

Reflections

Uncertainty and Decision Making in Climate Change Economics

Geoffrey Heal* and Antony Millner†

Introduction

The issue of climate change is beset with uncertainties, many of which are only partially captured by our existing analytical tools. This “Reflections” explores these uncertainties, some of which are rooted in the underlying science while others stem from an imperfect understanding of how the socioeconomic system will evolve and cope with climatic shifts. Although robust general conclusions about the nature and impacts of climate change can be drawn from the existing scientific and economic literatures, in many cases our knowledge of the climate problem is ambiguous and our uncertainty is deep. In the words of Knight (1921), we are in a world of “uncertainty” (unknown probabilities) rather than “risk” (known probabilities). How then should we make policy decisions in such an informational environment?

There have been several surveys of uncertainty in environmental economics, and climate change in particular. These include Heal and Kristrom (2002) and Pindyck (2007), as well as a 2011 symposium in this journal (Nordhaus 2011; Pindyck 2011; Weitzman 2011) that examined some of the questions raised by fat-tailed probabilities of climate catastrophes.¹ These reviews focus on issues that have become the bread and butter of economic approaches to uncertainty: irreversibility, discounting, and the consequences of the standard expected utility approach to representing uncertainty.

This article takes a different approach, arguing that the expected utility framework for decision making under uncertainty may be of limited use for analyzing the climate problem (see also Kunreuther et al. 2013). We review the decision theory literature, explore alternative decision frameworks that may be more appropriate, and investigate how they can be applied to the

*Columbia Business School, Columbia University; e-mail: gmh1@columbia.edu.

†Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science; e-mail: a.millner@lse.ac.uk.

A.M. acknowledges support from the Centre for Climate Change Economics and Policy, funded by the UK's Economic and Social Research Council.

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

Review of Environmental Economics and Policy, volume 8, issue 1, 2014, pp. 120–137
doi:10.1093/reep/ret023

issue of climate policy. In particular, 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? We argue that some of the most important uncertainties in climate economics fall into the last category: they are inherently unpredictable.

The remainder of the article is organized as follows. First we discuss sources of uncertainty and their relative importance for climate policy. We decompose the sources of uncertainty into scientific and socioeconomic components, with 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, how those societies will respond, and how we should value these impacts. Next we examine decision frameworks that may be more appropriate in the context of deep uncertainty, discuss how they can be applied to the climate problem, and review some examples of such applications. We conclude with a brief summary and discussion of the challenges and opportunities for economists to develop approaches to uncertainty and decision making in climate economics.

Sources of Scientific Uncertainty

Advances in the science of climate change detection and attribution over the past twenty years have made it clear beyond a reasonable doubt that human emissions of greenhouse gases (GHGs) are largely responsible for the observed changes in global temperature during the twentieth century (Intergovernmental Panel on Climate Change [IPCC] 2007; Stone et al. 2009). Thus as we continue to pump carbon dioxide (CO₂) and other GHGs into the atmosphere, the climate will change further. What is less clear is exactly how much and how quickly the global climate will respond to these changes in the composition of the atmosphere. The scientific community has constructed models that attempt to quantify the relationship between future climate and GHG emissions, but it is still highly uncertain. This section describes the nature of uncertainty in scientific projections of climate change and its relevance for policy choices.

Uncertainty in Climate Sensitivity

Climate science can provide some of the quantitative information we need about climate change, but it is not as precise as we would like. Consider the scientific community's estimates of the climate sensitivity parameter that measures the average surface warming resulting from a doubling of the atmospheric concentration of CO₂, *in equilibrium*. This parameter tells us how sensitive global average temperature is to changes in CO₂ concentration in the long run, and it is perhaps the most studied and most frequently quoted summary statistic in all of climate science.²

Figure 1 presents estimates of the probability distribution for climate sensitivity from several recent studies. Clearly there is much disagreement among these studies. To identify the sources

²Increasingly, scientists are focusing on other summary statistics, partly because the uncertainty in climate sensitivity has remained essentially unchanged over several decades (Knutti and Hegerl 2008), and partly because other measures, such as transient climate response—the global mean temperature change that is realized at the time of CO₂ doubling if CO₂ concentrations are increased by 1 percent per year—provide more dynamic 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).

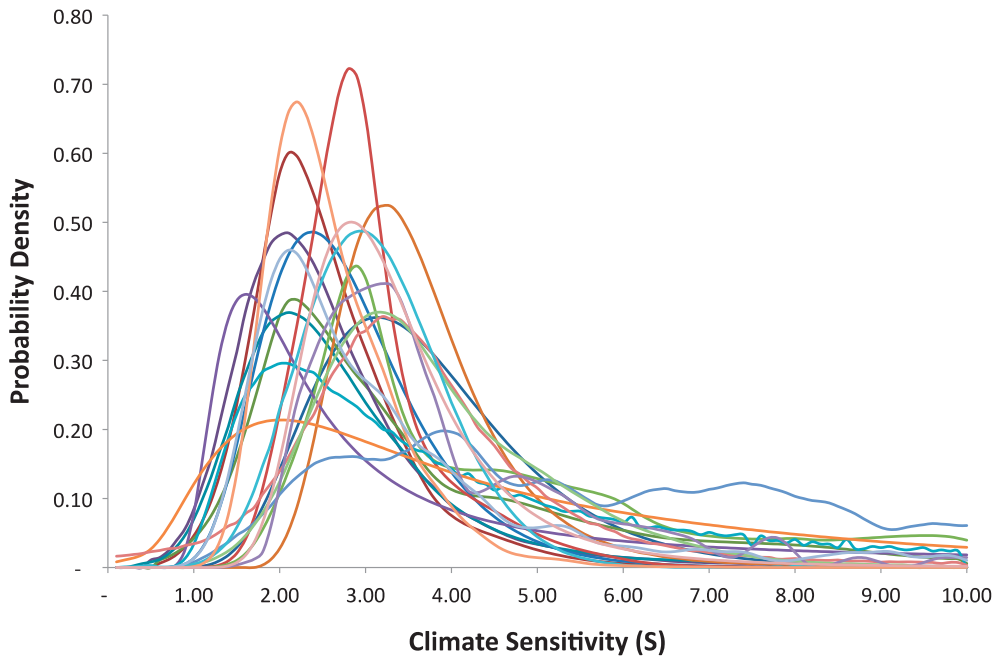


Figure 1 Estimates of the probability distribution for climate sensitivity in degrees Celsius. Reproduced from Millner Dietz and Heal (2013b) with permission.

