

Uncertainty and Ambiguity in Environmental Economics: Conceptual Issues

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

¹Columbia University

²London School of Economics and Political Science

May 25, 2017

1 Introduction

Uncertainty is ubiquitous in environmental economics. This is inevitable: we study the interactions between socio-economic systems and biogeochemical systems, and in general neither of these is fully understood. As a consequence, our grasp of their interactions is necessarily rather limited. Climate change is a good case study: the scientific community understands some aspects of the behavior of the climate system well, but others poorly. We are certainly no better off, and often worse off, when it comes to our understanding of economic systems. And we are, as we will argue below, particularly weak on the interactions between the two. Biodiversity loss is another important problem for which our lack of knowledge is striking. We are in the midst of a mass extinction comparable to those of pre-history, an event which will transform the world around us, and one that scientists suspect will be greatly detrimental to human well-being. Yet we have little formal understanding of why biodiversity matters to us or of how to model the economic consequences of its loss.

The prevalence of uncertainty in our field has long been recognized, and has led to some seminal papers. The work of Arrow & Fisher (1974) and Henry (1974) on the option value associated with uncertainty, learning, and irreversibility are well-known examples. Because

*gmh1@columbia.edu

†a.millner@lse.ac.uk

of this, there are already several surveys of uncertainty in environmental economics. Mäler & Fisher (2005), Heal & Kriström (2002), Pindyck (2007), and Aldy & Viscusi (2014) survey the literature up to early this century. Rather than repeating what has gone before, in this chapter we review new ideas that were only beginning to gain currency in environment economics when these surveys were written. These ideas touch on deep conceptual questions about how uncertainty can be modeled, and which decision criteria should be applied when objective probabilistic information about the consequences of policy choices is not available.¹

The early literature on decision under uncertainty is based on the expected utility model of von Neumann & Morgenstern (1944). They axiomatized preferences over ‘lotteries’, in which the probabilities of alternative outcomes are objectively known (e.g. the toss of a coin or the roll of a die). Their theory was powerfully extended by Savage (1954), who developed a sophisticated theory of choice under ‘subjective’ uncertainty. He showed that if agents obey certain primitive axioms, which do not presuppose the existence of probabilities, they should act as if they are maximizing a subjective expected utility functional. Subjective probabilities in the Savage framework capture ‘degrees of belief’ even when no objective information exists, as in answering the question “what is the probability of life elsewhere in the universe?” This development provided a basis for applying the expected utility model in a wide variety of contexts, both as a positive model of behaviour (later largely refuted (Kahneman et al., 2000)), and as a benchmark normative theory of rational decision-making. Since then however important developments in decision theory have challenged the applicability of expected utility theory in situations characterized by ‘deep’ uncertainty, or ambiguity. New models of rational decision-making designed for informationally poor environments have been developed, and are beginning to filter into applied economics, particularly finance and macroeconomics. Their recent applications in environmental economics are our subject matter here.

Our treatment of uncertainty in environmental applications will be motivated by two leading examples: climate change and biodiversity loss. We argue that in these cases uncertainty is sufficiently far reaching that standard decision-making tools such as expected utility theory may no longer capture important aspects of our uncertainty preferences. Richer models of decision-making, which allow us to express lack of confidence in our information, may be desirable in these cases.

¹A previous review paper of ours (Heal & Millner, 2014a) covers some of the same ground, but in considerably less detail.

1.1 Uncertainty and Climate Policy

Climate policy choices must be made in the face of several sources of uncertainty. At the highest level we can classify them into uncertainty about climate science and uncertainties about the socio-economic consequences of climate change. These categories can be further subdivided: scientific uncertainty is often resolved into internal variability, model uncertainty and emissions uncertainty, and socio-economic uncertainty into positive uncertainties about the magnitude of climate damages, opportunity costs of mitigation, and future mitigation costs, and normative disagreements² about the welfare framework that should be used to evaluate policy options.

We review the origins of scientific uncertainty first. Internal variability arises because climate models are highly non-linear, and exhibit sensitive dependence on initial conditions [“chaotic behavior”]. Small errors in the specification of initial conditions in model runs can lead to significant differences in predicted outcomes. Since the current state of the climate system is known imperfectly (our observation network is sparse, and measurement instruments introduce errors of their own), global climate models need to be run many times with a variety of initial conditions in order to build up an ensemble of projections that reflect the possible future states of the climate within a given model.

Model uncertainty just reflects the fact that there is a lot that is still unknown about the physics of several important processes in the climate system. An important example is the so-called cloud radiative feedback, which is a major source of uncertainty in the response of the climate system to changes in GHG concentration (Stephens, 2005; Zickfeld et al., 2010). Changes in temperature affect cloud formation, and since cloud cover affects the emission and absorption of solar radiation, this has a feedback effect on temperature itself. The magnitude of this feedback is not well understood (although it is likely to be positive (Dessler, 2010)), and different modeling groups represent cloud physics, and other poorly understood physical phenomena, in different ways.³

Finally, emissions uncertainty simply refers to the fact that we do not know the quantity of greenhouse gases that will be emitted over the coming decades. Climate models are thus usually forced with a variety of emissions scenarios, leading to a further expansion in the range of predicted outcomes. All of these types of scientific uncertainty are depen-

²Normative disagreements are usually seen as conceptually distinct from empirical uncertainties. However, recent work by philosophers has somewhat blurred this distinction (MacAskill, 2016).

³Technically, cloud formation is a sub grid process, i.e. it occurs on a spatial scale smaller than the spatial resolution of climate models. So cloud feedbacks are often put into the models through ‘reduced form’ parameterizations.

dent on scale: uncertainties are largest on small spatial scales, and reduce significantly at continental and global scales.

Uncertainty about the response of economic and social systems to climate change is probably greater than scientific uncertainty - though both are large in some absolute sense. Scientific uncertainty can be reduced by refining climate models and testing their ability to reproduce historical data, in many cases stretching back hundreds of thousands of years. Clearly there is no comparable data set that will allow us to understand the response of future social and economic systems to climatic change. Our current forms of economic and social activity are at most a few hundred years old, and during that period the global climate has been remarkably stable by historical standards.

Socio-economic uncertainties can be roughly broken into two categories: model uncertainty again – including uncertainty about parameter values in a given model, and uncertainty about the structural relationships that govern the evolution of the climate-economy system – and disagreements about values. The importance of model uncertainty is highlighted by the substantial variations in the predictions of different integrated assessment models (IAMs). IAMs combine simplified scientific models of climate change with macroeconomic growth models, estimates of the dynamic costs of mitigation, and damage functions that quantify the impact of climate change on economic activity. Different models represent these components, and many other factors, in different ways, leading to a very large distribution of policy recommendations. These models are playing an increasingly important role in climate policy analysis. For example, an interagency committee of the US government used three IAMs to inform its choice of the social cost of carbon for regulatory cost-benefit analysis (Greenstone et al., 2013). Yet it is unclear what empirical basis we have for believing that any given IAM captures the evolution of the climate-economy system over the coming centuries. The reliance of IAMs on untested structural assumptions, and the strong dependence of their outputs on arbitrary modeling choices, has caused some to suggest that they “have crucial flaws that make them close to useless as tools for policy analysis.” (Pindyck, 2013).

To illustrate some of the empirical difficulties IAMs face, consider the problem of specifying a damage function – a function that translates climatic changes into changes in aggregate economic output. The climate impacts literature (IPCC 5AR WG3 2014) generally breaks the socio-economic impacts of climate change into several components. These include, for example, the impact of rising sea levels on coastal properties and infrastructure (Yohe et al., 1996), the impact of heat on food crop yields (Lobell et al., 2011), the effects

of higher temperatures on labor productivity (Heal & Park, 2013), the effect on crime and conflict (Hsiang et al., 2013) and the effect of climate change on health through temperature stress and the spread of disease vectors. Reviews of the recent micro-econometric literature on climate damage estimates can be found in Dell et al. (2014); Houser et al. (2015); Carleton & Hsiang (2016). But there is as yet no systematic attempt to aggregate and monetize these damages for the world as a whole, so we are in the dark about the global economic costs of climate change. In practice, IAMs use reduced form damage functions to quantify climate impacts, but as Pindyck (2013) trenchantly notes, their “descriptions of the impact of climate change are completely ad hoc, with no theoretical or empirical foundation.” For example, it has become conventional, following Nordhaus (1994), to assume an inverse quadratic relationship between damages and temperature. Yet this assumption has no empirical or theoretical justification, and has great consequences for IAM outputs, especially when it comes to estimating the costs of extreme climatic changes. Other IAM components – e.g. their representation of long-run technological change – are subject to similar concerns (Millner & McDermott, 2016).

A second type of socio-economic uncertainty is not really uncertainty at all, but rather disagreement about values. The values that are chosen for the parameters of intertemporal social welfare functions are key inputs to IAMs, and are the subject of debate and, on occasion, controversy. The most prominent examples are the pure rate of time preference, which discounts the wellbeing of future generations, and the elasticity of the marginal utility of consumption, which captures aversion to intergenerational consumption inequalities. Both of these parameters express distributional value judgments that have been hotly debated. These disagreements are fundamental, and can have a huge impact on the policy recommendations IAMs produce. Quoting Pindyck (2013) again,

“Nordhaus (2008) finds that optimal abatement should initially be very limited, consistent with an SCC (social cost of carbon) around \$20 or less, while Stern (2007) concludes that an immediate and drastic cut in emissions is called for, consistent with an SCC above \$200. Why the huge difference? Because the inputs that go into the models are so different. Had Stern used the Nordhaus assumptions regarding discount rates...he would have also found the SCC to be low. Likewise, if Nordhaus had used the Stern assumptions, he would have obtained a much higher SCC.”

