sanctity that makes in incommensurate with the mundane goods that make up the quality of life. This is not correct. What makes it good for a person to live longer is that she will receive further good things in her life if she does. The further good things she will receive are goods of the mundane sort that make up her qualify of life. So the value of a life is just made up of mundane goods. Consequently, there is no incommensurability between the value of the quantity of life and the value of the quality of life. In principle, they can be measured on the same scale.

However, though this is so in principle, it is hard to put them on the same scale in practice. So far, we have the value of the quality of life measured in terms of people's incomes, which is an amount of money. We have the value of the quantity of life measured in terms of the number of lives saved, or years of life, or adjusted years of life. We have to aggregate the two. To do that, we need a rate of exchange between money and quantity of life: a value of life in terms of money, or of money in terms of life. But no correct rate of exchange can be found.

A rate of exchange is nevertheless in common use in cost-benefit analysis. It is generally based on people's willingness to pay to extend their lives. In its 1995 report, the IPCC considered how a cost-benefit analysis of climate change might be conducted, and it used a willingness-to-pay rate of exchange there.[9] But the use of a rate of exchange faces a deep difficulty. What one person is willing to pay to extend her life will be different from what another person is willing to pay to extend hers. Moreover, this is so even if the two people are rational, and are accurately promoting their own best interest. Generally, a poorer person will be willing to pay less to extend her life than a richer person will, because the poorer person has more urgent alternative uses for her money. In terms of money, then, a poorer person's life appears to be worth less than a richer person's. This conclusion was incorporated into the 1995 IPCC report, so that Indian lives were valued much less than American lives.

It is genuinely true that a poorer person's life is worth less *in terms of money* than a richer person. To put the very same thing another way: *in terms of life*, money to a poorer person is worth more than money to a richer person. This is exactly what one would expect. Money is indeed worth more to a poorer person than to a richer person. An extra dollar brings more benefit to a poorer person than it does to a richer one. On the other hand, it is reasonable to assume as an approximation that life is worth the same to a poorer person as it is to a richer one. It follows that the value of money in terms of life is greater to a poorer person than to a richer person.

This makes it impossible to establish a single rate of exchange between money and life. This is not simply a difficulty for willingness to pay as a way of determining a rate of exchange. There simply is no rate of exchange at all. The problem is that our measures of the quality of life and of the quantity of life are mutually incompatible. We are measuring quality of life by means of income, an amount of money. Doing so implies that the value of money to each person is the same. We are measuring the quantity of life in terms of the number of lives or years of life. Doing so implies that the value of a life or of a year of life to each person is the same. These are contradictory assumptions. So long as our separate measures are founded on these incompatible bases, we shall not be able to aggregate them together.

Poverty and length of life: double counting?

A further question arises over aggregating together the measures of the quality of life and quantity of life. Climate change will kill people. It will do so by various means: more people will die in heat waves; more will die in storms and other bad weather; more will be killed by tropical diseases, which will increase their range. It will also kill people through poverty. It will make people in some areas of the world poorer than they presently are, and poverty tends to shorten people's lives. So some of the killing done by climate change will result from a reduction in income, which is the measure we use for the quality of life. This raises the question of whether our measures of the quality of life and of the quantity of life are truly independent of each other.

Compare a different case. Imagine that, rather than quantity of life, we were dealing with a different value: nutrition. Suppose we measured quality of life by means of income and we separately had some measure of people's nutrition. We should not try in some way to aggregate these two measures together to reach an overall value, because nutrition is part of quality of life. It is one of the things that people's income is spent on. In measuring the quality of life, income should already incorporate that value. We might worry that it similarly incorporates the value of the quantity of life, or at least a part of that value.

This would be a serious worry if the length of a person's life was, like nutrition, just one of the various values that a person spends her income on. But this is not so. A person's income and the length of her life are generally fairly independent causally. There is a causal connection, in that having a low income tends to shorten a person's life, but the connection is weak. The length of a person's life is determined by many factors besides the person's income.