of the variability in these climate sensitivity estimates, it is helpful to examine how they are produced. Climate sensitivity is usually estimated by running global climate models using hypothetical emissions scenarios. These models are very complex, and there is a great deal of uncertainty about the values of their many parameters. To address this parametric uncertainty, scientists run the models with many different parameter values and weight the predicted values of climate sensitivity using a measure of the likelihood of that run's parameter values. There are many ways of assigning likelihoods to different parameter configurations. One is to assume that all sets of parameter values are equally likely. Another is to ask climate experts for their prior probabilities concerning the model parameters and use them to weight the model runs. Almost all of the more recent studies weight parameters based on how well they reproduce historical climate observations. Many studies arrive at their estimates by using Bayesian methods to combine experts' prior probabilities with historical observations. All of these techniques are used by the studies in Figure 1, although not always in the same way.

Even if we restrict our attention to those studies that base their estimates of climate sensitivity on historical data, there is still a lot of variability in the estimates: some studies use paleoclimate data, some use the cooling from the aerosols released by large volcanic eruptions, and some use the instrumental record over the twentieth century. 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 because there is no single agreed-upon best climate model for estimating climate sensitivity. Instead, given that there is so much about the climate system that we do not understand, there are many models, each with its own representation of physical, chemical, and biological processes. Although these models have many features in

common, they are not statistically independent. Thus we cannot treat the studies in Figure 1 as independent estimates of some underlying distribution. Moreover, because the models are all calibrated on related data, we cannot objectively quantify *how* the models depend on one another. Finally, it is important to point out that it is impossible to obtain the key data that would allow us to determine which model is “best.” To do so would require having an exact copy of the earth, exposing it to increased CO₂ emissions, and ranking models according to how well they predicted this fictitious planet’s climatic response. All of these complications make it difficult to interpret the information in Figure 1 in a standard probabilistic framework without making some strong subjective judgements about how to combine the climate sensitivity estimates from the different studies (Knutti 2010; Knutti et al. 2010).

Despite these methodological issues, it would be incorrect to say that the set of distributions in Figure 1 provide us with *no* information about climate sensitivity. For example, the distributions indicate that it is very unlikely that climate sensitivity is less than 1°C. In addition, a lot of the weight in most of the distributions is in the 2 to 4.5°C range, which is the IPCC’s official “likely” range for climate sensitivity. Finally, a lot of the disagreement between the estimates occurs in the upper tails of the distributions. That is, the data do not limit the estimates of the high end of climate change well at all (Allen et al. 2006; Roe and Baker 2007), which means that we understand little about the likelihood of worst case outcomes.

Contributors to Uncertainty over Time

Climate sensitivity is a very coarse measure of the magnitude of the climate problem. On its own, it does not tell us much about when and where changes in temperature are likely to occur, information that is crucial for policymakers. The main contributors to the uncertainty in temperature projections differ by the lead time and geographic region of the prediction. Hawkins and Sutton (2009) decompose the uncertainty in temperature projections into three parts: initial condition uncertainty (i.e., variations in model predictions due to small changes in how they are initialized), emissions scenario uncertainty (i.e., uncertainty about the path of future GHG emissions), and model uncertainty (i.e., uncertainty about which climate model captures the dynamics of the climate system). Figure 2 shows how these factors contribute to uncertainty about global mean temperature over time.

For short-term predictions, that is, less than about twenty years into the future, initial condition and model uncertainty are dominant. The fact that initial condition uncertainty plays an important role for shorter term predictions should come as no surprise. The climate system is chaotic, and the evolution of chaotic systems, by definition, shows great sensitivity to initial conditions. Hawkins and Sutton (2009) show that initial condition uncertainty is especially important for regional predictions, making substantial contributions to total uncertainty for predictions of sixty years or more. At the global level, for projections twenty to fifty years into the future, the dominant source of uncertainty is model uncertainty (i.e., the variation in predictions across the set of climate models). Beyond about fifty years, emissions scenario uncertainty takes over. That is, the policies we ultimately choose drive the uncertainty in the long run.

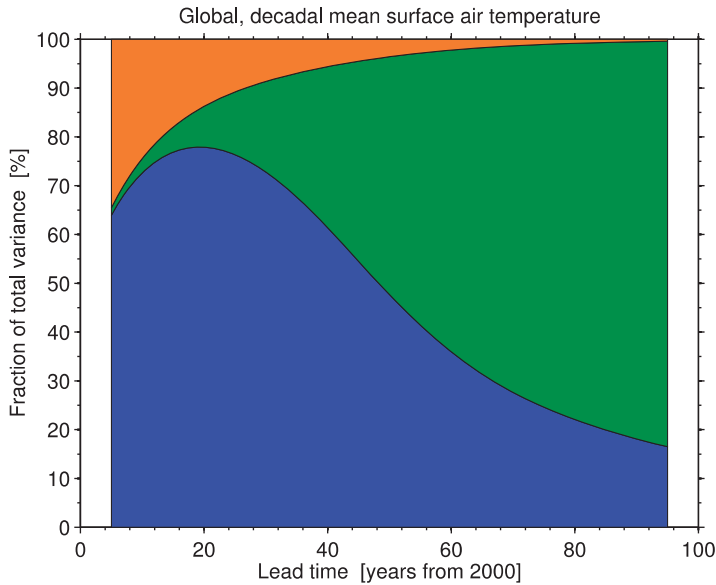


Figure 2 Sources of uncertainty in temperature predictions as a function of lead time. *Source:* Hawkins and Sutton (2009). ©American Meteorological Society. Used with permission.

Notes: The top, middle, and bottom areas denote initial condition uncertainty, emissions scenario uncertainty, and model uncertainty, respectively.

Spatial Dependence of Uncertainty

The effect of CO₂ concentrations on global average temperatures is one of the *best* understood relationships in climate change science. In a hierarchy of increasing uncertainties, climate sensitivity and global mean temperatures are somewhere near the top. However, if we shift our attention to smaller spatial scales, impacts that are localized in time, and less well understood climate variables (e.g., precipitation), uncertainties explode.