Disagreements on the ethical principles that should govern the measurement of intergenerational social welfare are unlikely to be resolved any time soon. As with many primitive

ethical questions, reasonable people can reasonably disagree on this issue. This suggests that methods for aggregating diverse views on social preferences into some compromise judgement could be useful tools for achieving a measure of consensus, allowing us to move beyond an ethical impasse (Heal & Millner, 2014b; Millner, 2016; Millner & Heal, 2017). Nevertheless, we acknowledge that the aggregation of opinions on social preferences raises new normative questions – there is no free lunch. In the remainder of this section we set ethical disagreements aside, and focus on the empirical uncertainties that climate policy choices must contend with.

Because of the substantial gaps in our knowledge about the magnitude and consequences of climate change, scientific and socio-economic uncertainties are not readily quantifiable by objective probability density functions (PDFs) that all reasonable people will agree on. It is common to find several PDFs for relevant variables in the scientific literature, or several expert opinions that need not be in the form of PDFs at all (Kriegler et al., 2009). By way of example, consider scientific estimates of the equilibrium climate sensitivity (ECS), perhaps the central summary statistic in climate change science. The ECS is the equilibrium increase in global mean temperatures that would occur if the concentration of atmospheric CO₂ were doubled. It is thus a coarse measure of the sensitivity of the global climate to changes in CO₂ concentrations. Figure 1 shows twenty PDFs for the ECS from the scientific literature. The differences between the studies are largely due to differences in the climate models that are used by different modeling groups, but are also in part due to the use of different data sets and statistical methodologies. Although there are significant differences between these estimates, all are generated by intellectually respectable modeling groups.

Given the substantial variation in scientific estimates of even highly aggregated summary statistics like the ECS, how should the analyst proceed? Should she seek to aggregate different PDFs into a single composite distribution and use decision techniques that rely on the existence of such a distribution, or should she work with the ambiguity inherent in the information available to her? Naively combining multiple PDFs is generally not advisable as different models are not independent estimates of a “true” underlying distribution. For example, although the climate models produced by different modeling groups have important differences, they are generally calibrated at least in part on the same data and based on the same physical principles and sets of equations. In order to objectively combine PDFs from multiple models of the same phenomenon into a single probability estimate it is necessary to account for the dependencies between models, and have some measure




Figure 1: Scientific estimates of equilibrium climate sensitivity

of their relative predictive performance. When meaningful verification data are available this can be a very productive strategy. For example, Nate Silver’s FiveThirtyEight blog used precisely such a methodology to aggregate many electoral polls into a single summary probability forecast which was more accurate than any individual poll. A lot of the skill in this exercise involves using the historical performance of each poll to design a weighting scheme for aggregating polls.⁴

Can we do the same thing for integrated assessment models of climate change? In principle there is nothing preventing us from evaluating the performance of these models on historical data sets (Millner & McDermott, 2016), however for some important model components – e.g. the damage function – this exercise is unlikely to reveal much information about which of the models is a better match to reality. One reason for this is that there has been only a small amount of warming in the roughly 150 years since the industrial revolution began, so finding a climate change signal in overall growth outcomes is very difficult.⁵ A second reason is that many of the crucial assumptions in IAMs are exogenously specified scenarios, rather than structural relationships that are subject to empirical verification. In the DICE model (Nordhaus & Sztorc, 2013) this applies to the trajectories of key time series such as population change and abatement costs. In the PAGE (Hope, 2006) and FUND (Tol, 1997) models the entire baseline sequence of global GDP is put in by hand. A third, and more fundamental, reason is that past predictive performance may not be a good indicator of future predictive performance in the case of climate policy. Unlike the fundamental physical principles that underlie climate models, which derive their authority from a myriad of successful applications across vastly different spatial and temporal scales, the structural assumptions that underpin our models of economic growth and technical progress have had only patchy predictive success in other economic applications. They are enormously valuable *explanatory* tools (e.g. much of economic growth theory uses models to *explain* historical differences in growth outcomes between countries), but their utility as quantitative *predictive* tools is far less established. Thus, even if we were able to calibrate our models so that they replicate growth outcomes over the previous century reasonably well, we do not have high confidence that we have identified the true dynamics that govern long-run changes in the global economy. Calibration does not imply out-of-

⁴See Knutti et al. (2017) for a recent discussion of weighting climate models.

⁵It is well known that temporary temperature shocks are statistically associated with changes in both the level of economic output, and the pace of economic growth (Dell et al., 2012; Heal & Park, 2015). However these relationships are deduced using data from relatively short time periods (50 years), over which the climate was approximately stationary. They are at best indicative of what might happen if the *distribution* of shocks changes.

sample predictive power. Millner & McDermott (2016) provide further discussion of the difficulty, and importance, of attempting to confirm the structural modeling assumptions that IAMs rely on.

The lack of availability of unique objective PDFs poses methodological questions for the economic analysis of climate policies. In practice the analyst must choose a framework for policy analysis that represents the state of our knowledge (or lack thereof) in as complete and honest a manner as possible. The tendency in the economic literature on climate change has been to adopt the expected utility framework for analyzing climate policies. In order to apply this framework to the climate problem it is necessary to make subjective judgements about how to transform all the known limitations and ambiguities in the scientific literature on climate science and impacts into a unique mapping between policy choices and probability distributions over outcomes. This process invariably does violence to the available data, mixing largely arbitrary subjective judgements about how to combine alternative models with more objective (conditional) probabilities that arise when fitting models to data. There are however alternative decision tools, involving “multiple priors,” that explicitly recognize that we generally do not have the quality of information needed to define a unique objective PDF. We investigate the possibility of applying these approaches in the climate context below.

1.2 Uncertainty and Biodiversity

We understand generally that the loss of biodiversity has an economic cost, which can be thought of as the loss of various contributions that biodiversity makes to the economy. We probably don’t understand fully these contributions, but they include the value of biological resources in bioprospecting, the insurance value of biodiversity, its value in enhancing productivity of natural ecosystems, and its value as the origin of all domesticated plants and animals (see e.g. Heal (2016)).

Consider the first of these: many important pharmaceutical products are derived from plants and insects, and pharmaceutical companies routinely scan extracts from such sources as possible drug leads. Aspirin is probably the most famous plant-derived drug, occurring naturally in the bark of willow trees. After decades and billions of dollars of research, the pharmaceutical industry has not come up with anything that is clearly better than aspirin as a painkiller and anti-inflammatory. In addition, aspirin reduces the risks of heart attack and stroke, and recent research (Drew et al., 2016) suggests strongly that regular use of aspirin reduces the risk of a range of common cancers. A more contemporary example

is Glucobay, a drug sold by Bayer that lowers blood glucose levels in diabetics and is in great demand in view of the growing menace of diabetes. Its key ingredient is a natural sugar called Acarbose, which reduces the absorption of glucose into the bloodstream. In a US patent application Bayer revealed that a bacterial strain that originates from Kenya's Lake Ruiru had genes that enable the synthesis of Acarbose, and subsequently confirmed that this was being used to manufacture Acarbose. In the two decades since 1990 Bayer has sold at least Euro 4 billion of Glucobay. A rather different example of the value of biodiversity comes from the development of the polymerase chain reaction (PCR), central to the amplification of DNA specimens for analysis, used in forensic tests for criminal investigations and in many processes absolutely central to the biotechnology industry. The PCR technique, which takes a minute sample of DNA and multiplies it manyfold so that there is enough to conduct extensive chemical tests, requires an enzyme resistant to high temperatures. A bacterium *Thermus aquaticus* containing such an enzyme was discovered in the Lower Geyser Basin of Yellowstone National Park, and has since been found in similar habitats around the world. The enzyme derived from it is now central to the rapidly growing biotechnology industry. It is not much of an exaggeration to say that the biotechnology industry could not have taken off without an obscure bacterium found only in a few hot springs. Several research groups have developed models of the value of biodiversity in bioprospecting, including Simpson et al. (1996), Rausser & Small (2000), and Costello & Ward (2006). Each of these gives very different distributions of the value of biodiversity conservation for bioprospecting.

In economic terms, genetic variation is a resource, something we can work with and develop, because it provides a pool of within-species differences on which we can draw when seeking to develop new varieties better adapted to particular places or tasks. Genetic variation is in fact a form of natural capital, allowing us to develop new varieties with valuable properties. Developing crop varieties resistant to disease is an insurance application for genetic diversity within a species. Different examples of the same species have different degrees of susceptibility to any particular disease, allowing us to breed varieties resistant to diseases or to conditions such as drought or heat. If a farmer plants a single crop variety and a disease to which it is susceptible strikes, the entire crop is destroyed. If instead he plants varieties differing in their disease susceptibility, there is some insurance against complete crop loss. The Irish potato famine of the nineteenth century illustrates the hazards of growing a single variety of potato, in that case *Solanum tuberosum*, vulnerable to a variety of potato blight then rampant in Europe.

