Climate Ethics and Climate Economics: Risk, Uncertainty and Catastrophe Scenarios

Conveners

Simon Beard (Centre for the Study of Existential Risk)

Kai Speikerman (London School of Economics)

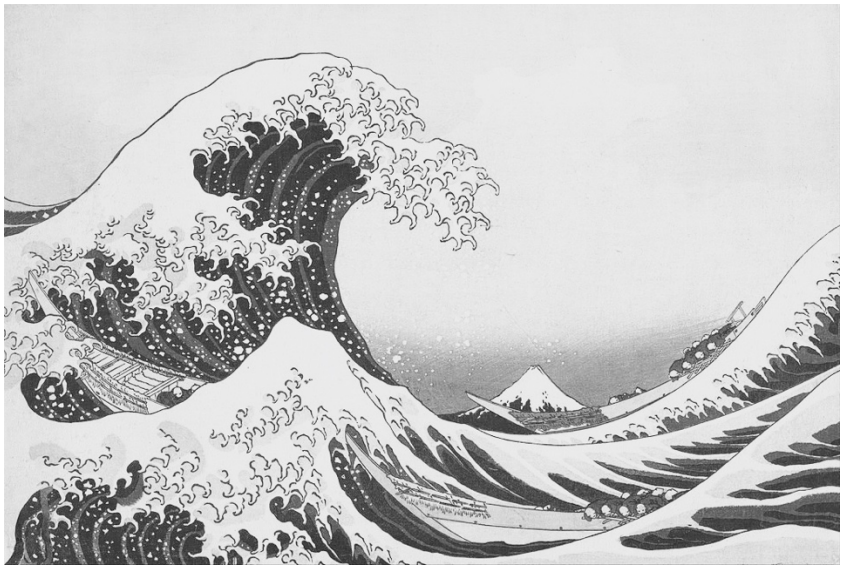

May 8th – 10th 2017

Workshop Report

Contents

Workshop Summary

Economic analyses of the effects of global warming often focus on the scenarios analysts believe are most likely to materialize. However, in recent years, both economists and moral philosophers have paid increasing attention to the possibility that climate change could prove much worse than we expect. However, it is not clear just how likely such possibilities are. Analysts and decision makers thus confront the problem of decision not only under risk, but under uncertainty. How should we respond?

One keynote speaker—Hilary Greaves (Oxford University) addressed the problem of 'cluelessness' about the consequences of our actions, and its implications for the effective altruism movement. The other, Doyne Farmer (Oxford), discussed the prediction of the effects of future forms of technology. John Halstead (Oxford) and Matthew Rendall (Nottingham) presented papers on decision making when we are uncertain which empirical or moral theories are true, while Elizabeth Baldwin (Oxford) and Kieran Marray (Oxford) examined high-stakes decision making with limited information. Mariam Thalos (Utah) discussed under which circumstances precautionary policy is appropriate rather than the maximization of expected utility, and Iñaki San Pedro (University of the Basque Country) addressed the problem of motivation in dealing with catastrophic risks. Tina Sikka (Newcastle) related feminist epistemology to scientific uncertainty, and Eike Düvel (Graz) examined moral and political problems connected with stranded carbon assets.

Coupled with the workshop were two public lectures by our keynote speakers. Opening the conference, Doyne Farmer argued that compared with the study of natural processes like climate change, we actually have a poor understanding of our collective effects on our own societies, and that greater resources could and should be devoted to understanding it. At its close, Hilary Greaves examined the connection between population growth and climate change, and argued that while it may well be a problem, the relationship is not a simple or linear one.

Summaries of papers presented

Full versions of many of these papers can be downloaded from http://www.nottingham.ac.uk/climateethicseconomics/papers/workshop-5.aspx

Keynote: Hilary Greaves (University of Oxford) – Cluelessness

Decisions, whether moral or prudential, should be guided at least in part by considerations of the consequences that would result from the various available actions. For any given action, however, the majority of its consequences are unpredictable at the time of decision. Many have worried that this leaves us, in some important sense, clueless.