Masson and Knutti (2011) have analyzed the dependence of uncertainty in climate predictions on spatial scale. Their results (see Figure 3) show how disagreement between models grows with reductions in spatial scale for both temperature (left panel) and precipitation (right panel). The parameter λ (on the horizontal axis) indicates the scale at which the data from models are averaged—the higher the value of λ , the more the fine spatial detail in the models is lost, and the more the model predictions at each geographic location look like spatial averages over regions of size λ . The more spatial averaging is applied, the more the models agree with one another. Masson and Knutti show that the models' ability to reproduce past observations has a similar dependence on spatial scale.

These results concern only predictions of climate variables themselves. If one continues down the ladder of uncertainty to specific localized predictions of variables that directly affect economic and natural systems (e.g., water availability, growing conditions, sea level rise, and extreme events), matters get even more vague. This suggests that when considering various climate policy options, whether they be mitigation or adaptation, uncertainty and its management are going to be major factors.

Taking a broad view of the scientific models, it is clear that our knowledge is most certain for coarse aggregate statistics such as climate sensitivity, that the main contributors to uncertainty

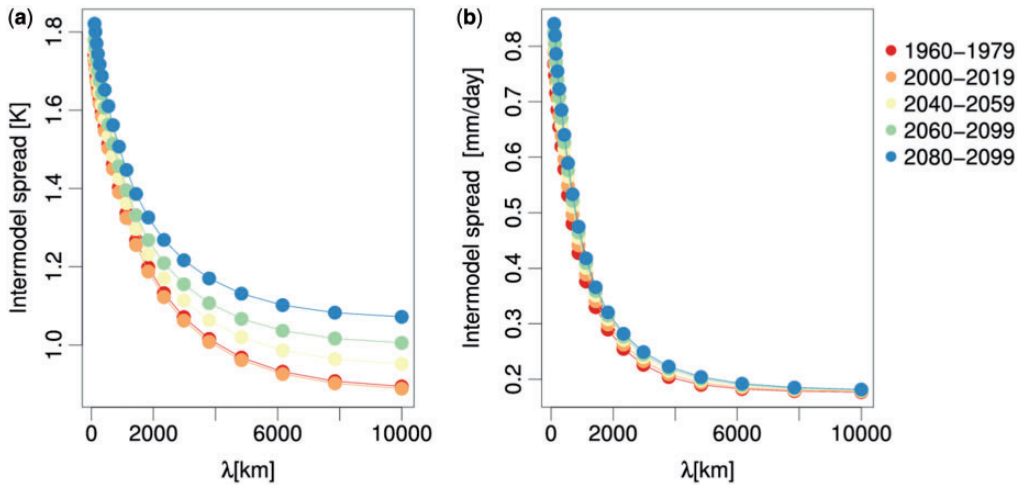


Figure 3 The spatial dependence of uncertainty in climate model predictions. *Source:* Masson and Knutti (2011). ©American Meteorological Society. Used with permission.

Notes: λ represents the scale of spatial averaging of model predictions. Panel a shows the spread, over the suite of climate models used by the IPCC in its fourth assessment report of spatially averaged predictions of changes in surface temperature for several twenty-year time periods. Panel b shows the same for precipitation.

in predictions are different for different lead times, and that uncertainty increases rapidly as we move from predicting aggregate variables to predicting spatially and temporally localized variables. Is there a chance that this uncertainty will be reduced dramatically in the near future? Despite the considerable advances in methods, data, and modeling that have been made over the past few decades, this seems unlikely (Knutti and Sedlacek 2013). Indeed, most uncertainties have not been significantly reduced over the last two decades (Knutti and Hegerl 2008). This suggests that the passage of time is unlikely to radically improve the data available to decision makers. Thus the climate policies we choose today and in the near future will have to be based on our current uncertain estimates.

Sources of Socioeconomic Uncertainty

Sources of socioeconomic uncertainty can be broken down into empirical uncertainty—i.e., lack of knowledge about climate impacts and future technologies—and normative disagreements about the values of key welfare parameters. In the balance of this section we review the empirical uncertainties and their interactions with scientific uncertainties, and then we discuss how differences in value judgments complicate economic analysis of climate policy.

Introducing Socioeconomic Uncertainty

Even if all scientific uncertainty were resolved, we would still face major uncertainties stemming from the socioeconomic dimensions of climate change. Suppose that all scientists agreed on both a climate model and the impacts of climate change on sea level rise, fresh water availability, and so on. Even in this case they would not be able to forecast the future climate because this

requires knowing future emissions. We are far from being able to forecast future emissions because they depend on whether technological change provides us with new ways of reducing GHG emissions, and the policies chosen, which are themselves difficult to forecast.

Even if future emissions were known, the future would *still* be unknown in many important ways. Take the case of sea level rise, for example. Let's suppose we had an accurate forecast of sea levels for the next century. How would societies react to rising seas? By protecting settlements or moving them? Would movements occur in a peaceful and organized manner, or would there be strife and dislocation? In a similar vein, let's consider agriculture. Suppose we know how much future food outputs will fall in response to climate stress. What would the economic consequences of this decline be? Would the poor carry the burden or would it be shared? Another example concerns the fate of small island states whose land is within 10 feet of sea level: it is highly likely that these islands will be uninhabitable by the end of this century. How will the global community react? Will it address the refugee problem, and what compensation will be offered? Clearly the answers to these types of questions will affect the welfare costs and distributional impacts of climate change. Thus even the full resolution of scientific uncertainties would leave huge residual uncertainties about the costs of climate change.

Estimates of Uncertainty in Socioeconomic Impacts

There are a limited number of econometric studies of the impacts of climate on socioeconomic systems (Burke et al. 2009; Dell, Jones and Olken 2008; Deschenes and Greenstone 2007; Fisher Hanemann and Schlenker 2012; Mendelsohn Nordgaus and Shaw 1994; Schlenker Hanemann and Fisher 2005; Schelnker and Lobell 2010), with their main focus being on agriculture, although there has also been some work on conflict and economic growth. The uncertainties associated with the regressions in these studies are measures of impact uncertainties: all of the studies use cross-section or panel data to estimate how climatic conditions affect economic outcomes (given today's technologies). By and large, the findings suggest that climate change is likely to be damaging, with negative impacts on agricultural yields, the frequency of civil conflicts, and economic growth rates (see the survey by Dell, Jones, and Olken 2013).