Genetic variation within species has historically been the source of almost all agricultural progress. It took us from hunter-gatherers to farmers, and in the 20th century it allowed us to increase food production to match the increase in world population from 1 to 7 billion. But today's food crops are genetically homogeneous, with the same varieties of most major crops grown worldwide, so the within-species variation we have drawn on in the past is disappearing. Institutions such as the International Rice Research Institute (IRRI) maintain seed banks to supplement the diversity we have in the fields; the IRRI has been critical in cases like the outbreak of the grassy stunt virus, which destroyed much of the Asian rice crop and was resistant to all attempts to neutralize it. A previously noncommercial variety of rice resistant to the grassy stunt virus—extinct in the wild—was found at the IRRI and crossbred with commercial varieties. This prevented further drastic crop losses, and showed that species diversity may provide our only protection against disastrous new diseases.⁶ Brock & Xepapadeas (2003) is a pathbreaking quantitative study of the economic value of genetic diversity in plants, and models the contribution that genetic diversity can make to the yield of an optimally-managed agricultural system. It establishes a framework within which such values can be computed, but the work of extending the model beyond a simple example, establishing parameter values and producing quantitative estimates, remains to be carried out. In the meantime we have only some very approximate estimates to work with (Heal, 2016).⁷

We conclude that in both of our case studies – climate change and biodiversity loss – we face deep uncertainties that are not readily described by unique, scientifically rigorous, objective, probability distributions. For some key aspects of climate change - such as the equilibrium climate sensitivity - we have a range of PDFs to work with, but for others we have little more than expert's guesses. And for biodiversity loss, we have a range of analytical estimates for the value of biodiversity in bioprospecting, but seldom (if ever)

⁶If the population of a species is reduced, then the genetic variation within it is lessened too. Smaller populations typically have less variation, less potential for innovative new varieties and more risk of inbreeding. So even the partial loss of a population, well short of extinction, can have an economic cost.

⁷Tilman et al. (2001) demonstrated the connection between biodiversity and ecosystem productivity: he planted similar plots of land with different varieties of grassland plants—some with many species, some with fewer. Each plot was planted with the same mix year after year and each year the experimenters noted the amount of available nutrients that the plants took up and the amount of biomass grown. Biomass, the total dry weight of the plants, is also a measure of the amount of carbon from the atmosphere that is photosynthesized into carbohydrate. Over about twenty years, the more species-diverse plots performed 270% better than the less-diverse ones. Furthermore, the plots that were more diverse were also more robust in the face of weather fluctuations. Subsequent research has confirmed these findings and the centrality of biodiversity to the functioning of natural systems. As with the previous case, we have no systemic models of the value of diversity in this context.

any objective likelihood information. For the agricultural value of biodiversity we have an analytical framework but little in the way of empirical estimates. In order to make decisions based on such information, we arguably need an uncertainty calculus that is suited to incomplete and qualitative information about the likely consequences of our actions.

2 Alternatives to Expected Utility

In this section we develop formal alternatives to the standard expected utility approach to decision under uncertainty. Before delving into mathematical details however, we discuss intuitive conceptual arguments that are often used to motivate a departure from expected utility theory when decision makers are faced with deep uncertainty.

2.1 Probabilities and confidence

Probabilities may either denote relative frequencies, as in a game of chance, or subjective degrees of belief. In the context of climate change, biodiversity, and many environmental applications, the interpretation generally has to be the latter, as we have argued above. One of the main issues we want to discuss is whether it is reasonable to require that decision makers' beliefs about the consequences of environmental policies with highly uncertain consequences always be represented by *unique* PDFs.

Formalization of choice under uncertainty in economics began with the work of von Neumann & Morgenstern (1944), who assumed that uncertainty can be characterized by an objectively known set of probabilities that is agreed by all observers. Their analysis thus builds on the relative frequency interpretation of probability, showing that if decision makers' preferences over lotteries obey certain axioms they will act as if maximizing their expected utility. In later work that has since become the mainstream approach (Savage, 1954) probabilities are not assumed *a priori*, but derived from a primitive set of axioms on preferences over 'acts', i.e. maps between states and outcomes. Savage proves a representation theorem implying that anyone whose behavior satisfies these axioms must behave as if she maximizes some subjective expected utility functional. We repeat: the existence of a complete subjective probability distribution follows from Savage's axioms, which do not presuppose the existence of probabilities. Proponents of the subjective expected utility paradigm identify rational choice with consistency with the Savage axioms. The crucial axiom in Savage's approach is the so-called "sure thing principle," which we discuss in

detail below. This, plus the other axioms, necessarily implies that agents must behave as if they have unique PDFs over all possible outcomes and can assign probabilities to any events, no matter how little information they have to draw on.

Despite the appeal of Savage's approach, scholars have always questioned whether it is reasonable to adopt a definition of rationality that requires rational agents to be able to assign a unique probability to absolutely any event. As Gilboa et al. (2009) [GPS] argue,

“Will the US president six years hence be a Democrat? The Bayesian approach requires that we be able to quantify this uncertainty by a single number; we should be able to state that our subjective belief for this event is, say, 62.4% or 53.7%. Many people feel that they do not have sufficient information to come up with such an accurate probability estimate. Moreover, some people feel that it is more rational not to assign a probabilistic estimate for such an event than to assign one. Choosing one probability number in the interval $[0,1]$ would be akin to pretending that we know something that we don't.”

They go on to argue that

“The Bayesian approach is lacking because it is not rich enough to describe one's degree of confidence in one's assessments. For any probability question it requires a single probability number as an answer, excluding the possibility of replies such as “I don't know” or “I'm not so sure”. A paradigm of rational belief should allow a distinction between assessments that are well-founded and those that are arbitrary.”

If we accept this critique, it follows that we must drop one or more of Savage's axioms when defining ‘rational’ choice. One possibility, of course, is to drop the completeness axiom and accept that preferences over uncertain prospects may be incomplete, with decision-makers (DMs) having no preferences over some sets of alternatives. Bewley (1986) adopts this approach, replacing completeness with an inertia assumption and accepting that DMs may pronounce some pairs of alternatives to be “incomparable”. His approach, like those based on ambiguity theory that we will develop further below, implies that DMs work with multiple probability distributions and prefer one alternative to another if and only if it gives a greater expected utility for all probability distributions in some set of distributions (see also Galaabaatar & Karni (2013)). An alternative to dropping the completeness axiom is to drop Savage's “sure thing principle,” and this is the approach we investigate in detail below.

2.2 Formal Development

All of the decision-making frameworks we shall talk about are developed axiomatically. Rules of behavior that are believed to be compelling desiderata of rational choice are posited, and a mathematical representation of behaviour that is consistent with these rules is deduced. In all cases there are some essential axioms whose role is technical, and other axioms that embody the intuitive essence of the approach. For example, all frameworks have a largely technical axiom requiring preferences to be continuous in some sense. Also most of them have an axiom implying that preferences are non-trivial, i.e. the decision-maker does not rank all alternatives equally. Clearly we need such an assumption for the problem to be interesting. In the review of alternative decision theories that follows we will focus on the axioms that embody the intuitive essence of each approach, and neglect the technical axioms that are needed for the mathematics, but which play a secondary role in giving the framework its distinguishing characteristics.

The von Neumann Morgenstern approach to decision-making under uncertainty is developed in the context of preferences over **lotteries**, a lottery being a list of outcomes, and a probability of occurrence for each item on the list:

Definition 1. A simple lottery l is a list of N exclusive and exhaustive outcomes $1, \dots, N$ with associated probabilities (p_1, p_2, \dots, p_N) , $\sum_n p_n = 1, p_n \in [0, 1]$, where p_n is the probability of outcome n occurring.

Probabilities are assumed to be objectively given and can be interpreted as the relative frequencies of outcomes in repeated experiments. It is assumed that preferences over simple or compound lotteries (lotteries over lotteries) depend only on the outcomes and their probabilities, and not in any way on the process used to arrive at these outcomes and probabilities. So we take the set of alternatives to be the set of all simple lotteries L over the set of outcomes. If there are three alternatives, the set of lotteries is just the set of numbers in R^3 that are non-negative and sum to one: this is the triangle joining the points $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$. Agents are assumed to have a preference ordering \succeq over L . This ordering is, as we indicated above, assumed to be continuous.⁸ The key assumption in developing this framework is an independence assumption:

⁸The preference relation \succeq on the space of lotteries L is continuous if for any $l, l', l'' \in L$ the sets

$$\{a \in [0, 1] : al + (1 - a)l' \succeq l''\} \subset [0, 1]$$

and

$$\{a \in [0, 1] : l'' \succeq al + (1 - a)l'\} \subset [0, 1]$$

Definition 2. The preference relation \succeq on the space of simple lotteries L satisfies the independence axiom if for all $l, l', l'' \in L$ and $a \in (0, 1)$ we have $l \succeq l' \iff al + (1 - a)l'' \succeq al' + (1 - a)l''$.

In words, if we mix two lotteries in the same proportions with a third one, preferences between the two resulting mixtures should be the same as preferences between the two original lotteries.

