

Centre for
Climate Change
Economics and Policy

An ESRC Research Centre



Grantham Research Institute on
Climate Change and
the Environment

Should climate policy account for ambiguity?

Geoffrey Heal and Antony Millner

August 2015

Centre for Climate Change Economics and Policy
Working Paper No. 229

Grantham Research Institute on Climate Change and
the Environment

Working Paper No. 202

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council. Its second phase started in 2013 and there are five integrated research themes:

1. Understanding green growth and climate-compatible development
2. Advancing climate finance and investment
3. Evaluating the performance of climate policies
4. Managing climate risks and uncertainties and strengthening climate services
5. Enabling rapid transitions in mitigation and adaptation

More information about the Centre for Climate Change Economics and Policy can be found at: <http://www.cccep.ac.uk>.

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and the Global Green Growth Institute. It has nine research programmes:

1. Adaptation and development
2. Carbon trading and finance
3. Ecosystems, resources and the natural environment
4. Energy, technology and trade
5. Future generations and social justice
6. Growth and the economy
7. International environmental negotiations
8. Modelling and decision making
9. Private sector adaptation, risk and insurance

More information about the Grantham Research Institute on Climate Change and the Environment can be found at: <http://www.lse.ac.uk/grantham>.

This working paper is intended to stimulate discussion within the research community and among users of research, and its content may have been submitted for publication in academic journals. It has been reviewed by at least one internal referee before publication. The views expressed in this paper represent those of the author(s) and do not necessarily represent those of the host institutions or funders.

Should climate policy account for ambiguity?

Geoffrey Heal¹ and Antony Millner*²

¹Columbia University

²London School of Economics and Political Science

August 10, 2015

Abstract

Climate change is fundamentally an ‘out-of-sample’ problem – our available information does not tightly constrain predictions of the consequences of rapid increases in greenhouse gas concentrations. Moreover, the fact that we haven’t observed much warming up to the present makes it very difficult to validate scientific and economic models of the medium to long-run consequences of climate change. We have many models, each based on roughly plausible assumptions about the dynamics of the climate-economy system, but we do not have *confidence* in our ability to select between them.

Traditional approaches to decision under uncertainty do not permit a decision maker’s confidence in her information set to influence her choices. They require us to combine probability distributions from different models into a single summary distribution, even if we are not confident of our ability to discern which model captures reality best. Since all probabilities are created equal in this framework, this summary distribution is treated the same as a distribution that arises from a single well validated model. The decision framework forces us to make subjective probability judgments, but decision makers do not distinguish them from probabilities derived by constraining models with data.

*Heal: gmh1@columbia.edu; Millner: a.millner@lse.ac.uk. This paper was prepared for the conference ‘Developing the next generation of economic models of climate change’, at the Heller-Hurwicz Institute, University of Minnesota, Sept. 2014. Support from the Grantham Foundation for the Protection of the Environment and the ESRC through the Centre for Climate Change Economics and Policy is gratefully acknowledged.

We suggest that approaches to decision under uncertainty that allow us to work with probabilities of different quality without combining them into a summary distribution may provide an attractive set of tools for analysis of climate policies. We set out the conceptual arguments for a departure from expected utility theory, give examples of alternative approaches to decision making under uncertainty, and discuss several common objections to them. We then suggest practical applications of these tools to integrated assessment models of climate policy, and investigate what difference they might make to policy recommendations.

1 Introduction

Uncertainty is ubiquitous in the science and economics of climate change. Climate science gives us a lot of robust broad-brush information about how increased greenhouse gas concentrations might affect climatic conditions at the global scale, but detailed empirical predictions are highly uncertain. It is routine in climate science for our state of knowledge to be represented by a set of probability distributions over key parameters or climatic outcomes. Each of these distributions is an internally consistent estimate, given some assumed model of the climate system. Unfortunately, choosing which of these models is ‘right’ or ‘better’ is a very difficult task – there is too much we don’t understand about the climate system, and too few meaningful opportunities for model verification, to make this choice entirely objective. The science is also the part of the problem we understand best. Our models of the long run evolution of the global economy and technological change, which are crucial for deciding on abatement policy, have less established predictive track records than their scientific counterparts. The policy recommendations that the integrated assessment community generates thus rely on detailed quantitative predictions from models that combine uncertain science with even more uncertain economics. The necessity of making these policy recommendations, e.g. prescribing a value for the Social Cost of Carbon, is clear. However, if the modeling exercise is to retain its scientific credentials it should attempt to capture the full scope of our uncertainty about the costs and benefits of climate policy. How should this be done? This is the topic of this paper.

Modeling uncertainty with a view to informing policy choices raises a number of complex issues. The default approach to modeling rational decision under uncertainty in economics is to assume that decision-makers [DMs] seek to choose an action that maximizes their expected utility. This requires the analyst to specify a utility function that captures society’s preferences or values, and a unique probability density function [PDF] over the space of

policy relevant consequences, which captures our beliefs. One of the main issues we want to discuss is whether it is necessarily ‘rational’ for degrees of belief about the consequences of climate policy to be represented by a unique PDF. If not, how should we represent uncertainty in our models when making policy recommendations?

