

NBER WORKING PAPER SERIES

UNCERTAINTY AND DECISION IN CLIMATE CHANGE ECONOMICS

Geoffrey Heal
Antony Millner

Working Paper 18929
<http://www.nber.org/papers/w18929>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2013

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Geoffrey Heal and Antony Millner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Uncertainty and Decision in Climate Change Economics
Geoffrey Heal and Antony Millner
NBER Working Paper No. 18929
March 2013
JEL No. D81

ABSTRACT

Uncertainty is intrinsic to climate change: we know that the climate is changing, but not precisely how fast or in what ways. Nor do we understand fully the social and economic consequences of these changes, or the options that will be available for reducing climate change. Furthermore the uncertainty about these issues is not readily quantified and expressed in probabilistic terms: we are facing deep uncertainty or ambiguity rather than risk in the classical sense, rendering the classical expected utility framework of limited value. We review the sources of uncertainty about all aspects of climate change and resolve these into various components, commenting on their relative importance. Then we review decision-making frameworks that are appropriate in the absence of quantitative probabilistic information, including non-probabilistic approaches and those based on multiple priors, and discuss their application in climate change economics.

Geoffrey Heal
Graduate School of Business
616 Uris Hall
Columbia University
New York, NY 10027-6902
and NBER
gmh1@columbia.edu

Antony Millner
Grantham Research Institute
London School of Economics and Political Science
Houghton St
London, WC2A 2AE
United Kingdom
A.Millner@lse.ac.uk

Uncertainty and Decision in Climate Change Economics

Geoffrey Heal^{*1} and Antony Millner^{†2}

¹Columbia Business School, Columbia University

²Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science

March 27, 2013

Abstract

Uncertainty is intrinsic to climate change: we know that the climate is changing, but not precisely how fast or in what ways. Nor do we understand fully the social and economic consequences of these changes, or the options that will be available for reducing climate change. Furthermore the uncertainty about these issues is not readily quantified and expressed in probabilistic terms: we are facing deep uncertainty or ambiguity rather than risk in the classical sense, rendering the classical expected utility framework of limited value. We review the sources of uncertainty about all aspects of climate change and resolve these into various components, commenting on their relative importance. Then we review decision-making frameworks that are appropriate in the absence of quantitative probabilistic information, including non-probabilistic approaches and those based on multiple priors, and discuss their application in climate change economics.

Key words: Climate Change, Uncertainty, Expected Utility, Multiple Priors, Maximin

JEL codes: D81, Q54

1 Introduction

The climate problem is beset with uncertainties, many of which are only partially represented by our current analytical tools. Our aim is to explore these uncertainties, some of which are rooted in the underlying science, and some of which stem from imperfect understanding of how the socio-economic system will evolve over the coming centuries, and cope with climatic shifts. While there are robust broad-brush messages to be gleaned from the scientific and economic literatures, in many cases our detailed knowledge is ambiguous, and our uncertainty is ‘deep’

^{*}gmh1@columbia.edu

[†]a.millner@lse.ac.uk

– we are in a world of ‘uncertainty’ rather than ‘risk’ in the words of Knight (1921). How then should we make policy decisions in such an informational environment?

We will argue that the standard von Neumann-Morgenstern Expected Utility framework for decision making under uncertainty is of limited use for the climate problem. We thus explore alternative decision frameworks that are arguably more applicable and investigate how they can be applied to the issue of climate policy. Alternatives to expected utility maximization have been developed extensively in the last two decades, and there are several examples of their application to the climate problem.

There have been several surveys of uncertainty in environmental economics, and climate change in particular, including Heal & Kriström (2002) (which now serves to show how far the subject has moved on) and Pindyck (2007). A set of papers in this journal in 2011 has also examined some of the questions raised by fat-tailed probabilities of climate catastrophes.¹ These reviews focus on questions that have become the bread and butter of economic approaches to uncertainty – irreversibility, discounting, and the consequences of the standard bayesian approach to representing uncertainty. These issues are unquestionably important, but we take a different approach. We ask: what do we know about the climate problem, can we meaningfully characterize our knowledge with objective probabilities derived from data, and what will, in all likelihood, remain unquantifiable? Some of the most important uncertainties in climate economics, we argue, fall in the latter category.

We decompose the sources of uncertainty into scientific uncertainty and socio-economic uncertainty, the former arising from our incomplete knowledge of the climate system, and the latter from our imperfect understanding of the impacts of climate change on human societies and of how those societies will respond. Scientific uncertainty is in turn decomposed into internal variability, arising from the dependence of outcomes of complex climate models on initial conditions, model uncertainty arising from the fact that we don’t know what the ‘right’ climate model is, and emissions pathway uncertainty that naturally arises from our inability to predict future greenhouse gas emissions. Socio-economic uncertainty is broken down into positive or empirical uncertainty, arising from lack of knowledge about how socio-economic systems work, and normative uncertainty arising from uncertainty or disagreement over key welfare parameters such as the pure rate of time preference and the elasticity of the marginal utility of consumption.

Following this review of the sources of uncertainty and their relative importance, we turn to decision frameworks that can be used in the context of deep uncertainty, discuss how these can be applied to the climate problem, and review examples of such applications.

¹See also Millner (2013) for a comprehensive analysis of the role of fat tails and catastrophes in climate economics.

2 Scientific uncertainty

Human emissions of greenhouse gases are largely responsible for the observed changes in global temperature over the 20th century. Advances in the science of climate change detection and attribution over the past 20 years have placed this proposition beyond reasonable doubt (IPCC, 2007; Stone et al., 2009). So it's clear that as we continue to pump CO₂ and other greenhouse gases into the atmosphere, we can expect the climate to change further. What is much less clear is just how much, and how quickly, the global climate will respond to these changes in the composition of the atmosphere. This quantitative information is of primary importance for the purposes of informing a policy response to the climate problem – knowing that climate change is happening is one thing, but to know how best to respond to it, we need to know when, where, and by how much the climate is going to change.

Climate science can provide some answers to these questions, but they are perhaps not quite as precise as we would like. As an example of this, consider the scientific community's estimates of the famed climate sensitivity parameter. Climate sensitivity measures the average surface warming that results from a doubling of the atmospheric concentration of CO₂, *in equilibrium*. It thus tells us how sensitive global average temperature is to changes in CO₂ concentration in the long run, and is perhaps the most studied, and often quoted, summary statistic in all of climate science². Figure 1 plots a set of recent estimates of the probability distribution for climate sensitivity.

There is a lot of disagreement between the studies. To understand where the variability in these estimates comes from, it helps to understand a little bit about how they are produced. Climate sensitivity is usually estimated by running global climate models with hypothetical emissions scenarios. These models can be very complex, and there is a lot of uncertainty about the values of their many parameters. So scientists run the models in lots of different configurations, and weight the predicted value of climate sensitivity from each of these model runs according to some measure of the likelihood of that run's parameter values. There are many ways of assigning likelihoods to different model configurations. One is to assume that all model runs are equally likely. Another is to ask experts for their priors on the model parameters, and use these to weight the model runs. Almost all of the more recent studies use the historical instrumental record to score model configurations according to how well they reproduce past climate observations. Many studies also use Bayesian methods to combine expert priors with historical observations to arrive at their estimates. All of these techniques are used by the studies in Figure 1, but not always in the same way.