In this paper, I distinguish between 'simple' and 'complex' possible sources of cluelessness. In terms of this taxonomy, the majority of the existing literature on cluelessness focusses on the simple sources. I argue, contra James Lenman in particular, that these would-be sources of cluelessness are unproblematic, on the grounds that indifference-based reasoning is far less problematic than Lenman (along with many others) supposes.

However, there does seem to be a genuine phenomenon of cluelessness associated with the 'complex' sources; here, indifference-based reasoning is inapplicable by anyone's lights. This 'complex problem of cluelessness' is vivid and pressing, in particular, in the context of Effective Altruism. This motivates a more thorough examination of the precise nature of cluelessness, and the precise source of the associated phenomenology of discomfort in forced-choice situations. The latter parts of the

paper make some initial explorations in those directions.

Keynote: Doyne Farmer (Oxford Martin School) - How Predictable is Future Technology?

Recently it has become clear that many technologies follow a generalized version of Moore's law, i.e. costs tend to drop exponentially, at different rates that depend on the technology. Here we formulate Moore's aw as a correlated geometric random walk with drift, and apply it to historical data on 53 technologies. We derive a closed form expression approximating the distribution of forecast errors as a function of time. Based on hind-casting experiments we show that this works well, making it possible to collapse he forecast errors for many different technologies at different time horizons onto the same universal distribution. This is valuable because it allows us to make forecasts for any given technology with a clear understanding of the quality of the forecasts. As a practical demonstration we make distributional forecasts at different time horizons for solar photovoltaic modules, and show how our method can be used to estimate the probability that a given technology will outperform another technology at a given point in the future.

John Halstead (University of Oxford) - Moral Uncertainty and Climate Change

Some climate economists, most notably Martin Weitzman, have argued for strong action against climate change on the basis that unmitigated climate change brings a non-negligible probability of a catastrophic outcome. Those who ignore these catastrophic tail risks and focus instead on the costs of the most likely levels of warming will tend to advocate for more moderate action on climate change. Many of those especially worried about the tail risks of climate change make some variant of the following claims:

- *Empirical Claim*: The probability of extreme global warming sufficient to destroy, or otherwise undermine the long-run potential of, human civilisation is >~1%.
- *Moral Claim*: Destroying, or otherwise undermining the long-run potential of, human civilisation would be extremely, perhaps infinitely, bad; the costs would swamp many times over all levels of warming that would not destroy, or undermine the long-run future of, human civilisation.

One of Weitzman's arguments, which has received very little attention in the literature, appears to run as follows:

1. The Empirical Claim and the Moral Claim entail that strong action on climate change has extremely high, perhaps infinite, expected utility, whereas rival reasonable theories entail its expected utility is bounded and much lower.
2. Rational policymakers ought to have non-negligible credence in the Empirical Claim and the Moral Claim.
3. Rational choice is a function of the degree of credence one has in different normative theories, and the stakes of the choice according to the theories in which one has credence.
4. Therefore, rational climate change policy should be guided by the Empirical Claim and the Moral Claim, rather than by reasonable rival theories.

Call this the *Metanormative Climate Change Argument* (MCC).

The question of how it is rational to act in the face of moral uncertainty has only recently received significant attention. My aim is assess the soundness of the MCC. If it is sound, the MCC shows that some of the arguments made by opponents of tail risk-based arguments for strong action on climate change face an insurmountable burden of proof. However, it remains very unclear whether the MCC succeeds because many of its premises are questionable. It is unclear whether metanormativism is true, unclear whether MITE is true, and unclear whether the requisite empirical assumptions about climate policy hold.

Matthew Rendall (University of Nottingham) - Betting on Theories: The Nordhaus-Weitzman Debate

Martin Weitzman's (2009) paper 'On Modelling and Interpreting the Economics of Catastrophic Climate Change' prompted a high-profile debate with William Nordhaus. Weitzman had estimated a roughly five percent chance of more than ten degrees of warming in the next two centuries without strong action, and about a one percent chance of more than twenty. Nordhaus argued that for such a catastrophe to materialize, the climate would have to prove improbably sensitive to greenhouse gas emissions, the damages from climate change would have to be improbably high, *and* scientists and governments would have to remain improbably complacent in the face of danger signs. Weitzman's reply was noteworthy. Suppose, he wrote, that a policymaker found his claims or another economist's less alarming ones

equally credible. The policy maker should accord his conclusions greater weight in decision making, because the consequences of ignoring his advice would be far more serious if it proved to be right. Weitzman was effectively advocating Pascalian wagering in response to expert disagreement.