The legitimacy of the expected utility approach to decision under uncertainty derives largely from Savage’s Subjective Expected Utility theory (SEU) (Savage, 1954). In his theory the existence of a unique PDF that summarizes the DMs beliefs follows from an attractive set of axioms: anyone whose preferences satisfy these axioms must behave as if she maximizes expected utility for *some* utility function and PDF. Unlike the more familiar von Neuman-Morgenstern axioms, in which a probabilistic representation of beliefs is assumed *a priori*, the existence of subjective probabilities is *derived* in the Savage result. This makes SEU a much more powerful and impressive result, which may, at least in principle, be applied to any decision problem. An axiom that is crucial in delivering the SEU representation is Savage’s so-called “sure thing principle,” which we discuss in detail below. This, plus the other axioms, necessarily implies that DMs must assign unique probabilities to any event, regardless of the quality of the information they have available to inform these judgments, if they are to respect the conception of rationality encoded in Savage’s axioms.

In the last few decades scholars have begun to question whether it is reasonable to adopt normative axioms that force us to assign a unique probability to absolutely any event, regardless of the information we have at hand to inform such a judgement. As Gilboa, Postlewaite, and Schmeidler [GPS] argue (Gilboa et al., 2009),

“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 say 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.”

This critique of the Savage approach captures a fundamental tension between what GPS call ‘internal’ and ‘external’ coherence in rational decision under uncertainty. The Savage axioms deliver a form of rationality that emphasizes internal coherence – each axiom taken in isolation seems an attractive feature of rational choice. However all forms of normative reasoning must be evaluated by the sensibility of their prescriptions when applied to relevant problems (external coherence), as well as by the abstract appeal in their axiomatic foundations. This back and forth between evaluating abstract principles and their applied prescriptions is a crucial part of refining normative frameworks, and defining their domain of applicability. It has been dubbed ‘reflective equilibrium’ by the philosopher John Rawls (Rawls, 1971). Uniquely among the decision sciences (including sub-disciplines in mathematics, statistics, computer science, and philosophy), economics has historically under emphasized the external coherence component of rationality (with some notable exceptions), and identified ‘rational’ choice with internal coherence. We argue below that this perspective may be too narrow, especially for applied policy problems as profoundly uncertain as climate change.

If we wish to move beyond SEU, we need to drop one or more of the Savage axioms. One possibility is to accept that preferences over uncertain prospects may be incomplete, with DMs being unable to rank some pairs 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” (see also Galaabaatar & Karni (2013)). 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. While this approach may be the most ‘honest’ in some very poor quality informational environments, it is very conservative with respect to our willingness to rank policy alternatives. Even if 999 out of 1000 plausible PDFs yield the same ranking of two policies, Bewley does not permit us to say anything about which policy is preferred. An alternative to dropping the completeness axiom is to drop the “sure thing principle”, and this is the approach whose application to climate change we investigate below. But first we need to understand more about the nature of our uncertainty about climate change.

2 The nature of climate uncertainty

There are several sources of empirical uncertainty in climate economics: at the coarsest level we can classify them into scientific uncertainties – how will increases in greenhouse gas (GHG) concentrations affect climate variables over time – and socio-economic uncertainty – what will the consequences of these changes be for society? We discuss these sources of uncertainty in detail in Heal & Millner (2014), and summarize the main points below.

The uncertainty in scientific predictions of the response of the climate system to increases in GHG concentrations can be resolved into three components: internal variability, model uncertainty and emissions uncertainty. Internal variability arises because climate models are complex non-linear dynamical systems, and thus exhibit sensitive dependence on initial conditions, or chaos. Small differences in the initialization of a model can lead to significantly different model trajectories over several decades. 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.¹ Emissions uncertainty arises because anthropogenic greenhouse gas emissions drive climate change projections in all models, and future emissions pathways are unknown, as they depend on our own future policy choices. The relative contribution of each of these uncertainties to the overall uncertainty in scientific projections depends on the temporal scale of the prediction. Generally speaking, internal variability and model uncertainty dominate in the short to medium run, while emissions uncertainty dominates in the long run. In addition, all uncertainties explode

¹Technically, 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.

when we decrease the spatial scale of the prediction. For example, we know much more about changes in global mean temperature than we do about changes in mean temperature at regional, country, or county scales.

Socio-economic uncertainty – uncertainty about how economic and social systems will respond to climate change – is probably greater than scientific uncertainty, though both are large in some absolute sense. A very broad classification of socio-economic uncertainties breaks them into empirical and normative components. Empirical uncertainties abound in every important process that determines the consequences of climate change for the global economy. These can be roughly summarized as follows: we don't know how damaging climate change will be for economic activities (damage function), we don't know how well we will adapt to these changes, we don't know what mitigation technologies will be available (abatement costs), we don't know how wealthy we will be (overall economic progress), and we don't know the opportunity costs of investments in climate mitigation. Each of these is a crucial empirical determinant of optimal climate policies. We have some knowledge about all of them, but we are arguably lacking anything that comes close to e.g. an empirically validated 'damage function' or a plausible model of long-run technical change. In contrast to climate science, where the physical principles that underpin the models have been verified in countless related applications, the representation of these important economic processes in integrated assessment models is largely *ad hoc*, and based on analytical convenience and precedent in the literature, rather than empirically demonstrated predictive power.