²Increasingly, scientists are calling for a focus on other summary statistics, partly because the uncertainty in climate sensitivity has remained essentially unchanged over several decades (Knutti & Hegerl, 2008), and partly because other measures, such as Transient Climate Response – the global mean temperature change that is realised at the time of CO₂ doubling if CO₂ concentrations are increased by 1% per year – carry more dynamical information about the near-term response of the climate system to increases in GHG concentrations. Other alternatives include peak warming from a cumulative emissions target (Allen et al., 2009).




Figure 1: Estimates of the probability distribution for climate sensitivity. Reproduced from Millner et al. (2012), data collated by Meinshausen et al. (2009).

Even if we restrict attention to studies that use historical data to constrain their estimates of climate sensitivity, there is a lot of variability in the estimates. Some use paleoclimate data, some use the cooling from the aerosols released by large volcanic eruptions, and some use only the instrumental record over the 20th century. Finally, and most importantly, even if we pooled all the data, and used a single consistent weighting scheme to evaluate model predictions, there would *still* be a lot of variability in estimates of climate sensitivity. This is because there is not a single ‘best’ climate model that everyone agrees should be used to estimate climate sensitivity. There are many models, each with its own representation of physical, chemical, and biological processes. There is so much we don’t understand about the details of the climate system that it is natural that there be a suite of competing approaches. The models have many features in common, but they are also different in important ways. The salient point is that they are not independent – we cannot treat the studies in Figure 1 as independent draws from some equally likely set of distributions. Because the models are all calibrated on related data, we also cannot objectively quantify *how* the models depend on one another. Finally, it is important to be aware that the key verification data that would allow us to decide between models is impossible to get. We would need to have a copy of the earth, force this planet with increased CO₂ emissions, and rank models according to how well they predicted this fictitious planet’s climatic response. Given these methodological

issues, it is difficult to know how to interpret the information in Figure 1 in a standard probabilistic framework without making some strong subjective judgements about how to combine estimates from different studies (Knutti et al., 2010; Knutti, 2010).

It would however be much too strong to say that there is *no* information in the set of distributions in Figure 1. It's clearly very unlikely that climate sensitivity is less than 1°C . It's also clear that a lot of the weight in most of the distributions is in the range $2 - 4.5^{\circ}\text{C}$ – the official IPCC 'likely' range for climate sensitivity. Finally, a lot of the disagreement between the estimates occurs in the upper tails of the distributions – the data don't constrain the 'high end' of climate change well at all (Allen et al., 2006; Roe & Baker, 2007). Thus, although the science provides us with some robust broad-brush messages, the detailed empirical predictions are certainly ambiguous – there is not a single, objectively arrived at, gold standard model that everyone agrees provides the answer to the climate sensitivity conundrum.

Climate sensitivity is a very coarse measure of the magnitude of the climate problem – on its own, it doesn't tell us much about when and where changes are likely to occur. The main contributors to uncertainty in temperature projections differ by the lead time of the prediction, and even by geographical region. Hawkins & Sutton (2009) decompose the contributions to total uncertainty into three components – initial condition uncertainty, model uncertainty, and emissions scenario uncertainty. Figure 2 shows how this decomposition plays out for predictions of global average temperature. For short lead times – i.e. less than about 20 years into the future – most of the uncertainty derives from uncertainty in the initial conditions of the models and model uncertainty. The fact that initial condition uncertainty plays an important role for shorter term predictions should not surprise us. The climate system is chaotic, and chaotic systems, by definition, exhibit sensitive dependence on initial conditions. As Hawkins & Sutton (2009) show, initial condition uncertainty is even more important for regional predictions, making substantial contributions to total uncertainty for lead times of 60 years or more. Back at the global level, for projections 20–50 years into the future, the dominant source of uncertainty is 'model' uncertainty, i.e. the spread in predictions across the set of climate models. Beyond about 50 years, emissions scenario uncertainty takes over – the policies we end up choosing drive the uncertainty in the long-run.

The effect of CO_2 concentrations on the global average temperatures we've discussed above is probably one of the *best* constrained relationships in climate change science. In a hierarchy of increasing uncertainties, climate sensitivity and global mean temperatures are somewhere near the top. As one moves down the ladder to smaller spatial scales, impacts that are localized in time, and less well understood climate variables (e.g. precipitation), uncertainties explode. Masson & Knutti (2011) have analyzed how the uncertainty in climate predictions depends on spatial scale. Figure 3 shows one of their results. The parameter λ on the horizontal axis indicates the scale at which the data from models are averaged – the higher is λ the more of the fine spatial detail in the models is lost, and the more model predictions at each geographic




Figure 2: Sources of uncertainty in temperature predictions as a function of lead time, from Hawkins & Sutton (2009). The orange, green, and blue areas denote uncertainty derived from internal variability, the emissions pathway, and model uncertainty respectively.

location look like spatial averages over regions of size λ . Clearly, the more spatial averaging is applied, the more the models agree with one another. Masson & Knutti show that the models' ability to reproduce past observations has a similar dependence on spatial scale.

These results are still only for predictions of climate variables themselves. If one keeps on going down the hierarchy of uncertainty to specific localized predictions of things that actually matter to economic and natural systems – water availability, growing conditions, sea level rise, and extreme events – matters get even more murky. So when it comes to thinking through climate policy options, whether mitigation or adaptation, uncertainty and its management are clearly going to be major considerations.

On the whole, our impression is that the climate science community has been impressively self-reflective about the meaning and interpretation of their models, and have devoted considerable energy to thinking through the challenges of characterizing uncertainty. High profile journals frequently publish discussions of what can be learned from climate models and how they should be interpreted (e.g. Oreskes et al., 1994; Smith, 2002; Frame et al., 2007; Stainforth et al., 2007; Knutti, 2008, 2010), and new approaches to characterizing uncertainty and ranking and combining models (e.g. Tebaldi et al., 2005; Murphy et al., 2007; Knutti et al., 2010; Palmer, 2012). Despite this careful soul-searching, and considerable advances in methods, data, and modeling over the past few decades, it is fair to say that scientific uncertainty is unlikely to be much reduced in the near future (Knutti & Sedlacek, 2012).