This paper explores how to conduct such wagering. Policy analyses often implicitly follow the 'maxiprobability method'. The analyst selects the theory she believes most credible, and recommends the policies optimal if it is right. This seems like common sense, but is seriously misleading if one ignores the respective payoffs. Instead, we should assign subjective probabilities to the likelihood that competing theories are true, and use them to weight the expected value of competing policies' outcomes. This has become a popular approach among moral philosophers examining how agents should decide when they are unsure of which moral theory to follow. It is equally appropriate, however - and indeed less problematic - when applied to uncertainty about the truth of descriptive theories. Weitzman was right - when in doubt, we *should* often bet on the more catastrophic prognosis. A brief last section discusses how we can do this without its driving us mad.

Elizabeth Baldwin (University of Oxford) - Choosing in the Dark: Incomplete Preferences, and Climate Policy

I consider decision-making when the stakes are high, but information is poor, and outcomes may be far from our experience. My leading example is climate change. We do not know the probabilities of diverse outcomes; we disagree about societal impatience, risk, and inequality aversion. I offer a simple model of "justifiable acts", providing maximal agreement between decision theories, and facilitating quantification of the remaining disagreement. When this disagreement is large, I characterise the choice situation as "dismal". I demonstrate that the question of climate policy is "dismal". This modelling framework undermines certain straightforward relationships between quantities, which have been claimed in a recent literature. Seeing climate policy evaluations as "dismal" illuminates how subjective much of the literature on climate change economics really is, and so how poor a guide to policy this literature may form. It generalizes Weitzman's (2009) "dismal theorem", giving a broader view, which shows that it may be unnecessary or unwise to focus on highly unlikely events.

Kieran Marray (University of Oxford) - Dealing with Uncertainty in Ethical Calculations of Existential Risk

Taking Bostrom's definition of existential risk, this paper will argue that either we are not able to make guesses from our temporal perspective as to the certainty of things which we predict, or the current approach to assessing the gravity of existential risk needs qualification.

Firstly, the paper looks at the current approach of using expected-value utility as a guide to ranking different sources of potential existential risk. It then points out two flaws with this approach that stem from the issue of uncertainty. The first of these is the equal inclusion of risks which would be set into motion by some change in human action, such as AI risk, with other existential risks which would occur simply from a continuation of an ongoing situation, like climate risks. Using a variation of Parfit's famous thought experiment in 'Reasons and Persons' it will show that uncertainty over ability to prevent the risk governs whether one should consider the first type of existential risk at all in such calculations. If not, then they simply should not be allowed to occur, so do not need to be considered in the calculations.

Consequently, the paper then shows that the current approach does not correctly incorporate the two main types of uncertainty which occur when making the necessary predictions into its calculations. These are the certainty of the probability estimate being correct and the certainty of the expected outcome actually occurring if the event indeed happens. This means that such calculations are not reliable, even as a rough guide for assessing existential risk, as they simply assume that the predictions contained within them are themselves reliable. An example of a rough heuristic is then proposed to correct for this. This is not intended as definite, but merely as a starting point, which is of greater use than the current types of calculation. It is that the probability measure and potentially the utility measure in the expected utility calculations should be presented as a range, which increases as certainty about each decreases, so that the potential variation in the outcome of the calculation can be shown. e.g instead of $E(U) = p(x) \cdot U(x)$ as the current approach dictates, it would be $E(u) = [a \geq p(x) \geq b] \cdot U(x)$ or $E(u) = [a \leq p(x) \leq b] \cdot [c \leq U(x) \leq d]$ (where a and b are some probabilities and c and d are some utility/disutility levels, the range of which is dictated respectively by each's level of certainty).

Finally, the paper looks at the fact that any attempt to guess certainty is in itself inherently uncertain due to the weaknesses of assessing the certainty of modelling and so on from a single temporal perspective. Disregarding any further weighting due to the slippery slope into infinite regress which would occur if this were used, it instead suggests a temporally mutable approach and mind-set as the best way of normatively applying this type of heuristic.