This axiom has a sharp geometric interpretation. As we noted above, with three alternatives a simple lottery can be represented by a point in R^3 in the triangle $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$. More generally in the case of N alternatives it is a point in the $N - 1$ -dimensional simplex $\Delta = \{p \in R_+^N : \sum_n p_n = 1\}$. The independence axiom implies that preferences over lotteries can be represented by parallel straight lines (planes, hyperplanes depending on the dimension) on the simplex. This means that preferences are linear in the probabilities, and given this we can fairly readily prove the following:

Proposition 1. *[von Neumann Morgenstern Expected utility theorem.] Suppose that the preference relation \succeq on L satisfies the continuity and independence properties. Then \succeq admits a representation in the expected utility form, that is we can assign numbers u_n to each outcome $1, \dots, N$ in such a manner that for any two lotteries $l = (p_1, \dots, p_N)$, $l' = (p'_1, \dots, p'_N)$ we have*

$$l \succeq l' \iff \sum_n u_n p_n \geq \sum_n u_n p'_n.$$

We now contrast this result with the Savage approach to decision under uncertainty. Unlike von Neumann and Morgenstern, Savage does not assume an objectively given set of probabilities over outcomes. In his approach probabilities are subjective, and are a consequence of choice behavior that is consistent with primitive behavioral axioms, in a manner similar to the approach developed by De Finetti (1937). The Savage result thus works where von Neumann-Morgenstern's doesn't, i.e. where no objective probabilities are available. This is clearly a very important advance, particularly in the contexts that we have discussed above.

Primitive concepts in Savage's framework are **states** and **outcomes**. The set of states S is an exhaustive list of all scenarios s that might unfold. Knowing which state occurs resolves all uncertainty. An **event** is any subset $A \subset S$. The set of outcomes is X ,

are closed. So the set of combinations of l, l' that are at least as good as l'' is a closed set, as is the set of combinations that are no better than l'' . As in the deterministic case, this requirement rules out lexicographic preferences where the agent places all emphasis on the probability of one particular outcome - for example on the risk of death being zero.

with typical member $x \in X$. An outcome specifies everything that affects the chooser's well-being.

The objects of choice are acts, which are functions from states to outcomes, and are denoted by $f \in F, f : S \rightarrow X$. The state s that will be realized after an act f has been selected is uncertain, but the agent does know that if the state s occurs then the outcome will be $f(s)$ if she chooses act f .

Acts whose payoffs do not depend on the state of the world s are constant functions in F . We will use the notation $x \in F$ to indicate the constant function in F whose outcome is always equal to $x \in X$. Suppose f, g are two acts and A is an event: then we define a new act by

$$f_A^g(s) = \begin{cases} g(s), & s \in A, \\ f(s), & s \in A^c \end{cases}$$

where A^c denotes the complement of the set A . Intuitively this composite act coincides with f everywhere except on A , where it coincides with g . Within this framework Savage's two key axioms are the following:

Axiom P2 (Sure Thing Principle) Preferences between two acts f, g depend only on events where the values of f and g differ:

Let A be an event, and suppose $f = g$ on A^c , so that f and g differ only on A . Now consider two new acts f', g' that are equal to f, g respectively on A , and equal to one another on A^c , i.e.

$$\begin{aligned} f'(s) &= f(s), & g'(s) &= g(s) & \text{for } s \in A \\ f(s) &= g(s), & f'(s) &= g'(s) & \text{for } s \in A^c \end{aligned}$$

Then P2 requires:

$$f \succeq g \Leftrightarrow f' \succeq g'.$$

An intuitive interpretation of P2 is provided by GPS. Consider the following four bets:

1. If horse A wins you get a trip to Paris, and otherwise you get a trip to Rome
2. If horse A wins you get a trip to London and otherwise a trip to Rome
3. If horse A wins you get a trip to Paris and otherwise a trip to Los Angeles
4. If horse A wins you get a trip to London and otherwise a trip to Los Angeles

Clearly 1 and 2 are the same if A loses, as are 3 and 4. Generally one's choices between these options will depend on the value one assigns to alternative outcomes and beliefs about their likelihood of occurring. Presumably however, the chance of A winning is the same in each case, so the choice between 1 and 2 depends on one's preferences between Paris and London. The same is true for 3 and 4. Axiom P2 thus requires consistency between preferences over the pairs of bets 1, 2 and 3, 4: $1 \succeq 2 \Leftrightarrow 3 \succeq 4$. If two acts are equal on a given event, it does not matter what they are equal to. So it doesn't matter if when the horse loses you get Rome or LA. It is hard to argue with this axiom viewed in isolation. Nevertheless, when taken together with the rest of Savage's axioms, it requires agents to be able to assign subjective probabilities to any events however unlikely or unusual, as we noted in the introduction.

The second axiom we highlight gives further insight into how Savage's primitive axioms for choice between acts gives rise to a probabilistic preference representation:

Axiom P4. For every $A, B \subset S$ and every $x, y, z, w \in X$ with $x \succ y$, $z \succ w$,

$$y_A^x \succeq y_B^x \Leftrightarrow w_A^z \succeq w_B^z$$

Here is an interpretation of this axiom. We have four alternatives, x, y, z, w with $x \succ y$ and $z \succ w$. And we have two subsets of the set of all states, two events, A and B . We look at y_A^x , which is y modified to be x on the event A . x is preferred to y so this is an improvement on y . The axiom says that we prefer y_A^x to y_B^x (that is, prefer y upgraded on A to x to y upgraded on B to x) if and only if we prefer w similarly upgraded to z . Why would we prefer to upgrade y and w on A rather than on B ? Presumably because we think A is more likely than B . So Axiom P4 encodes some notion of the relative 'likelihood' of the two events A and B . This axiom clearly plays an important role in introducing subjective probabilities into the Savage formalism.

There are other essentially technical assumptions, including an axiom roughly analogous to continuity, but it is the two above that give Savage's theory its main characteristics. With this set of assumptions Savage proves

Proposition 2. *[Savage] Assume that X is finite. Then \succeq satisfies the Savage axioms if and only if there exists a probability measure μ on states S and a non-constant utility function $u : X \rightarrow R$ such that for every $f, g \in F$,*

$$f \succeq g \Leftrightarrow \int_S u(f(s)) d\mu(s) \geq \int_S u(g(s)) d\mu(s)$$

Furthermore μ is unique and u is unique up to positive linear transformations.

Savage’s theory has been the economists’ workhorse since the 1960s, and is the default approach to choice under uncertainty in environmental economics and in economics generally. But as we noted above, there are reasons to think that it is limited in important ways, and in particular has a limited capacity to reflect some important aspects of the information available to us on environmental problems. So we next turn to the other candidates for addressing these issues, referred to generally as “multiple priors” models. They are models of decision-making in which the decision-maker recognizes the ambiguity inherent in the problem she wishes to solve and instead of working with a single PDF, she works with all of the PDFs consistent with the information available to her.

A prominent approach is that of Gilboa & Schmeidler (1989), who work within the same conceptual framework as Savage but of course use a different set of axioms. In particular, they drop Savage’s second axiom, the sure thing principle.⁹ Their axiom set contains several technical axioms as usual, including a continuity axiom. In addition it includes an independence axiom similar to that used by von Neumann and Morgenstern, and an explicit invocation of uncertainty aversion:

Axiom GS5. Independence: For every act $f, g \in F, \forall$ constant $h \in F, \forall \alpha \in (0, 1)$,

$$f \succeq g \Leftrightarrow \alpha f + (1 - \alpha) h \succeq \alpha g + (1 - \alpha) h$$

Axiom GS6. Uncertainty Aversion. For every $f, g \in F, \forall \alpha \in (0, 1)$

$$f \sim g \Rightarrow \alpha f + (1 - \alpha) g \succeq f$$

The axiom GS5 is similar to but weaker than the von Neumann Morgenstern independence axiom, because the axiom applies only to *constant* third alternatives h . The uncertainty aversion axiom GS6 is unique; it implies that if we are indifferent between f and g then we regard any strictly convex combination of the two as at least as good as either. This sounds a lot like risk aversion – a mixture of two acts, which hedges some of the uncertainty associated with each act, is weakly preferred to each act itself. With this set of assumptions, and other technical conditions, GS prove the following result:

⁹Amartya Sen, John Rawls, and others have argued that axioms can only be fully evaluated when one knows their implications – they are reasonable only if they have reasonable implications. Arguably, alternatives to the Savage axioms are attractive in situations where imprecise probabilities seem to occur naturally because they do not require us to behave as if we know things we don’t. We discuss this more fully below.

Proposition 3. [Gilboa-Schmeidler] *The preference \succeq satisfies the above axioms if and only if there exists a closed convex set of probabilities Δ and a non-constant function $u : X \rightarrow R$ such that for every $f, g \in F$*

$$f \succeq g \Leftrightarrow \min_{p \in \Delta} \int_S u(f(s)) dp(s) \geq \min_{p \in \Delta} \int_S u(g(s)) dp(s)$$

C is unique and u is unique up to a positive linear transformation.