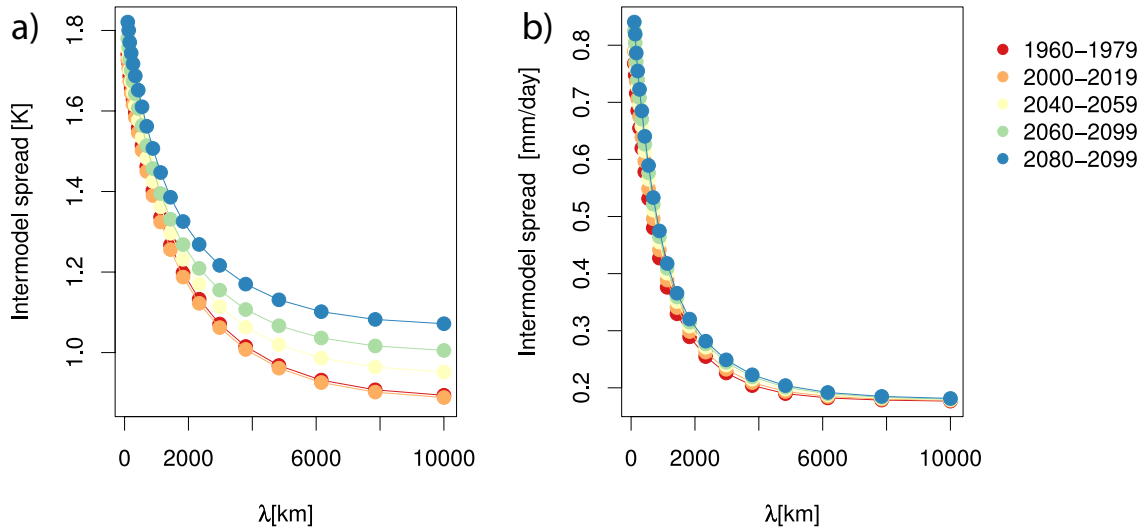


Figure 3: The spatial dependence of uncertainty in climate model predictions, from Masson & Knutti (2011). λ represents the scale of spatial averaging of model predictions. Panel a) shows the spread over the CMIP3 suite of climate models used by the IPCC, of spatially averaged predictions of changes in surface temperature, for several 20 year time periods. Panel b) shows the same for precipitation.

3 Socio-economic uncertainty

Even if all scientific uncertainty were resolved, we would still face major uncertainties stemming from the socio-economic dimensions of climate change. To understand this point, suppose for the sake of argument that all scientists agreed on a climate model, and also on the impacts of climate change on all such matters as sea level rise, fresh water availability, etc. Even in this extreme case they would not be able to forecast future climate because this requires knowing future emissions - climate models operate on emissions and map them into outcomes. Clearly we are far from being able to forecast future emissions, depending as they do on whether technological change provides us with ways of reducing greenhouse gas emissions that are not currently open to us, and on policies chosen, both of which are hard to forecast.

Even if we did have a believable forecast of future emissions, the future would still be unknown in many important ways. Take the case of sea level rise, and suppose we had an accurate forecast of future sea levels for the next century. How will societies react to rising seas? By protecting settlements or moving them? Will movements occur in a peaceful and organized manner, or will this involve strife and dislocation? In a similar vein, think about agriculture, and suppose we know how much future food outputs will fall in response to climate stress. What will be the economic consequences of this decline? Will the poor carry the burden or will the shortfall be shared? Another illustration is given by the fate of small island states whose land is within ten feet of sea level: it is highly likely that these will be uninhabitable

by the end of this century. How will the global community react to this, how will it address the refugee problem and what compensation will be offered? Clearly the answers to questions such as these will affect the welfare costs of climate change and its distributional impacts, so that even the full resolution of scientific uncertainties would leave huge residual uncertainties about the costs of climate change.

The conclusion is that in addition to obvious uncertainties stemming from incomplete scientific knowledge, there are uncertainties associated with future policies - which will affect emissions - and with future reactions to these policies and to an altered climate. The uncertainties about policies go deep: as we shall see below, there are uncertainties and disagreements about parameters such as discount rates that should be used in evaluating policy choices.

Scientific, impacts, and policy uncertainties of course interact, and accounting for their interaction is essential for deducing plausible estimates of the possible effects of climate change on social and economic systems. Such an 'end-to-end' approach is rare in empirical work, an interesting recent exception being the paper by Burke et al. (2011). They take seven influential studies of the economic impact of climate change, each of which gives estimates of the average impact of climate change on a particular sector - generally agriculture - and error bounds around this reflecting the uncertainty of the regressions carried out. The seven studies are Mendelsohn et al. (1994), Schlenker et al. (2005), Deschenes & Greenstone (2007), Fisher et al. (2012), Schlenker & Lobell (2010), Burke et al. (2009), and Dell et al. (2008). The first five papers all look at the impact of climate on agriculture (the first four in the US and the last in Africa), Burke et al. (2009) look at its impact on conflict in Africa, and Dell et al. (2008) study the relationship between temperature shocks and growth across countries. These papers all use a limited number of different climate models to generate predictions about impacts, in fact most use just one, whereas the IPCC uses a suite of more than 20 models for its predictions of climate change.

The differences between climate models reflects scientific uncertainty, whereas the uncertainty associated with a regression in one of these seven studies is a measure of impact uncertainties: they all use cross section or panel data to estimate how climatic conditions affect economic outcomes. By re-estimating the impact estimates with a full range of climate models and the uncertainties associated with these, Burke et al. show that incorporating climate model uncertainty into the impacts estimates increases the 95% confidence intervals around central estimates of outcomes by as much as a factor of five, and also changes the central estimates. They also argue that failing to allow fully for scientific uncertainty gives underestimates of the thickness of the tails of the error distributions. This plus the increased confidence intervals greatly increases the probability of extreme outcomes. Economists have recently focused on the importance of extreme outcomes in policy evaluation (Weitzman, 2009; Millner, 2013), so this change may have a disproportionate impact on policy choices.

Studies of climate impacts, such as those cited above, are necessarily based on existing tech-

nologies. They show how we have adapted to different climate regimes with the technologies and resources available to us. But in the context of climate we hope, and with some confidence expect, that new technologies will be available within a couple of decades, perhaps sooner. In particular investors are actively developing carbon-free energy sources, with wind and solar power now being competitive with fossil energy, in the case of wind, and close to competitive in the case of solar. Another dramatic illustration of the power of technological change is the rapid growth of hydraulic fracturing in the US, pushing down the prices of natural gas from \$12 per mmBTU to less than \$4, and in the process replacing coal and reducing US greenhouse gas emissions. Coal, which provided 52% of US electric power in 2008, now (2013) provides only 38%, and its share is falling under pressure from gas and wind. Such a change would have been hard to anticipate as recently as ten years ago. This illustrates the importance of technological uncertainty: there are many possible developments that could make it easier to reduce or to respond to climate change. Technological forecasting is a notoriously difficult field.

In addition to uncertainty about technologies that affect the climate, such as energy production, there is of course uncertainty about the future course of technologies in general. We can't forecast how total factor productivity will move, and this is a major determinant of overall economic growth. How much our economies grow - how rich we will be in the future - determines how many resources we will have to adapt to an altered climate, and sets the benchmark from which we measure deviations as a result of a more hostile climate. These technological uncertainties are particularly complex because they are partly exogenous, but also to some degree endogenous, in that they are influenced by policies, though we are not sure what the nature of these influences is. Many current integrated assessment models treat the growth rate of total factor productivity as a parameter, chosen once based on some general reference to historical growth rates, but we have no assurance that past technological progress is a good predictor of future trends (Gordon, 2012).