Scientific and socioeconomic uncertainties interact, and accounting for their interaction is essential for developing plausible estimates of the effects of climate change on social and economic systems. Such an end-to-end approach is rare in empirical work. An interesting exception is the article by Burke et al. (2011) that examines the seven studies of the economic impact of climate change cited earlier. Most of these studies use just one climate model to generate predictions about impacts. However, the IPCC uses more than twenty models for its predictions of climate change. In order to reflect the variation in scientific estimates, Burke et al. (2011) reestimate the socioeconomic impacts in these seven studies using a full range of climate models. They find that this increases the 95% confidence intervals around the central estimates of outcomes by as much as a factor of five and also changes the central estimates themselves. They also argue that failing to account fully for scientific uncertainty results in underestimates of the thickness of the tails of the error distributions, which, in conjunction with the increased confidence intervals, greatly increases the probability of extreme outcomes. Economists have recently focused on the importance of extreme outcomes in policy evaluation (Millner 2013;

Weitzman 2009), suggesting that the possibility of such extreme outcomes may have a disproportionate impact on policy choices.

Technological Uncertainty

Technology determines the scale and carbon intensity of our economic activities. This means that technical change will alter the energy mix of the economy in the future, and thus to a large extent determine how much CO₂ we emit. A good example of how technical change affects the carbon intensity of economic activity can be seen in the rapid growth of hydraulic fracturing in the United States, which between 2002 and 2012 pushed down the price of natural gas from \$12 to less than \$4 per mmbTU, thus replacing coal and reducing US GHG emissions (Energy Information Agency 2013). Such a change would have been hard to imagine even ten years ago. This illustrates the importance of technological uncertainty and the fact that there are many potential developments that could make it easier to reduce GHG emissions or adapt to the consequences of climate change.

Technical change also influences the level of economic output and the growth path of the economy as a whole. Thus technology affects both our demand for energy and how rich we will be in the future. Much of the extraordinary economic growth that has occurred since the Industrial Revolution can be traced to the productivity enhancing effects of technological innovations (Phelps 2013). But we cannot be sure that past progress is a good predictor of future trends (Gordon 2012). In general, we have no theory of technical change that can reliably forecast long-run economic growth rates. This is unsettling because the growth rate of the economy strongly influences the discount rate, which has been shown to be a major driver of climate policy recommendations, as reflected in the debate between Stern and the critics of his 2006 report (Nordhaus 2007; Weitzman 2007).

Uncertainty and Disagreement about Discount Rates

Discount rates allow us to aggregate costs and benefits that are distributed across time. This is a vital task for assessing climate policies because they usually have short-run costs but long-run benefits that accrue mainly to future generations. Calculating discount rates requires us to account for both empirical uncertainties and normative disagreements about the values of welfare parameters.

Empirical uncertainties in the discount rate derive largely from uncertainty about the future growth rate of the global economy. The rate of economic growth enters the discount rate through the Ramsey formula, which breaks the discount rate r into two components: one that depends on how much weight we give to the welfare of future generations (the pure rate of time preference δ) and one that depends on aversion to inequality in consumption between generations (the product of the elasticity of the marginal utility of consumption η and the consumption growth rate g): $r = \delta + \eta g$. Hence if there is uncertainty about future economic growth rates g , then there is also uncertainty about how to discount future changes in consumption (Arrow et al. 2012; Gollier 2012) and thus the benefits of climate policy.

Weitzman (2001) presents a novel methodology for arriving at the discount rate. He notes that different economists recommend different discount rates and suggests that we view these rates as independent estimates of the true underlying rate, which is uncertain. More specifically,

he suggests combining these estimates by taking a weighted average of their implied discount factors, with the weights being the frequency of the individual discount rates in the population. As the Ramsey formula shows, however, people's estimates of the discount rate combine empirical forecasts of the growth rate g with their personal welfare parameters (i.e., their rate of time preference and elasticity of marginal utility). Differences in welfare parameters reflect fundamental *disagreements* about how to measure welfare and are distinct from empirical *uncertainties*. This subtle distinction has important consequences for how we compute discount rates. In particular, if people have different preferences we need to aggregate them in some way in order to identify social preferences that are representative of the population overall.

Gollier and Zeckhauser (2005) and Heal and Millner (2013) pursue this point further and consider how groups should discount the future when different people have different pure rates of time preference. That is, how can we combine these different value judgments into a representative discount rate? One option would be to view this as a social choice problem, and consider reconciling these disagreements about values through a process such as voting. In this case, the outcome would likely be that the median rate of time preference should be used to compute the discount rate. Another option would be to use a more traditional economic approach that seeks to maximize a weighted sum of the diverse agents' welfare, and at the same time asking whether there is a representative agent who can stand in for the group. It can be shown that such a representative agent's rate of time preference declines with time (and approaches the lowest rate in the population asymptotically). Gerlagh and Liski (2012) and Heal and Millner (Forthcoming) show that a declining pure rate of time preference has a substantial effect on climate policy choices.

In summary, the discount rate is determined by empirical factors that are difficult to forecast and welfare parameters whose values are likely to be contested.

Challenges of Modeling Uncertainty

When it comes to modeling uncertainty, most economists instinctively reach for the expected utility framework of von Neuman and Morgenstern (1944), which assumes a well-defined set of possible states of the world, some exogenously given probability distribution over these states, and preferences over uncertain choices that can be represented by expected utilities. We will argue that there are several characteristics of the climate problem that may make it a poor match for the expected utility paradigm.

The first challenge to representing uncertainty in climate models is that we do not have a good picture of all the possible states of the world. One very concrete example of this is the possible disruption of the thermohaline circulation in the Atlantic Ocean, potentially a serious consequence of climate change. Although such a consequence is now believed to be unlikely, it was not even recognized as a possibility until the 1990s (Broecker 1997), suggesting that even today there may be consequences of climate change that we have not considered.