This result does not require Savage’s “sure thing principle”, nor does it need his axiom P4 about probabilities. What this result says is the following: for each act, look at the probability distribution in the set Δ that yields the lowest expected utility for that act, and evaluate the act using these probabilities. Then choose the act with the largest evaluation, i.e. the act with the largest minimum expected utility across the set of probabilities Δ . What is important from our perspective is that the idea of multiple priors emerges naturally from this framework. Just as primitive axioms over acts yield subjective probabilities in the Savage framework, so the GS axioms require beliefs to be described by a *set* of probabilities. In addition, the GS axioms require decision-makers to respond to the multiplicity of priors in a specific way, i.e. by focussing on the worst prior for each act. This result is remarkable in that it uses the same basic machinery as the Savage result, and a simple (indeed simpler) set of axioms, to arrive at a radically different representation of agents’ beliefs that does not collapse all uncertainty into a single probability distribution.

Obviously one can conceive of many other ways of reacting to a multiplicity of priors. A more flexible alternative to GS that is widely used is that proposed by Klibanoff et al. (2005) (henceforth KMM). The KMM result is in some senses less ‘primitive’ than the Savage and GS results, although it too can be seen as a consequence of a set of axioms on preferences over Savage acts, albeit over an enriched state space that includes lotteries over states. The idea is relatively simple. KMM assume a set Δ of ‘objective’ von Neuman-Morgenstern lotteries p over the state space. Preferences over acts, conditional on a lottery p being realized, are represented by a vNM Expected Utility functional, with utility function U . They then assume a set of ‘second order acts’, which map *lotteries* (not states) into consequences. Preferences over second order acts are also assumed to have a subjective expected utility representation, with some subjective probability $\mu(p)$ on lottery p , and some utility function Φ over consequences. Now given an objective PDF p , an act f induces a lottery p_f over the set of consequences, whose value can be captured by a certainty equivalent $c_f(p)$. These certainty equivalents take values in the space of outcomes,

and are an increasing function of $E_p U(f(s))$. Since there are many possible priors p , each act f generates a lottery over the outcomes $c_f(p)$. Preferences over such ‘second order’ acts have a subjective expected utility representation with utility Φ by assumption, so we must be able to represent preferences over acts by subjective expected utilities over certainty equivalents. This is effectively the content of their result:

Proposition 4. [*Klibanoff, Marinacci, and Mukherji*]: *There exists a set of PDFs Δ with generic element p , a utility function $U : S \rightarrow R$, a ‘second order’ utility function $\Phi : U \rightarrow R$, and second order subjective probabilities $\mu(p)$, such that for all $f, g \in F$,*

$$f \succeq g \iff \int_{p \in \Delta} \Phi \left(\int_{s \in S} U(f(s))p(s)ds \right) \mu(p)dp \geq \int_{p \in \Delta} \Phi \left(\int_{s \in S} U(g(s))p(s)ds \right) \mu(p)dp$$

The subjective probability $\mu(p)$ on PDFs p is unique, and Φ is unique up to positive affine transformations.

Just as in conventional expected utility theory, if Φ is concave, this implies that the decision maker is averse to the spread in expected utilities across the set of PDFs Δ , i.e. $\Phi'' < 0$ implies

$$\begin{aligned} \int_{p \in \Delta} \Phi \left(\int_{s \in S} U(f(s))p(s)ds \right) \mu(p)dp &\leq \Phi \left(\int_{p \in \Delta} \left(\int_{s \in S} U(f(s))p(s)ds \right) \mu(p)dp \right) \\ &= \Phi \left(\int_{s \in S} U(f(s)) \left(\int_{p \in \Delta} \mu(p)p(s)dp \right) ds \right) \end{aligned}$$

The expectation over s on the right hand side of this inequality is just a standard compound lottery, so that the effective probability on state s is $\int_{p \in \Delta} \mu(p)p(s)dp$. Thus, if Φ is concave, the decision maker would always prefer it if his ‘subjective’ uncertainty $\mu(p)$ could be made ‘objective’ – he is averse to subjective uncertainty over conditionally objective PDFs in the set of priors Δ . In more common terminology, he is ambiguity averse.

Speaking loosely and intuitively, the KMM approach involves applying a von Neumann Morgenstern framework twice, once to each of the individual models or distributions p and then secondly over the set of expected utilities that emerges from these distributions. Concavity of the U functions, the von Neumann Morgenstern utilities, would reflect the normal aversion to risk, and concavity of the Φ function would reflect aversion to not knowing the expected utility, which is only known conditional on a particular distribution. The KMM approach converges to the GS approach as the second order utility function Φ becomes more and more concave. A KMM decision-maker who is infinitely averse to

uncertainty about expected utilities is a GS decision-maker.

We have now reviewed the traditional expected utility model and modern alternatives to it that allow decision-makers to express lack of confidence in their probabilistic beliefs. In particular, these new tools allow a distinction to be drawn between objective probabilities and subjective judgements, and are naturally suited to dealing with decision problems in which several plausible probabilistic descriptions of the world present themselves to us. Arguably, these features of the new decision models make them well suited to environmental problems such as climate change and biodiversity, as we seldom have objective probabilistic knowledge of the consequences of policy choices in these applications. But is it ‘rational’ to dispense with the appealing axioms that underpin Savage’s subjective expected utility theory? Perhaps violation of the sure thing principle (for example) is too high a price to pay in order to have a decision framework that admits a distinction between objective probabilities and arbitrary subjective judgements? Next we digress slightly and try to put this movement away from the expected utility framework into historical context. The key issue we must address is the extent to which a given axiom system (in this case the Savage axioms) can be thought of as providing a universal and unassailable definition of ‘rationality’. While there can be no ultimate answer to this question – it is in large part a matter of definition – the history of science does suggest that placing one axiom system on a pedestal can sometimes be a barrier to progress.

2.3 Is ambiguity aversion rational?

Not everyone agrees on whether the models of decision under ambiguity developed above should be used in normative applications. As we have emphasized, the EU approach requires that agents be able to assign a precise probability to any event. In many areas of environmental economics this seems to do violence to the facts. For example, we don’t know the equilibrium climate sensitivity, and we have multiple PDFs over its possible values, which we cannot sensibly condense into a single PDF. In other areas, such as the cost of biodiversity loss, our knowledge is even less precise. One response to this is to adopt a decision framework which takes multiple probabilistic descriptions of reality as its raw material, and which allows us to distinguish between subjective judgements and conditionally objective probabilities (i.e. probabilities that arise from constraining models with data, *assuming* that the model is a ‘correct’ description of the world). The approaches we have discussed above achieve this at the cost of dropping Savage’s “sure thing” principle. While some eminent decision theorists have argued that EU theory is not

a universally applicable standard of rational choice (Binmore, 2009; Gilboa et al., 2008), others strenuously adhere to the Savage framework (Sims, 2001; Al-Najjar & Weinstein, 2009; Al-Najjar, 2015). Our view of this debate is that, much like the debate about the appropriate value of welfare parameters, it is something that reasonable people can reasonably disagree about. Nevertheless, it is important to be clear what the arguments for and against the normative status of ambiguity aversion are. We now discuss some of these.

It is sometimes asserted that models of decision under ambiguity are merely descriptive theories of people's behaviour in the face of ambiguity, as was memorably described by Ellsberg (1961). Briefly, in one version of Ellsberg's choice experiment a DM is asked whether she would prefer betting on a red or black ball being drawn from urn 1, which contains 50 red and 50 black balls, or from urn 2 which contains 100 balls, each of which is red or black, but whose proportions are not known. Most people are indifferent between betting on red or black in both Urn 1 and Urn 2. However, when asked whether they would prefer to bet on a red ball being drawn from urn 1, or a red ball being drawn from urn 2, many people prefer to bet on urn 1. The important point here is that according to the Savage approach, being indifferent between red and black in Urns 1 and 2 reveals that the DM must have a subjective probability that red and black are equally likely in *both* urns. Since the subjective probabilities and payoffs are the same for bets on both urns, from the perspective of the Savage axioms these two urns are indistinguishable, and it is irrational to prefer betting on a red ball being drawn from Urn 1 to a red ball being drawn from Urn 2. And yet, many people do prefer betting on the Urn with known composition to that with unknown composition, i.e. they are ambiguity averse.

An ability to explain behaviour in Ellsberg's choice experiments is certainly an attractive feature of the decision theoretic tools developed above, but this is clearly not a reason to believe that these tools have *normative* legitimacy. A compelling argument against the normative validity of the behaviour Ellsberg observes was provided by Raiffa (1961). He asks us to choose a colour in Urn 2 by flipping a fair coin, say choose red if the coin lands on heads. Clearly our chance of matching the ball drawn from Urn 2 using this strategy is now an 'objective' 50% – the coin flip is either 'right' or 'wrong'. Raiffa thus uses a randomization device (the coin) to turn an ambiguous situation (unknown probabilities) into a risky one (known probabilities). According to Raiffa, this shows that preferring to bet on Urn 1 is irrational. We do not deny the force of this argument in the context of Ellsberg style choice problems. It is however important to realize that the argument turns

on the high degree of symmetry in the stylized Ellsberg examples. This symmetry gives rise to a natural prior on the space of possible probability distributions that describe the composition of Urn 2, i.e. that all distributions are equally likely. Raiffa's coin is simply a device for bringing this symmetry to the fore. But, as David Schmeidler (quoted in Gilboa (2009)) says: 'Real life is not about balls and urns.' What is the natural prior for the probability of war in the South China Sea in 2050? Of course, there is none.