The second class of socio-economic 'uncertainties' are not really uncertainties at all, but rather disagreements about values. This arises because the policy recommendations that flow from economic models of climate change are heavily driven by value judgments. These judgments are reflected in our models by the objective functions we assume planners seek to optimize. The appropriate values of the parameters of these objective functions – e.g. the infamous pure rate of time preference, which determines our attitude to the welfare of future generations – are the subject of heated debate, both within and outside the economics profession. At a higher level of abstraction, the correct objective function itself is also heavily debated: should it have multiple attributes (Stern & Persson, 2008; Heal, 2009), should we disentangle risk aversion from consumption smoothing motives (Epstein & Zin, 1989), should we adopt discounted utilitarianism or some other criterion for intertemporal choice (Asheim, 2010), how should we value climate catastrophes and population change (Millner, 2013), how should we account for inequality (Schelling, 1995), and of course, how should we make decisions under uncertainty?

None of the empirical uncertainties we have discussed, be they scientific or socio-economic, is readily quantifiable by a unique PDF that everyone can agree represents our current state of knowledge. How should an economist faced with such pervasive uncertainty proceed if she wishes to be an ‘honest broker’, and build models that accurately reflect the nature of our information? Should she seek to derive a single summary PDF over policy consequences and use decision techniques that rely on the availability of such a distribution, or should she work with the ambiguity inherent in the information available to her? In order to combine PDFs from multiple models of the same phenomenon into a single probability estimate it is desirable to account for the dependencies between models, and have some measure 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 approximately independent 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.

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, however this exercise is unlikely to reveal much information about which of the models is likely to be 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.² 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 (e.g. Nordhaus & Sztorc, 2013) this applies to the trajectories of key time series such as total factor productivity, population change, and abatement costs. In the PAGE (e.g. Hope, 2006) and FUND (e.g. 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

²It 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.

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-sample predictive power.

Although these issues are prominent when thinking about the empirical validity of the economic components of IAMs, the scientific components of these models are not immune to these concerns. The climate models that provide PDFs over key scientific parameters in integrated assessment models are also difficult to verify. For example, as Figure 1 illustrates, different models provide substantially different estimates of the crucial climate sensitivity parameter (the equilibrium increase in global average surface temperature from a doubling of CO₂ concentration), which features prominently in all IAMs. It is again very difficult to objectively weight these different estimates, as they are calibrated on related datasets, and share similar, but not identical, modeling assumptions. Teasing out the dependencies between the different estimates is a practical impossibility.

Our inability to validate models of the consequences of climate policy poses methodological questions. The decision analyst has to choose a framework within which the information available can best be integrated into the analysis. As we have noted, the default expected utility approach requires us to summarize our knowledge in a single PDF over policy consequences. For climate applications this raises an obvious question: where is this PDF going to come from? Any such summary distribution will require us to aggregate information of varying quality from a variety of largely unvalidated models. This will be a heavily subjective exercise – we have low *confidence* in our ability to discern between models. An alternative to this approach is to recognize that we do not have the quality of information needed to define a unique PDF. Instead of forcing our information into a probabilistic strait-jacket, we can work with decision frameworks that are sensitive to the difference between subjective judgments and ‘conditionally objective’³ model probabilities. We investigate the formal foundations of these approaches in detail in the next section.

³We use this term to denote probabilities that arise within a given model, conditional on the assumption that the model is empirically ‘correct’.




Figure 1: Estimates of the PDF of Climate Sensitivity from the scientific literature.

3 Theories of decision under uncertainty

A common misconception about non-expected utility theories of decision under uncertainty is that they all have purely *positive* aspirations, i.e. they attempt to describe, rather than prescribe, behavior. This is certainly true of some of these theories (see Machina (1987) for a review). However, normative alternatives to SEU have always been a part of decision theory. For example, the statistician Abraham Wald advocated a ‘Maxmin’ approach to decision under deep uncertainty. Arrow & Hurwicz (1977) later axiomatized a set of preferences for choice under complete ignorance, and Wald’s prescription is a special case of their representation. Savage himself also saw limitations to his SEU theory, advocating a ‘minimax regret’ criterion in low information environments. All these decision criteria do not require any likelihood information, a great advantage if none is available, but a distinct limitation when the analyst does have likelihood information at her disposal, albeit not in the form of a unique PDF. Adopting these criteria avoids the need to turn ambiguous information into a PDF but at the potentially high cost of throwing valuable information

away.

