

Session 2 – Risk, Probability and the Precautionary Principle in Scientific Scepticism

A Small Chance of Disaster

JOHN BROOME

Corpus Christi College, Oxford OX1 4JF, UK.

E-mail: john.broome@philosophy.ox.ac.uk

Expected utility theory tells us how we should make decisions under uncertainty: we should choose the option that leads to the greatest expectation of utility. This may, however, not be the option that is likely to produce the best result – that may be the wrong choice if it also creates a small chance of a great disaster. A small chance of disaster may be the most important consideration in decision making. Climate change creates a small chance of disaster, and some authors believe this to be the most important consideration in deciding our response to climate change. To know whether they are right, we need to make a moral judgement about just how bad the disaster would be.

Expected utility theory tells us the right way to manage uncertainty: we should act in accordance with the axioms of the theory. I believe expected utility theory is intuitively attractive to most people. Although psychologists have shown that empirically we do not conform to it, I don't think this is because we find the theory unattractive as an account of what we should do. If you explain it to people, they get it. We don't apply it correctly because it's difficult to apply.

So what does this theory say? Its old-fashioned version, which was invented in the seventeenth century by Daniel Bernoulli, says that you ought to act so as to maximize the mathematical expectation of value or goodness. (I treat 'value' and 'goodness' as synonymous.) When you are faced with a choice among alternative things to do, each will have a range of possible outcomes, and you should choose the alternative whose possible outcomes have the greatest mathematical expectation of goodness. (The mathematical expectation is the weighted average goodness of the possible outcomes, where each outcome is weighted by the probability of its occurrence.)

The modern version, which was invented in the twentieth century by Frank Ramsey, says you should maximize expected utility rather than expected goodness, where *utility* is an artificial notion that may be distinct from goodness. The main point of making the distinction is to allow for 'risk aversion' about goodness: to recognize that it may be rational to prefer a less risky option to a riskier one, even if it has a lower expectation of goodness. As it happens, for complicated reasons, I think Bernoulli got it right, and we should maximize expected goodness. But I don't intend to dwell on this, or on the distinction between value and utility.

Expected utility theory is intuitively attractive, but justifying it is another matter. The only sound justification I know is to support the axioms one by one. And this justification is not incontrovertible. There are some genuine reasons for doubting the theory. I'm not planning to dwell on justification either. Even if expected utility theory is not the whole truth, it is nearer to the truth than the alternatives.

Expected utility theory tells us that what matters in making the right decision is not necessarily what is likely to happen. In particular, if there is something that is not likely to happen, but will be very bad if it does happen, it may well be more important than what is likely to happen. It may dominate the calculation of expected utility. It is not likely that your house will catch fire, but if it does, the result will be very bad. Even discounted by the small chance of its happening, this bad possibility may outweigh the cost of a fire extinguisher. If it does, you should buy a fire extinguisher.

This is obvious. However, at least in the area of climate change, scientists seem focused on telling us what is likely to happen. The most recent IPCC report says what is likely to happen, what is very likely to happen, what is more likely than not to happen, and so on. Its predictions for increasing temperature and sea level concentrate on the modes of the distributions: which is to say the most likely outcomes. But, as mentioned, the most likely outcome may not be the important thing. What's unlikely to happen may be more important.

I sympathize with the scientists. They have only the data they have, and data inevitably tells you more about the centre of a distribution than about its edges. They are telling us what they know, but that unfortunately is not what we need to know about.

Furthermore, in the last IPCC report, the scientists did actually do better in one respect. The report does tell us about the whole probability distribution of the key parameter in climate change, which is something called 'climate sensitivity'. This is a measure of the sensitivity of temperature to increasing concentrations of greenhouse gas. It is defined as the number of degrees the atmosphere will eventually warm by, if the concentration is doubled above its pre-industrial level and held at the doubled level for ever. It gives you some idea of what actual increase in temperature we can expect, since we are on course to double the concentration within a few decades.

Different groups of scientists come up with different distributions. The modes of their distributions tend to be between one degree and three degrees Celsius. But a conspicuous feature is that all the distributions are asymmetrical. More than half of them give a greater than 5% probability to climate sensitivity's being more than 6.7°C.

We should be worried about this. It means that, if we don't do something serious about limiting climate change, warming of 6°C or 10°C or more is a real possibility.

Temperatures like those will alter the geography of the world. At the peak of the last ice age, when an ice-cap sat on Wales, the Earth was about 5°C colder than now. It has not been 10°C warmer for tens of millions of years. That temperature will eventually melt most of Antarctica, which will give us a sea level about 70 m above the present one. Farming will be much harder, and water much scarcer. It will be quite impossible for the Earth to sustain the billions of people we have now. There will be a collapse of our population, and quite possibly we will become extinct.

This is unlikely to happen. But as I say, an unlikely chance of something very bad may be more important than what is likely to happen. We need to think about whether this unlikely event should be our main concern.

This gives me another reason for sympathizing with the scientists. Given the forces of climate denial that are arrayed against them, they do not want to appear alarmist. They could easily be portrayed as extremists or lunatics. So they have to appear calm. To say that what we really have to worry about is the extinction of humanity would leave them too exposed.