But there are also cases where the connection is closer. Think about a person who has just enough income to keep her alive. She eats only to survive. In this case, we might think that survival is the only value that her income achieves. If that is so, we should not count income and length of life as different values to be aggregated together. However, it is not necessarily so. It depends on the details of our theory of value. For instance, we might think that assuaging hunger is itself a good thing, even if it goes no further than what is needed for survival. If so, then two values are achieved by income: assuaging hunger and extending life. There is a lot of room for debate about this.

Until that debate is settled, we cannot be confident that there is no serious overlap between our measure of income and our measure of the quantity of life. If we try to aggregate these two measures together, we may be double-counting to some extent.

The value of a person's existence

The possibility that humanity may become extinct raises different values. It seems natural to assume that our extinction would be a dreadfully bad thing. But we need to ask what would be wrong with it. Some of the things that would be wrong with it are of sorts I have already considered. First, the means by which extinction will happen will be through terrible damage to the environment. While it is happening, the quality of people's lives will be very poor, and people's lives will be shortened. They will suffer a reduction in both the quality and quantity of their lives. I have already considered harms of those sorts.

Second, humanity is a species. If the preservation of species is valuable, presumably the preservation of humanity is valuable to a high degree. So the extinction of humanity is a loss of another value I have already considered.

The difficult question is whether extinction leads to a further loss of a sort I have not yet considered. If humanity becomes extinct, there will be no later generations of human beings.

Huge numbers of people who otherwise would have existed will actually never exist. A great deal of life that otherwise would have been lived will actually not be lived. The total quantity of human life will be enormously diminished. Is this a loss? We have been assuming life is valuable. But that was in a different context. We were then thinking about events that shorten people's lives. We are now thinking about an event that will prevent the existence of people in the first place.

We now have to compare the value of two possible outcomes. In one of them a whole lot of people exist who do not exist in the other. (I mean 'exist' in a timeless fashion; everyone who ever lives, at any time, exists in my sense.) Does the fact that these people exist count in favour of the outcome where they do exist? Is it a value we should take into account? Many people have the intuition that it does not. They think that having a person in the world is not in itself valuable. They think that, once we have a person in the world, it is valuable to extend her life, but it is not valuable in the first place to have a person in the world. They value extending life, but not creating it. They think the existence of a person is in itself of neutral value – neither good nor bad. As the philosopher Jan Narveson succinctly put it: 'We are in favour of making people happy, but neutral about making happy people'.[10] Call this 'the intuition of neutrality'.

The intuition of neutrality is built into much of our practice of valuation. For example, when the National Institute for Clinical Excellence evaluated fertility treatment, it included the happiness of the parents among the benefits of successful treatment, but it did not include the existence of the baby.[11] Here is another example. When the life of a young person is saved, it is likely that she will later have a child of her own. This child is a person who would never have existed had her parent not been saved. But when transport economists assess the value of saving a person's life, they never include the existence of her future child as a benefit.[12] This is presumably because they do not count it as valuable in itself.

If the intuition of neutrality is correct, the extinction of humanity will be much less of a catastrophe that it might seem at first. It will prevent a lot of people from existing who otherwise would have existed. But their existence would not have been valuable in any case. That is what the intuition of neutrality tells us.

Actually, the intuition of neutrality has to be false. It cannot be consistently fitted into any theory of value. I shall not go into the argument here.[13] Its conclusion is that there must be a particular quality of life such that adding a person whose life has that quality is neutral – neither good nor bad. Call this 'the neutral level' of life. Adding a person with a better quality of life than neutral is a good thing, and adding one with a worse quality of life is a bad thing. I am sorry to say I cannot say what the neutral level is.

Given that the intuition of neutrality if false, the extinction of humanity might be a very great disaster indeed. It would prevent the existence of huge numbers of future people, and the existence of each one of them might well have been a good thing.