In recent decades decision theorists have developed criteria with the ability to use incomplete or degenerate probabilistic information. These criteria have rigorous axiomatic foundations which are analogous to the axioms that underpin SEU theory. Of course, the axiom sets they employ are slightly different to SEU, as are the representation results they achieve. A familiarity with these axiomatic results is necessary if we are to discuss the tradeoffs between the different approaches, and understand how they differ. To this end, we review two of the most popular models: the “MaxMin Expected Utility” model due to Gilboa & Schmeidler (1989) [GS], and the “Smooth Ambiguity” model due to Klibanoff Marinacci and Mukherji (2005) [KMM]. In order to put these more recent developments in perspective we begin by discussing Savage’s axioms.

3.1 Formal Development

The primitive concepts in Savage’s theory of decision under uncertainty are **states** and **outcomes**. The set of states $s \in S$ is an exhaustive list of all scenarios that might unfold. Knowing which state occurs resolves all uncertainty. An **event** is any subset $A \subset S$. The set of outcomes is X , 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 acts are denoted $f \in F, f : S \rightarrow X$. The state is uncertain and so not known when the act is chosen, but the actor does know that if the state is s then the outcome is $f(s)$.

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 and g are two acts, A is an event, and A^c is the complement of A . Then we define a new act f_A^g by

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

Intuitively this is f but replaced by g on A . Within this framework Savage uses the following axioms:

Axiom P1. Preferences \succeq are a complete transitive relation on F .

Axiom P2 (‘**Sure thing principle**’). Preferences between two acts f, g depend only on the values of f, g where they differ.

Suppose f, g are equal if A does not occur, that is on A^c , so they differ only on A . Alter f, g only on A^c to get f', g' such that $f(s) = f'(s), g(s) = g'(s), s \in A$. So there is no change on A . They are still equal off A , though not necessarily to f, g : $f(s) = g(s), f'(s) = g'(s), s \in A^c$.

Then P2 requires:

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

This is sometimes written $f \succeq_A g$, f preferred or indifferent to g given A .

Axiom P3. If you take an act that guarantees an outcome x on an event A and you change it on A from x to another outcome y , the preference between the two acts should follow the preference between the two outcomes. Formally let f_A^x be an act that produces x for every state in A . A null event is, roughly, one that is thought to be impossible. Formally an event E is null if whenever two acts yield the same outcome off E they are ranked as equivalent. For every act $f \in F$, every non-null event $A \subset S$, and $x, y \in X$,

$$x \succeq y \Leftrightarrow f_A^x \succeq f_A^y$$

Here x, y can also be interpreted as the acts that yield x, y respectively in every state. This is a monotonicity assumption. Another interpretation is that rankings should be independent of the events with which they are associated (note the requirement that this hold for every non-null event A).

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

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

This is an axiom about probabilities: presumably $y_A^x \succeq y_B^x$ means that you think event A is more likely than event B .

Axiom P5. There are f, g such that $f \succ g$.

Axiom P6. For every $f, g, h \in F$ with $f \succ g$ there exists a partition of S [a collection of pairwise disjoint events whose union is S] denoted $\{A_1, A_2, \dots, A_n\}$ such that for every i

$$f_{A_i}^h \succ g \text{ \& } f \succ g_{A_i}^h$$

This is roughly like a continuity assumption, but it is hard to state continuity in Savage's framework. With this set of assumptions Savage proves

Theorem 1. [Savage] Assume that X is finite. Then \succeq satisfies P1 to P6 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.

Remarkably, Savage proves the existence of **both** the utility function **and** the subjective probabilities that go into the expected utility function. Recall that von Neumann and Morgenstern take probabilities as given, and prove that there exists a utility function with the requisite properties. So, at least in principle, Savage's theorem applies where von Neumann-Morgenstern's doesn't, i.e. where there are no 'given' probabilities.

Gilboa and Schmeidler work within the same framework but of course use a different set of axioms:

Axiom 1. We have a complete transitive ordering \succeq over F .

Axiom 2. Continuity: For every $f, g, h \in F$ if $f \succ g \succ h$ then there exist $\alpha, \beta \in (0, 1)$ such that

$$\alpha f + (1 - \alpha) h \succ g \succ \beta f + (1 - \beta) h$$

Axiom 3. Monotonicity: For every $f, g \in S$, $f(s) \succeq g(s) \forall s \in S \Rightarrow f \succeq g$

Axiom 4. Nontriviality: There exist $f, g \in X : f \succ g$

Axiom 5. Independence: For every $f, g \in F, \forall$ 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 6. Uncertainty Aversion. For every $f, g \in F, \forall \alpha \in (0, 1), f \sim g \Rightarrow \alpha f + (1 - \alpha) g \succeq f$.

This axiom can be understood intuitively as a 'preference for hedging'. If we have two equivalent acts, we weakly prefer a convex combination of them that 'hedges' the uncertainty in either one, to each act on its own.