Iñaki San Pedro (University of the Basque Country) - No Need for Catastrophism

Catastrophism, understood as any discourse or argumentation that has catastrophe at its core, is commonplace in the context of climate change. This approach sees catastrophe as a good reason to take such and such decisions or such and such action, regardless of whether the actual chances of catastrophic climate events taking place are high or low. Even if we might, should or must be hoping for the best, we must expect the worst and action needs to be taken on this basis.

In this paper, I argue against catastrophism as an effective way to influence individuals' decisions and their capacity to take action aimed at diminishing or mitigating the impact of climate change. In particular, I argue that, at least when it comes to individuals, catastrophism may not be an effective motivation for decision making and action.
I articulate my claims making use of the idea of possible worlds, on the one hand, and assuming, on the other hand, that catastrophe is fundamentally an agent based notion. Something is catastrophic for an agent in an environment if there is a dramatic radical change in the present state of affairs such that the environment is no longer in equilibrium with the agent, and re-adaptation on the part of the agent is almost impossible, at least in the short term. Also, I will take it that the capacity of taking real action on the part of an agent is directly linked to her awareness of both the present state of affairs and the significance of changes, big or small, in it.

My claim is that such catastrophes correspond to possible worlds which are located epistemically too far from an agent's present perspective for her to develop awareness as regards the situation. In contrast, considering non-catastrophic scenarios - i.e. closer possible worlds - will put the agent in a better epistemological position to gain such awareness, which is ultimately necessary for taking action. So, while worrying about possible catastrophes is important for various reasons, taking the possibility that a catastrophe is going to take place as the reason to take action in order to help prevent climate change may not be as effective as intended. Reference to non-catastrophic scenarios involving smaller but significant changes in our present state of affairs may instead prove more useful.

Eike Düvel (University of Graz) - Duties to Strand Assets and Compensatory Claims

Avoiding dangerous climate change will require us to leave a majority of carbon fuel (oil, gas, coal) reserves in the ground – reserves states and companies made plans and investments to extract. The current valuation of the companies holding these assets does not reflect the renewed political pledges to phase out carbon fuels rather earlier than later.

Recently, the issue of the risk of stranded assets has gotten a lot of attention, with divestment campaigns being increasingly successful in convincing investors to lower their exposure to carbon assets. Instead of focusing on the risks of stranded assets facing owners and investors of carbon fuel companies, we ask under what circumstances it is permissible for governments and the international community to implement policies and enter in agreements that will exacerbate the risk of assets being stranded or even mandate reserves to be left in the ground.

It is usually assumed that putting a price on carbon will lead to a phase-out of carbon fuel extraction. We argue that there are practical, political, and justice-related reasons to focus instead directly on responsibilities to leave carbon fuels in the ground. Furthermore, we sketch a 2-stage argument outlining what our responsibilities towards future generations affected by climate change are and how to allocate responsibilities to strand carbon assets. Finally, we consider some of the implications of a stranding approach, among them the question of compensation in the case of wrongfully stranded assets.

Mariam Thalos (University of Utah) - Precaution-first frameworks for decision analysis

Expected-Utility (EU) analyses are premised on applications in a context in which the decision in question, or similar such decisions, are faced again and again. But some decisions are never repeated, indeed some are unrepeatable, if only for the fact that some outcomes that result from decisions in such contexts are irreversible. So, EU is not well-founded in such contexts. This is the context in which one makes decisions that affect global climate.

What sort of decision analysis is appropriate in such contexts? Some researchers propose a Precautionary Principle, that we exercise an aversion to risk when there are large stakes. The trouble with this concept is that there is almost always a small risk associated with large stakes - for example, there are non-negligible risks to life associated with crossing busy streets - but we do not always feel it is appropriate to exercise extraordinary precaution. Ordinary precaution seems to be appropriate in such contexts, so why not in the iconic cases of risks to environment or climate?