But to recognize this point we must give up the intuition of neutrality, and that will require us to revise our thinking about a great many things. NICE will need to revise its assessment of fertility treatment. When a baby results from treatment, the existence of the baby cannot be assumed to have neutral value. It will have some value (positive or negative). Transport economists will need to alter the value they assign to saving a young person on the roads. Furthermore, these alterations in value are likely to be large. A new child added to the world's population is likely herself to have children later, and in due course a whole line of descendants. This is a whole succession of people who would not have existed had the new child not been created. So the effect on the world's population (taking all times together) of adding a person is likely to be large. So giving up the intuition of neturality will force on us

some radical rethinking.

Moreover, the assessment of climate change will have to face up to an entirely new problem. As well as its other effects, climate change will certainly alter the population of the world. It will force people to migrate in huge numbers, and this cannot happen without altering people's decisions about having children. The intuition of neutrality allows us to ignore the effect on population, because the intuition is that it is neutral. But the intuition is false. Any change in population will be either a good thing or a bad thing. Moreover, the change in population, taking all times together, is likely to be large, even in comparison to the number of people who are directly killed by climate change.

To summarize: the intuition of neutrality is attractive but nevertheless false. Given that it is false, it may well be that the extinction of humanity has to be counted as an extraordinarily bad event. But recognizing that the intuition is false will also force us to reconsider radically the values we attach to all the effects of climate change. This is a serious problem for the ethics of climate change.

A final note on discounting

Discounting future lives – counting lives that come later in time as less valuable than ones that come earlier – would reduce the severity of this problem. It takes a long time for the population of the world to grow or shrink. That means that the big changes will lie far in the future. Discounting will therefore diminish their significance relative to values in the present.

However, as a solution to the difficulties over valuing population, discounting would be entirely ad hoc. Discounting is not itself anything to do with population issues. Discounting offers a solution only because it happens to take a long time for populations to grow and shrink, and from a theoretical point of view this is just a contingency.

In any case, it is not really plausible that a future life is worth less than present life. If we are to discount future lives, we have two options. We may do it in either a temporally neutral or a temporally relative way. The difference is this. Temporal neutrality means that the value of things is not affected by the temporal perspective from which we view them. Temporal relativity means that values may be relative to the perspective from which they are viewed. Either sort of discounting is implausible.

It is a surprising at first that there can be temporally neutral discounting, but it is indeed possible. A neutralist thinks that a life saved in 2020 is worth less than a life saved in 2010, and this is so at whatever time the valuation is done. Later lives are worth less than earlier lives, from whatever perspective we view them. The neutralist discounter will have to think that a life saved in 1066 was worth a huge amount more than a life saved in 1943, so the killing at the Battle of Hastings was enormously worse than at the Battle of Stalingrad. I know of no inconsistency in this view (even if the discounting is not exponential). The problem with it is that it is utterly implausible.

So most people who discount the future are temporal relativists. They mostly discount both forward and backward in time. So, from the perspective of 2006, a life saved in 2020 is worth less than a life saved in 2010, and also a life saved in 1990 is worth less than a life saved in 2000. But from the perspective of 1990, a life saved in 1990 is worth more than a life saved in 2000. Relativism suffers from a general problem, when a single agent (such as a government) occupies different perspectives in sequence. It makes for incoherence. The agent may do something today, judging it correctly to be the very best thing to do, and yet find in a year's time that is was not the best thing to do. Indeed, the agent may be able to predict today that, what she does will not be the best thing to do from the perspective of a year from now, and yet still judge today that it is the best thing to do. If she does something

today that she ought to do, in a year's time it may be the case (and predictably so) that she ought not to have done it. It is doubtful that this even makes sense.[14]

So either neutralist or relativist discounting is hard to accept.