Philosophers are less constrained, because they are expected to be insane anyway. And interestingly, economists are becoming increasingly concerned with this small chance of what they call ‘catastrophe’. Martin Weitzman is the economist most associated with this view.

This brings me to the second elementary lesson to draw from expected utility theory. The first was that we have to be concerned with more than what is likely to happen. The second is that we have to weigh up the values of things. It hardly takes expected utility theory to tell us that; it is obvious anyway. Making the right judgements about such things as climate change, and almost every other aspect of policy, involves weighing some good things against other good things. In the case of climate change we are constantly told that we, the current generation, have to make sacrifices for the sake of creating a better life for future generations.

It is not strictly true that sacrifices are required from the present generation to solve the problem of climate change. Climate change is caused by what economists call an ‘externality’: those who emit greenhouse gas, and benefit from doing so, do not bear the cost of their emissions. The costs are borne by all the people around the world who suffer their harmful effects. It is an elementary conclusion of economics that an externality causes what is called ‘Pareto inefficiency’, which is defined as a situation in which some people can be made better off without anyone’s being made worse off. It follows directly that no sacrifice is required from anyone to eliminate the effects of the externality.

However, although sacrifices are not strictly required, the evidence from economics is that the best way to solve the problem is for the present generation to make sacrifices. To know whether this is so, and if it is what sacrifices we should make, we need to weigh our potential sacrifices against the benefits they will bring to people in the future. This is just what economists such as Nicholas Stern and William Nordhaus have done. In doing so they are answering a moral question. The question of what sacrifices we make for other people is obviously a matter of morality.

This is obvious, but again it seems not to be recognized in the debates about climate change. Take the concept of ‘dangerous’, which is enshrined in the UN Framework

Convention on Climate Change: 'The ultimate objective of this Convention ... is to achieve ... stabilization of greenhouse gas concentrations in the atmosphere at a level that would prevent dangerous anthropogenic interference with the climate system.' What's dangerous interference? Plainly, it is interference that brings with it risks we should not tolerate. This is a matter of valuation: of judging the bad things that may happen, and their probability of happening, and setting this badness against the cost – another loss of value – of doing something about it. Dangerousness is not something that can be assessed by science alone. It requires moral judgement.

And yet the way the UNFCCC proceeds is to ask scientists to determine what concentrations of greenhouse gases would be dangerous. This is not scientists' business. I often find the view expressed (by economists and by scientists) is that ethics should be kept out of climate change. But the question of what we should do about climate change is a moral question.

I'm going to mention one particular place where we need ethics. The possible catastrophe that worries Weitzman and others involves a collapse of our human population, and perhaps our extinction. Weitzman claims the small risk of this catastrophe is the thing we should really worry about. But this is so only if it dominates the calculation of expected utility. And that is so only if the badness of the catastrophe is dominant even when it is discounted by the small chance of it occurring. We cannot decide whether Weitzman is right without thinking about how bad a catastrophe would actually be. Weitzman assumes we can put no bound on its badness. But that is clearly false. We are a finite species living on a finite planet. There has to be a finite limit on the badness of anything that can happen here.

Even so, you might think a catastrophe would obviously be so enormously bad that it should dominate our calculations. But actually that is not obvious. Very bad things are going to be caused by the climate change that is likely to happen. In particular, according to reports of the World Health Organization, tens of millions of people will die. Compare the badness of these tens of millions of deaths with the expected badness of billions of people's dying in a catastrophe. The latter does not greatly outweigh the former if there is, say, only a one in a hundred chance of the catastrophe's happening. Billions divided by 100 is tens of millions. So the killing expected in a catastrophe does not dominate the killing done by likely climate change.

If a catastrophe should really dominate our thinking, it will not be because of the people it kills. There will be other harms, of course. But the effect that seems the most potentially harmful is the huge number of people whose existence might be prevented by a catastrophe. If we become extinct within the next few thousand years, that will prevent the existence of tens of trillions of people, as a very conservative estimate. If those non-existences are bad, then this is a consideration that might dominate our calculations of expected utility.

To many people, the thought of our extinction seems at first appallingly bad. But many of us think differently when we look at things from another direction. Is the non-existence of people a bad thing? Most of us would answer 'No'. Most of us intuitively think that adding to the population of the world is not a good thing in itself, which means that not adding to the population is not a bad thing. Most of us think that having extra

people in the world is ethically neutral. We value improving the lives of the people there are, but we don't value adding extra people. I call this 'the intuition of neutrality'. If it's your view, on the face of it you shouldn't mind the extinction of humanity. Extinction simply prevents the existence of lots of people. If you think that is bad, in the face of the intuition of neutrality, you at least have something to explain.

All I want to say about this is that there is some work to do. This is a job of valuation: how bad is a collapse of population or extinction? It is a job for moral philosophy. Moral philosophers have in fact been thinking about the value of population for some decades now, and it has proved to be a particularly intractable problem. It now appears to matter practically for our decision-making about climate change.

About the Author

John Broome is the White's Professor of Moral Philosophy at the University of Oxford. He was previously Professor of Philosophy at the University of St Andrews, and before that Professor of Economics at the University of Bristol. He has worked on value theory, and now principally works on rationality, reasoning and normativity. His books include *Weighing Goods* (1991), *Weighing Lives* (2004), *Climate Matters* (2012) and *Rationality Through Reasoning* (2013).