A second challenge is that there is not a generally agreed-upon objective probability distribution over the possible consequences of climate change. Our discussion of Figure 1 demonstrates that this is a problem even for one of the best understood quantities in climate science: the climate sensitivity parameter. It is thus difficult to know how to construct the objective probability distribution that the von Neumann-Morgenstern framework requires. Given these

difficulties, we are led to ask more fundamental questions: Can a unique probability distribution be used to represent our uncertain knowledge, and does rational policy analysis require such a distribution as an input?

Subjective Expected Utility

The traditional answer to these questions draws on a fundamental result of decision theory. This is Leonard Savage's theory of subjective expected utility (SEU) (Savage 1954), which is even more important than von Neuman and Morgenstern's result but is perhaps less familiar to applied economists. 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). Savage's results are much more general than those of von Neuman and Morgenstern. He considers preferences over "acts" (which economists might interpret as policies) that map states of the world into outcomes. For example, the "Buy an Umbrella" act may map the "rain" state into the "happy to be dry" outcome, and the "sun" state into the "irritated to be carrying an unneeded umbrella" outcome. Savage states plausible axioms for preferences over acts, and he 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 basic choice axioms. In contrast to the von Neuman-Morgenstern framework, SEU theory does not assume the existence of unique probabilities. Rather, they are derived endogenously in the result. Thus if one accepts Savage's axioms, the problem of modeling 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!

Challenges to the SEU Framework

Although SEU theory provides a neat one-size-fits-all solution to the problem of decision making under uncertainty, is it the right solution in all situations? This issue 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 (Keynes 1921; Knight 1921). The most famous challenge to the SEU theory's claim to define rational choice came from Daniel Ellsberg's hypothetical choice experiments (Ellsberg 1961), which revealed that when we have no information on which to base decisions, we may violate one of Savage's key axioms: the "sure thing principle."³ 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

³See Gilboa (2009) for a detailed discussion of Savage's axioms and the Ellsberg thought experiments.

landing on “heads”). Many of us are ambiguity averse, preferring to bet on known risks rather than uncertain prospects,⁴ even if, from a subjective probability standpoint, these two situations are indistinguishable (Camerer and Weber 1992; Slovic and Tversky 1974).

Just because many people are ambiguity averse, it does not follow that ambiguity aversion is rational and should be a part of normative policy analysis. In fact, many economists interpret Ellsberg’s (1961) results as a contribution to behavioral economics, grouping it with other behavioral “mistakes” (e.g., Sims 2001). Ellsberg, however, viewed his work in a normative light (Ellsberg 2001). Subsequently, eminent decision theorists have argued that SEU theory is not universally applicable (Binmore 2009), that subjective probabilities cannot always capture our state of knowledge (Gilboa et al. 2008), and that 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 values of welfare parameters, it is something about which reasonable people can reasonably disagree.⁵ There is unlikely to be a knockdown argument that determines whether ambiguity aversion is a behavioral anomaly or a rational response to uncertainty, and both positions contain valuable insights into how we might want to make choices under uncertainty.⁶ From this perspective, the problem of decision making under uncertainty takes on a new social choice dimension. That is, there are many different legitimate preferences, of which some are consistent with SEU theory and some are not, and policy analysis should pay attention to them all. This suggests that it is important to explore alternatives to the standard SEU paradigm because these will form part of any comprehensive approach to designing climate policy in the face of uncertainty.

Alternatives to Expected Utility

Over the past two decades, there has been an explosion of activity in formal approaches to decision making under uncertainty, beginning with the seminal work of Schmeidler (1989) and Gilboa and Schmeidler (1989).⁷ This section reviews several alternatives to the standard expected utility approach to decision making under uncertainty, explaining the ideas underlying them and investigating whether they provide an appropriate framework for policy choice in the climate context.⁸ Where possible, we also provide examples of how they have been applied in climate change economics. There are two types of approaches. One assumes no probabilistic

⁴Millner et al. (2013a) have adapted Ellsberg’s experiments to show that some climate scientists’ beliefs about the value of climate sensitivity cannot be captured by subjective probabilities.

⁵A special issue of *Economics and Philosophy*, 25, no. 3 (2009), is devoted to an exchange of views on this topic.

⁶This 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 (Berlin 1969). We just have to live with the fact that different, equally legitimate, values may come into conflict. This lesson seems to be just as applicable to questions of rationality.

⁷Etner et al. (2012) provides a comprehensive review of these approaches.

⁸In our discussion, “model” and “distribution” are used interchangeably. That is, for us, a model “maps” policies (e.g., GHG emissions) to distributions over outcomes. For example, Figure 1 represents twenty such “models,” each of which maps a doubled CO₂ concentration to a probability distribution over the equilibrium climate sensitivity.

information at all, and the other, known as the multiple priors approach, recognizes that many probability distributions are consistent with what we know and seeks to work with all of them.

Nonprobabilistic Approaches

The first approaches we examine— α -maxmin, and minmax regret—are nonprobabilistic, which means they require no likelihood information. This is a plus in situations where such information is not available, but a minus if it is, because these approaches ignore information that could be valuable in policy selection. This suggests that such approaches may be more appropriate for studies on a local scale where climate predictions are least informative. Such small-scale studies are natural in the context of adaptation to climate change, where location-specific policies are often needed.

Maxmin decision rules

The α -maxmin decision rule ranks policies according to a weighted average of their worst and best outcomes. The parameter α determines how much weight the decision maker gives to a policy's worst outcome, with $1-\alpha$ being the weight given to the policy's best outcome. This parameter could be viewed as representing aversion to a lack of knowledge, with all the weight on the minimum ($\alpha = 1$) implying extreme concern for ambiguity and all the weight on the maximum ($\alpha = 0$) implying a lack of concern about ambiguity. The α -maxmin criterion was derived axiomatically by Arrow and Hurwicz (1977) as a criterion for choice under "ignorance" (i.e., when no information is available). It is the least informationally demanding of the decision criteria. No probabilistic information is needed to implement this approach; all we need is a ranking of outcomes.