There are some arguments in its favour. The best one is this. We can use resources to save some life today, or we can invest them at interest in order to save a quantity of life in the future. Because of the interest, these resources will save more life in the future. If life in the future is worth as much as life today, we should therefore invest the resources rather than use them today. Indeed, that is what we should do at any time. We should always save up your resources to save future life, and never use them to save present life. The consequence is that life-saving should be indefinitely postponed. That is a paradoxical conclusion. The paradox arises because we value future and present lives the same. So, the argument goes, to avoid it we have to discount future life.

This is a mistaken argument.[15] Over time, life (a life, or a life-year, or a daly) retains a constant value. But produced commodities fall in value over time, because present commodities can be converted through the productive system into a greater quantity of commodities in the future. If everything was in equilibrium, the price of life in terms of commodities would increase over time. The marginal quantity of life saved would become more expensive in terms of commodities. That is to say, the marginal quantity of life would cost more to save. As time goes on, we would save quantities of life that are progressively more expensive to save. We would keep people alive in progressively more expensive ways, at the margin. This equilibrium does not require all life-saving to be indefinitely postponed. It is not paradoxical at all. I do not know whether we are close to this equilibrium or not, but we easily could be. We certainly spend progressively more on the marginal life.

In conclusion, I see no convincing grounds for discounting future lives.

Notes

1. John Rawls, in *A Theory of Justice*, pp. 22–30, strongly emphasizes this point.
2. The grounds of expected utility theory, and this particular interpretation of it, are set out in my *Weighing Goods*, chapters 4–6.
3. In his 'Specimen theoriae novae'.
4. Incommensurability is described in my *Weighing Lives*, chapter 12.
5. See UNCED, *Rio Declaration*.
6. Pascal, *Pensées*, §233.
7. See Christopher Murray, 'Rethinking DALYs'.
8. See my *Weighing Lives*, section 18.3.
9. See Pearce *et al*, 'The social costs of climate change'.
10. Narveson, 'Utilitarianism and new generations'.
11. NICE, *Fertility*, p. 6.
12. For example, M. W. Jones-Lee, *The Economics of Safety and Physical Risk*.
13. The full argument appears in my *Weighing Lives*, chapters 10–12. A brief version is in my 'Should we value population?'.
14. See my *Weighing Lives*, section 4.3.
15. The reason why is spelt out more fully in my 'Discounting the Future'.

References

Bernoulli, Daniel, 'Specimen theoriae novae de mensura sortis', *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 5 (1738), translated by Louise Sommer as 'Exposition of a new theory on the measurement of risk', *Econometrica*, 22 (1954), pp.

23-36.
Broome, John, ‘Discounting the future’, *Philosophy and Public Affairs*, 23 (1994), pp. 128–56.
Broome, John, ‘Should we value population?’, *Journal of Political Philosophy*, 13 (2005), pp. 399–413. To be reprinted in a book belonging to the series *Philosophy, Politics and Society*, edited by James Fishkin and Robert Goodin.
Broome, John, *Weighing Goods*, Blackwell, 1991.
Broome, John, *Weighing Lives*, Oxford University Press, 2004.
Jones-Lee, M. W., *The Economics of Safety and Physical Risk*, Blackwell, 1989.
Murray, Christopher J. L., ‘Rethinking DALYs’, in *The Global Burden of Disease*, edited by Christopher J. L. Murray and Alan D. Lopez, Harvard University Press, 1996.
Narveson, Jan, ‘Utilitarianism and new generations’, *Mind*, 76 (1967), pp. 62–72.
National Institute for Clinical Excellence, *Fertility: Assessment and Treatment for People With Fertility Problems*, February 2004.
Pascal, Blaise, *Pensées*, J. M. Dent, 1932.
Pearce, D. W. *et al*, ‘The social costs of climate change: greenhouse damage and the benefits of control’, in *Climate Change 1995, Volume III: Economic and Social Dimensions of Climate Change*, by the Intergovernmental Panel on Climate Change, Cambridge University Press, 1996, pp. 179–224.
Rawls, John, *A Theory of Justice*, Harvard University Press, 1972.
UNCED, *Rio declaration on environment and development*, 1992.