The key normative argument for deviating from subjective expected utility theory is that it forces us to treat situations that are manifestly different from an informational perspective as if they were identical. Savage's theory admits no distinction between arbitrary subjective judgements and objective probabilistic knowledge (e.g. known relative frequencies of events). In practice it seems desirable for policy analysts to be able to make use of multiple conflicting probabilistic models of how the world works, while acknowledging that their policy recommendations rely on subjective judgements about how to weight alternative models. These judgements clearly do not have the same character as objective probabilities. The models of decision under ambiguity described above provide such tools.

A more technical set of objections to giving up the Savage axioms concerns the extensions of ambiguity averse decision rules to dynamic contexts (e.g. Al-Najjar & Weinstein, 2009). Extending static decision frameworks to dynamic contexts requires us to specify how information sets are updated, and how preferences at different nodes of decision trees are related. This aspect of intertemporal decision-making is pinned down by invoking additional axioms: dynamic consistency (conditional plans, once made, are not revised due to the mere passage of time), and consequentialism (preferences depend only on nodes of the decision tree that are reachable from the current node). There is a tight connect between the sure thing principle and dynamic consistency (Epstein & Le Breton, 1993), implying that ambiguity averse decision rules cannot be made to respect dynamic consistency on a universal domain of decision trees. However, dynamic consistency may be respected for any particular decision tree. Epstein & Schneider (2003) and Mukerji (2009) pursue this approach in the GS and KMM models respectively. See Mukerji (2009); Siniscalchi (2009); Machina & Siniscalchi (2014) for further details and discussion.

It is very likely that some readers may feel that giving up the elegant Savage axioms, in particular the sure thing principle, is too high a price to pay for the additional epistemic nuance that the ambiguity models provide. That is a legitimate normative perspective, although ultimately not one we agree with. As we have argued above, axioms cannot be evaluated in isolation, but must be judged by their implications when combined with other

axioms. If an axiom system leads to conclusions we are not prepared to accept, it must be revised – a process that the philosopher John Rawls refers to as *reflective equilibrium* (Rawls, 1971). While the appeal of the Savage axioms is undeniable, their implications are, to our eyes, unpalatable in situations in which information is imprecise or incomplete. In our view, the Savage axioms are best seen as a recipe for translating arbitrary choice problems into choices over bets that look like coin tosses or rolls of a die. This is an elegant and powerful trick, but the sleight of hand comes at a high price. In executing this recipe we are required to behave as if we knew things (i.e. the relative frequencies of events) that we often do not.

For some commentators the Savage (or worse, von Neumann-Morgenstern) axioms provide a *definition* of rational behaviour in the face of uncertainty in all situations. We would argue, however, that concepts of rationality can (and should) evolve, as our understanding of the implications of axiom systems deepens. The history of science demonstrates that adopting an inflexible view of the nature of reality based on preconceived notions of mathematical elegance can be a barrier to progress. The same may well be true of normative theories of rationality. To illustrate, consider the following example:

It is an elementary fact of geometry, known to schoolchildren the world over, that the sum of the angles in a triangle is equal to 180 degrees. This fact is a direct consequence of the five axioms of Euclid's Elements:

- 1) Any two points can be connected by a straight line.
- 2) Any line segment can be extended indefinitely in a straight line.
- 3) Given any straight line segment, a circle can be drawn having the segment as radius and one endpoint as its center.
- 4) All right angles are equal to one another.
- 5) Given any straight line and a point not on it, there exists exactly one straight line which passes through that point and never intersects the first line.

These axioms are so self-evident that doubting them was seen as a symptom of madness by learned minds for almost 2000 years after Euclid. The Elements became the gold standard of rational thought throughout that period. Nevertheless, as the centuries rolled by it was occasionally suggested that the fifth axiom, the so-called 'parallel postulate', seemed less satisfying than the others. It was more complex, and some argued that it should be possible to derive it as a consequence of the other four more primitive axioms. Several attempts were made, but no one could find a proof. Then, in the 18th century, a Jesuit priest named Giovanni Girolamo Saccheri adopted a novel approach. His strategy

was simple: assume that the parallel postulate does not hold, and derive a contradiction. It was known that the parallel postulate is equivalent to the fact that the angles of a triangle sum to 180 degrees, so Saccheri considered two cases: either the angle sum of a triangle is greater, or less than 180 degrees.

Saccheri set about deriving the consequences of these assumptions, which he was sure would lead to a self-evident contradiction. He first considered the case where the angle sum exceeds 180 degrees, and quickly derived a contradiction with the second axiom – the assumption implied that straight lines would have to be finite. He saw this as a victory, as the second axiom was seen as self-evidently true. Next he considered the case where the angle sum of a triangle is less than 180 degrees. He derived proposition after proposition based on this flawed initial assumption, but could not find a contradiction. Finally, exasperated, he asserted that ‘the hypothesis of the acute angle is absolutely false; because it is repugnant to the nature of straight lines’.

Today we recognize Saccheri’s many propositions not as a failed attempt to vindicate the parallel postulate and Euclid’s Elements, but as an inadvertent first step towards a new, non-Euclidean, geometry. The sum of the angles of a triangle can be less, or indeed greater, than 180 degrees. In so-called elliptical geometries, the angle sum is greater than 180, all parallel lines eventually meet, and just as Saccheri discovered, straight lines cannot be extended indefinitely. In hyperbolic geometries, the angle sum is less than 180, and there are many lines that pass through the same point and are parallel to a given line. These non-Euclidean geometries are now a part of the mathematical mainstream, and have a myriad of applications in diverse areas of science. Most famously, non-Euclidean geometry lies at the heart of Einstein’s General Theory of Relativity.




Figure 2: Illustration of non-Euclidean geometries

There are two lessons we wish to draw from this remarkable episode in the history of science, one intellectual, the other, for want of a better word, sociological.

The intellectual lesson is this: although relaxing the parallel postulate leads to a new and deeper understanding of geometry, the old Euclidean geometry is still a vital part of our analytical toolkit. It turns out that all smooth geometries are approximately Euclidean on small enough spatial scales. If we zoom in to the surface of a sphere or a hyperboloid, the curvature of the surface becomes negligible, and we reclaim the old Euclidean geometry to a very good approximation. The message of this fact is that even if one axiom system is superseded by another, this does not render the original axioms useless. For some problems accounting for the non-Euclidean effects of surface curvature is crucial (e.g. in applications of General Relativity to cosmology, astrophysics, and the technologies underpinning Global Positioning Systems), while for others, we can safely rely on our tried and trusted Euclidean tools. Suggesting that ambiguity aversion may be a normatively legitimate stance in no way invalidates the expected utility approach in applications where objective probabilities are available, or as a parsimonious modeling tool in descriptive applications.

The sociological lesson is this: history teaches us that viewing a single axiom system as the universal standard of truth or rationality can be a barrier to progress. Had the mathematicians who followed Euclid been less dogmatic in their adherence to the Euclidean axioms, non-Euclidean geometry, and its many applications, might have been discovered hundreds of years earlier, perhaps paving the way for further advances in other areas. While Euclid's axioms concern the nature of consistent geometrical spaces, the parallels to our discussion of consistent theories of rational choice are obvious.

3 Application to Environmental Policy Choices

The next step is to show how these new ideas can be applied to the analysis of environmental choices. For this we develop a simple static model of the optimal choice of an environmental policy, and show how the multiple priors framework can be applied. We then briefly review more sophisticated environmental applications of the decision frameworks discussed above from the literature.

3.1 A simple analytical model

We focus on a class of problems in which there are several alternative models of the relationship between a policy choice that the DM has to make, and the outcome that is of

value to the DM. The DM does not know which of these models is correct. She can however assign subjective second-order probabilities to each model being correct.

There are many possible interpretations of this framework: in keeping with our earlier discussions we could, for example, be thinking of an integrated assessment model of climate policy choice where the equilibrium climate sensitivity is unknown and there are several alternative distributions over its possible values, each corresponding to a different underlying climate model. A different interpretation could be that we are concerned with modeling the allocation of resources to biodiversity conservation. The consequences of biodiversity loss are unknown and there are again multiple models of how biodiversity affects human welfare.

We assume that each alternative model can be identified with the value of a key parameter S of the underlying system - this could be the equilibrium climate sensitivity or the cost of biodiversity loss. A model is a distribution $p_m(S)$ over this parameter, indexed by m , and the set of such distributions of S is Δ . The level of the chosen policy variable – think for example of greenhouse gas abatement or tropical forest conservation – is a , and the utility level associated with this level depends on the parameter S and is denoted $U(a, S)$.

For each model or distribution $p_m(S)$, the expected utility associated with policy level a is given by $EU_m(a) = \int U(a, S) p_m(S) dS$. The correct distribution is of course not known. We assume that the policy maker has KMM preferences. Thus, there exists a second order probability distribution over the set of distributions Δ , with the weight on distribution $p_m(S)$ given by μ_m . So the DM’s overall objective is to choose policy a so as to

$$\max_a \sum_m \Phi [EU_m(a)] \mu_m \tag{1}$$

where Φ is a concave function. We can think of the probabilities $p_m(S)$ within each expected utility calculation as scientific probabilities from some model of a physical or economic process, and the μ_m as subjective judgments about how to weight alternative models.