The special case of $\alpha = 1$ is called the maxmin rule. It was suggested long before Arrow and Hurwicz's work (Wald 1945). This rule requires us to pick the policy whose worst possible outcome (min) is the least bad (max). An interesting application of this rule to climate change is provided by Woodward and Bishop (1997), who simulate alternative policies in Nordhaus's DICE integrated assessment model of climate policy (Nordhaus 2008) for two alternative climate scenarios: a catastrophic scenario and a scenario 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 rule is to assume the catastrophic scenario to be correct and optimize for this, altering course later if the base case emerges as the correct worldview.

Minmax regret

The minmax regret rule, advocated by Savage (1954), has many similarities to the maxmin rule. The regret of an action $a(s)$ in state s is given by the difference between the value of the best possible outcome in that state and the actual outcome. Minmax regret instructs us to choose the action we regret least over all states of the world. The key here is that rather than focusing on the worst outcome, you focus on how much better off you could have been (this is the "regret") and try to come as close as possible to doing as well as possible over all states of the world. The focus on missed opportunities, rather than worst cases, makes this decision rule less conservative than the maxmin rule. This approach has been used in a number of computational

studies of decision making that seek to find policies that are robust to a broad range of future scenarios (e.g., Lempert et al. 2006).

Multiple Priors Approaches

A second class of alternatives to the expected utility framework is based on the idea that if the world is not kind enough to provide us with a unique probability distribution over states, then there must be many different probability distributions that are consistent with what we know: that is, we have multiple priors. The idea is to use all of these distributions and weight them in some way. Because decision rules based on multiple priors make use of likelihood information, they are naturally more complex than the nonprobabilistic approaches discussed in the previous section. This section summarizes the maxmin expected utility approach, the smooth ambiguity approach, and the multiplier preference approach to decision making with multiple priors.

Maxmin expected utility approach

The oldest of the multiple prior approaches, and in many ways the simplest, is the maxmin expected utility rule (Gilboa and Schmeidler 1989). The first step in implementing this approach is to identify all distributions (or “priors”) that are consistent with what is known. Policies are then ranked according to their lowest expected utility over this set of possible distributions. This means that for each policy we find the distribution that gives it the worst expected outcome and evaluate policies based on these worst cases. Thus this approach focuses on “bad distributions.” As with all of the multiple prior approaches we will review here, there is no assumption of a unique probability distribution, and we can have incomplete probabilistic information that is compatible with many different distributions. Figure 1 illustrates such a situation for climate sensitivity: the data there are consistent with many but not all probability distributions. With this approach, if you do not know what the probabilities are, you ask “what probabilities are consistent with what I do know?”, and then use all of these probabilities.

What this approach does not do, however, is *rank* the possible distributions in any way; it simply focuses on the worst case. Thus this approach reflects extreme ambiguity aversion—we are very strongly averse to the spread in expected utilities across the set of different probability distributions. In practice, however, we may not want to be so cautious and instead may prefer to reflect the fact that some distributions are more plausible than others. We next discuss the smooth ambiguity model, an approach that allows us this freedom.

Smooth ambiguity model

In the smooth ambiguity model (Klibanoff Marinaci and Mukherji 2005) there are again many possible distributions consistent with what we know, but in this case we assign a subjective weight to each distribution. We then use each of the possible distributions to evaluate a policy, and we combine these evaluations into a single value using the subjective weights of the distributions and a flexible measure of ambiguity aversion.

The smooth ambiguity approach allows us to use all of the likelihood information contained in the possible distributions. But in order to do this we need to specify weights for each distribution, and these weights reflect information that is very different from the distributions

themselves. The weights are subjective judgments, whereas the distributions are generally considered to be “objective” (i.e., informed by data). The smooth ambiguity model recognizes the necessity of such subjective judgments, in much the same way as the standard SEU framework. However, unlike the SEU framework, the smooth ambiguity model allows us to treat these two kinds of information very differently. That is, by introducing ambiguity aversion, this model preserves the distinction between subjective and objective judgments.

Millner, Dietz, and Heal (2013b) present a model that illustrates the maxmin expected utility and smooth ambiguity approaches. They assume multiple scientific models (such as those in Figure 1), each of which can be used to evaluate a policy. Each individual model, in conjunction with a policy choice, predicts a distribution of consumption levels, enabling us to compute an expected utility for that policy-model combination.

MDH find that the introduction of ambiguity aversion can increase or decrease the optimal level of GHG abatement: the key is whether models that show a high payoff from abatement (in expected utility terms) also have a high marginal payoff. Their dynamic simulations show that for the DICE model, aversion to ambiguity about the distribution of climate sensitivity increases the welfare benefits of mitigation, in some cases dramatically. In the extreme case, where aversion to ambiguity is very large, the model reduces to the maxmin expected utility model.

Multiplier preferences approach

The final multiple priors approach we consider is the multiplier preferences approach, introduced by Hansen and Sargent (2007) and used increasingly widely in macroeconomics. Macroeconomists have recognized that it is not possible to know the true evolution of the economy and that as a result they must contend with multiple possible models. In this sense, the challenges they face are similar to those faced by economists concerned about climate change. Under the Hansen and Sargent (2007) approach, the analyst chooses a best guess distribution but is not sure that this distribution is correct, and so allows for possible model misspecification, again by allowing for a set of possible distributions. To evaluate a policy, the analyst computes its expected utility for each of these distribution and adds to this a penalty term that depends on how “close,” in a statistical sense, a given distribution is to the best guess. The further a distribution is from the best guess, the bigger the penalty. Policies are then evaluated by taking the minimum over all distributions of this sum of the expected utility and penalty terms. Thus since distributions that are far away from the best guess are penalized, there is a trade-off between identifying the policy that maximizes the minimum expected utility over all distributions and identifying a policy that does well under the best guess distribution.

The multiplier preferences approach makes sense if there is a clearly preferred model that we can use to describe our uncertainty but we are not sure about 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 for the “best” model, then this approach is not so attractive. 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, but it is not clear that it is appropriate for climate change. As with maxmin expected utility, this approach ignores likelihood information and simply ranks priors by their closeness to our best guess.

Xepapadeas (2012) has applied the Hansen-Sargent multiplier preferences approach to climate problems. He considers a model in which output produces a pollutant, which could be a GHG, according to a stochastic process, but we are not sure exactly which process it is. By applying the Hansen and Sargent framework, Xepapadeas solves for optimal emissions policies that account for model misspecification and demonstrates a “cost of robustness”: the larger the set of distributions to which our policies have to be robust, the lower their expected payoffs.