Uncertainty about future economic growth is reflected in uncertainty about one of the key parameters of climate economics, the consumption discount rate, the rate at which future increments to consumption should be discounted. Under conditions of certainty this is given by the Ramsey formula: the consumption discount rate ρ equals the pure rate of time preference (δ) plus the elasticity of the marginal utility of consumption (η) times the rate of growth of consumption (g). So if we are uncertain about how fast we will grow in the future, we are uncertain about how to discount future changes in consumption. In the context of uncertainty about growth rates the Ramsey formula is amended by the addition of another term. If the consumption growth rate is normally distributed with mean μ and variance σ^2 then the formula becomes³

$$\rho = \delta + \eta\mu - \frac{1}{2}\eta^2\sigma^2.$$

³See Traeger (2009); Gollier (2012) for much more complete accounts of discounting under uncertainty.

The new third term is a precautionary effect: it lowers the discount rate and makes society save more in the face of uncertainty. If, as is generally more likely, consumption growth rates are serially correlated, then we no longer have a constant consumption discount rate. Gollier (2008) shows that in this case, and with a utility function that shows constant relative risk aversion, the consumption discount rate will fall over time. An extensive literature addresses the term structure of discount rates in this case (for a review see Arrow et al. (2012)).

Weitzman (2001) has a different uncertainty-based argument about why the discount rate should not be constant. He notes that different people recommend different discount rates - he conducted a survey of 2000 PhD economists to elicit their suggested discount rates for the long-term future - and suggests that we see these different rates as independent estimates of a true underlying rate. He suggests combining these estimates by taking a weighted average of their implied discount factors, the weights being probabilities attached to the individual discount rates. Weitzman draws these from a Gamma distribution, titling his paper “Gamma Discounting.” In this case it is readily shown that the discount rate declines over time, asymptoting to the lowest of all the individual discount rates.

There is another way of looking at this. Weitzman sees different discount rates as independent estimates of a true underlying rate: Heal & Millner (2013) suggest instead that part of the heterogeneity in discount rate estimates may reflect the fact that different people have different pure rates of time preference, i.e. different fundamental value judgments. The question is then how to combine these different value judgments into a representative discount rate. One possibility is to think of this as a social choice problem, and consider reconciling these disagreements by a process such as voting - in which case the outcome would likely be the median discount rate. Another is to follow a more traditional economic approach and seek to maximize a weighted sum of the diverse agents’ utility integrals, asking at the same time whether there is a representative agent who can stand in for the group, an approach pioneered by Gollier & Zeckhauser (2005). Heal and Millner generalize their approach by endogenizing the group’s income stream, and show that the group’s rate of impatience is declining, and approaches the lowest rate asymptotically. Thus a heterogeneous group of time-consistent agents can only be represented by a social planner with a declining pure rate of time preference. A declining pure rate of time preference has a big effect on climate policy choices, as demonstrated by e.g. Gerlagh & Liski (2012). In general, the discount rate is a major determinant of future policies, as highlighted by the debate between Stern and the critics of his 2006 report (Nordhaus, 2007; Weitzman, 2007), so uncertainty over discount rates translates into uncertainty about future policies.

In summary, there are multiple sources of uncertainty that affect the evaluation of climate change policies. One is scientific uncertainty, as reflected in uncertainty about the values of the physical parameters in a given climate model, and also in differences between climate models themselves. Scientists refer to this as model uncertainty. For a given emissions scenario the

different models predict different outcomes. In addition there is within-model uncertainty in climate predictions arising from the fact that the models display a high degree of sensitivity to initial conditions: running the same model from two slightly different initial conditions, all else unchanged, can lead to predictions that differ significantly over several decades. Over and above these categories of scientific uncertainty, the socio-economic part of the system gives rise to uncertainties. We can't forecast the operation of the economic system and how it will respond to climate change, can't forecast the policies that will be chosen in the future, don't know what technologies will emerge, and can't even be sure of the parameters that will govern our evaluations of the options. In the language of classical welfare economics, there are positive (descriptive, empirical) and normative (evaluative) uncertainties.

4 Representing Uncertainty

When it comes to modeling uncertainty, most economists instinctively reach for the expected utility framework associated with the names of von Neuman and Morgenstern (vNM). We assume a well-defined set of possible outcomes, some exogenously given probability distribution over these, and preferences over lotteries that can be represented by expected utilities. This is the paradigm that we all learn in graduate school. It is hopefully clear by now that the uncertainties associated with climate change do not fit this model well.

The first problem is that we do not always have a clear set of possible states of the world. A very concrete example is the possible disruption of the thermohaline circulation in the Atlantic (ending or weakening of the Gulf Stream in plain English). This is a possible consequence of climate change, and a very serious one, though one now believed to be unlikely. But it was not even recognized as a possibility until the 1990s (Broecker, 1997). This is an illustration of learning about the set of possible outcomes, and suggests that even today there may be consequences of climate change that we are not considering.

In addition, it should be abundantly clear that we do not have a generally agreed objective probability distribution over what are currently thought to be the possible consequences of climate change. So where are we supposed to get our probability distribution from? More generally, are probabilities capable of representing our uncertain knowledge, and does rational policy analysis require them as inputs?

There is a traditional answer to this question that draws on a fundamental result of decision theory that is even more important than vNM's results, but is perhaps less well known to applied economists. This is Leonard Savage's theory of Subjective Expected Utility (SEU) (Savage, 1954). Savage's results are much more general than those of vNM. He considers preferences over abstract mathematical objects known as 'acts' that map states into outcomes. For example, the 'Umbrella' act may map the 'rain' state into the 'happy to be dry' outcome, and the 'sun' state into the 'irritated to be carrying unneeded umbrella' outcome.

Savage states plausible axioms for preferences over different acts (e.g. an ‘Umbrella’ act vs. a ‘No Umbrella’ act), and proves that preferences that obey his axioms can be represented by expected utilities over a *subjective* probability distribution over the possible states. The remarkable thing about Savage’s result is that it derives utilities *and* subjective probabilities from primitive choice axioms. Unlike vNM, he does not need to put probabilities in by hand. Thus, if one accepts Savage’s axioms, the problem of representing uncertainty, and making rational decisions when faced with it, is solved in one fell swoop. Even if our information is ambiguous, in fact even if we have no information at all, it is rational to represent our knowledge with some subjective probability distribution, and choose the action that maximizes expected utility given some utility function. From this perspective, it does not matter where our probabilities come from, so long as there are some! The SEU framework provides the theoretical justification for the Bayesian methods economists have usually used to model the uncertain effects of climate change (e.g. Weitzman, 2009).

SEU theory provides a neat one-size-fits-all solution to the problem of decision under uncertainty. But is it clear that it’s the right solution in all situations? This has been hotly debated since the theory’s inception. Savage himself saw limits to his theory, cautioning us to use it only in ‘small worlds’, in which it is possible to ‘look before you leap’ (Binmore, 2009). Even before Savage, eminent theorists questioned the ability of the probabilistic framework to capture our knowledge when information is incomplete, inconsistent, or nonexistent (Knight, 1921; Keynes, 1921). The most famous challenge to SEU’s claim to define rational choice was mounted by Daniel Ellsberg, in a series of hypothetical choice experiments (Ellsberg, 1961). These show that when we have no information to base decisions on, we may wish to violate one of Savage’s key axioms – the ‘Sure thing principle’⁴. In essence, Ellsberg’s results show that many people treat objectively known risks (e.g. the chance of a fair coin landing on ‘heads’) very differently from ‘unknown’ uncertainties (e.g. the chance of a coin with an unknown bias landing on ‘heads’). Many of us are ambiguity averse, and prefer to bet on known risks rather than uncertain prospects⁵, even if, from a subjective probability standpoint, these two situations are indistinguishable (Slovic & Tversky, 1974; Camerer & Weber, 1992).