The first order conditions for (1) are

$$\sum_m \hat{\mu}_m(a^*) \left. \frac{dEU_m}{da} \right|_{a=a^*} = 0 \tag{2}$$

where a^* is the optimal abatement level and the $\hat{\mu}_m(a)$ are “ambiguity-adjusted” second

order probabilities:

$$\hat{\mu}_m(a^*) = \frac{\Phi'(EU_m(a^*)) \mu_m}{\sum_n \Phi'(EU_n(a^*)) \mu_n} \quad (3)$$

In the case of ambiguity neutrality (i.e. $\Phi'(x)$ is a constant), the first order condition is

$$\left. \frac{d(\sum_m \mu_m EU_m)}{da} \right|_{a=a^*} = 0$$

i.e. uncertainty over which ‘model’ is correct is treated as a standard compound lottery, which reduces the set of priors Δ to a single prior given by $\sum_m \mu_m p_m(S)$. When $\Phi'(x)$ is concave however (i.e. Φ' is decreasing), (3) shows that the decision-maker’s first order condition effectively places more weight on models that have low expected utilities, through the ambiguity adjusted second order weights. Indeed, all else equal, an increase in the concavity of Φ always places more weight on models with low expected utility models in the first order conditions. If $\Phi'(x)$ is very steeply decreasing (i.e. Φ is strongly concave), then priors that predict really poor outcomes will receive most of the weight in the first order condition. In the limit of very large concavity, only the prior with the lowest expected utility enters the first order condition, and we recover the GS model of ambiguity aversion.

It is important to realize that the ambiguity adjusted second order weights (3) that appear in the first order condition depend on the endogenous variable a . Thus, the fact that an increase in ambiguity aversion places more weight on models with low expected utilities in the first order condition does not imply that increasing ambiguity aversion has an unambiguous effect on optimal policies. Conditions that allow us to sign the effect of increased ambiguity aversion on policy choice are provided in a result by Millner et al. (2013):

Proposition 5. *Suppose that $\frac{d^2 EU_m}{da^2} < 0$ for all m , and assume that for every fixed value of a the sequences $\{EU_m(a)\}$ and $\{\frac{dEU_m}{da}\}$ are anti-comonotonic [comonotonic] in m . Then an increase in ambiguity aversion increases [decreases] the optimal level of the policy variable a .*

Figure 2 gives examples of anti-comonotonic and comonotonic sequences of expected utilities and marginal expected utilities. Policy intensity (e.g. conservation, greenhouse gas abatement) a is plotted horizontally and a set of expected utilities EU_m is plotted vertically for a set of three distinct models. Anti-comonotonicity means that for each policy value the models with high expected utility have low derivatives of expected utility with respect to policy, and vice versa. When the anti-comonotonicity condition holds an increase in

policy will reduce the spread of expected utilities across models, so that a rise in aversion to ambiguity will increase optimal abatement.




Figure 3: Comonotonicity [left] and anti-comonotonicity [right]

This simple model illustrates the basic workings of ambiguity aversion in static optimization models. Applying such models to environmental problems, in particular in more sophisticated dynamic contexts, is a topic of on going research. We next review some of the recent work in this area.

3.2 Applications in the literature

There is an emerging, but not yet extensive, literature on ambiguity in environmental economics; the idea is more firmly established in macroeconomics and finance. Most of the applications to dynamic environmental problems have focussed on some aspect of climate policy. In an early paper, Lange & Treich (2008) survey the effects of ambiguity aversion on optimal policy choices in a simple two period discrete time model motivated by climate policy. Using a simple two period model, Chambers & Melkonyan (2017) show that the subjective expected utility, maxmin expected utility and incomplete expected utility theories can lead to radically different policy recommendations. Millner et al. (2013) develop the KMM framework used in the previous section and apply it to the ranking of exogenous dynamic climate abatement policies using Nordhaus's DICE integrated assessment model, and the set of distributions over the ECS depicted in Figure 1. They find that ambiguity aversion can have a substantial positive effect on the welfare benefits of greenhouse gas abatement if damage functions are steep. Lemoine & Traeger (2016) analyzes the consequence of ambiguity aversion in a climate model in which there are tipping points (regime

changes in a complex dynamical system) whose locations are unknown. In their model ambiguity aversion increases the optimal tax on CO2 emissions, but only a little. Asano (2010) uses the Gilboa-Schmeidler approach to study the optimal timing of climate policy in a model that incorporates irreversibility, model uncertainty, and strong ambiguity aversion.

A slightly different formulation of ambiguity aversion, known as ‘multiplier preferences’, has been developed by Hansen & Sargent (2001), and is becoming increasingly popular in macroeconomics. Macroeconomists have long recognized that it is often not possible to know the ‘correct’ model of the 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 or biodiversity loss. Under the Hansen-Sargent approach, the analyst chooses a ‘best guess’ model or distribution, but is not sure that this distribution is correct. So the analyst allows for possible model misspecification and evaluates policy options according to her best guess as well as other distributions that are ‘similar’ to this guess. Distributions that are far from the best guess distribution are penalized by adding a non-negative cost term to the DMs objective function, which depends on a measure of the distance between the distribution and the best guess. Policy makers are assumed to have strong ambiguity aversion, and thus evaluate each action using the distribution that is most pessimistic for that action across the set of models in consideration, as in the Gilboa-Schmeidler framework.

The multiplier preferences approach makes sense if there is a clear preferred model but we are not sure about some of its details. However, if there are several rather different models that are all serious candidates for the ‘best’ model, then this approach is not so attractive. That is, 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 environmental issues, climate change in particular. However there is a return for this assumption – a tractable framework that allows us to identify analytical solutions to complex dynamic optimization problems, at least in the case of linear-quadratic models.

Xepapadeas (2012) and Athanassoglou & Xepapadeas (2012) apply the Hansen-Sargent multiplier preference approach to climate problems. They consider models in which output produces a pollutant 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 dynamic optimization problem permits analytical solutions. Vardas & Xepapadeas (2010) use this approach to study biodiversity management under model uncertainty.

The effects of ambiguity and ambiguity aversion on social discount rates, essential inputs to the cost-benefit analysis of marginal environmental projects, have also been studied. Traeger (2014) studies the consequences of allowing for ambiguity about consumption growth rates or policy payoffs on the choice of a social rate of discount. Gierlinger & Gollier (2017) study the related question of how ambiguity affects equilibrium interest rates in a Lucas tree model, finding a term structure that is qualitatively different from the ambiguity neutral case.

4 Conclusions

We have argued that profound uncertainty is endemic to the field of environmental policy. Arguably, traditional approaches to decision-making under uncertainty based on expected utility maximization are out of their depth in this area, as they force us to act as if we know things (i.e. unique probabilities) that we know we do not. Alternative decision frameworks based on the explicit recognition of multiple priors are naturally suited to the information we in fact possess about the consequences of environmental policy choices. These frameworks allows us to reflect our lack of confidence in our ability to discern between alternative probabilistic views of the world. The virtue of this approach is that it explicitly acknowledges the limitations of our knowledge, and allows us to exhibit preferences that are sensitive to this lack of knowledge. Applying these tools to environmental policy choice leads to new policy conclusions, which would be lost if all available information were forced into a single PDF.

References

- N. I. Al-Najjar (2015). ‘A Bayesian Framework for the Precautionary Principle’. *The Journal of Legal Studies* **44**(S2):S337–S365.
- N. I. Al-Najjar & J. Weinstein (2009). ‘The ambiguity aversion literature: a critical assessment’. *Economics and Philosophy* **25**(Special Issue 03):249–284.
- J. E. Aldy & K. Viscusi (2014). ‘Environmental Risk and Uncertainty’. In M. M. a. K.

- Viscusi (ed.), *Handbook of the Economics of Risk and Uncertainty*, vol. 1, pp. 601–649. North-Holland.
- K. J. Arrow & A. C. Fisher (1974). ‘Environmental Preservation, Uncertainty, and Irreversibility’. *The Quarterly Journal of Economics* **88**(2):312–319.
- T. Asano (2010). ‘Precautionary Principle and the Optimal Timing of Environmental Policy Under Ambiguity’. *Environmental and Resource Economics* **47**(2):173–196.
- S. Athanassoglou & A. Xepapadeas (2012). ‘Pollution control with uncertain stock dynamics: When, and how, to be precautionous’. *Journal of Environmental Economics and Management* **63**(3):304–320.
- T. Bewley (1986). ‘Knightian decision theory: Part I’. *Cowles foundation discussion paper* (807).
- K. Binmore (2009). *Rational Decisions*. Princeton University Press.
- W. A. Brock & A. Xepapadeas (2003). ‘Valuing Biodiversity from an Economic Perspective: A Unified Economic, Ecological, and Genetic Approach’. *The American Economic Review* **93**(5):1597–1614.
- T. A. Carleton & S. M. Hsiang (2016). ‘Social and economic impacts of climate’. *Science* **353**(6304):aad9837.
- R. G. Chambers & T. Melkonyan (2017). ‘Ambiguity, reasoned determination, and climate-change policy’. *Journal of Environmental Economics and Management* **81**:74–92.
- C. Costello & M. Ward (2006). ‘Search, bioprospecting and biodiversity conservation’. *Journal of Environmental Economics and Management* **52**(3):615–626.
- B. De Finetti (1937). *La prévision : ses lois logiques, ses sources subjectives*. Annales de l’Institut Henri Poincaré. Paris: Institut Henri Poincaré.
- M. Dell, et al. (2012). ‘Temperature Shocks and Economic Growth: Evidence from the Last Half Century’. *American Economic Journal: Macroeconomics* **4**(3):66–95.
- M. Dell, et al. (2014). ‘What Do We Learn from the Weather? The New Climate–Economy Literature’. *Journal of Economic Literature* **52**(3):740–798.