I refer to frameworks focused on the scope of extraordinary precaution as *precaution-first* frameworks and argue that precaution-first analyses should be construed, not as antithetical to EU, but rather as giving articulation to where EU is appropriately applied. These are *routine* or *quotidian* cases of decision-making, calling only for *ordinary* caution, not the extraordinary kind. Circumscribing their appropriate application is the task of a precaution-first framework.

What are the necessary elements of the precaution-first framework? There are, I am convinced, at least three. First, there must be an irreversibility condition. Second, the framework has to articulate the terms under which tradeoffs in stakes can be conducted, whether in conditions calling for extraordinary caution or otherwise. And finally, the framework must address the question of potential incomparabilities in the system of values (preferences) being employed. Each of these elements is associated with a cluster of questions that a precaution-first framework should address. This paper focuses on the reasons for inclusion of these elements, and makes a beginning at structuring the precaution-first framework for decisions in public policy contexts.

Tina Sikka (University of Newcastle) - Feminist Epistemology: A Reevaluation of Uncertainty, Risk, and Climate Disruption

One of the most fundamental tensions in science communication, as it relates to climate disruption, is the how to deal with and reconcile the complexities that characterize contemporary climate science, one the one hand, with the popularized and politicized conception of science as acceptable only when there is 100% unanimity. In this paper, I take up the subject of uncertainty and science through the lens of feminist epistemology. Specifically, my objective is to reconcile a pluralist, situated, and value laden science rooted in feminist empiricism with a practice of science and scientific modeling that contains a large degree of uncertainty, and the need for public and political confidence in climate science. This is particularly necessary in a context where the calculus of risk, insurance, reparations, equity, trust and justice in regard to climate change are, as Ulrick Beck argues, not only unstable, but constituted by incalculability, non-compensatibility, and de-localization.
In pursuit of this, I discuss what climate modeling and science is meant to do, how to deal with the uncertainties inherent in their findings; how feminist science intersects with these uncertainties in particular ways; and the productive conclusions that can be drawn from these intersections. This establishes a path through which to reconcile scientific uncertainty with epistemological pluralism and, in doing so, takes care of a problematic line of argumentation, which asserts that a pluralistic science endangers scientific consensus and that uncertainties in science cannot be productively reconciled. These are distinct claims which, if misunderstood, can have the effect of policy paralysis and distorted public perceptions of climate change. When coupled with Beck's conception of risk, a productive reevaluation of how to deal with and, most importantly, understand climate science can be realized.

Summaries of public lectures

Hilary Greaves (University of Oxford) - Overpopulation: A driver of Climate Change?

It is often remarked that the significant drivers of climate change include not only high and rising levels of fossil fuel use per person, but also high and rising human population size. The logic behind this remark appears at first sight to be simple: climate change is driven by emissions, and total emissions are equal to per-capita emissions multiplied by population, so of course (one might think) higher population will lead to more climate change. I will argue that given a proper understanding of the physics of climate change, this simple argument is flawed. High population may indeed be damaging for reasons related to climate change, but if so, the reasons for this are more subtle; I will outline what they might be.

This lecture can be viewed at http://cser.org/event/hilary-greaves/

Doyne Farmer (Oxford Martin School) - Collective awareness: A vision of a new economics and how it could reduce risk

Science gives us a collective awareness that turns unknown unknowns into probabilities, and helps us deal with risks and avoid catastrophic scenarios. It is worth distinguishing three levels of collective awareness, that involve understanding the external environment, our effect on the environment, and our collective effect on ourselves. My lecture will focus on the hardest of these — our collective effect on ourselves — and on economics in particular. The economy underpins almost everything we do, and economic fluctuations cost the world many tens of trillions of dollars, yet the budget for polar research is greater than that for economics. Why is there no large scale effort to better understand the economy? I will argue that our lack of making a serious effort and our lack of progress is due to fundamental problems with the current culture of economics, and macroeconomics in particular. I will present an alternative vision of the economics of the future, with a much stronger emphasis on our ability to simulate the world. This will give us a better day-to-day understanding of the economy, but most importantly, it will allow us to better use science to think about the big problems in our future, such as climate change, the digital economy, and the overarching changes to human existence that the bio, info, nano and cognitive technologies of the future will bring.