Conclusions

This article has reviewed the sources of uncertainty about climate change, methods of representing them, and alternatives to the expected utility approach to decision making under uncertainty. What can we conclude from this survey? First, it is clear that uncertainty is here to stay. Although we are constantly making incremental progress, it seems unlikely that our knowledge about the scientific and socioeconomic consequences of climate change will radically improve over the next few decades. However, considerably more could be done to effectively characterize our uncertainty about socioeconomic impacts. End-to-end analyses, which integrate uncertainty in both climate science and climate impacts estimates, should be an important part of such efforts. Point estimates of impacts based on projections from a single climate model are not particularly informative. We also have to recognize, however, that many of the most important empirical uncertainties, such as those 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 unknowable. We simply do not have credible tools for making such predictions.

Second, although the nature and extent of these uncertainties could lead to policy paralysis, we believe this would be a big mistake. We have more than enough information about climate change to understand that it is a serious problem that requires immediate policy attention: uncertainty does not imply ignorance. Our job as economists is to assess the welfare consequences of climate policies as best we can and provide analytical tools that can be used to inform policy choices. However, the methods that we use to evaluate small-scale projects with well-defined, short-run consequences are unlikely to be appropriate for a problem as global, long run, and uncertain as climate change. This offers us an opportunity to develop an inclusive and creative approach to policy analysis that incorporates a variety of viewpoints. As illustrated by the debate about discounting, the decision tools and welfare parameters we use to evaluate policy can have as much of an effect on policy recommendations as the data themselves. This is no less true when thinking about how to incorporate uncertainty into our models: different approaches can lead to radically different policy recommendations.

The survey of the decision theory literature we have presented here indicates that there are many promising approaches to decision making under deep uncertainty that provide new methods for representing our information about the nature and impacts of climate change, and allow us to make policy recommendations that account for the limitations of our knowledge. When combined with empirical inputs that represent both what we do and do not know, such tools could provide useful policy guidance. It is important to recognize, however, 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

- Allen, M., N. Andronova, B. Booth, S. Dessai, D. Frame, C. Forest, J. Gregory, G. Hegerl, R. Knutti, C. Piani, D. Sexton, and D. Stainforth. 2006. Observational constraints on climate sensitivity. In: *Avoiding dangerous climate change*, ed. H. Schellnhuber, W. Cramer, N. Nakicenovic, T. Wigley, and G. Yohe, 406. Cambridge, UK: Cambridge University Press.
- Allen, M. R., D. J. Frame, C. Huntingford, C. D. Jones, J. A. Lowe, M. Meinshausen, and N. Meinshausen. 2009. Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature* 458 (7242): 1163–66.
- Arrow, K., and L. Hurwicz. 1977. An optimality criterion for decision-making under ignorance. In *Studies in resource allocation; processes*, 482. Cambridge, UK: Cambridge University Press.
- Arrow, K. J., M. L. Cropper, C. an Gollier, B. Groom, G. M. Heal, R. G. Newell, W. D. Nordhaus, R. S. Pindyck, W. A. Pizer, P. R. Portney, T. Sterner, R. S. J. Tol, and M. L. Weitzman. 2012. How should benefits and costs be discounted in an intergenerational context? Resources for the Future Discussion Paper 12 (53).
- Berlin, I. 1969. *Four essays on liberty*. Oxford: University Press.
- Binmore, K. 2009. *Rational decisions*. Princeton, New Jersey: Princeton University Press.
- Broecker, W. S. 1997. Thermohaline circulation, the Achilles heel of our climate system: Will man-made CO₂ upset the current balance? *Science* 278 (5343): 1582–88.
- Burke, M., J. Dykema, D. Lobell, E. Miguel, and S. Satyanath. 2011. Incorporating climate uncertainty into estimates of climate change impacts, with applications to U.S. and African agriculture. NBER Working Paper 17092.
- Burke, M. B., E. Miguel, S. Satyanath, J. A. Dykema, and D. B. Lobell. 2009. Warming increases the risk of civil war in Africa. *Proceedings of the National Academy of Sciences* 106 (49): 20670–74.
- Camerer, C., and M. Weber. 1992. Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty* 5 (4): 325–70.
- Dell, M., B. Jones, and M. Olken. 2008. Climate change and economic growth: Evidence from the last half century. NBER Working Paper 14132.
- . 2013. What do we learn from the weather? The new climate-economy literature. NBER Working Paper 19578.
- Deschenes, O., and M. Greenstone. 2007. The economic impacts of climate change: Evidence from Agricultural output and random fluctuations in weather. *American Economic Review* 97 (1): 354–85.
- Ellsberg, D. 1961. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 75 (4): 643–69.
- . 2001. *Risk, ambiguity and decision*. New York: Gerland Publishing, Inc.
- Energy Information Agency. 2013. Today in Energy. at <http://www.eia.gov/todayinenergy/detail.cfm?id=13731> (accessed November 10, 2013).
- Etner, J., M. Jeleva, and J.-M. Tallon. 2012. Decision theory under ambiguity. *Journal of Economic Surveys* 26 (2): 234–70.
- Fisher, A. C., M. Hanemann, and W. Schlenker. 2012. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. *American Economic Review* 102 (7): 3749–60.
- Gerlagh, R., and M. Liski. 2012. Carbon prices for the next thousand years, CESifo Working Paper 3855.
- Gilboa, I. 2009. *Theory of decision under uncertainty*. New York: Cambridge University Press.
- Gilboa, I., A. Postlewaite, and D. Schmeidler. 2009. Is it always rational to satisfy Savage's axioms? *Economics and Philosophy* 25 (3): 285–296.
- Gilboa, I., A. W. Postlewaite, and D. Schmeidler. 2008. Probability and uncertainty in economic modeling. *Journal of Economic Perspectives* 22 (3): 173–88.