With this set of assumptions they prove the following result:

Theorem 2. [Gilboa-Schmeidler] A preference satisfies the above axioms if and only if there exists a closed convex set of probabilities C and a non-constant function $u : X \rightarrow R$ such that for every $f, g \in F$

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

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

GS don't need Savage's "sure thing principle", nor do they need his axiom P4 about probabilities: they do however need an analogue of von Neumann Morgenstern's independence axiom, axiom 5, and axiom 6 about aversion to uncertainty.

What their theorem is saying, is: look at the probabilities that give the worst possible expected utility for each act, evaluate the acts according to these probabilities, and choose the act that gives the largest of these 'worst' expected utilities. What is important from our perspective is that the idea of multiple priors emerges naturally from their framework. If you believe that the GS axioms are an acceptable conception of rationality, your beliefs must be represented not by a unique PDF, but a set of PDFs which cannot be compounded into a single summary distribution.

The focus on the 'minimum' expected utility over the set of plausible priors in the GS result can be seen as conservative with respect to our willingness to make subjective judgements between different PDFs. It allows us probabilistic information over the set of states given a model, but not over the set of PDFs or models. It might be desirable to have a decision framework that allows us the freedom to make subjective judgements between different PDFs, but treats these as different from the 'conditionally objective' information captured by the PDFs themselves. This is what is achieved by the Smooth Ambiguity model of KMM.

The KMM representation is in some senses less 'primitive' than the SEU 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 'conditionally objective' von Neumann-Morgenstern lotteries π over the state space. Preferences over acts, conditional on a lottery π 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 $\mu(\pi)$ on lottery π , and some utility function Φ over consequences. Now given an objective PDF π , an act f induces a lottery π_f over the set of consequences, whose value can be captured by a certainty equivalent $c_f(\pi)$. These certainty equivalents take values in the space of outcomes, and are an increasing function of $E_\pi U(f(s))$. Since there are many possible π , each act f generates a lottery over the outcomes $c_f(\pi)$. 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:

Theorem 3. [*Klibanoff, Marinacci, and Mukherji*]: *There exists a set of PDFs Δ with generic element π , a utility function $U : S \rightarrow R$, a ‘second order’ utility function $\Phi : U \rightarrow R$, and second order subjective probabilities $\mu(\pi)$, such that for all $f, g \in F$,*

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

The subjective probability $\mu(\pi)$ on PDF π 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

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

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 $\int_{\pi \in \Delta} \mu(\pi)\pi(s)d\pi$. Thus, if Φ is concave, the decision maker would always prefer it if his ‘subjective’ uncertainty $\mu(\pi)$ could be made ‘objective’ – he is averse to subjective uncertainty over conditionally objective PDFs. In more common terminology, he is ambiguity averse.

There are two novel elements here: second order probabilities, reflecting the degree of belief in each of the priors, and a concave function expressing aversion to uncertainty about the expected utility that will come from any choice of policy. In the pure expected utility world, once a policy is chosen then the expected payoff is known. With multiple priors this is not true: given a choice of policy the expected utility depends on which model turns out to be correct, so there is still uncertainty and it is reasonable to posit aversion to this, reflected by the concavity of the function used in computing the expectation of the expected utilities.

3.2 Evaluating the theories

We now have alternatives to the traditional expected utility model which seem more consistent with the kind of information that occurs naturally when evaluating climate policies. They provide a natural way of using multiple different PDFs in making a decision, without having to reduce them down to a single summary PDF. However, this comes at the price of dropping the “sure thing principle,” Savage’s axiom P2. It is important to understand exactly what this axiom requires of us. GPS provide an intuitive example. Consider the following four bets

1. If horse A wins you get a trip to Paris, and otherwise you get 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 your choice will depend on preferences and beliefs or probabilities, but presumably the chance of A winning is the same in each case, so the choice between 1 and 2 depends on your preferences between Paris and London. The same is true for 3 and 4, and axiom P2 requires $1 \succeq 2 \Leftrightarrow 3 \succeq 4$. 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 A loses you get Rome or LA. It is hard to argue with this axiom viewed in isolation, but nevertheless it does have, when taken together with the rest of Savage’s axioms, the implication that you will be able to assign probabilities to any events whatever the quality of information available to you.

There is no better illustration of just how strong a constraint the sure thing principle is on what is considered a ‘rational’ decision than the famous examples provided by Ellsberg (1961). Violation of the sure thing principle in this case – in which decision makers are forced to choose between bets in the absence of any information whatsoever – is commonly observed, and is moreover persistent even after subjects are explained how their choices violate P2 (Slovic & Tversky, 1974). Despite their beauty and insight, it would be a mistake to think that the case for deviating from SEU rests *only* on these examples (see also Mukerji, 2009). The Ellsberg examples are set up with an artificially high degree of symmetry – we certainly do not view adherence to the sure thing principle in this case as irrational! Rather, they should be seen as a very clean illustration of exactly what the sure thing principle requires of us in informationally poor environments. The case for normative

deviations from SEU, especially in the climate context, rests much more on the importance of the ‘external’ coherence requirements of rational choice, as we discuss below.