Of course, the fact that many people are ambiguity averse does not mean that ambiguity aversion is rational, and should form part of normative policy analysis. Many economists interpret Ellsberg’s results as a contribution to positive economics, grouping it with other behavioral ‘mistakes’ such as the Allais paradox. It’s clear however that Ellsberg himself viewed his work in a normative light (Ellsberg, 2001). Subsequently eminent decision theorists have argued that SEU is not universally applicable (Binmore, 2009), that subjective probabilities cannot always capture our state of knowledge (Gilboa et al., 2008), and that

⁴See Gilboa (2009) for a detailed discussion of Savage’s axioms and the Ellsberg thought experiments.

⁵Millner et al. (2013) have adapted Ellsberg’s experiments to show that some climate scientists’ beliefs about the value of climate sensitivity cannot be captured by subjective probabilities.

ambiguity aversion can be rational when our information is of poor quality (Gilboa et al., 2009). Our view of this debate is that, much like the debate about the appropriate value of the pure rate of time preference, it is the kind of thing reasonable people can reasonably disagree about⁶. There is unlikely to be a knock-down argument in favor of either position, and both contain valuable insights about how we might wish to make choices under uncertainty⁷. If we acknowledge this, the problem of decision making under uncertainty contains a new social choice dimension – there are many different legitimate preferences, some obey SEU, some do not, and policy analysis should give some weight to all. Given this position, it is clearly going to be important to explore alternatives to the standard SEU paradigm, as these will form part of any comprehensive approach.

5 Decision making under ‘deep’ uncertainty

There has been a resurgence of activity in formal approaches to decision under uncertainty over the past two decades, beginning with the seminal work of Schmeidler (1989) and Gilboa & Schmeidler (1989). There is now a large literature, with many different decision rules being proposed. Following in the decision theoretic tradition, all of them are rigorously derived from preferences that satisfy primitive choice axioms. Etner et al. (2012) provides a comprehensive review of approaches. Some argue that the plurality of new approaches to choice under uncertainty is reason enough to fall back on SEU – how do we know which one to use? This seems a perverse piece of historicism. Rather, we advocate an exploration of different approaches.

Table 1 lists a few of the most promising approaches for applications. These can be divided into non-probabilistic approaches, which require no likelihood information at all, and multiple priors approaches, in which information is characterized by a set of possible probability distributions, which are not reducible to a single distribution.

Our next step is to review these various criteria, explaining the ideas underlying them and investigating whether they could provide a framework for thinking about uncertainty in the climate context. In our analysis we will use the words model and distribution interchangeably: a model for us is a mapping from policies – e.g. greenhouse gas concentrations – to distributions over outcomes. Figure 1 represents twenty such models, each of which maps a doubled CO₂ concentration to a probability distribution over the equilibrium climate sensitivity.

The first three criteria – maxmin, α -maxmin, and minmax regret – are all non-probabilistic

⁶A recent special issue of *Economics and Philosophy* (Vol. 25, No. 3, 2009) is devoted to an exchange of views on this topic.

⁷This is true even though these insights may be mutually inconsistent. The tension between the twin virtues of ‘consistency’ (SEU) and epistemic sensitivity (ambiguity approaches) parallels similar tensions between fundamental values in ethics, e.g. between freedom and equality. The philosopher Isaiah Berlin famously argued that this is a fundamental property of value systems. We just have to live with the fact that different, equally legitimate, values may come into conflict. This lesson seems just as applicable to questions of rationality.

Name	Reference	Rule
<i>Non-probabilistic approaches</i>		
Maxmin	Wald (1949)	$\max_{a \in \mathcal{A}} [\min_{s \in \mathcal{S}} a(s)]$
α -Maxmin	Arrow & Hurwicz (1977)	$\max_{a \in \mathcal{A}} [\alpha \min_{s \in \mathcal{S}} a(s) + (1 - \alpha) \max_{s \in \mathcal{S}} a(s)]$
Minmax Regret	Savage (1954)	$\min_{a \in \mathcal{A}} [\max_{s \in \mathcal{S}} \{(\max_{a' \in \mathcal{A}} a'(s)) - a(s)\}]$
<i>Multiple priors approaches</i>		
Maxmin EU	Gilboa & Schmeidler (1989)	$\max_{a \in \mathcal{A}} [\min_{p \in \mathcal{P}} \mathbf{E}_p U(a(s))]$
Smooth Ambiguity	Klibanoff et al. (2005)	$\max_{a \in \mathcal{A}} \mathbf{E}_{p \in \mathcal{P}} \Phi(\mathbf{E}_p U(a(s)))$
Variational preferences	Maccheroni et al. (2006)	$\max_{a \in \mathcal{A}} [\min_{p \in \mathcal{P}} \{\mathbf{E}_p U(a(s)) + C(p)\}]$

Table 1: Alternatives to Expected Utility Theory: \mathcal{S} and \mathcal{A} denote sets of states (s) and actions (a) respectively. An action $a : s \rightarrow a(s)$ is a map from states s to outcomes $a(s)$. \mathcal{P} is set of probability distributions (p), and \mathbf{E} denotes the expectation operator. Each rule is explained further in the text.

and so require no likelihood information. This is a plus in a context where such information is not available, but a minus if it is, because they discard information that could add value to the policy selection process. This suggests that they may be more appropriate for studies that are local in scale, as climate predictions are most uninformative on small spatial scales. Such studies occur naturally in the context of adaptation to climate change, where location-specific policies are often needed.

The Maxmin criterion is the least demanding, in the sense that it needs least information of all. We pick the strategy whose worst possible outcome (min) is least bad (max). To implement this we just need a ranking of outcomes. No probabilistic information is needed. An interesting application of this criterion to climate change is provided by Woodward & Bishop (1997), who simulate alternative policies in Nordhaus’s DICE model in the face of two alternative climate scenarios, a catastrophic scenario and one that corresponds to Nordhaus’s base case. They allow for the possibility of learning and revising policy choice in 2085, and find that the best policy according to the maxmin criterion is to assume the catastrophic scenario to be correct and optimize for this, altering course later if the base case emerges as the correct world view.

The α -maxmin criterion generalizes Maxmin: it ranks alternatives $a(s)$ according to a linear mixture of their minimum and maximum outcomes (see Table 1). The parameter α determines how much weight the decision maker gives to the worst outcome of a policy, with $1 - \alpha$ being the weight on its best outcome. This parameter could be seen as representing aversion to a lack of knowledge: all the weight on the min implies extreme caution and vice versa. This rule was derived axiomatically by Arrow & Hurwicz (1977) as a criterion for choice under ‘ignorance’, a situation in which no information is available at all.