- Gilboa, I., and D. Schmeidler. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18 (2): 141–53.
- Gollier, C. 2012. *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton, New Jersey: Princeton University Press.
- Gollier, C., and R. Zeckhauser. 2005. Aggregation of heterogeneous time preferences. *Journal of Political Economy* 113 (4): 878–96.
- Gordon, R. J. 2012. Is U.S. economic growth over? Faltering innovation confronts the six headwinds. NBER Working Paper 18315.
- Hansen, L. P., and T. J. Sargent. 2007. *Robustness*. Princeton, New Jersey: Princeton University Press.
- Hawkins, E., and R. Sutton. 2009. The potential to narrow uncertainty in regional climate predictions. *Bulletin of the American Meteorological Society* 90: 1095–1107.
- Heal, G., and B. Kristrom. 2002. Uncertainty and climate change. *Environmental and Resource Economics* 22 (1–2): 3–39.
- Heal, G., and A. Millner. 2013. Discounting under disagreement. NBER Working Paper 18999.
- . Forthcoming. Agreeing to disagree on climate policy.
- Intergovernmental Panel on Climate Change (IPCC). 2007. Summary for policymakers. In: *Climate change 2007: The physical science basis*. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, ed. S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. Averyt, M. Tignor, and H. Miller. Cambridge, UK, and New York: Cambridge University Press.
- Keynes, J. M. 1921. *A treatise on probability*. Macmillan and Co.
- Klibanoff, P., M. Marinaci, and S. Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* 73 (6): 1849–92.
- Knight, F. 1921. *Risk, uncertainty, and probability*. New York: Houghton Mifflin.
- Knutti, R. 2010. The end of model democracy? *Climatic Change* 102 (3): 395–404.
- Knutti, R., R. Furrer, C. Tebaldi, J. Cermak, and G. A. Meehl. 2010. Challenges in combining projections from multiple climate models. *Journal of Climate* 23 (10): 2739–58.
- Knutti, R., and G. C. Hegerl. 2008. The equilibrium sensitivity of the Earth's temperature to radiation changes. *Nature Geoscience* 1 (11): 735–43.
- Knutti, R., and J. Sedlacek. 2013. Robustness and uncertainties in the CMIP5 climate model projections. *Nature Climate Change* 3: 369–73.
- Kunreuther, H., G. Heal, M. Allen, O. Edenhofer, C. B. Field, and G. Yohe. 2013. Risk management and climate change. *Nature Climate Change* 3: 447–50.
- Lempert, R. J., D. G. Groves, S. W. Popper, and S. C. Bankes. 2006. A general, analytic method for generating robust strategies and narrative scenarios. *Management Science* 52 (4): 514–28.
- Masson, D., and R. Knutti. 2011. Spatial-scale dependence of climate model performance in the CMIP3 ensemble. *Journal of Climate* 24 (11): 2680–92.
- Mendelsohn, R., W. D. Nordhaus, and D. Shaw. 1994. The impact of global warming on agriculture: A Ricardian analysis. *American Economic Review* 84 (4): 753–71.
- Millner, A. 2013. On welfare frameworks and catastrophic climate risks. *Journal of Environmental Economics and Management* 65 (2): 310–25.
- Millner, A., R. Calel, D. A. Stainforth, and G. MacKerron. 2013a. Do probabilistic expert elicitation capture scientists' uncertainty about climate change? *Climatic Change* 116 (2): 427–36.
- Millner, A., S. Dietz, and G. Heal. 2013b. Scientific ambiguity and climate policy. *Environmental and Resource Economics* 55 (1): 21–46.
- Nordhaus, W. D. 2007. A review of the Stern review on the economics of climate change. *Journal of Economic Literature* 45 (3): 686–702.
- . 2008. *A question of balance*. New Haven & London: Yale University Press.
- . 2011. The economics of tail events with an application to climate change. *Review of Environmental Economics and Policy* 5 (2): 240–57.
- Phelps, E. 2013. *Mass flourishing: How grassroots innovation created jobs, challenge, and change*. Princeton, New Jersey: Princeton University Press.
- Pindyck, R. S. 2007. Uncertainty in environmental economics. *Review of Environmental Economics and Policy* 1 (1): 45–65.

- . 2011. Fat tails, thin tails, and climate change policy. *Review of Environmental Economics and Policy* 5 (2): 258–74.
- Roe, G. H., and M. B. Baker. 2007. Why is climate sensitivity so unpredictable? *Science* 318 (5850): 629–32.
- Savage, L. J. 1954. *The foundations of statistics*. Wiley and Sons.
- Schlenker, W., M. Hanemann, and A. C. Fisher. 2005. Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach. *American Economic Review* 95 (1): 395–406.
- Schlenker, W., and D. B. Lobell. 2010. Robust negative impacts of climate change on African agriculture. *Environmental Research Letters* 5 (1): 014010.
- Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica* 57 (3): 571–87.
- Sims, C. 2001. Pitfalls of a minimax approach to model uncertainty. *American Economic Review* 91 (2): 51–54.
- Slovic, P., and A. Tversky. 1974. Who accepts Savage's axiom? *Behavioral Science* 19 (6): 368–73.
- Stone, D. A., M. R. Allen, P. A. Stott, P. Pall, S.-K. Min, T. Nozawa, and S. Yukimoto. 2009. The detection and attribution of human influence on climate. *Annual Review of Environment and Resources* 34 (1): 1–16.
- von Neumann, J., and O. Morgenstern. 1944. *Theory of games and economic behaviour*. Princeton, New Jersey: Princeton University Press.
- Wald, A. 1945. Statistical decision functions which minimize the maximum risk. *The Annals of Mathematics* 46 (2): 265–80.
- Weitzman, M. L. 2001. Gamma discounting. *American Economic Review* 91 (1): 260–71.
- . 2007. A review of the Stern review on the economics of climate change. *Journal of Economic Literature* 45 (3): 703–24.
- . 2009. On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 91 (1): 1–19.
- . 2011. Fat tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy* 5 (2): 275–92.
- Woodward, R. T., and R. C. Bishop. 1997. How to decide when experts disagree: Uncertainty-based choice rules in environmental policy. *Land Economics* 73 (4): 492–507.
- Xepapadeas, A. 2012. Ambiguity and robustness in international pollution control. In *Climate change and common sense: Essays in honour of Tom Schelling*, ed. R. W. Hahn, and A. Ulph. Oxford: Oxford University Press.