What are the advantages of the the GS and KMM results? The GS result is quite remarkable – it uses axioms that seem, at least to us, to have a good claim to be seen as criteria of rational choice. A common critique of the GS axioms is that they need to introduce an attitude to uncertainty explicitly into the axiom system via their Uncertainty Aversion axiom. The Savage axioms do not take a position on the decision maker’s attitude to uncertainty (he could be risk loving, for example), whereas the GS axioms do. While this is clearly a difference between the two, it seems to us to have limited power as a means for deciding which of the two axiom sets offers a more attractive conception of rationality. In all normative applications of the Savage axioms some form of aversion to uncertainty is assumed. This happens via the choice of the utility function in the Savage framework, and at the level of the axioms in the GS framework. So normative choices about the decision maker’s attitude to uncertainty are being made either way in applications, and it is not clear that choices about parameters given a preference representation should carry less weight than choices about the axioms that determine a preference representation. Both are normative choices, and they constrain the set of decisions that are considered ‘rational’, albeit in rather different ways.⁴

What is clear is that both the GS and KMM results give more importance to the ‘external’ coherence of the choice framework than the Savage result. In the GS system this flows from the uncertainty aversion axiom, and in the KMM framework it is very explicitly built in via the *a priori* assumption of two different kinds of probabilities – subjective weights on the set of PDFs, and conditionally objective probabilities. The KMM result seems to us to be to choice under uncertainty what the von Neumann-Morgenstern result is to choice under risk. Both these results take the representation of the decision maker’s information set as given – this is the set of 1st and 2nd order probabilities in KMM, and the set of lotteries in vNM. They then derive representations of preferences,

⁴A more technical set of objections concerns the extensions of the GS representation 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. This aspect of inter-temporal decision-making is pinned down by invoking additional axioms, in particular dynamic consistency (conditional plans, once made, are not revised due to the mere passage of time). Insisting on this requirement in the the standard SEU setting implies that beliefs must be updated in accordance with Bayes’ rule (Epstein & Le Breton, 1993). Epstein & Schneider (2003) demonstrate that imposing dynamic consistency in the GS model requires Bayes’ rule to hold on a prior by prior basis, however their representation result holds only if the set of priors in the GS model satisfies a technical condition called ‘rectangularity’. Mukerji (2009) notes that the dynamic extension of the KMM model (Klibanoff et al., 2009) requires no such restriction on the set of priors.

given these *a priori* representations of the decision maker's information. The assumed separation between subjective and objective probabilities in KMM means that it places a lot more importance on external coherence than the GS result. In GS the multiple priors representation of beliefs is derived from primitive axioms (internal coherence), just like in the Savage framework, as opposed to being put in by hand.

Amartya Sen (and many others) has persuasively argued that axioms can only be fully evaluated when one knows their implications - they are reasonable only if they have reasonable implications. Moreover, what counts as a 'reasonable' implication in one situation may be unreasonable in another. There is a natural human tendency to want to have a single 'one size fits all' solution to the problem of decision making under uncertainty. A more nuanced approach to deciding what counts as rational choice might however require us to abandon the pursuit for a 'universal' theory of decision-making. Instead, it would require us to carefully delineate *when* one set of criteria should take precedence over another.

It could be that the outcome of such a deliberation process leads some people to reaffirm the sure-thing principle and SEU theory. Such individuals might acknowledge the difficulty of specifying probabilities in informationally poor environments, but feel that this is simply the price we must pay for consistency. We must make hard choices, and specify these probabilities anyway. This is a defensible view, even if it is difficult to swallow when it comes to policy applications that aspire to provide objective and scientific input to public decision making. Alternatively, we might acknowledge that the domain of application of SEU theory and the sure thing principle should be limited. It is fine for private decisions, or decisions with a marginal impact on society's welfare, and it certainly applies when probabilities arise naturally in applications, and where subjective judgements don't have a radical impact on choices. But when there are many plausible models that are difficult to validate and make substantially different predictions, and the future of humanity is at stake, it seems at best bold, and at worst dishonest, to adopt a framework that forces us to act as if we know things that we do not.

4 Application to Climate Policy

The next step is to show how these new ideas can be applied to climate policy choices. We develop a simple one-period model of the optimal choice of an abatement level, and show how the multiple priors framework can be applied here.

We assume that all uncertainty is about the equilibrium climate sensitivity, and that the

set of distributions over this is P , with π_m a typical distribution. The level of greenhouse gas abatement is a , and the utility level associated with this abatement level depends on the state and is denoted $U(a, S)$. For each distribution or prior $\pi_m(S)$, the expected utility associated with abatement a is $EU_m(a) = \int U(a, S) \pi_m(S) dS$. The correct prior distribution is of course not known, but in the KMM approach the decision-maker has second order probabilities over these given by $\mu_m = \mu(\pi_m)$. So the overall objective is to choose abatement a so as to

$$\text{Max}_a \sum_m \phi[EU_m(a)] \mu_m \quad (1)$$

where ϕ is a concave function. We can think of the probabilities $\pi_m(S)$ as scientific probabilities and the μ_m as subjective judgments, so we are taking two expectations in succession in this maximization, one according to scientific probabilities and one according to subjective assessments of the plausibility of different priors or models.