Minmax regret has many similarities to these two criteria, but requires a more demanding assessment of outcomes. If s is a state, and $a'(s)$ a policy choice, $\max_{a' \in \mathcal{A}} a'(s)$ is the best policy choice conditional on s being the state, and $a(s)$ is the value of choosing policy a if

the state is s , then the regret of policy a in state s is $\max_{a' \in A} a'(s) - a(s)$. Max regret is the maximum of this over all states for a given policy, and minmax regret involves choosing a policy to minimize this maximum.

Informationally this is more demanding than the other two criteria, requiring that we rank not just alternatives but also the differences between them: regret is a difference and we need to maximize and minimize this. The motivation here is that rather than focusing on the worst outcome, you want to focus on how much better you could have done (this is regret) and want to come as close as possible to doing as well as possible over all states of the world. The probabilistic criteria are naturally all more complex. All of them are based on the idea of multiple priors: if the world isn't kind enough to give us a unique probability density function over states, then there must be many different prior probability distributions that are consistent with the data. The idea is to use all of them, and to weight them in some way. Maxmin expected utility places all weight on the least favorable prior, a probabilistic equivalent of the maxmin criterion. Multiplier preferences weight distributions by how close they are to our "best guess", and in the smooth ambiguity model we assign subjective weights to the alternative priors: it has us make subjective judgments about their relative merits.

The oldest of the probabilistic approaches, and in many ways the simplest, is the Maxmin Expected Utility criterion (Gilboa & Schmeidler, 1989). In implementing this, the first step is to identify all priors consistent with what is known. Then one act (policy) is better than another if the consequence of that act, evaluated at the least favorable prior (the prior which gives it the lowest expected utility), is better than the worst consequence of the other. So we are looking for the policy whose worst outcome over all priors is the best. The focus is on bad priors. This is just Maxmin but in prior space. As with all the multiple prior approaches that we will review there is no assumption of a unique PDF. We can have incomplete probabilistic information that is compatible with many different priors. Figure 1 shows such a situation for climate sensitivity: the data there are consistent with many but not all PDFs for its value. When you feel "I don't know what the probabilities are" you ask "what probabilities are consistent with what I do know?" and work with all of these.

What this approach is not doing, is ranking the priors in any way: all priors have equal status, and we focus on the worst. This reflects extreme ambiguity aversion – why in general should we not take the best case or the average case, and why should we not believe some priors to be more plausible than others and reflect this in our choices? There is no way of parameterizing ambiguity aversion here.

These observations lead naturally to a discussion of the smooth ambiguity model Klibanoff et al. (2005): again there are many possible priors or models consistent with what we know. However in this case we don't necessarily treat them all equally. We assign weights to each of these priors. For every prior distribution we find the expected utility associated with each act or policy, conditional on that prior being correct. So for each policy we now have an expected

utility according to each prior. We then take the expected utility of all of these, using the subjective weights over priors and a second order utility function whose arguments are the expected values according to the individual priors or models. The best policy is now the one that maximizes the expectation of this second order utility (Φ in Table 1) over the weights on models. The concavity of the second order utility function is a measure of ambiguity aversion: if this function is linear there is no aversion to ambiguity and the second order process is just evaluating a compound lottery, taking the expectation of expected utilities. If on the other hand the second order utility function is concave, this reflects ambiguity aversion. In the limit as the degree of ambiguity aversion goes to infinity this approach reduces to maxmin over priors.

The novelty here is that we rank the priors, which we didn't in the maxmin EU case. The smooth ambiguity model permits us to use all of the likelihood information in a set of priors. In order to do so we need to specify weights on each prior. These weights reflect a very different kind of information to the probabilities that arise within each model. They are subjective judgments, whereas the model probabilities can be thought of as 'objective', i.e. informed by data. The smooth ambiguity model recognizes the necessity of such subjective judgments, in much the same way that the standard Bayesian framework does. The difference however is that it allows us to treat these two kinds of knowledge very differently. Whenever the second order utility function is nonlinear, the smooth ambiguity model drives a wedge between the 'subjective' and the 'objective' – we are not able to reduce an information set that contains both subjective and objective elements to a single distribution that is treated as if it contained only objective information.

Lastly we have multiplier preferences, introduced by Hansen & Sargent (2001), and increasingly widely used in macroeconomics. Macroeconomists have finally recognized that they don't know what the true macro model is and that they have to contend with multiple possible models. In this sense they are in a similar position to economists concerned about climate change. The inspiration in the Hansen-Sargent case is that you have an idea of what the true model/distribution is but are not sure. So you want to allow for possible misspecification and you evaluate the policy options not only according to your main distribution but also according to others nearby. Policy makers are assumed to be strongly ambiguity averse and so evaluate actions by the distribution that gives the lowest value to them. But distributions away from your main one are penalized by adding on a cost term that depends on how different they are from the best guess distribution. This addition is a penalty to the distribution as we are ranking by the minimum. These preference are a special case of variational preferences, formalized by Maccheroni et al. (2006) and Strzalecki (2011).

This approach makes sense if you have a clear preferred model but are not totally sure of some of its details – in other words if there is a possibility of model misspecification. However if there are several rather different models that are all serious candidates then it is not so

attractive, indeed it is not applicable. It may be a good approach for macroeconomists who think they more or less know the true model and only need to consider small variations around it: it is not so clear that it fits the environmental scenario so well. As with Maxmin EU this approach is discarding likelihood information. It is just ranking priors by their closeness to our best guess. And the fact that it works with the minimum over possible distributions means that extreme ambiguity aversion is built in. But as Hansen & Sargent (2007) show we get a return for these costs – a tractable framework that permits analytical solutions, at least for linear-quadratic models.

All of these frameworks have been deployed in the analysis of climate policies. A model that can be used to illustrate two of the approaches, maxmin expected utility and smooth ambiguity, is that of Millner, Dietz, and Heal (2012) (MDH). A simplified version assumes two models, each a map from GHG concentrations to a distribution over variables that matter to us, which we summarize as consumption. Assuming GHG concentrations are determined by policies, a model is also a map M_i from policies P to a distribution F_i over consumption C , i.e. $M_i : P \rightarrow F_i(C|P)$ where $F_i(C|P), i \in \{1, 2\}$ is a distribution over C conditional on the policy P . Consumption is evaluated by a utility function $U(C)$. Models 1 and 2 are candidates for being the correct model, but we cannot be sure which is right. The maxmin EU approach says that each model is a distribution F_i over outcomes and we evaluate a policy according to the model that gives it the lowest expected utility, so if $V(P)$ is the value of policy P then $V(P) = \min_i EU(C)F_i(C|P)$ where E is the expectation operator. According to the maxmin EU rule the best policy is that for which this min is greatest. Note that we are not attaching any weights to the models here.