This lecture can be viewed at http://cser.org/event/collective-awareness-a-vision-of-a-new-economics-and-how-it-could-reduce-risk-wprof-doyne-farmer/

Should we care about the worst-case scenario when it comes to Climate Change?

This article, written by Simon Beard, draws on a number of themes presented at the workshop and was first published on Huffington Post on September 27th 2017

After two of the most damaging hurricanes in history affected the Gulf of Mexico just a few days apart, the impact of climate-induced catastrophes is finally getting some attention. However, in truth, a few hurricanes, even ones that cause hundreds of billions of dollars' worth of damage, are far from the worst things climate change could cause.

Let's face it, climate change could be 'beyond catastrophic'. In the worst-case scenario, climate change could end human civilization once and for all, taking most of Earth's species with us.

At present, countries around the world have committed to holding global warming at less than 1.5 degrees Celsius above pre-industrial levels - "not yet a geophysical impossibility" according to a study published last week in Nature. This means around twice as much warming as we have already experienced. However, even if all countries honour these commitments, that may not be enough to stave off climate catastrophe.

For one thing, while we can aim to limit climate change to such levels, it's rather like playing 'pin the tail on the donkey.' Global climate systems are incredibly complex, and understanding the long-term impact of emissions is near impossible. For instance, the last time the earth's atmosphere contained as much Carbon Dioxide as it does now, global surface temperatures were 6 degrees warmer, and sea levels were 30 meters higher. It is estimated that under current plans to reduce emissions we still have a 10% chance of causing 6 degrees of warming or more, levels commonly associated with a real threat of civilizational collapse.

However, even if we do limit global surface temperature rise to 1.5-degrees, this still may not avoid catastrophe. For instance, hurricanes, like Harvey and Irene, are governed by oceanic temperature, not surface temperature, and these can be poorly correlated. Other equally important features of the climate include the availability of fresh water and the acidity of our oceans. Even if we manage to keep global temperatures to within 1.5 degrees Celsius of pre-industrial levels; if this is the only thing we achieved, then we could still face many other catastrophic global climate changes.

When I tell people about these kinds of problems, the response I usually get is 'well then, we're doomed!' However, this is a mistake. For one thing, all this complexity and uncertainty could also mean that we are safer than we feared. The climate may respond less to greenhouse gas emissions than our models predict and global temperature rises may be less dangerous than we thought. For another, even if things do turn out for the worst, it is not like we are powerless to do anything about it!

I work at the Centre for the Study of Existential Risk, a University of Cambridge research group working to prevent human extinction. We believe that humanity is a lot more resilient, creative and productive than many people give us credit, and that, so long as we do not go extinct, we have a promising future ahead of us. The challenge we face right now is making sure we survive to see it!

For us, avoiding global catastrophes is not just a good subject for science fiction, it is the key global challenge of our age.

The effects of climate change are not linear. For one or two degrees of warming, they amount to 'only' a few percentage points drop in global GDP, a few hundred thousand extra deaths each year and a few hundred million climate refugees. For six degrees of climate change, it becomes reasonable to talk about the total collapse of the global economy, billions of deaths and the prevention of trillions of future lives.

Seen this way, averting climate catastrophe should be our number one goal in mitigating and adapting to climate change. That certainly involves adopting the toughest policies towards reducing global greenhouse gas emissions, because even a 10% chance of a climate catastrophe is way too high. And it means studying the potential for these catastrophes much more so that we understand how likely they are and how to prevent them.

However, it also means not getting too attached to other climate goals, such as keeping global temperatures to within 1.5 degrees of pre-industrial levels. If, and when, we break through such targets, it will be a black day indeed in human history. However, uncertainty about the long-term impact of our emissions means that these goals may already have been broken, and to lose hope and give up on trying to prevent the greater danger of human extinction would be worse.

This workshop was sponsored by the Economic and Social Research Council as part of a series on 'Climate Ethics and Climate Economics' and is part for the Templeton World Charity Foundation funded project 'Evaluating Extreme Technological Risks'

The workshop took place on May 8th, 9th and 10th 2017 and was hosted by the Centre for the Study of Existential Risk at the University of Cambridge