The first order conditions for problem 1 are

$$\sum_m \hat{\mu}_m(a^*) \frac{dEU_m}{da} \Big|_{a=a^*} = 0 \quad (2)$$

where a^* is the optimal abatement level and the $\hat{\mu}_m$ are ‘‘ambiguity-adjusted’’ second order probabilities:

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

The first order conditions say that a weighted sum over models of marginal expected utility of abatement should be zero, where the weights are the ambiguity-adjusted second order probabilities. These are the second order probabilities μ_m multiplied by the marginal expected utilities associated with them. This reweighting implies that more weight is given to priors that predict low utilities, as they will have high values of ϕ' . If $\phi' \rightarrow \infty$ as $EU \rightarrow 0$ then priors that predict really poor outcomes will in the limit receive all the weight: the decision-maker’s focus will be on the really bad outcomes.

Within this model Millner et al. (2013) prove the following result:

Proposition 1. *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 abatement.*

Figure 2 gives examples of anti-comonotonic and comonotonic sequences of expected




Figure 2: Comonotonicity [left] and anti-comonotonicity [right]

utilities and marginal expected utilities. Abatement a is plotted horizontally and EU_m vertically for a set of three distinct models. Anti-comonotonicity means that for each value of abatement the models with high expected utility have low derivatives of expected utility with respect to abatement, and vice versa. The key point here is that when the anti-comonotonicity condition holds then an increase in abatement will reduce the spread of expected utilities across models, so that a rise in aversion to ambiguity will increase optimal abatement.

While this simple static model can provide the basic intuition for how the spread in expected utilities across a set of possible priors drives the effects of ambiguity aversion on policy choice, it is seriously deficient as a guide to how these effects might play out in detailed fully dynamic models. The study of the effects of ambiguity aversion in dynamic models of climate policy is still in its infancy. Nevertheless, several early studies exist. Millner et al. (2013) use a simplification of the dynamic version of the KMM model (Klibanoff et al., 2009) to evaluate exogenous dynamic abatement policies in the DICE integrated assessment model, accounting for the ambiguity in the set of scientific estimates of climate sensitivity represented in Figure 1. They show that estimates of the welfare benefits of CO₂ stabilization policies in DICE can be highly sensitivity to ambiguity preferences, but only if climate damages are sufficiently steep at high temperatures. This analysis only considers ambiguity over scientific parameters, and not abatement costs or damage functions. This limitation is partially addressed by more recent work (Drouet et al., 2015), which evaluates exogenous CO₂ stabilization targets using a variety of ambiguity averse decision criteria, taking multiple estimates of abatement costs and damage functions into account.

The authors again argue that the choice of decision criterion can have a material affect on policy evaluation and selection.

While these analyses provide estimates of how ambiguity aversion might affect the welfare evaluation of exogenous policy proposals, they do not inform us about optimal policy choice. Fully optimizing models that account for ambiguity can be technically challenging, especially if they attempt to accurately reflect empirical likelihood information from climate science and impact studies. Invariably there is a tradeoff to be made between fidelity to the data and the tractability and elegance of the modeling. As noted above, one possible way around this is to employ decision criteria that do not require likelihood information. This is the approach adopted in e.g. Xepapadeas (2012); Athanassoglou & Xepapadeas (2012), who use the robust control methods of Hansen & Sargent (2007) to identify optimal policies in a simple climate-economy model. While these elegant methods pay clear tractability dividends, they neglect much of the information at our disposal (see Heal & Millner (2014) for further discussion of the virtues and limitations of robust control methods). Fully dynamic optimizing models that are based on the KMM framework (and hence are able to retain likelihood information), are under development by several authors. It is clear that this is a fruitful area for future research. The authors pursuing these models are doing so for explicitly normative reasons. They believe that ambiguity averse decision rules enable us to better reflect our lack of confidence in our modeling assumptions than the standard SEU paradigm, and moreover that this lack of confidence is rationally decision relevant. We hope that we have laid out some of the reasons why this is not a misguided view.

5 Conclusion

We conclude with a parable from the history of science, which we believe may serve as something of a cautionary tale.

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 banal 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 either criminal stupidity, or madness, by all 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, the Jesuit priest 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 absurd 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 is so 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-




Figure 3: Triangles in non-Euclidean geometries. Encyclopaedia Britannica, 1997.

Euclidean geometries are now a part of the mathematical firmament, 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.

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 different axiom systems have different domains of applicability. For some problems accounting for the non-Euclidean effects of surface curvature is crucial (e.g. in applications in cosmology, astrophysics, and the technologies underpinning Global Positioning Systems), while for others, we can safely rely on our tried and trusted Euclidean tools.

The sociological lesson is this: history teaches us that viewing a single axiom system as the universal standard of 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 earlier advances in other areas.