In the smooth ambiguity approach we introduce subjective weights on alternative models, which we will denote w_i . In effect, we place weight w_i on the event that if we choose policy P it will give us expected utility $EU(C)F_i(C|P)$. We combine the expected utilities from the two models through a second order utility function Φ and evaluate policies by $w_1\Phi(EU(C)F_1(C|P)) + w_2\Phi(EU(C)F_2(C|P))$. So the optimal policy maximizes this expectation over expected utilities. MDH analyze the comparative statics of a generalized form of this model and then simulate a dynamic version using Nordhaus’s DICE integrated assessment model to evaluate how the ambiguity in climate sensitivity shown in Figure 1 affects the welfare benefits of CO₂ abatement policies. They show that the introduction of ambiguity aversion can increase or decrease the optimal level of greenhouse gas abatement: the key is whether models that show a high payoff from abatement (in expected utility terms) also have a high marginal payoff. If they do, ambiguity aversion lowers abatement and otherwise it increases it. Their dynamic simulations show that for the DICE model, aversion to scientific ambiguity increases the welfare benefits of mitigation, in some cases dramatically. In a more complex model, Lemoine & Traeger (2012) use the smooth ambiguity framework to study the impact of possible climate tipping points and regime changes on optimal greenhouse gas

policies, assuming that knowledge of such tipping points is ambiguous.

The Hansen-Sargent multiplier preference approach has been applied to climate problems by Xepapadeas and coauthors. Xepapadeas (2012) considers a model in which output produces a pollutant – which could be a GHG – according to a stochastic process, taken to be a Brownian motion. Possible misspecification of this process is captured, as in the macroeconomic work, by adding a weighted error term to the stochastic process, and then conducting a min-max expected utility calculation. By the magic of the linear-quadratic model assumption, this complex problem permits analytical solutions.

6 Conclusion

What can we conclude from this discussion? A first important point that we need to internalize is that uncertainty is here to stay. Although we are making incremental progress all the time, it is unlikely that we will learn anything over the next few decades that will radically improve our state of knowledge about the likely consequences of climate change in the medium to long run. This is probably true both of the science and of the socio-economic consequences of climate change, although in the case of the latter, considerably more could be done to characterize our current knowledge more comprehensively. End-to-end analyses, which integrate uncertainty in both climate science and climate impacts estimates, are an important part of this approach – point estimates of impacts based on projections from a single climate model are all but meaningless. But we also have to recognize that many of the most important empirical uncertainties, for example about future abatement technologies, the long-run growth path of the global economy, and the adaptability of future societies to an altered climate, are fundamentally irreducible. We simply do not have credible tools for making such predictions. We can be confident that the future will not look like the past, but we don't know how it will differ.

This proliferation of uncertainties can give rise to a temptation towards policy paralysis, and methodological conservatism – we may be tempted to stick our heads in the sand and deal with easier problems, and to rely on methods that sweep many of the uncertainties under the carpet. This would be a big mistake. We know more than enough about climate change to know that it is a serious problem that requires immediate policy attention: uncertainty does not imply ignorance. As economists, our job is to characterize the welfare consequences of such policies as best we can, and provide analytical tools that can inform policy choices. These decision tools must fit the task – the methods that we are used to using for evaluating small social projects with well defined, short-run, consequences may be less appropriate for a problem as global, long-run, and uncertain, as climate change. This is an opportunity for an inclusive, creative approach to policy analysis, that should aim to represent a variety of legitimate approaches – all methods and opinions should be considered, and nothing dismissed

out of hand. As the debate about discounting has shown, the decision tools and welfare parameters use to evaluate policy have as strong an effect on policy recommendations as empirical inputs. None of us is qualified to make such choices in isolation – they are political and philosophical, as much as economic, in nature.

The survey of the decision theory literature we have offered above shows that there are many promising approaches to decision making under ‘deep’ uncertainty that may allow us to represent the nature of our information with greater fidelity, and make policy recommendations that account for the ambiguity in our knowledge. Such tools, when combined with empirical inputs that represent both what we do and don’t know comprehensively, could provide useful policy guidance. But we must also recognize that the world is under no obligation to agree with our models. We must expect the unexpected, and heed Mark Twain’s warning: “It ain’t what you don’t know that gets you into trouble. It’s what you know for sure that just ain’t so”.

References

- M. Allen, et al. (2006). ‘Observational constraints on climate sensitivity’. In H. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, & G. Yohe (eds.), *Avoiding dangerous climate change*, p. 406. Cambridge University Press, Cambridge, UK.
- M. R. Allen, et al. (2009). ‘Warming caused by cumulative carbon emissions towards the trillionth tonne’. *Nature* **458**(7242):1163–1166.
- K. Arrow & L. Hurwicz (1977). ‘An optimality criterion for decision-making under ignorance.’. In *Studies in resource allocation processes*, p. 482. Cambridge University Press, Cambridge, UK.
- K. J. Arrow, et al. (2012). ‘How should benefits and costs be discounted in an intergenerational context?’. *Resources for the Future Discussion Paper* **12**(53).
- K. Binmore (2009). *Rational Decisions*. Princeton University Press.
- W. S. Broecker (1997). ‘Thermohaline Circulation, the Achilles Heel of Our Climate System: Will Man-Made CO₂ Upset the Current Balance?’. *Science* **278**(5343):1582–1588.
- M. Burke, et al. (2011). ‘Incorporating Climate Uncertainty into Estimates of Climate Change Impacts, with Applications to U.S. and African Agriculture’. *NBER working paper no.17092*.
- M. B. Burke, et al. (2009). ‘Warming increases the risk of civil war in Africa’. *Proceedings of the National Academy of Sciences* **106**(49):20670–20674.

- C. Camerer & M. Weber (1992). ‘Recent developments in modeling preferences: Uncertainty and ambiguity’. *Journal of Risk and Uncertainty* **5**(4):325–370.
- M. Dell, et al. (2008). ‘Climate Change and Economic Growth: Evidence from the Last Half Century’. Working Paper 14132, National Bureau of Economic Research.
- O. Deschenes & M. Greenstone (2007). ‘The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather’. *The American Economic Review* **97**(1):354–385.
- D. Ellsberg (1961). ‘Risk, Ambiguity, and the Savage Axioms’. *The Quarterly Journal of Economics* **75**(4):643–669.
- D. Ellsberg (2001). *Risk, Ambiguity and Decision*. Routledge.
- J. Etner, et al. (2012). ‘Decision Theory Under Ambiguity’. *Journal of Economic Surveys* **26**(2):234–270.
- A. C. Fisher, et al. (2012). ‘The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment’. *American Economic Review* **102**(7):3749–3760.
- D. J. Frame, et al. (2007). ‘Probabilistic climate forecasts and inductive problems’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857):1971–1992.
- R. Gerlagh & M. Liski (2012). ‘Carbon prices for the next thousand years’. *CESifo Working Papers* (3855).
- I. Gilboa (2009). *Theory of Decision under Uncertainty*. Cambridge University Press, 1 edn.
- I. Gilboa, et al. (2009). ‘Is It Always Rational to Satisfy Savage’s Axioms?’. *Economics and Philosophy* **25**(3):285–296.
- I. Gilboa, et al. (2008). ‘Probability and Uncertainty in Economic Modeling’. *Journal of Economic Perspectives* **22**(3):173–188.
- I. Gilboa & D. Schmeidler (1989). ‘Maxmin expected utility with non-unique prior’. *Journal of Mathematical Economics* **18**(2):141–153.
- C. Gollier (2008). ‘Discounting with fat-tailed economic growth’. *Journal of Risk and Uncertainty* **37**(2-3):171–186.
- C. Gollier (2012). *Pricing the Planet’s Future: The Economics of Discounting in an Uncertain World*. Princeton University Press.