- A. E. Dessler (2010). ‘A Determination of the Cloud Feedback from Climate Variations over the Past Decade’. *Science* **330**(6010):1523–1527.
- D. A. Drew, et al. (2016). ‘Aspirin and colorectal cancer: the promise of precision chemoprevention’. *Nature Reviews Cancer* **16**(3):173–186.
- D. Ellsberg (1961). ‘Risk, Ambiguity, and the Savage Axioms’. *The Quarterly Journal of Economics* **75**(4):643–669.
- L. G. Epstein & M. Le Breton (1993). ‘Dynamically Consistent Beliefs Must Be Bayesian’. *Journal of Economic Theory* **61**(1):1–22.
- L. G. Epstein & M. Schneider (2003). ‘Recursive multiple-priors’. *Journal of Economic Theory* **113**(1):1–31.
- T. Galaabaatar & E. Karni (2013). ‘Subjective Expected Utility With Incomplete Preferences’. *Econometrica* **81**(1):255–284.
- J. Gierlinger & C. Gollier (2017). ‘Do interest rates decline when there is ambiguity about growth? <https://c6ac9fe8-a-62cb3a1a-s-sites.googlegroups.com/site/jgierlingerecon/GierlingerGollier2017.pdf>’. *Working Paper* .
- 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.
- M. Greenstone, et al. (2013). ‘Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation’. *Review of Environmental Economics and Policy* **7**(1):23–46.
- L. P. Hansen & T. J. Sargent (2001). ‘Robust Control and Model Uncertainty’. *The American Economic Review* **91**(2):60–66.

- G. Heal (2016). *Endangered Economies: How the Neglect of Nature Threatens Our Prosperity*. Columbia University Press.
- G. Heal & B. Kriström (2002). ‘Uncertainty and Climate Change’. *Environmental and Resource Economics* **22**(1-2):3–39.
- G. Heal & A. Millner (2014a). ‘Uncertainty and Decision Making in Climate Change Economics’. *Review of Environmental Economics and Policy* **8**(1):120–137.
- G. Heal & J. Park (2013). ‘Feeling the Heat: Temperature, Physiology & the Wealth of Nations’. Working Paper 19725, National Bureau of Economic Research.
- G. Heal & J. Park (2015). ‘Goldilocks Economics? Temperature stress and the direct impact of climate change’. *Review of Environmental Economics and Policy*, forthcoming .
- G. M. Heal & A. Millner (2014b). ‘Agreeing to disagree on climate policy’. *Proceedings of the National Academy of Sciences* **111**(10):3695–3698.
- C. Henry (1974). ‘Investment Decisions Under Uncertainty: The “Irreversibility Effect”’. *The American Economic Review* **64**(6):1006–1012.
- C. Hope (2006). ‘The marginal impact of CO₂ from PAGE2002: An integrated assessment model incorporating the IPCC’s five reasons for concern’. *Integrated Assessment* **6**(1).
- T. Houser, et al. (2015). *Economic Risks of Climate Change: An American Prospectus*. Columbia University Press.
- S. M. Hsiang, et al. (2013). ‘Quantifying the Influence of Climate on Human Conflict’. *Science* **341**(6151):1235367.
- D. Kahneman, et al. (2000). *Choices, values, and frames*.
- P. Klibanoff, et al. (2005). ‘A Smooth Model of Decision Making under Ambiguity’. *Econometrica* **73**(6):1849–1892.
- R. Knutti, et al. (2017). ‘A climate model projection weighting scheme accounting for performance and interdependence’. *Geophysical Research Letters* .
- E. Kriegler, et al. (2009). ‘Imprecise probability assessment of tipping points in the climate system’. *Proceedings of the National Academy of Sciences* **106**(13):5041–5046.

- A. Lange & N. Treich (2008). ‘Uncertainty, learning and ambiguity in economic models on climate policy: some classical results and new directions’. *Climatic Change* **89**(1):7–21.
- D. Lemoine & C. P. Traeger (2016). ‘Ambiguous tipping points’. *Journal of Economic Behavior & Organization* **132**(B):5–18.
- D. B. Lobell, et al. (2011). ‘Climate Trends and Global Crop Production Since 1980’. *Science* **333**(6042):616–620.
- W. MacAskill (2016). ‘Normative Uncertainty as a Voting Problem’. *Mind* **125**(500):967–1004.
- M. J. Machina & M. Siniscalchi (2014). ‘Ambiguity and Ambiguity Aversion’. In K. Viscusi & M. Machina (eds.), *Handbook of the Economics of Risk and Uncertainty*, vol. 1, pp. 729–807. North-Holland.
- K.-G. Mäler & A. Fisher (2005). ‘Environment, Uncertainty, and Option Values’. In K.-G. Mäler & J. R. Vincent (eds.), *Handbook of Environmental Economics*, vol. Vol. 2, pp. 571–620. Elsevier.
- A. Millner (2016). ‘Non-paternalistic Social Discounting’. *Working Paper*, <http://personal.lse.ac.uk/millner/files/non-paternalism.pdf> .
- A. Millner, et al. (2013). ‘Scientific Ambiguity and Climate Policy’. *Environmental and Resource Economics* **55**(1):21–46.
- A. Millner & G. Heal (2017). ‘Discounting by committee’. *Working paper* .
- A. Millner & T. K. J. McDermott (2016). ‘Model confirmation in climate economics’. *Proceedings of the National Academy of Sciences* **113**(31):8675–8680.
- S. Mukerji (2009). ‘Foundations of ambiguity and economic modelling’. *Economics and Philosophy* **25**(Special Issue 03):297–302.
- W. D. Nordhaus (1994). *Managing the global commons : the economics of climate change*. MIT Press, Cambridge, Mass ; London. William D. Nordhaus; Includes bibliographical references (p. [199]-204) and index.
- W. D. Nordhaus & P. Sztorc (2013). ‘DICE 2013: Introduction and User’s Manual’. Tech. rep.

- R. S. Pindyck (2007). ‘Uncertainty in Environmental Economics’. *Review of Environmental Economics and Policy* **1**(1):45–65.
- R. S. Pindyck (2013). ‘Climate Change Policy: What Do the Models Tell Us?’. *Journal of Economic Literature* **51**(3):860–872.
- H. Raiffa (1961). ‘Risk, Ambiguity, and the Savage Axioms: Comment’. *The Quarterly Journal of Economics* **75**(4):690–694.
- G. C. Rausser & A. A. Small (2000). ‘Valuing Research Leads: Bioprospecting and the Conservation of Genetic Resources’. *Journal of Political Economy* **108**(1):173–206.
- J. Rawls (1971). *A theory of justice*. Harvard University Press.
- L. J. Savage (1954). *The Foundations of Statistics*. Wiley and Sons.
- R. D. Simpson, et al. (1996). ‘Valuing Biodiversity for Use in Pharmaceutical Research’. *Journal of Political Economy* **104**(1):163–185.
- C. A. Sims (2001). ‘Pitfalls of a Minimax Approach to Model Uncertainty’. *The American Economic Review* **91**(2):51–54.
- M. Siniscalchi (2009). ‘Two out of three ain’t bad: A comment on ‘The ambiguity aversion literature: A critical assessment’’. *Economics & Philosophy* **25**(3):335–356.
- G. L. Stephens (2005). ‘Cloud Feedbacks in the Climate System: A Critical Review’. *Journal of Climate* **18**(2):237–273.
- D. Tilman, et al. (2001). ‘Diversity and Productivity in a Long-Term Grassland Experiment’. *Science* **294**(5543):843–845.
- R. Tol (1997). ‘On the optimal control of carbon dioxide emissions: an application of FUND’. *Environmental Modeling and Assessment* **2**(3):151–163.
- C. P. Traeger (2014). ‘Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity’. *Economic Theory* **56**(3):627–664.
- G. Vardas & A. Xepapadeas (2010). ‘Model Uncertainty, Ambiguity and the Precautionary Principle: Implications for Biodiversity Management’. *Environmental and Resource Economics* **45**(3):379–404.

- J. von Neumann & O. Morgenstern (1944). *Theory of Games and Economic Behaviour*. Princeton University Press.
- 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.
- G. Yohe, et al. (1996). ‘The economic cost of greenhouse-induced sea-level rise for developed property in the United States’. *Climatic Change* **32**(4):387–410.
- K. Zickfeld, et al. (2010). ‘Expert judgments about transient climate response to alternative future trajectories of radiative forcing’. *Proceedings of the National Academy of Sciences* .