In the present context it might be argued that the von Neumann Morgenstern and Savage axioms are relevant to problems where a unique PDF can plausibly be agreed upon. Alternatives that admit multiple priors might well be more appropriate for problems

such as climate change where it is implausible that a unique PDF over the consequences of policies can be identified with anything approaching an objective scientific method.⁵ Perhaps the best we can do is to build many plausible representations of the dynamics of the climate-economy system, and make decisions that reflect our lack of confidence in our ability to discern between them. Adherence to the sure thing principle prevents us from adopting such an approach, which may be more suited to the decision problem at hand.

References

- N. I. Al-Najjar & J. Weinstein (2009). ‘The ambiguity aversion literature: a critical assessment’. *Economics and Philosophy* **25**(Special Issue 03):249–284.
- K. Arrow & L. Hurwicz (1977). ‘An optimality criterion for decision-making under ignorance.’. In *Studies in resource allocation processes*, p. 482. Cambridge University Press, Cambridge, UK.
- G. B. Asheim (2010). ‘Intergenerational Equity’. *Annual Review of Economics* **2**(1):197–222.
- 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).
- M. Dell, et al. (2012). ‘Temperature Shocks and Economic Growth: Evidence from the Last Half Century’. *American Economic Journal: Macroeconomics* **4**(3):66–95.
- A. E. Dessler (2010). ‘A Determination of the Cloud Feedback from Climate Variations over the Past Decade’. *Science* **330**(6010):1523–1527.
- L. Drouet, et al. (2015). ‘Selection of climate policies under the uncertainties outlined in IPCC AR5’. *Nature Climate Change*, forthcoming .

⁵Gilboa et al. (2010) provide a helpful definition of ‘objective’ and ‘subjective’ uncertainties. If a DM is objectively uncertain he can convince everyone that his ranking of acts is rational. If he is subjectively uncertain, no one can convince him that his ranking of acts is irrational.

- 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.
- L. G. Epstein & S. E. Zin (1989). ‘Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework’. *Econometrica* **57**(4):937–969.
- T. Galaabaatar & E. Karni (2013). ‘Subjective Expected Utility With Incomplete Preferences’. *Econometrica* **81**(1):255–284.
- I. Gilboa, et al. (2010). ‘Objective and Subjective Rationality in a Multiple Prior Model’. *Econometrica* **78**(2):755–770.
- I. Gilboa, et al. (2009). ‘Is It Always Rational to Satisfy Savage’s Axioms?’. *Economics and Philosophy* **25**(3):285–296.
- I. Gilboa & D. Schmeidler (1989). ‘Maxmin expected utility with non-unique prior’. *Journal of Mathematical Economics* **18**(2):141–153.
- L. P. Hansen & T. J. Sargent (2007). *Robustness*. Princeton University Press.
- G. Heal (2009). ‘The economics of climate change: a post-Stern perspective’. *Climatic Change* **96**(3):275–297.
- G. Heal & A. Millner (2014). ‘Uncertainty and Decision Making in Climate Change Economics’. *Review of Environmental Economics and Policy* **8**(1):120–137.
- G. Heal & J. Park (2015). ‘Goldilocks Economics? Temperature stress and the direct impact of climate change’. *Review of Environmental Economics and Policy*, forthcoming .
- 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).

- P. Klibanoff, et al. (2009). ‘Recursive smooth ambiguity preferences’. *Journal of Economic Theory* **144**(3):930–976.
- M. J. Machina (1987). ‘Choice Under Uncertainty: Problems Solved and Unsolved’. *Journal of Economic Perspectives* **1**(1):121–154.
- A. Millner (2013). ‘On welfare frameworks and catastrophic climate risks’. *Journal of Environmental Economics and Management* **65**(2):310–325.
- A. Millner, et al. (2013). ‘Scientific Ambiguity and Climate Policy’. *Environmental and Resource Economics* **55**(1):21–46.
- S. Mukerji (2009). ‘Foundations of ambiguity and economic modelling’. *Economics and Philosophy* **25**(Special Issue 03):297–302.
- W. Nordhaus & P. Sztorc (2013). ‘DICE 2013: Introduction and User’s Manual’. Tech. rep.
- J. Rawls (1971). *A theory of justice*. Harvard University Press.
- L. J. Savage (1954). *The Foundations of Statistics*. Wiley and Sons.
- T. C. Schelling (1995). ‘Intergenerational discounting’. *Energy Policy* **23**(4–5):395–401.
- P. Slovic & A. Tversky (1974). ‘Who accepts Savage’s axiom?’. *Behavioral Science* **19**(6):368–373.
- G. L. Stephens (2005). ‘Cloud Feedbacks in the Climate System: A Critical Review’. *Journal of Climate* **18**(2):237–273.
- T. Sterner & U. M. Persson (2008). ‘An Even Sterner Review: Introducing Relative Prices into the Discounting Debate’. *Review of Environmental Economics and Policy* **2**(1):61–76.
- R. Tol (1997). ‘On the optimal control of carbon dioxide emissions: an application of FUND’. *Environmental Modeling and Assessment* **2**(3):151–163.
- 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.

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*