- C. Gollier & R. Zeckhauser (2005). ‘Aggregation of Heterogeneous Time Preferences’. *Journal of Political Economy* **113**(4):878–896.
- R. J. Gordon (2012). ‘Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds’. *National Bureau of Economic Research Working Paper* (18315).
- L. P. Hansen & T. J. Sargent (2001). ‘Robust Control and Model Uncertainty’. *The American Economic Review* **91**(2):60–66.
- L. P. Hansen & T. J. Sargent (2007). *Robustness*. Princeton University Press.
- E. Hawkins & R. Sutton (2009). ‘The potential to narrow uncertainty in regional climate predictions’. *Bulletin of the American Meteorological Society* **90**:1095–1107.
- G. Heal & B. Kriström (2002). ‘Uncertainty and Climate Change’. *Environmental and Resource Economics* **22**(1-2):3–39.
- G. Heal & A. Millner (2013). ‘Discounting under disagreement’. *Working paper* .
- IPCC (2007). ‘Summary for Policymakers’. In S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. Tignor, & H. Miller (eds.), *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- J. M. Keynes (1921). *A treatise on probability*. Macmillan and Co.
- P. Klibanoff, et al. (2005). ‘A Smooth Model of Decision Making under Ambiguity’. *Econometrica* **73**(6):1849–1892.
- F. Knight (1921). *Risk, uncertainty, and profit*. Houghton Mifflin, New York.
- R. Knutti (2008). ‘Should we believe model predictions of future climate change?’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **366**(1885):4647–4664.
- R. Knutti (2010). ‘The end of model democracy?’. *Climatic Change* **102**(3):395–404.
- R. Knutti, et al. (2010). ‘Challenges in Combining Projections from Multiple Climate Models’. *Journal of Climate* **23**(10):2739–2758.
- R. Knutti & G. C. Hegerl (2008). ‘The equilibrium sensitivity of the Earth’s temperature to radiation changes’. *Nature Geoscience* **1**(11):735–743.
- R. Knutti & J. Sedlacek (2012). ‘Robustness and uncertainties in the CMIP5 climate model projections’. *Nature Climate Change* (doi:10.1038/nclimate1716).

- D. Lemoine & C. P. Traeger (2012). ‘Tipping points and ambiguity in the Economics of Climate Change’. *NBER working paper no. 18230*.
- F. Maccheroni, et al. (2006). ‘Ambiguity Aversion, Robustness, and the Variational Representation of Preferences’. *Econometrica* **74**(6):1447–1498.
- D. Masson & R. Knutti (2011). ‘Spatial-Scale Dependence of Climate Model Performance in the CMIP3 Ensemble’. *Journal of Climate* **24**(11):2680–2692.
- M. Meinshausen, et al. (2009). ‘Greenhouse-gas emission targets for limiting global warming to 2C’. *Nature* **458**(7242):1158–1162.
- R. Mendelsohn, et al. (1994). ‘The Impact of Global Warming on Agriculture: A Ricardian Analysis’. *American Economic Review* **84**(4):753–71.
- A. Millner (2013). ‘On welfare frameworks and catastrophic climate risks’. *Journal of Environmental Economics and Management* **65**(2):310–325.
- A. Millner, et al. (2013). ‘Do probabilistic expert elicitations capture scientists’ uncertainty about climate change?’. *Climatic Change* **116**(2):427–436.
- A. Millner, et al. (2012). ‘Scientific Ambiguity and Climate Policy’. *Environmental and Resource Economics, forthcoming* (10.1007/s10640-012-9612-0):1–26.
- J. M. Murphy, et al. (2007). ‘A methodology for probabilistic predictions of regional climate change from perturbed physics ensembles’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857):1993–2028.
- W. D. Nordhaus (2007). ‘A Review of the Stern Review on the Economics of Climate Change’. *Journal of Economic Literature* **45**(3):686–702.
- N. Oreskes, et al. (1994). ‘Verification, Validation, and Confirmation of Numerical Models in the Earth Sciences’. *Science* **263**(5147):641–646.
- T. N. Palmer (2012). ‘Towards the probabilistic Earth-system simulator: a vision for the future of climate and weather prediction’. *Quarterly Journal of the Royal Meteorological Society* **138**(665):841–861.
- R. S. Pindyck (2007). ‘Uncertainty in Environmental Economics’. *Review of Environmental Economics and Policy* **1**(1):45–65.
- G. H. Roe & M. B. Baker (2007). ‘Why Is Climate Sensitivity So Unpredictable?’. *Science* **318**(5850):629–632.
- L. J. Savage (1954). *The Foundations of Statistics*. Wiley and Sons.

- W. Schlenker, et al. (2005). ‘Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach’. *The American Economic Review* **95**(1):395–406.
- W. Schlenker & D. B. Lobell (2010). ‘Robust negative impacts of climate change on African agriculture’. *Environmental Research Letters* **5**(1):014010.
- D. Schmeidler (1989). ‘Subjective Probability and Expected Utility without Additivity’. *Econometrica* **57**(3):571–587.
- P. Slovic & A. Tversky (1974). ‘Who accepts Savage’s axiom?’. *Behavioral Science* **19**(6):368–373.
- L. A. Smith (2002). ‘What might we learn from climate forecasts?’. *Proceedings of the National Academy of Sciences of the United States of America* **99**(Suppl 1):2487–2492.
- D. Stainforth, et al. (2007). ‘Confidence, uncertainty and decision-support relevance in climate predictions’. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857):2145–2161.
- D. A. Stone, et al. (2009). ‘The Detection and Attribution of Human Influence on Climate’. *Annual Review of Environment and Resources* **34**(1):1–16.
- T. Strzalecki (2011). ‘Axiomatic Foundations of Multiplier Preferences’. *Econometrica* **79**(1):47–73.
- C. Tebaldi, et al. (2005). ‘Quantifying Uncertainty in Projections of Regional Climate Change: A Bayesian Approach to the Analysis of Multimodel Ensembles’. *Journal of Climate* **18**(10):1524–1540.
- C. P. Traeger (2009). ‘Recent Developments in the Intertemporal Modeling of Uncertainty’. *Annual Review of Resource Economics* **1**(1):261–286.
- A. Wald (1949). ‘Statistical Decision Functions’. *The Annals of Mathematical Statistics* **20**(2):165–205.
- M. L. Weitzman (2001). ‘Gamma Discounting’. *The American Economic Review* **91**(1):260–271.
- M. L. Weitzman (2007). ‘A Review of The Stern Review on the Economics of Climate Change’. *Journal of Economic Literature* **45**(3):703–724.
- M. L. Weitzman (2009). ‘On Modeling and Interpreting the Economics of Catastrophic Climate Change’. *Review of Economics and Statistics* **91**(1):1–19.

- R. T. Woodward & R. C. Bishop (1997). 'How to Decide When Experts Disagree: Uncertainty-Based Choice Rules in Environmental Policy'. *Land Economics* **73**(4):492–507.
- A. Xepapadeas (2012). 'Ambiguity and Robustness in International Pollution Control'. In R. W. Hahn & A. Ulph (eds.), *Climate Change and Common Sense: Essays in Honour of Tom Schelling*. Oxford University Press.